AVERAGE TYPICAL MISSION AVAILABILITY: A FREQUENCY MANAGEMENT METRIC

Charles H. Jones, PhD
412TW/ENTI
Edwards Air Force Base
charles.jones@edwards.af.mil

ABSTRACT

One approach to improving spectrum usage efficiency is to manage the scheduling of frequencies more effectively. The use of metrics to analyze frequency scheduling could aid frequency managers in a variety of ways. However, the basic question of what is a good metric for representing and analyzing spectral usage remains unanswered. Some metrics capture spectral occupancy. This paper introduces metrics that change the focus from occupancy to availability. Just because spectrum is not in use does not mean it is available for use. A significant factor in creating unused but unusable spectrum is fragmentation. A mission profile for spectrum usage can be considered a rectangle in a standard time versus frequency grid. Even intelligent placement of these rectangles (i.e., the scheduling of a missions spectrum usage) can not always utilize all portions of the spectrum. The average typical mission availability (ATMA) metric provides a way of numerically answering the question: Could we have scheduled another typical mission? This is a much more practical question than: Did we occupy the entire spectrum? If another mission couldn't have been scheduled, then the entire spectrum was effectively used, even if the entire spectrum wasn't occupied.

KEYWORDS

Frequency Management, Spectrum, Metrics, Frequency Scheduling, Fragmentation

INTRODUCTION

In its daily incarnation, frequency management is the process of scheduling blocks of frequencies and time in a non-interfering fashion. The unit of scheduling is a mission profile, which is a single contiguous block of frequencies over a contiguous period of time. Thus, geometrically, a mission profile can be considered a rectangle in a time vs. frequency grid. Figure 1 shows such a grid with four mission profiles scheduled. A fifth potential mission that conflicts with one of the already scheduled missions is also shown. This conflict is geometrically illustrated by the overlap of the rectangles.

A significant concern in frequency scheduling is that of spectrum fragmentation. (This is virtually identical to the concept of disk fragmentation on a computer.) The placement of mission profiles can easily lead to a fragmentation of the domain into pieces that are not usable. For example, the potential new mission in figure 1 can not be scheduled anywhere without conflicting with the missions already scheduled. Within this example, a simple solution is to move the left-most mission down a notch or two and everything fits just fine. However, real world applications may have tens or hundreds of missions so that deciding what rectangles to move may not always be immediately obvious. In fact, many people spend all day trying to deconflict frequency schedules.

There are a variety of reasons why scheduling frequencies is hard. One reason is simply that scheduling is a mathematically difficult problem. This scheduling problem is an example of a dynamic bin packing problem. Technically, such problems are **NP**-hard, meaning that they are exponentially difficult and, in the worst case, computers could take years or centuries to find optimal solutions. Other realities that make frequency scheduling difficult include: not all radios are tunable; missions often require multiple frequency ranges; changes require coordination, time, and often have cascading effects; frequency isn't the only thing that has to be rescheduled; prioritizations are not always well established; and a variety of other reasons.

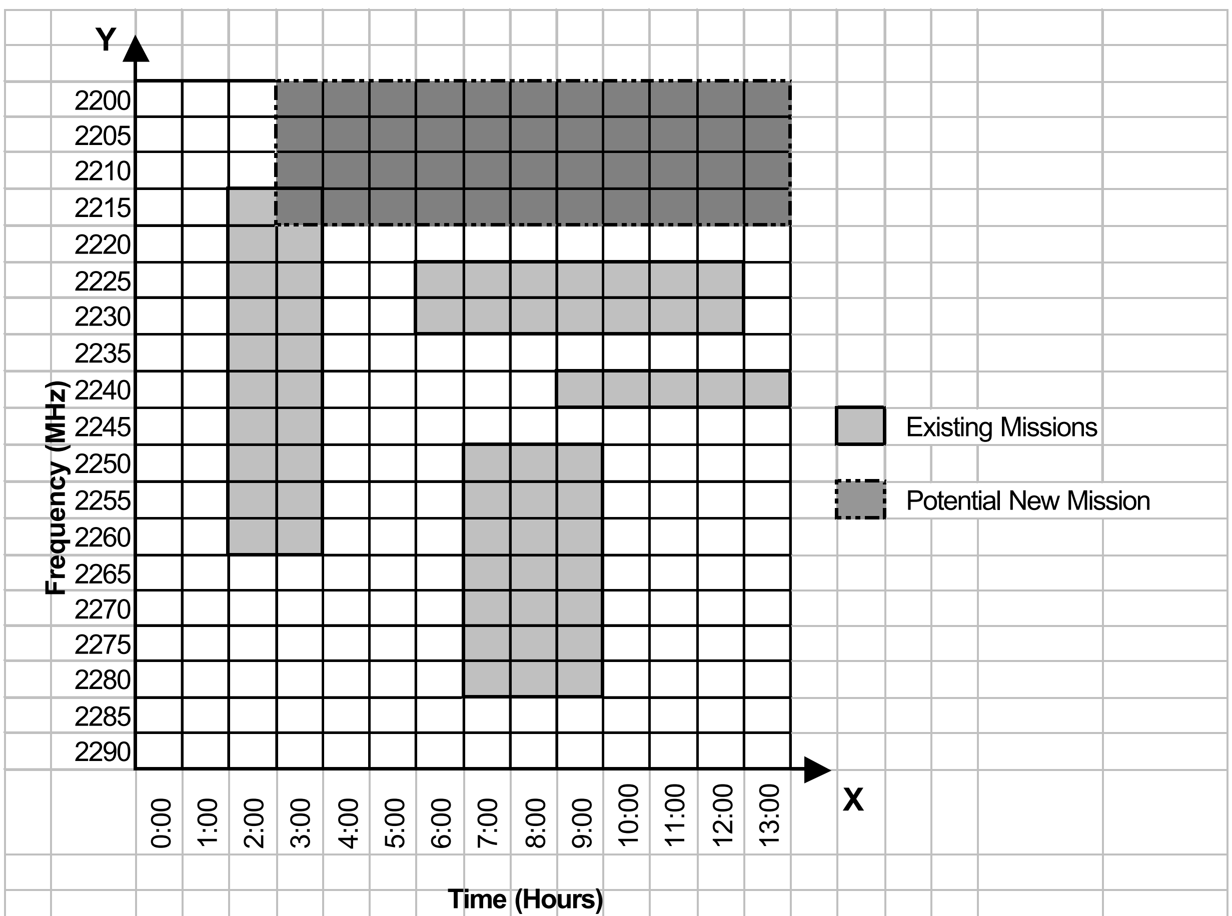

Figure 1 – Time vs. Frequency Grid with scheduled and conflicting missions.[1]

[1] The mission profiles in all examples are fictional. In particular, most of the profiles are larger than in practice. They are provided for illustration of the concepts only.

Because of the difficulties of scheduling, the reality is that the spectrum gets fragmented and not all of the spectrum is used. The questions addressed in this paper are: What metrics capture how well the scheduling was done considering the reality of fragmentation? How do we generate a numerical measure of whether or not another mission could have been scheduled? In other words, was there a *usable* portion of the spectrum *available*?

MISSION AVAILABILITY METRICS (FIXED TILE METHODS)

Focusing on the idea of whether a new mission profile will fit into a grid with already scheduled missions, one approach is simply to try every possible position and see if a conflict is created. That is, think of the mission profile as a fixed size, rectangular, tile and see if it fits into the jigsaw puzzle anywhere. Although, the question isn't just: Can a particular mission profile fit somewhere? The question is more about quantifying probabilities of scheduling missions in general. So a next level use of the fixed tile approach is to ask: How many places can a mission profile be placed? This allows the calculation of a probability by dividing the number of places a mission profile can be scheduled without conflict by the total number of possible placements.

Depending on what restrictions are applied, there are several metrics that can be defined.[2] The fixed tile approach can also be used to define spectral fragmentation.

1. Time Required Mission Availability (TRMA). The probability of scheduling a mission given a mission profile and a required start time, but not a required frequency.
2. Frequency Required Mission Availability (FRMA). The probability of scheduling a mission given a mission profile and a required frequency, but not a required start time.
3. Ad Hoc Mission Availability (AHMA). The probability of scheduling a mission given a mission profile and flexibility in both frequency and start time.
4. Average Typical Mission Availability (ATMA). Most users doing frequency scheduling have more than one typical mission profile. ATMA is the average of the AHMA for all typical mission profiles. (Although, certainly you could average any of these metrics.)
5. Percent Fragmentation. Define a portion of the frequency-time domain to be fragmented if that portion is not usable to schedule a typical mission profile. Percent fragmentation can then be calculated as the percent unusable out of the total domain or as the percent unusable out of the portion of the domain not already scheduled. Since this is defined for a specific mission profile, the final percent fragmentation is defined to be the average fragmentation for all typical mission profiles.

The following sections provide a detailed example of ATMA, which the author puts forth as the most useful of these metrics for answering the questions: Could another mission have been scheduled? and Has the spectrum been scheduled efficiently? ATMA requires the introduction of "typical" mission profiles. This allows individual users to customize some of the metrics to their particular environment by identifying what their most common mission profiles are. Another customizable aspect for all these metrics is the definition of both the total frequency available and

[2] These and the maximum availability metrics listed have all been formally defined by the author.

the portion of the day used for the calculations. In order to give an example of ATMA, an example of AHMA must be given first since ATMA is the average of AHMAs.

AD HOC MISSION AVAILABILITY (AHMA)

AHMA is the probability of scheduling a mission given a mission profile and flexibility in both frequency and start time. A supporting metric is an absolute count of the available (frequency, start time) pairs the mission can be scheduled at. Numeric interpretations of AHMA include:

1. AHMA > 0 means the mission can be scheduled for some (frequency, start time) pair.
2. AHMA = 1 means there are no other missions scheduled in the frequency and start time ranges.
3. The greater AHMA is, the more flexibility there is to schedule the mission.

Predetermined inputs to the algorithm

1. Required bandwidth
2. Required duration
3. Available frequency range
4. Available mission time range
5. Existing scheduled missions

Algorithm

```
Calculate earliest and latest start times from available mission time range and required duration.
Calculate lowest and highest frequency from available frequency range and required bandwidth.
Set dT to the minimum time increment (1 hour in this example.)
Set dB to the minimum bandwidth increment (5 MHz in this example.)

available count = 0
for start time = earliest start time to latest start time step dT
     for frequency = lowest frequency to highest frequency step dB
        if schedulable then
            Available count = Available count + 1
        end if
     end for
end for

number of available start times = (latest start time – earliest start time ) / dT
number of available frequencies = (highest frequency – lowest frequency) / dB
Number of (frequency, start time) pairs
    = number of available start times * number of available frequencies
AHMA = available count / number of (frequency, start time) pairs.
```

Note that AHMA = Average TRMA over all available frequencies
= Average FRMA over all available start times.

SPECIFIC EXAMPLE OF AHMA

Given these inputs to the algorithm:

1. Required bandwidth 15 MHz
2. Required duration 5 hrs
3. Available frequency range 2200 – 2295 MHz
4. Available mission time range 0000 – 1400
5. Existing scheduled missions (See figure 1.)

Then:

Earliest start time = 0000
Latest start time = 0900
Lowest frequency = 2200
Highest frequency = 2280
Number of available start times = 10 (using dT = 1 hr)
Number of available frequencies = 16
Number of available (frequency, start time) pairs = 160
Available Count = 33 (10 times at 2200 MHz, 1 time at 2235 MHz, 5 times at 2205 MHz, 5 times at 2210 MHz, and 3 times each at 2265, 2270, 2275, and 2280 MHz)
AHMA = 33/160 = 0.21 (or 21%)

The first four iterations of the algorithm for this example are illustrated in figure 2. Place a rectangular tile, 15 MHz by 5 hrs, into the grid at the earliest start time, 0000, and at the lowest frequency, 2200. The tile does not intersect any other scheduled missions, so increment the available count to 1. Then move the tile down 5 MHz (dB). The tile intersects an already scheduled mission so do not increment the available count. Continue moving the tile down 5 MHz at a time. The tile continues to intersect a scheduled mission until the lowest frequency used by that tile is 2765 MHz. At which point the available count starts to be incremented again. After the tile reaches 2280 MHz, the algorithm resets the frequency and increments the start time by 1 hour (dT). Then the inner iteration of walking the tile down the frequencies repeats. This continues until all the possible (frequency, start time) pairs have been checked for availability.

AVERAGE TYPICAL MISSION AVAILABILITY (ATMA)

This calculates the average AHMA for several typical mission profiles. This is a summary statistic that gives a one number estimate of the probability of scheduling a typical mission on an ad hoc basis. Numeric interpretations of ATMA include:

1. ATMA low (near 0) means a very low probability of scheduling an ad hoc mission. Also indicates a schedule that is very full or very fragmented. In other words, scheduling a mission would require major rework of existing scheduled missions.
2. ATMA high (near 1) means high probability of scheduling an ad hoc mission.
3. The greater ATMA is, the more flexibility there is to schedule a mission.

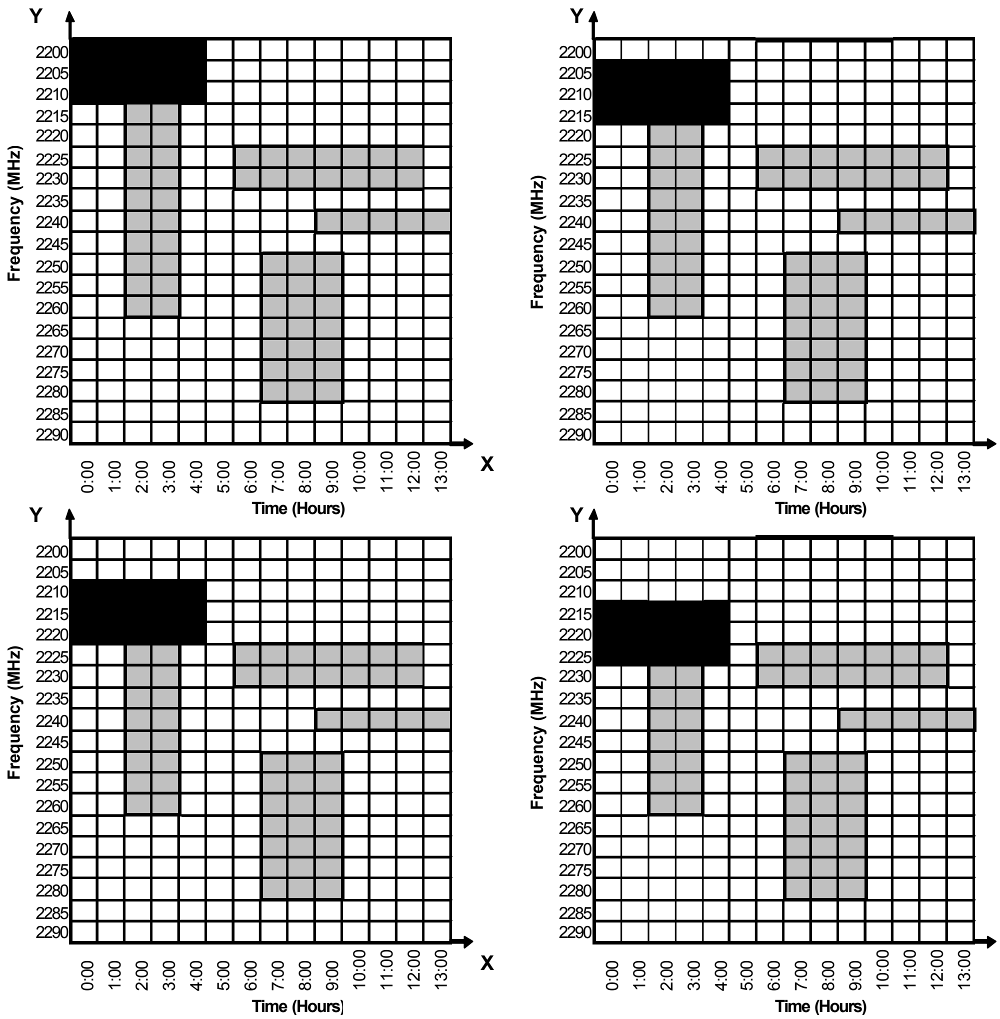

Figure 2 – Moving the fixed tile to calculate AHMA

<u>Predetermined inputs to the ATMA algorithm</u>

1. Predefined typical mission profiles $\{(b_i, d_i) : i \quad 1,..., n\}$, where the b_i are bandwidths in MHz and the d_i are durations in hours
2. Available frequency range
3. Available mission time range
4. Existing scheduled missions

Algorithm

$$\text{ATMA} = \frac{\sum_{i=1}^{n} AHMA((b_i, d_i))}{n}$$

SPECIFIC EXAMPLE OF ATMA

Given these inputs to the algorithm:

1. Predefined typical mission profiles {(5 MHz, 3 hrs), (15 MHz, 5 hrs), (20 MHz, 11 hrs)}
2. Available frequency range 2200 – 2295 MHz
3. Available mission time range 0000 – 1400
4. Existing scheduled missions (See figure 1.)

ATMA = (118/228 + 33/160 + 0/64) / 3 = 0.07

That is, there is only about a 7% chance of scheduling a typical mission.

COMPARISON TO PERCENT OF OCCUPANCY (PO)

In order to appreciate the significance of ATMA, it is useful to compare it to percent of occupancy (PO) which is based on how much of the spectrum is actually being used. Specifically, PO is calculated by dividing the amount of the spectrum used by the total available spectrum. In the example used previously there are four mission profiles: {(50, 2), (10, 2), (5, 5), (35, 3)}. These make a total occupancy of 100 + 20 + 25 + 105 = 250 MHz. There is a total available occupancy of 95 MHz x 14 Hours = 1330 MH. This results in a PO of 250/1330 = 0.19 or 19%.

A 19% occupancy gives the impression that 81% of the spectrum is available for use. But, as we can see from the ATMA calculations, there is only a 7% chance of scheduling a typical mission! This illustrates dramatically the effects of fragmentation on scheduling and the main thesis of this paper. *Just because spectrum is not occupied does not mean it is usable. Just because there is a low overall occupancy does not mean that the spectrum is being used inefficiently.*

MAXIMUM AVAILABILITY METRICS (VARIABLE TILE METHODS)

Instead of asking whether a particular mission profile can be scheduled it might be desirable to ask how large a profile can be scheduled. This question can be answered by using a similar method of looking at all possibilities, but this time, instead of using fixed-size tiles, the tiles are varied in size until they intersect existing scheduled missions. For example, to find the largest possible mission profile that can be scheduled at time 0:00 and 2290 MHz, start by placing a 1x1 tile in that location. If that doesn't intersect a scheduled mission, then try a 1x2 tile, and then a 1x3 tile, etc. until the tile does intersect. At that point, say at a 1x19 sized tile, then start increasing in the X direction and try a

2x19 tile. With a carefully crafted algorithm that checks all possible tiles, the largest such tile (or mission profile) can be found. In the example illustrated by Figure 2, the largest tile that can be placed at time 0:00 and 2290 MHz is a 7x6 tile representing a mission profile of 7 hours by 30 MHz. It is the largest mission profile in the sense of largest mission occupancy, which is calculated by multiplying the time by the bandwidth. In this case, the mission occupancy is 210 MHz hrs.

Just as with the fixed tile methods, there are several variant metrics that can be defined.

1. Maximum Available Duration (MAD). The longest possible duration a mission can be scheduled for a given start time and required bandwidth.
2. Maximum Available Bandwidth (MAB). The largest possible bandwidth a mission can be scheduled with for a given start time and required duration.
3. Maximum Available Mission Occupancy (MAMO). The mission profile with the largest mission occupancy that can be scheduled for a given start time and frequency.
4. Maximum MAMO. The maximum MAMO over a set of start times and frequencies.
5. Averages for each of these over a set of start times and frequencies can also be calculated.

USE OF THESE METRICS

There are always many ways to use these and other metrics. Here are some suggestions.

Availability Trend Analysis

Many of these metrics could be used for long term planning – in both time and frequency domains. For example, if average AHMA for evenings or weekends is extremely high while average AHMA during normal workdays is nearly 0, then the metrics would indicate the need for more flights scheduled during evenings and weekends. Similarly, comparisons of AHMA for L and S bands could indicate the need to migrate one way or the other.

Another use would be in analyzing future ability to support a large program. For example, consider the situation where a new program appears. It is likely they will be able to provide an estimate of spectrum requirements over the life of the project. There will be existing projects that will continue to operate over all or part of the time period of this new project. It would be possible to take the requests that these existing projects have made and forecast their future usage. Add these to the new project's estimates. Run scheduling algorithms over future time periods and see what kinds of availability metrics are returned. If general availability is low, there should be real concern that the new project could not be supported and that further efficiency efforts should be implemented.

Defense Against Further Spectrum Reduction

This is certainly one of the driving forces for introducing these metrics. Actual percent of occupancy has never been, and probably never will be, 100 percent for Test and Evaluation (T&E) purposes (except possibly for short periods of time.) This makes defending usage difficult. However, this doesn't change the fact that when the T&E community needs spectrum, either the spectrum is available or the mission doesn't happen. A proactive approach to spectrum defense is to present spectral usage data in a manner that captures the full complexity of the problem. Availability metrics are an attempt to do so.

As has been shown in the examples, availability may be very small even with a very small percent of occupancy. Although actual implementation of these metrics will provide the real answer, it is expected that most of the availability metrics are close to zero over normal operational times. This should help illustrate numerically that the T&E community is not wasting spectrum; that the spectrum is, in fact, being used efficiently within the working constraints of the discipline.

Real Time Reassignment Aids

To the person on the front line trying to schedule missions, the mission availability metrics could be useful, especially if the specific time and frequency choices were returned in addition to the counts. That is, instead of having to manually search for possible positions to schedule a requested mission, it would be quite useful if the schedulers could be provided with all possible choices. The maximum availability metrics could also be helpful, especially for ad hoc requests for projects that are willing to take what they can get.

Scheduling Algorithm Analysis

This may be wishful thinking, but it would be nice if the scheduling process could be automated or, at least, semi-automated. That is, are there algorithms that can help increase scheduling efficiency in both manpower and number of missions scheduled? Several of the metrics identified could be used for static analysis. That is, given a set of either past requests, fictional requests, or projected requests, it would be possible to run several algorithms and compare the resulting metrics.

DISCUSSION

To a certain extent, these metrics have been developed by starting with a couple of algorithmic ideas and defining the various metrics that come out of the process. However, these metrics also have been developed starting from the idea that, just because a portion of the spectrum is not occupied, doesn't mean that it is not used efficiently. A phrase that helps codify this is: "Use" is "denial to others." A simple example of this is that when you schedule a block of frequencies, no one else can use it whether you do or not. A more complicated example is fragmentation. If, all constraints considered, a fragmented but unused portion of the spectrum can not be used by a particular mission, that mission has been denied use of that portion of the spectrum. Hopefully, these ideas illustrate that a metric based solely on spectral occupancy is a somewhat uninformed metric.

Availability metrics are not currently in use. There is some potential for them to be incorporated into the Integrated Frequency Deconfliction System (IFDS), but a final decision has not been made. The frequency management community is engaged in a discussion of how to present spectral usage to the world. There are certainly other metrics or variants to be considered, but the author believes that the idea of availability, rather than occupancy, is a valid and strong candidate for capturing actual scheduling efficiency.