Incumbent Radio Systems in the Defense Advanced Research Projects Agency Spectrum Collaboration Challenge Test Bed

Kaun J. Yim, Kenneth R. McKeever, and Daniel R. Barcklow

ABSTRACT

The application of artificial intelligence and machine learning promises to usher in a new paradigm for emerging wireless communication systems. The goal of the Defense Advanced Research Projects Agency (DARPA) Spectrum Collaboration Challenge (SC2) was to push this new paradigm forward. However, legacy radio systems already in place, such as radars for weather monitoring, receivers for spectrum monitoring, and battlefield jammers, will remain in use for a long time. Therefore, intelligent radios must operate around and adapt to these legacy systems to avoid interfering with them. In support of DARPA's SC2, the Johns Hopkins University Applied Physics Laboratory (APL) designed and built a wireless research test bed, referred to as the Colosseum, where SC2 competitors could test and develop solutions to enable this new communications paradigm. A critical component of the Colosseum was its legacy radio emulators, referred to as Colosseum incumbents, that represented today's systems. These incumbents emulated the radio frequency (RF) behavior of existing real-world radio systems, serving as RF obstacles that SC2 competitors had to detect and work around while simultaneously administrating their own communications for maximum data throughput efficiency.

INTRODUCTION

As the number of wireless devices and the bandwidth of applications increase as technology advances, the radio frequency (RF) spectrum becomes a limited resource. As RF spectrum availability becomes scarcer, radios must share (intelligently and autonomously) the spectrum in the presence of legacy radio actors. If not, emerging communication systems will cause interference. For example, when Wi-Fi began using the 5-GHz band, Wi-Fi stations inadvertently interfered with Terminal Doppler Weather Radar (TDWR) stations, which serve the critical function of providing quantitative measurements for gust fronts, wind shear, microbursts, and other weather-related hazards.[1]

The Defense Advanced Research Projects Agency (DARPA) Spectrum Collaboration Challenge (SC2) competition provided a test bed, known as the Colosseum, where competitors staged their designs for radio spectrum sensing and interference avoidance in adverse RF environments. Incumbent systems, designed and developed by APL as part of the Colosseum, emulated the behavior of legacy systems (i.e., served as obstacles in the RF spectrum). Each incumbent system was hosted on a standard radio node (SRN) and interacted with competitor systems through the RF Emulation System in the Colosseum. (For more details on the Colosseum in general and on SRNs and the RF Emulation System, see the articles by Coleman et al., White et al., and Barcklow et al., respectively, in this issue.)

To test the competitors' designs, DARPA developed several scenarios that simulated real-world challenges today's wireless communications systems would face. (See the article by Coleman et al. in this issue for details on SC2 scenarios.) Each scenario contained at least one incumbent system, and information about the incumbent system(s) was not conveyed to competitors in the scenario description. However, competitors did have a connection to an incumbent monitoring system over the collaborative intelligent radio network (CIRN), and this system provided information on incumbent systems' performance and interference levels. With this information, competitors could opt to back off transmission power or avoid frequency bands when they detected an incumbent system during their testing or a formal challenge event (see the article by Coleman et al. for more information on the challenge events).

SC2 scenarios featured three types of incumbent systems: (1) passive receivers, (2) radars, and (3) noise jammers. Passive receivers had exclusive rights to a specific RF band within a scenario, and competitors were penalized for transmitting in a passive receiver's band. The radar incumbent included periodic transmissions from the incumbent radio. To achieve a maximum score, a competitor had to learn the duty cycle of the radar and refrain from self-transmissions during the radar's transmit cycle. Jammers, on the other hand, had a 100% duty cycle that competitors had to avoid completely. The shape and power of the jammers varied across the scenarios so that DARPA could evaluate the robustness of the competitor solutions.

REQUIREMENTS

As mentioned, incumbents served as RF obstacles in the scenarios. As such, each incumbent had unique design goals and requirements, but two requirements were common to all incumbent systems:

1. **Automated scheduled reconfiguration**—Incumbents had to support a schedule-based reconfiguration mechanism to change RF characteristics within a scenario.

2. **Violation reporting**—While each incumbent had its own unique violation criteria, they all followed the general rule that if the current measurement metric exceeded a violation threshold, a violation event occurred. Incumbents had to detect violation events caused by competitors and report them to the Colosseum scoring engine.

INCUMBENTS

Software Architecture

All incumbents were built using a Linux container on an SRN. Within the Linux container, the GNU Radio framework[2] and Universal Software Radio Peripheral (USRP) hardware driver (UHD)[3] defined incumbent RF behavior. These two architectural designs allowed incumbents to be integrated into Colosseum infrastructure and could be easily distributed/reproduced by their respective license agreements (e.g., GNU General Public License).

Collaborative Intelligent Radio Network

The CIRN was an out-of-band messaging channel that SC2 competitors could use to exchange RF performance metrics and RF actions with other competitors. While the CIRN was primarily meant for competitor-to-competitor coordination, incumbents also leveraged the CIRN to notify competitors of their current performance. The amount and specificity of information passed through the CIRN was dependent on the incumbent type. For example, the passive receiver incumbent type advertised its center frequency and occupied bandwidth, while others did not. Ultimately, competitors could fuse incumbent CIRN information to their RF sensing metrics to enhance their decision engines.

Types

The SC2 Colosseum encompassed three incumbent types: passive, active, and jammer. Each incumbent was designed with unique behavioral motivators to test competitors' radio performance:

1. **RF avoidance**—The incumbent advertised a frequency band that competitors had to avoid.

2. **RF power management**—The incumbent advertised a shared frequency band that competitors could leverage, but only if the competitor's aggregate interference stayed below a certain threshold.

3. **RF detection**—The incumbent did not advertise its transmitter and/or receiver characteristics (e.g., center frequency, occupied bandwidth). Therefore, at run time competitors had to detect the presence of an incumbent without any a priori knowledge and avoid the incumbent.

4. **Pattern recognition**—The incumbent transmitted a repeating frequency-hopping signal. If competitors could successfully determine the pattern, they could leverage the temporal and spectral gaps left by the incumbent for their own transmissions.

5. **Intelligent sacrifice**—The incumbent continuously occupied a frequency band, such that the remaining

Table 1. Incumbent types and their associated motivators

Type	Description	RF Avoidance	RF Power Management	RF Detection	Pattern Recognition	Intelligent Sacrifice
Passive	The incumbent emulated a fixed satellite service Earth station located 30 km from a spectrum collaboration site. The station operator agreed to share its frequency band with the collaboration site—but only if their aggregate interference stayed below the Earth station's protection threshold. The Earth station's transmit pattern was the same, but the protection threshold varied over time.	✓	✓	×	×	×
Active	When Wi-Fi dynamic frequency selection (DFS) was first deployed in the Unlicensed National Information Infrastructure (U-NII) band, there was significant interference to TDWR systems because Wi-Fi devices did not properly sense and share with TDWR.[1] Could competitors implement a better sense-and-share approach than Wi-Fi DFS did? The incumbent transmitted a periodic transmission pattern and competitors had to work around the incumbent.	×	×	✓	✓	×
Jammer	A malicious user set up jammer signals to interfere with a radio network. The incumbent was the jammer source—creating multiple jammer signal types such as tone, multi-tone, wideband noise, or a combination barrage. For added complexity, the jammer signals could be stationary or sweeping and could switch modes pseudo-randomly. Such a chaotic RF environment emphasized a competitor's RF agility to dynamically detect and avoid jamming signals to maintain reliable communications.	✓	×	✓	✓	✓

scenario bandwidth made it impossible for competitors to achieve 100% traffic delivery. Since traffic types had different scoring weights, competitors had to intelligently decide which traffic to sacrifice to maximize points.

Table 1 summarizes the incumbent types, including descriptions of each type and their associated motivators.

Passive Incumbent

The passive incumbent emulated a radio receiver that had dedicated access to a portion of the RF spectrum. SC2 scenarios using the passive incumbent required that the aggregate interference level from competitor radios remain below a violation threshold level. If the threshold level were crossed, the incumbent entered a violation state and competitors were notified. During a violation state, the Colosseum scorekeeper began deducting competition points. Figure 1 illustrates a passive incumbent's high-level decision flow.

To calculate the aggregate interference level (P_{agg}), the incumbent buffered the most recent N in-phase quadrature (IQ) receiver samples, which was then applied to the following equation:

$$P_{agg} = 10\log\left(\frac{1}{N}\sum_{i=0}^{N}\left(I_i^2 + Q_i^2\right)\right) \qquad dBFS, \qquad (1)$$

where N is the product of reporting interval and the sampling rate, which results in a power measurement in decibels full scale (dBFS) relative to the Colosseum's analog-to-digital conversion (ADC) quantization. For example, if the passive incumbent reports the aggregate interference level every 0.100 seconds with a sampling rate of 10 Msps, the incumbent uses the most recent 1 million IQ samples to determine P_{agg}.

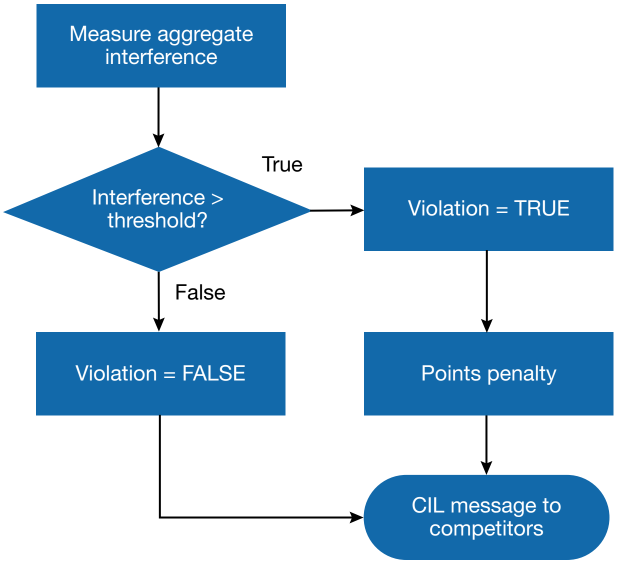
Measure aggregate interference
Interference > threshold?
True
False
Violation = TRUE
Violation = FALSE
Points penalty
CIL message to competitors

Figure 1. Passive violation decisioning. The passive incumbent continuously measured aggregate interference due to competitor transmissions. If the interference crossed a threshold, a violation was triggered and competition points were deducted.

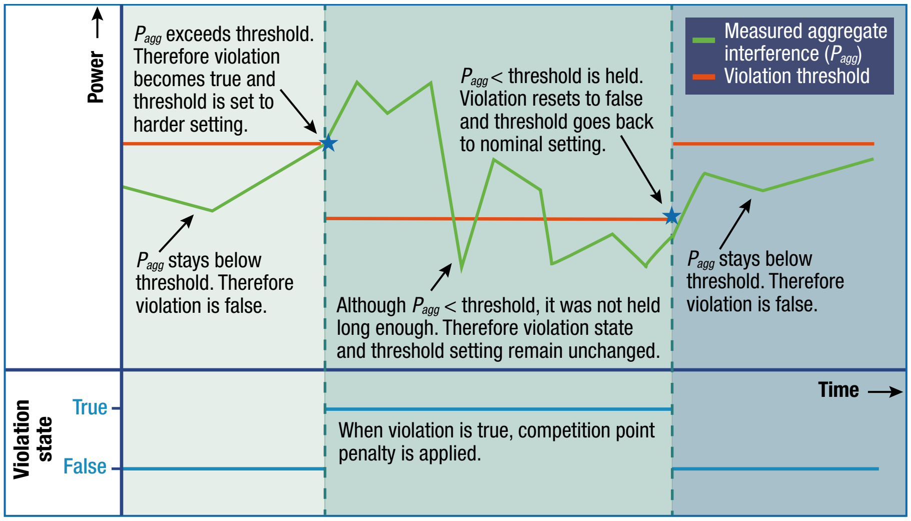
Power →
P_{agg} exceeds threshold. Therefore violation becomes true and threshold is set to harder setting.
P_{agg} < threshold is held. Violation resets to false and threshold goes back to nominal setting.
🟩 Measured aggregate interference (P_{agg})
🟥 Violation threshold
P_{agg} stays below threshold. Therefore violation is false.
Although P_{agg} < threshold, it was not held long enough. Therefore violation state and threshold setting remain unchanged.
P_{agg} stays below threshold. Therefore violation is false.
Time →
Violation state
True
False
When violation is true, competition point penalty is applied.

Figure 2. Passive incumbent example. Relationship between P_{agg} and violation state. The left-most and right-most columns show P_{agg} < *Threshold*; therefore, violation state evaluates to false and competition points are fully awarded. In the center column, P_{agg} > *Threshold*, violation state becomes true, which will result in a points penalty.

When competitors' aggregate interference exceeded the violation threshold, a persistent violation state was triggered. While in a violation state, competitors received a scoring penalty. To exit the violation state, competitors had to (1) lower their aggregate interference level so it fell below the violation threshold, and (2) maintain said aggregate interference level for a set duration. Only when these two criteria were met did the passive incumbent revert back to a non-violation state so that full scoring points could be awarded again. Additionally, the violation threshold used in the first criterion listed above was an even more challenging threshold than that normally used in a non-violation state. This incentivized competitors to avoid triggering violation events altogether—as it was harder to get out of a violation state than it was to avoid the violation state. Figure 2 illustrates an example competitor aggregate interference profile and how the passive incumbent reacts to it.

While the passive incumbent was operating, it periodically notified the competitors (over the CIRN Interaction Language [CIL]) whether the threshold was crossed. The results were reported in two message types: report and violation. Both messages contained the same information content, but report messages were delivered more frequently and could be used as a warning barometer by competitors (e.g., "I crossed the threshold, so I should lower my transmit power to get back below the threshold again."). However, if a competitor remained in violation of the threshold, a violation message was sent to notify the competitor that a violation event had occurred and therefore scenario points were deducted.

The passive incumbent was designed such that the threshold could change over time following a preconfigured schedule. If the aggregate interference level remained below a nominal threshold, the threshold would remain the same. If the aggregate interference level exceeded the threshold to the point of causing a violation, the threshold would change to a stricter threshold and stay there until the aggregate interference went back below the strict threshold. Once below the strict threshold for a set duration, the threshold would return back to the nominal threshold setting.

To ensure that the passive incumbent measured aggregate interference level accurately, a calibrated vector generator fed an incremental stair-step power profile into the passive incumbent's radio. If the interference-over-time measurement plot mirrored the signal generator's output profile, the passive incumbent was verified (Figure 3a). Additionally, to ensure that the passive incumbent was only capturing signals within its monitored spectrum, a tone sweep was fed into the incumbent's radio. For example, if the passive incumbent measured aggregate interference between 999.5 and 1000.5 MHz, and a tone sweep was generated from 995 to 1005 MHz, the incumbent reported zero aggregate interference while the tone swept from 995 to 999.5 MHz and 1000.5 to 1005 MHz. But while the tone swept from 999.5 to 1000.5 MHz, the passive incumbent reported a non-zero amount of aggregate interference (Figure 3b).

Active Incumbent

The active incumbent built on the passive incumbent by including a transmitter component that transmitted samples from a premade IQ file source. As it was for the passive incumbent, competitors still had to manage their transmissions to avoid aggregate interference threshold violations. But the new challenge posed by the active incumbent was that competitors had to detect and adapt to the incumbent's transmission pattern. If a competitor interfered with the reception of an incumbent's transmission, a violation state could be triggered. Since the active incumbent's transmission pattern was not advertised, competitors had to use real-time mechanisms to deduce the transmission pattern. The overall intent was for competitors to exploit both temporal and spectral gaps in the active incumbent's transmissions for their own communications, while still preserving the transmission integrity of the active incumbent.

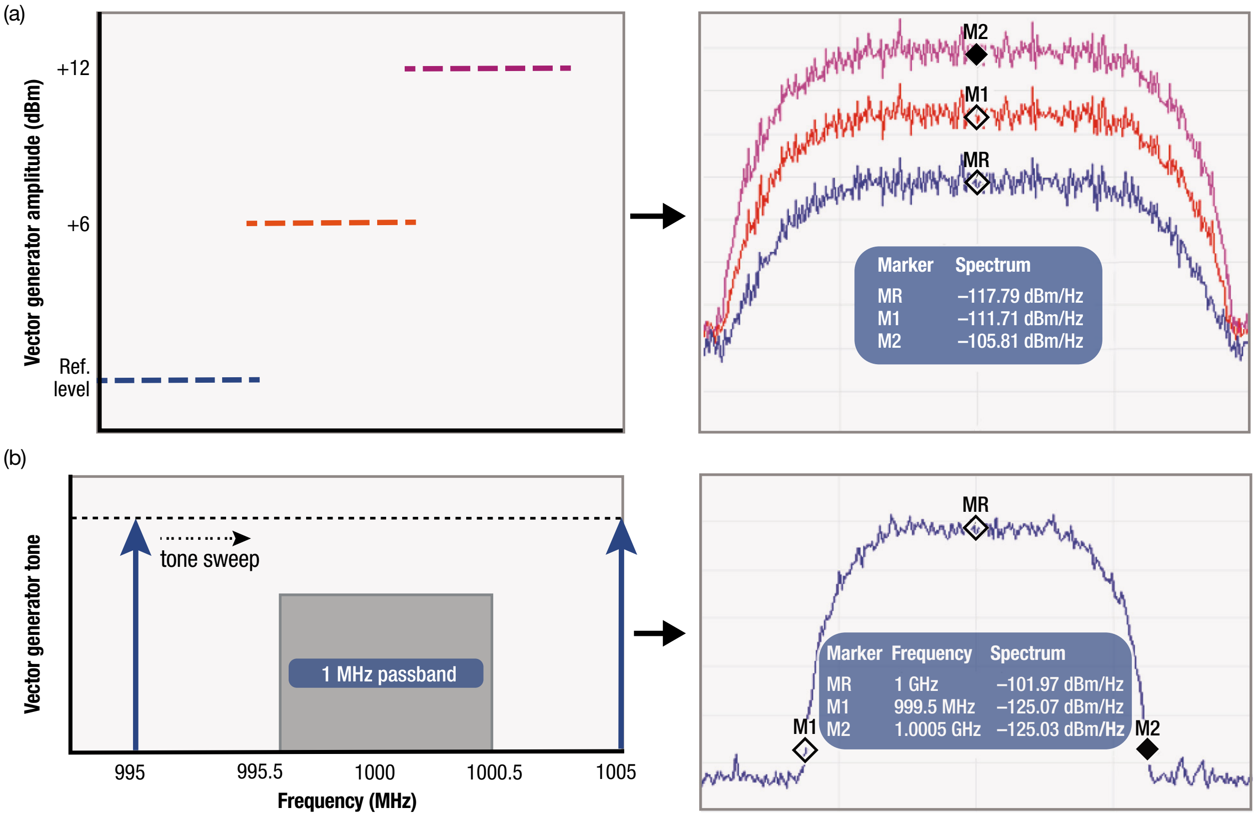

Figure 3. Passive incumbent validation. (a) Power measurement. Input power increased in a stair-step pattern, and the resulting measured power rose in the same amount. (b) Signal rejection. When the tone sweep was outside the passband, the signal was rejected, but when the tone sweep was inside the passband, the signal was kept.

The active incumbent resided on a single CIRN node—acting as both transmitter and receiver. The real-life counterpart of such a setup is a radar system. Because the active incumbent did not actually perform modulation/demodulation of the transmitted/received signal, it instead collected two sets of integrated power measurements to determine violation events (using the same power equation used in the passive incumbent, Equation 1). The first set simply measured the incumbent's own transmit power and nothing else (competitor transmissions were completely zeroized and masked away through wireless channel emulator channel taps—see the article by Barcklow et al. in this issue for details on the wireless channel emulator). If any power was present, the incumbent was transmitting; otherwise the incumbent was not transmitting. The second set measured competitor power and nothing else (incumbent transmissions were zeroized and masked away through wireless channel emulator channel taps). If competitor power exceeded the specified threshold, a violation event occurred; otherwise there was no violation. Finally, the violation decision was sent to competitors over the CIL. Figure 4 illustrates the event sequencing.

Since the active incumbent used the same violation decision mechanism as used in the passive incumbent, the example illustrated in Figure 4 also applies to the active incumbent. However, there were two operational differences. The first was the need to take account of the active incumbent's transmitting state. If the active incumbent was not transmitting, violation events were always disabled, allowing competitors to use the incumbent's frequency band however they wished. However, if the active incumbent was transmitting, a violation

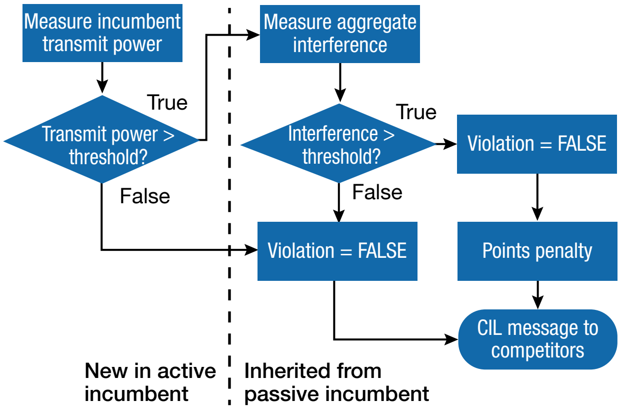
Measure incumbent transmit power
Transmit power > threshold?
True
False
Measure aggregate interference
Interference > threshold?
True
False
Violation = FALSE
Violation = FALSE
Points penalty
CIL message to competitors
New in active incumbent
Inherited from passive incumbent

Figure 4. Active violation decisioning. The active incumbent used the same violation decision as the passive incumbent, but it was in effect only when the incumbent was transmitting a signal.

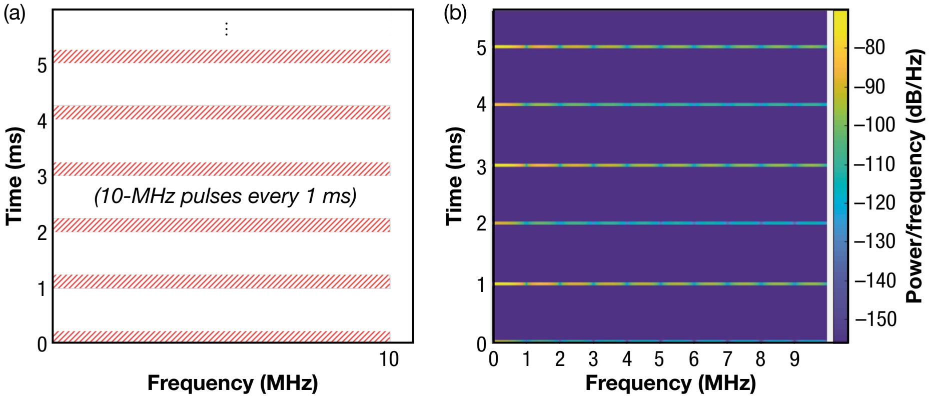

Figure 5. Active incumbent validation. (a) The theoretical design was a 10-MHz-wide pulse with a 1-ms repetition pattern. (b) The generated IQ samples (which were played back by the active incumbent) matched the theoretical design.

event was be triggered if the competitor's aggregate interference level exceeded the violation threshold. The second difference was the violation threshold. In the passive incumbent, if a competitor entered a violation state, they were challenged with a stricter threshold to be able to exit the violation state. However, in the active incumbent, there was only one violation threshold regardless of whether the competitor was in or not in a violation state.

For validating the functional aspects of the active incumbent, the same verification methods used in the passive incumbent could be reused to validate the active incumbent's receiver components. But since the active incumbent incorporated a transmitter component, additional verification methods were necessary. To ensure the transmitter was correctly transmitting the premade IQ file, an IQ recorder was connected to the active incumbent. The recorded IQ file was then analyzed using MATLAB. The active incumbent's transmitter was verified if the spectral characteristics (both in timing, magnitude response, and frequency response) met the scenario requirements. Figure 5 shows a validation of the active incumbent's transmitter playback of an IQ source and recording.

Jammer Incumbent

The jammer incumbent was transmit only, and its purpose was to inject varied jammer signals into the scenario to impede competitor RF performance. The jammer incumbent did not advertise its configuration, nor did it provide any real-time RF collision information to the competitors. Therefore, competitors were completely reliant on their own sensing mechanisms to detect and circumvent jammer effects.

The jammer incumbent supported three jammer waveforms (Figure 6): tone, additive white Gaussian noise (AWGN) band noise, and uniform random M-ary phase shift keying (M-PSK) modulation noise.

Since all jammer waveforms were generated in real time, multiple configuration options were available, as shown in Table 2. The jammer incumbent could support any number of waveforms and any combination of waveforms (each with their own configuration options). As a

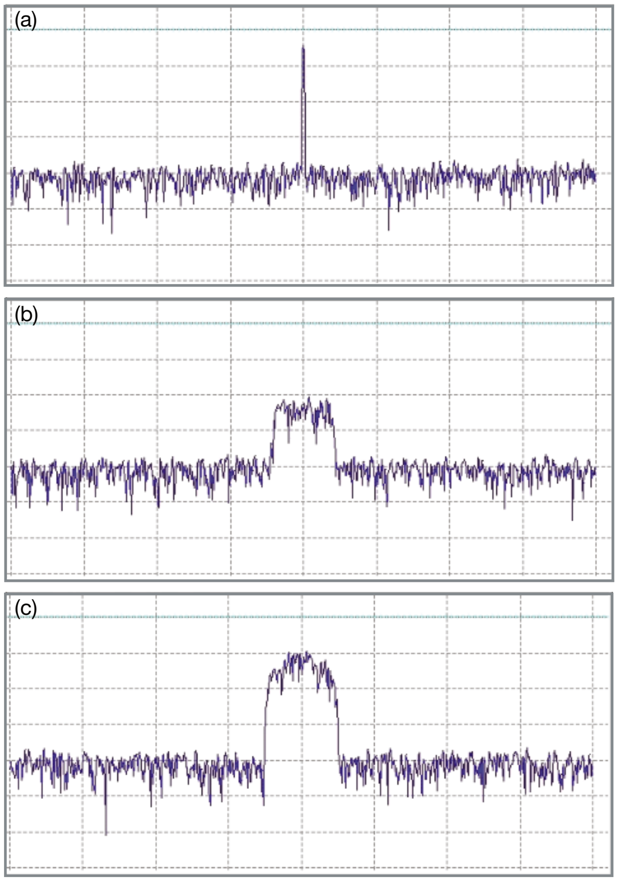

Figure 6. Jammer incumbent waveforms. (a) Tone. (b) AWGN band. (c) M-PSK modulation. Each waveform could be dynamically reconfigured at run time using the parameters specified in Table 2.

result, the jammer incumbent could create a plethora of complex jamming patterns that could stress competitors' learning algorithms. For example, Figure 7 illustrates the jammer response (orange) used during the second SC2 preliminary event competition.

Table 3 describes the stages in preliminary event 2. As shown in Figure 7, the red and green teams made no effort to avoid the jammer signals and simply stayed in their transmission bands throughout the scenario. As a result, during the jamming stages, this team's receivers were not able to close their links, which caused significant packet loss and ultimately prevented the team from earning points.

The blue team, on the other hand, used a detection-and-avoidance algorithm enabling it to (1) stay away from the red and green teams, and (2) interweave itself between gaps in the jammer signals to maintain communications (see Figure 8 for a closer view of stage 8).

Table 2. Jammer configuration options

Parameter	Description
Time start	Timestamp of when the jammer signal should start
Duration	Duration of how long the jammer signal should last
Amplitude	Relative dBFS power of the jammer signal
Center frequency	Center reference frequency of the jammer signal
Band transition width	Transition width between passband and stopband of the jammer signal
Band span	Width of the jammer band (3 dB down points)
Sweep span	Width of the jammer sweep
Sweep period	Periodicity of the jammer sweep; specifies the "speed" at which the signal sweeps from start to finish
Sweep direction	Specifies whether the sweep should move from left to right (ascending frequency) or right to left (descending frequency)

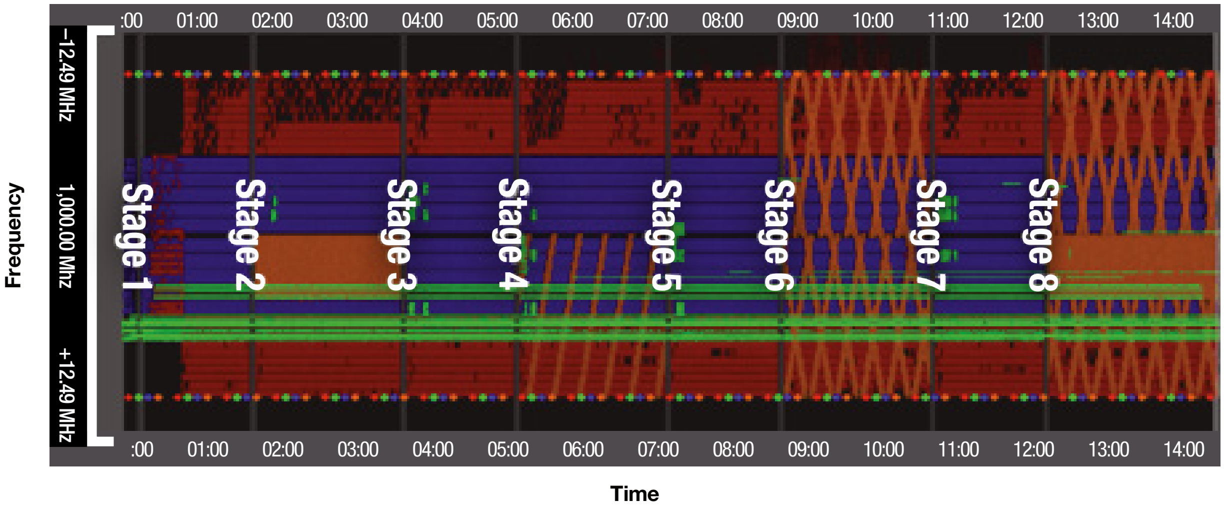

Figure 7. SC2 preliminary event 2 jammer incumbent scenario. Competitor radio activity (red, blue, and green) in the presence of jammer transmissions (orange). Red and green competitors lacked jammer detection and attempted to work brute-force through jammer effects, whereas the blue competitor had jammer detection and dynamically worked around jammer effects.

Table 3. Stages in SC2 preliminary event 2

Stage	Description
1	Jammer off
2	5-MHz wide stationary wideband AWGN
3	Jammer off
4	1-MHz sweeping wideband AWGN over 12.5 MHz
5	Jammer off
6	Four 1-MHz sweeping wideband AWGN over 12.5 MHz (creates a cross-hatch pattern)
7	Jammer off
8	Stage 4 and stage 6 combined

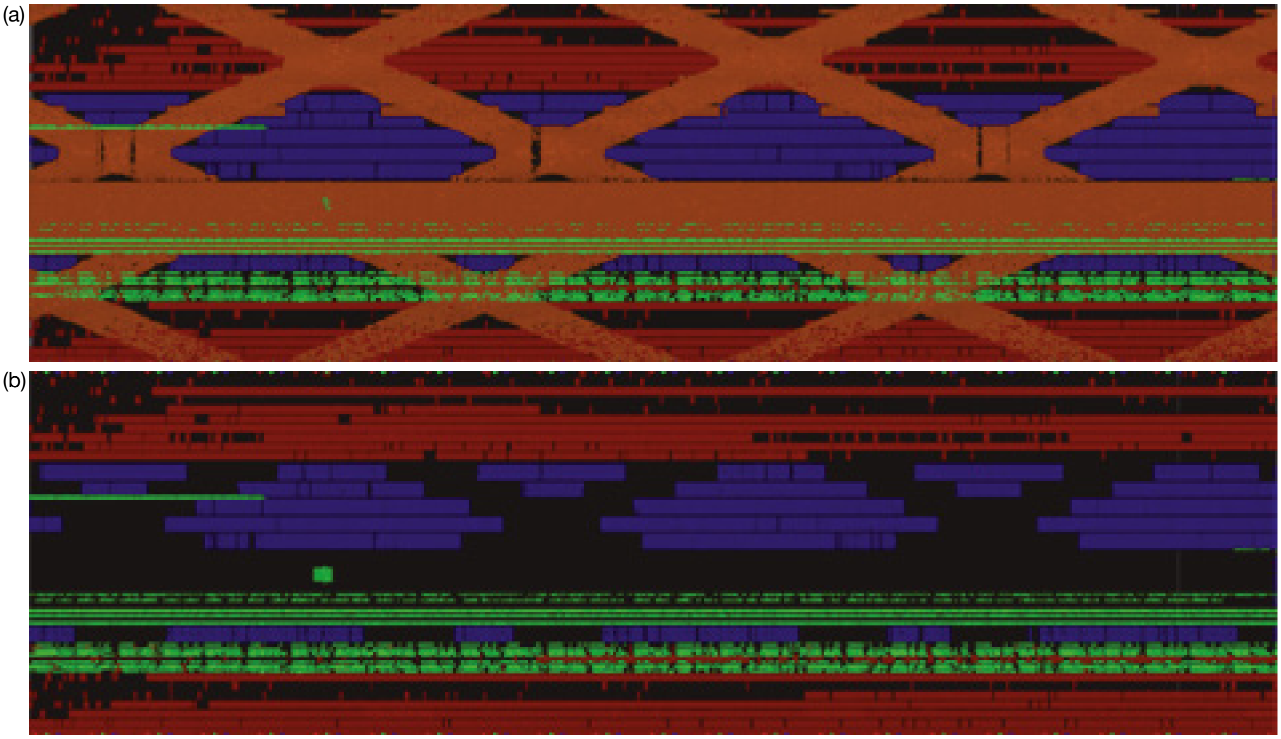

Figure 8. Stage 8 detailed view. (a) Jammer transmission graphic overlay on. (b) Jammer transmission graphic overlay off. The red and green teams ignored the presence of jammers and were unable to perform optimally, whereas the blue team interweaved in between the jammer and other competitor signals and was therefore able to maintain a radio link.

Despite the jammer sweeping across the entire frequency band, the blue team was still able to exploit these gaps to maintain communications and therefore had a higher uptime of reliable data transfer, earning it more competition points.

CONCLUSION

Incumbents served as obstacles for SC2 competitors, testing the robustness and resiliency of their radio designs in challenging RF environments. Passive incumbents were receive-only, and their goal was to ensure that competitors did not encroach on protected spectrum. Active incumbents were an extension of passive incumbents; they included a transmitter component and challenged competitors to detect and exploit temporal and spectral gaps for their own communications. Jammer incumbents were transmit-only and had the sole goal of creating an extremely contested RF environment to test competitor radio agility and recovery. By acting as non-competitor-controlled elements in the competition, incumbents forced competitors to react differently than they would have otherwise, helping them to push their architectures further.

ACKNOWLEDGMENTS: We thank Paul Tilghman (DARPA SC2 program manager) and Craig Pomeroy and Kevin Barone (Systems Engineering and Technical Assistance at DARPA) for their invaluable collaboration and support. We also thank the many APL SC2 contributors, whose names are listed on the inside back cover of this issue of the *Digest*. This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the US government.

REFERENCES

[1] "U-NII and TDWR Interference Enforcement." Federal Communications Commission. https://www.fcc.gov/general/u-nii-and-tdwr-interference-enforcement (last updated Aug. 22, 2019).

[2] "GNU Radio." https://www.gnuradio.org/ (accessed Aug. 22, 2019).

[3] "UHD." Ettus Research. https://kb.ettus.com/UHD (last modified Aug. 30, 2019).

Kaun J. Yim, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Kaun J. Yim is a software engineer in the Wireless Cyber Capabilities Group in APL's Asymmetric Operations Sector. He received a BS in computer engineering at Virginia Tech and an MS in electrical engineering at Johns Hopkins University. He has 15 years of experience, with an early-career focus on large-scale network design for military communication systems. Now at APL, his research interests and areas of focus include software-defined radios, digital communications, and wireless exploitation. For the SC2 Colosseum project, Kaun was the software lead for incumbent systems and co-developer for infrastructure systems. He also managed the Colosseum help desk portal. His email address is kaun.yim@jhuapl.edu.

Kenneth R. McKeever, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Kenneth R. McKeever is the acting assistant group supervisor of the Wireless Cyber Capabilities Group in APL's Asymmetric Operations Sector. He has a BS in electrical engineering from Penn State and an MS in electrical and computer engineering from Johns Hopkins University. His background and experience is in wireless security, communication systems theory, military and commercial communication technologies and protocols, signal processing, and computer networking. He has contributed to design, simulation, and both laboratory and field test and evaluation of various systems. As acting assistant group supervisor, Ken assists with directing and developing the group, including developing staff, prioritizing tasking, reaching out to and collaborating with other groups, and identifying and fostering new, impactful opportunities. He is a member of IEEE. His email address is kenneth.mckeever@jhuapl.edu.

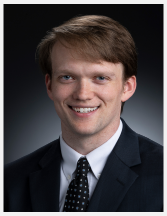

Daniel R. Barcklow, Asymmetric Operations Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Daniel R. Barcklow is a wireless communications engineer in APL's Asymmetric Operations Sector. He holds a BS in computer engineering from George Mason University and expects to complete an MS in electrical engineering at Johns Hopkins University in 2019. Daniel has experience in field-programmable gate array (FPGA) and software development for wireless channel emulation, wireless receiver design, embedded signal processing, FPGA development, RFNoC block development on Ettus X310 software-defined radios, script development to protect Red Hat Enterprise Linux systems against software exploits and to limit system vulnerabilities, and formal testing on large Java-based systems. Before joining the Lab in 2016, he interned at NIST and Micron Technology. His email address is daniel.barcklow@jhuapl.edu.