

Spectrum Collaboration Challenge

Preliminary Event 2 Scoring Procedures

November 21, 2018

FINAL



Defense Advanced Research Projects Agency
675 North Randolph Street
Arlington, VA 22203

Revision Summary

Revision	Date	Description
6	November 21, 2018	Refined file transfer specification; Declared Scoring Procedures as FINAL VERSION
5	November 7, 2018	Section 2.2 updated Payline acceptance threshold
4	October 11, 2018	Defined Tournament Scoring
3	September 28, 2018	Removed references to discrete performance thresholds
2	September 10, 2018	Clarification on Performance Thresholds, Individual Mandates, and Match Scores
1	April 26, 2018	(DRAFT) Initial Release for PE2

1 Overview

This document defines a set of scoring procedures for SC2 Preliminary Event #2 (PE2). These procedures apply to all participants in SC2. Further refinements and clarifications to this document will be released as needed. Prior to Preliminary Event #2, a final scoring document will be released. That document will be clearly marked FINAL, all others will be marked DRAFT. The Final Scoring Procedures document will carry the full authority of the rules in the SC2 Rules document as outlined in Section 8 of the SC2 Rules.

The Scoring procedures laid out in this and any future version of this document are subject to change at DARPA's sole discretion as outlined in the SC2 Rules Document Section 7.1.

2 Tournament Scoring

PE2 shall be composed of one or more rounds. A round is a collection of matches played in one or more scenarios. There are two types of rounds, a standard round and a payline round.

The Round Score, which aggregates Match Scores, is used to compute competition ranks.

2.1 Standard Round

A team's score in a standard round is the sum of the percentage of MOs achieved in each of the matches in that round.

$$\text{Round Score} = \sum_i \frac{\text{Match Score } i}{\text{Total MOs in Match } i} \times 100$$

At the conclusion of a standard round, rankings will be determined using the Round Score, and will be used to down-select the teams participating in the next round. Scores are reset at the beginning of each standard round. The final standard round scoring will determine the final rankings.

2.2 Payline Round

One round of the PE2 tournament, known as the payline round, will be used to determine the eligibility of a team to receive prize money. This round does not affect rankings, only prize eligibility.

The payline round is designed to test the ability of an ensemble to outperform a naïve division of spectrum resources. Each match first consists of one or more baseline stages. In each of these baseline stages the total spectrum bandwidth available will be equitably divided into contiguous, but non-overlapping allocations. Each team's network will receive its own allocation. A second set of stages, called payline stages, will permit all teams to use the total scenario bandwidth, without predetermined bandwidth allocations per network.

In each stage, the number of Individual Mandates (IMs, see Section 3) given to each team's network is the same, however each network is given a different share of total data traffic (as measured in bits-per-second, bps). Each network will have one baseline and one payline stage where they are given the largest portion of overall traffic, this is called the lead position.

To be eligible for PE2 prizes, a team's network must participate in at least one payline match where the number of IMs achieved in the payline stage, in which they occupy the lead position, **meets or** exceeds the corresponding baseline stage where they occupied the lead position. Additionally, the sum of all

other networks' achieved IMs in the payline stage must be equal to or exceed the sum of all achieved IMs in the corresponding baseline stage. Multiple payline matches will be conducted.

To maintain competition integrity, teams may not artificially lower their performance in the baseline stage to inflate their performance in the payline stage. Teams who do artificially reduce their performance during the baseline stage may be disqualified and determined ineligible for prizes.

3 Match Score

Each competitive PE2 match is divided into *stages*. A stage score measures the completion of objectives. A team's match score is the sum of its stage scores.

During a stage, each network is given a number of *individual mandates (IM)*. An IM specifies the required wireless communications performance parameters for a particular IP traffic flow. A network is said to achieve an IM when the required parameters for that traffic flow have been met continuously for at least as long as a period of time called the *steady-state period (SSP)*. *Mandated outcomes (MO)* are achieved when every network in the ensemble is achieving one or more IMs simultaneously. The value of the MO at any given time is the minimum number of IMs being achieved by any member of the ensemble.

The stage score is comprised of two parts, 1) the maximum value of MO, and 2) the earliest time at which the ensemble achieves that value during the stage.

In addition to IMs, scenarios have *gates*. A gate is a required behavior not associated with an individual IP traffic flow. A network cannot meet the required parameters of any of its IMs when a gate is being violated. Thus, the value of MO for the ensemble will be 0 until at least one SSP duration after correcting a gate violation.

The following sections detail the above outlined scoring procedure.

3.1 Stage Score

Measurement Period (MP): An interval of time in a stage over which performance parameters are assessed. The length of the measurement period (MP) for PE2 is 1 second. All times t referred to in the below are discretized to *MP*.

Gates: A gate is a required behavior not associated with an IP traffic flow. A gate is met if the network or ensemble complies with the required behavior.

Performance Thresholds (PT): A PT is a required behavior of an IP traffic flow (such as required throughput or latency). The PT is deemed to be met in a given MP if the associated parameter meets or exceeds the specified threshold in that MP, and if no gate is violated in that MP.

During a given MP, if no IP packets are offered, or if the parameters of a PT do not require the delivery of IP packets, the following conditions hold:

- The PT is deemed to be met if it was met in the previous MP.
- The PT is deemed to be not met if it was not met in the previous MP.

Steady State Period (SSP): A required duration of time over which IMs must be achieved to demonstrate that the outcome is non-transient; this value is common to all IMs in a stage.

SC2 PRELIMINARY EVENT #2, FINAL SCORING PROCEDURES

Individual Mandates (IM): An IM specifies one or more PTs for an IP traffic flow. An IM is deemed to be achieved at time t if the specified PTs have been satisfied for the preceding period of time at least as long as the SSP.

Mandated Outcome (MO) is defined as:

$$MO(t) \equiv \min_{\text{networks } n} (\text{number of IMs achieved by network } n \text{ at time } t)$$

Stage Score (SS) is defined as follows.

$$MAX_{MO} \equiv \max_t MO(t)$$

$$SST \equiv \operatorname{argmin}_t (MO(t) = MAX_{MO})$$

$$SS \equiv [MAX_{MO}, SST]$$

When comparing two different stage scores, SS_1 and SS_2 , comparison is performed primarily on the first element, MAX_{MO} , with higher values ranking higher. Comparison using the second element, SST , is only performed if SS_1 and SS_2 have the same MAX_{MO} value, with lower SST values ranking higher.

3.2 Individual Mandates

An *IP traffic flow* is a unique instance of IP traffic of a single application layer type to be delivered between source and destination IP addresses that are both in the team's network.

An Individual Mandate (IM) has two components:

- 1) A Flow Unique Identifier (flow UID) that identifies an instance of a specific traffic type, a source IP address, and a destination IP address.
- 2) Performance Thresholds (PT) for the IP traffic flow.

The PTs and their associated interpretation are given in Table 1. The below performance thresholds are computed by the scoring subsystem for each Measurement Period (MP) during the match.

Table 1 – Performance Thresholds and associated scoring interpretation.

Performance Threshold	Description
<i>max_latency_s</i>	The maximum allowed latency (in seconds) for any IP packet in the flow during the measurement period.
<i>min_throughput_bps</i>	The minimum allowed throughput for the flow during the measurement period.
<i>file_transfer_deadline_s</i>	The maximum allowed latency (in seconds) for 90% of the packets of a file transfer to occur.

Throughput is measured as bits per second of IP traffic delivered to the traffic generator sink at the destination node.

Packet latency is measured from when the traffic generator source provides a packet to the source node to when the traffic generator sink receives it from the destination node. A packet which is not delivered in less than *max_latency_s* is considered dropped, decreasing the measured throughput, even if it is eventually delivered.

File transfer latency is measured from when the traffic generator source provides a packet to the source node to when the traffic generator sink receives it from the destination node. Latency is measured on a per packet basis. 90% of packets offered in a MP must be delivered in order to meet the PT. A packet which is not delivered in less than *file_transfer_deadline_s* is considered dropped.

3.3 Gate: Incumbent Protection

An *Incumbent* is a radio spectrum user that must be afforded protection in the spectrum. In scenarios with Incumbents, protecting the Incumbent is a gate for the mandated outcomes.

An *informing Incumbent* is one that provides feedback via the collaboration network to inform other radio networks if its threshold has been violated. Informing incumbents may have a threshold set as a function of acceptable interference power, or have a threshold based on acceptable degradation of nominal throughput.

In either case, an informing incumbent will advertise both its current threshold (interference power in dBFS, or minimum throughput in bps) as well as its current measurement against the threshold (interference power in dBFS, or achieved throughput). These reports will be sent every *Reporting Period*. A *Violation Period* is comprised of a number of successive Reporting Periods. The current threshold will not change during a Violation Period. If the average of the reported measured values over the Violation Period exceeds the current interference power threshold, or is below the minimum throughput threshold (whichever is applicable), the Incumbent is deemed to have experienced interference during that entire Violation Period. This is referred to as an *Interfering Violation Period*.

The *incumbent protection gate (IPG)* is deemed to be met during an MP, if no Violation Period in or overlapping that MP is an Interfering Violation Period.

3.4 Gate: Transmission Only in Allowed Bandwidth

RF transmissions are only permitted within the allowed bandwidth of each scenario. Respecting the allowed bandwidth limit is an additional gate for the mandated outcomes.

A transmission outside the allowed bandwidth is determined according to the following equation, with variables as defined in Table 2:

$$PSD(f_{disallowed}) > \Gamma$$

Table 2: Variables for determining transmission outside of allowed bandwidth

Variable	Description	Value in PE2
----------	-------------	--------------

SC2 PRELIMINARY EVENT #2, FINAL SCORING PROCEDURES

$PSD(f_{disallowed})$	The power spectral density of any transmitting node at frequency $f_{disallowed}$ as observed by the observer node with 0 dB of digital attenuation (see Figure 1)	<i>variable</i>
Γ	Transmission energy threshold referenced to full scale (dBFs). "Full scale" refers to the maximum signal that can be represented by the observer A/D converter.	-141 dBFs/Hz

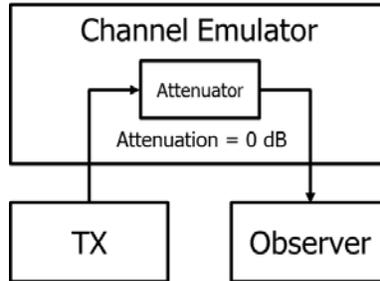


Figure 1 – Configuration of observer node

Multiple power spectral density observations will be taken of the transmissions of each node within a measurement period by an observer. Each PSD observation will average the power received over a time window. The length of the time windows is not specified in advance and may change during the course of a match.

The receiver gain of the observer node will be set according to Table 3 below. The gain of the closest center frequency to that of the scenario is used.

Table 3: Observe Node Receiver Gains

Center Frequency	Observer Receive Gain
1.0 GHz	7 dB
2.4 GHz	8 dB
5.8 GHz	15 dB

The *transmission bandwidth gate (TBG)* is deemed to be met during an MP if no transmissions beyond the allowed frequency bandwidth are detected within that MP.

3.5 Match Score

The match score for any given team is the sum of the Stage Scores.

$$S \equiv \sum_i SS_i = \left[\sum_i MAX_{MO_i}, \sum_i SST_i \right]$$

4 Feedback to networks

A team's network does not receive real-time feedback on scoring. Specifically, the system does not inform the network whether any IM or MO has yet been achieved in the current stage, the current value of the *SSP* timer, or whether applicable gates have been met.