EBC Spectrum Observatory Demo

1/17/2013

Contents

1. Script

1.1. Overview

The Microsoft Spectrum Observatory collects and presents frequency spectrum usage data in a way that interested parties can have an open discussion about current and future uses of allocated frequency spectrum. A new measurement station has been installed at the Executive Briefing Center in Brussels, Belgium to support this effort by collecting spectrum usage data.

1.2. Objectives

Illustrate that the Microsoft Spectrum Observatory can be used to present spectrum usage data that is correct, reliable, and relevant.

Demonstrate the thought process behind analyzing the presented data, understanding what the data means, and inferring valid conclusions.

Grow interest in the Spectrum Observatory with the hope that other parties will become interested in contributing their own data for others to view.

1.3. Background

Microsoft has been working on technologies that enable wireless devices to opportunistically use spectrum without interfering with TV broadcasts and other licensed operations for a number of years – in fact for more than a decade. Our trials and commercial pilots in the UK, Singapore and elsewhere have demonstrated the economic potential of this technology as well as answered most of the technical questions.

A large part of our success was our ability to measure the usage of wireless spectrum and use and share that information with regulators to better inform decisions on which spectrum bands are the most appropriate near term targets for expanded spectrum sharing. So building on these efforts, and Observatories in Washington and Seattle, we have now added a spectrum sensing node in central Brussels.

At the heart of this measurement station is a RFeye[1] device (manufactured by CRFS), which constantly monitors the power of transmitted signals across a frequency range of 30 MHz to 6 GHz, completing a scan every 3 seconds. Once the raw readings for power spectral density[2] are collected, they are uploaded to the Microsoft Spectrum Observatory (cloud-based tool) and made available for interactive viewing. A backend system has a number of pre-process stages to calculate the minimum, maximum, and average values and aggregate the data for each of the available time ranges (hourly, daily, weekly, and monthly).

---

[1] http://www.crfs.com/technology/rfeye/
[2] http://en.wikipedia.org/wiki/Spectral_density

To demonstrate the tool and demonstrate the usefulness of the data currently being collected by the Brussels measurement station, let's take a look at the UHF frequency band between 400 MHz and 1 GHz.

1) Open the Spectrum Observatory website: http://spectrum-observatory.cloudapp.net/
2) Click on the "Charts" tab at the top

This brings us to the first of four steps in a chart "wizard", which lets us choose which data to view. Currently the system contains data collected from four measurement stations located in Redmond, WA, Seattle, WA, Washington D.C, and the one located here in Brussels.

3) Select "Brussels" from the drop-down and click "Next"

The system currently supports three different visualizations of the data: spectrogram, occupancy chart, and power density chart. The spectrogram is a time-versus-frequency waterfall plot which allows us to easily see how the measured values change over time. The other two plots – occupancy and power density – provide views of the data aggregated over different time periods. In the case of occupancy, an estimate of the percentage of time that a particular frequency band is being used is shown. The power density chart displays an interactive view of the data that was measured, so let's start there.

4) Click on "Power Density"

Next we need to specify which frequency range we are interested in viewing. A number of predefined options are provided in the "Known Frequency Bands" tab as a convenient starting point, but since we specifically know the range of frequencies we wish to investigate, we can type in 400 and 1000 for the start and stop frequencies.

5) Click on the "Custom Frequency Range" tab
6) Enter "400" into the "Start Frequency" textbox
7) Enter "1000" into the "Stop Frequency" textbox
8) Click "Next"

In this next screen, we can select the range of time we are interested in viewing. The collected data is processed into useful time ranges – hours, days, weeks, and months – to make them easier to find. The Brussels station has actually been collecting data for a little more than a month now – since December 21, 2012 – so let's take a look at the current month, January 2013.

9) Select "Month" from the "Time Scale" dropdown
10) Leave the default selection of "Jan" for the "Month of" dropdown
11) Lave the default selection of "2013" for the year dropdown
12) Click "Next"

This screen allows us to confirm that this is the chart we are interested in viewing, so…

13) Click "Next"

This sends our request to the server where the data is aggregated and returned along with a new interactive chart so that we can explore the results.

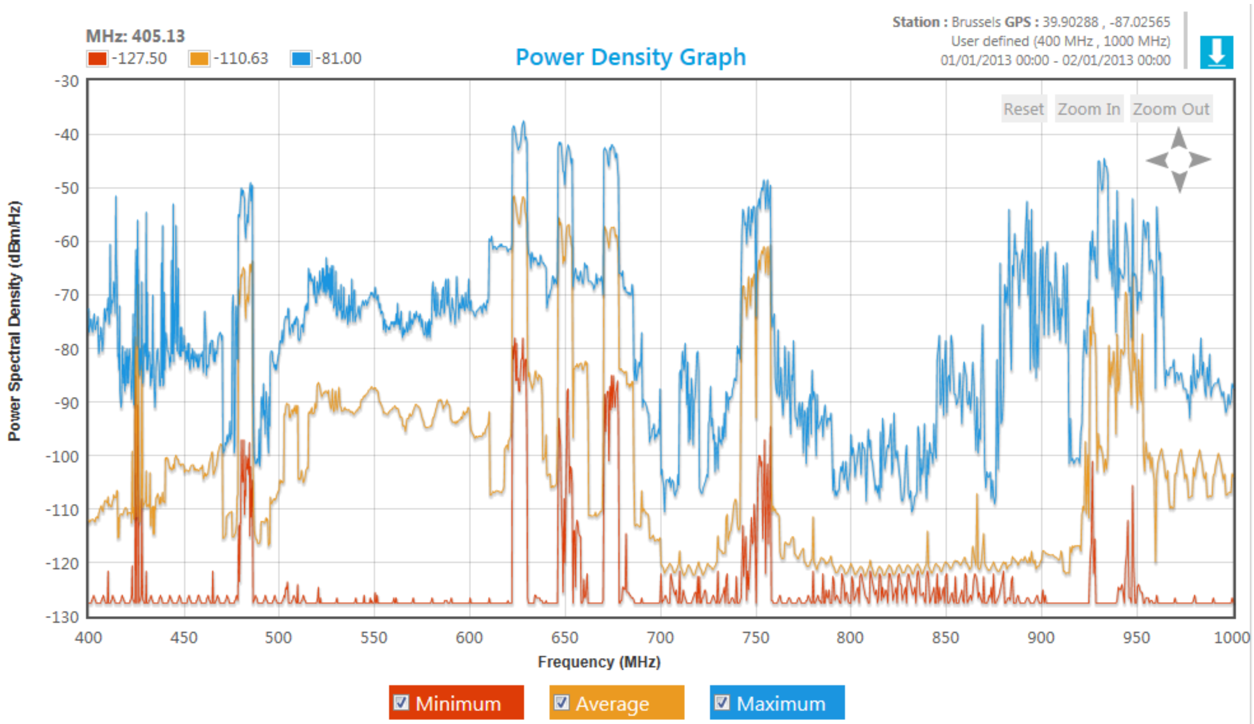

Figure 1 – Power Spectral Density as a function of frequency for a portion of the UHF band

This chart contains three traces – one each for the minimum, maximum, and average of the measured data.

As you can see, the Power Density Graph is interactive – we see a cursor move and show us the values in the upper left corner as the mouse is hovering (or tapped[3], for Microsoft Surface users). The chart can be zoomed with the mouse wheel and panned by dragging. For touch users, pinch-to-zoom works, as does touch-dragging for panning.

14) Demonstrate hovering, zooming, and panning
15) Click the Reset button in the top right corner of the chart

So what are we looking at here? Higher values of Power Spectral Density (towards the top of the chart) represent stronger transmissions (more power) at those frequencies. In this scale, -40 is fairly strong, while -120 is extremely faint (down in the noise floor). The y-axis is a logarithmic scale, so every change of 10dB that the signal goes down on the axis represents a reduction in power by 10x, so faint signals at -120 dBm are 100 million times[4] smaller.

---

[3] The Spectrum Observatory site is touch-enabled, but only on Internet Explorer 10.

[4] $10^{(-40+120)/10} = 10^{8}$ = 100 million

There are clearly some strong signals around 482, 625, 650, 672, 752, and 755 MHz. If we hover the cursor at the left edge, and then at the right edge, we can see that these "plateaus" are 8 MHz wide:

16) Hover at the right edge of the highest signal plateau – 630 MHz
17) Hover at the left edge of the highest signal plateau – 622 MHz

We can also see that not only is the maximum value quite high (-40), but so are the average and minimum values. This is indicative of signals that are consistently high across the aggregated time range, in this case a month. Can anyone take a guess what these high-power, always-on signals are? Yes, these are broadcast television, COFDM[5] DVB-T[6] signals to be specific. In fact, according to the center frequency of 626 MHz, this particular one is channel 40.

Just from looking at the chart, it seems that there are other interesting areas. The 880-915 MHz band has a slightly different signature than the TV stations – in this case, the maximum value is very high, while the average and minimum values are both very low. This seems to indicate that transmissions in this frequency band are very strong at times, but are not captured by the RFeye often enough to affect the average values, most likely from sporadic two-way communication. The BIPT frequency allocation chart[7] indicates that the 880-915 MHz band is allocated for GSM[8] and IMT-2000[9] - that is, 3G.

What about the low spots to the left (between 800 and 850 MHz)? Are these bands "unoccupied"? Let's take a deeper look.

The Microsoft Spectrum Observatory also provides another view of the power spectral density data where it calculates the percentage of time that a frequency band is occupied. Let's take a look.

18) Click on the "Create Chart" tab
19) Leave "Brussels" selected and click "Next"
20) Click on the "Occupancy" chart type
21) Leave the custom frequency range of 400-1000 and click "Next"
22) Leave "Month" for the "Time Scale" dropdown
23) Leave "Jan" and "2013" for the month selection, then click "Next" and "Next" again

The occupancy chart displays a two-dimensional line chart with frequency on the x-axis and occupancy percentage on the y-axis.

---

[5] http://en.wikipedia.org/wiki/Orthogonal_frequency-division_multiplexing
[6] http://en.wikipedia.org/wiki/DVB-T
[7] http://bipt.be/en/217/ShowContent/1057/Table/Table.aspx
[8] http://en.wikipedia.org/wiki/GSM
[9] http://en.wikipedia.org/wiki/3G

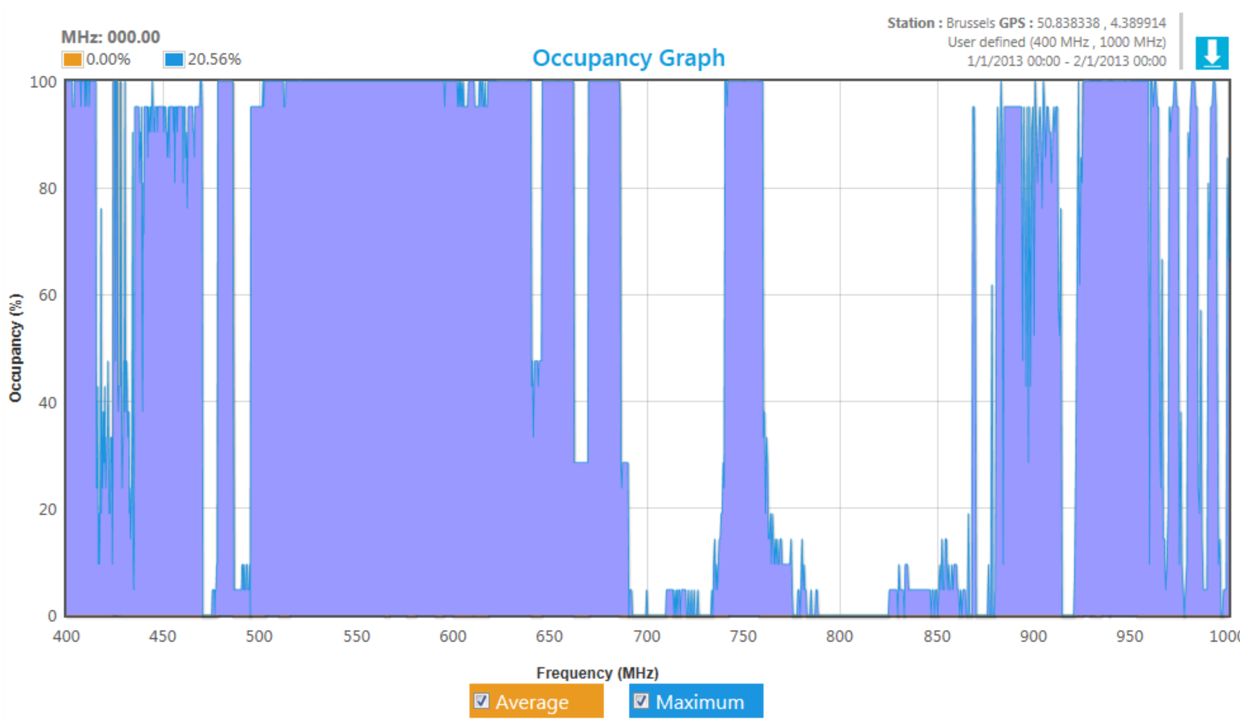

Figure 2 – Occupancy Chart of a portion of the UHF band

The occupancy percentage is calculated as the percentage of data points (average or maximum) over the specified time range that is greater than a -90 dBm/Hz threshold. So, for example, if a measurement at 203.45 MHz finds a power spectral density value of -80 dBm/Hz for the first 30 minutes of the hour and a value of -100 dBm/Hz for the last 30 minutes of the hour, then the occupancy percentage would be 50% since half of the values in the time range (1 hour) are above -90 dBm/Hz.

In this chart we can see that there are two large gaps – 690 – 740 MHz and 760 -860 MHz. What this chart is saying is that, over the past month, it is estimated that the signal in these ranges has risen above -90 dBm/Hz less than five percent of the time.

That being said, there are quite a few caveats - the associated occupancy percentage is an extremely course heuristic, for many reasons:

- Transmissions that are very brief pulses that are not sustained for long periods of time may never be measured by the RFeye, since it only scans each frequency about every 3 seconds.
- Guard bands are not accounted for. Some parts of the spectrum between allocated frequency bands are specifically allocated as "guard bands" to allow space between adjacent bands where there is intentionally no signal because the allocated bands on either side may have some amount of spillover.
- Transmitter power and receiver sensitivity vary greatly. Some applications intentionally operate at low signal levels and these will not be detected by the monitoring equipment. As an example, low-power indoor Wi-Fi signals are not realistically captured.

- Bands allocated for passive listening applications (such as radio astronomy) which do not transmit any information also will not be displayed on the occupancy chart.
- The measurement system is not calibrated. For example, the antenna gain or cable loss as a function of frequency is not taken into consideration. The measurement values are intended to be relative, not absolute.

Inferences about whether or not a particular frequency is actually occupied should be made only after careful study.

In conclusion, some parts of the spectrum are clearly occupied, while some are not. The Microsoft Spectrum Observatory provides an excellent tool for viewing real data and providing a starting point for further detailed investigation. We hope that this tool will provide a useful starting point for an ongoing conversation about innovation in the frequency allocation space to support new technologies like White Spaces and dynamic spectrum access.

2. Frequently Asked Questions

This section contains answers to a number of questions that may arise when using the Spectrum Observatory for the first time. They are arranged them into three general categories: data collection, data storage/processing, and data interpretation.

2.1. Data Collection

2.1.1 How much and how often does the RFeye sample data?

The RFeye device sweeps the entire frequency spectrum between 30 MHz and 6000 MHz every 3 seconds, though not all at once. Each sweep is comprised of around 20 smaller (contiguous) sweeps.

Each device produces about 60MB of raw data every hour, or 1.4GB per day.

2.1.2 Can anyone contribute data?

At the moment, no, but this is certainly something we would like to open up as a possibility in the future. If you are interested, please talk to us.

2.2. Data Storage and Processing

2.2.1 How much data is currently being stored?

We currently have about 3TB of raw data and about 1.4TB of processed data (hourly, daily, weekly, monthly). All raw data is stored in a local backup, while all of the processed data is stored in Azure blob storage.

2.2.2 How long does it take for all of the data to be processed?

Whenever a large backlog of data is bulk imported, it takes us about 12 hours using 40 Windows Azure virtual machines (VMs) to fully complete the processing of the last 6 months of data. However, if the data is loaded incrementally each day it only takes a few minutes.

2.3. Data Presentation

2.3.1 How far back in time is there data available for exploration?

For the Redmond, Seattle, and Washington DC measurement stations, we currently have data back to August 1, 2012. We have more raw data that needs to be uploaded and processed, and will be slowly adding it to the archive as we have time.

The Brussels station was installed more recently on December 21, 2012, so it has less data available.

2.3.2 You mentioned that the Spectrum Observatory site is touch-enabled, but why doesn't the Spectrum Observatory site work on my iPad?

The Spectrum Observatory is currently touch-enabled for Internet Explorer 10, so devices like the Microsoft Surface or other touch-enabled ultrabooks running Windows 8 will work, while other products which do not use Internet Explorer will not. That's not to say that other browsers may not be supported in the future – it was simply a technical limitation due to limited time and developer resources.

2.3.3 How were the Known Frequency Bands chosen?

We tried to group large sections of the spectrum into a few (< 20) groups. These are only intended as initial "jumping points" so that the user can see the data quickly. Then they can fine-tune what is interesting and create a custom frequency range that suits their needs.

2.3.4 How is occupancy calculated?

The occupancy percentage is calculated as the percentage of data points (average or maximum) over the specified time range that is greater than a -90 dBm/Hz threshold. So, for example, if a measurement at 203.45 MHz finds a power spectral density value of -80 dBm/Hz for the first 30 minutes of the hour and a value of -100 dBm/Hz for the last 30 minutes of the hour, then the occupancy percentage would be 50% since half of the values in the time range (1 hour) are above -90 dBm/Hz.

That being said, there are quite a few caveats - the associated occupancy percentage is an extremely course heuristic, for many reasons:

- Transmissions that are very brief pulses that are not sustained for long periods of time may never be measured by the RFeye, since it only scans each frequency about every 3 seconds.
- Guard bands are not accounted for. Some parts of the spectrum between allocated frequency bands are specifically allocated as "guard bands" to allow space between adjacent bands where there is intentionally no signal because the allocated bands on either side may have some amount of spillover.
- Transmitter power and receiver sensitivity vary greatly. Some applications intentionally operate at low signal levels and these will not be detected by the monitoring equipment. As an example, low-power indoor Wi-Fi signals are not realistically captured.
- Bands allocated for passive, receive-only applications (such as radio astronomy) which do not transmit any information also will not be displayed on the occupancy chart.
- The measurement system is not calibrated. For example, the antenna gain or cable loss as a function of frequency is not taken into consideration. The measurement values are intended to be relative, not absolute.

2.3.5 The numbers displayed in the Power Density Graph look too large/small. By my calculations, they should be X. Why is there a discrepancy?

The numbers presented represent the power spectral density values (in dBm/Hz) that are given to us by the RFeye device. They are intended to be relative values, not absolute, since there is no calibration for antenna gain (there are 2 different ones for higher/lower frequencies) and cable loss.

2.3.6 Can anyone download the data?

Yes. All of the chart options provide the ability to freely download the aggregated data in CSV format if additional analysis is required.

Currently the raw data (direct from the RFeye) is not available for download through the website, but if you’re interested in this, please talk to us.

2.3.7 What is the "Sign In" button for?

One of the more recent features that we've added to the Spectrum Observatory is the ability to save the view of a chart if you'd like to be able to look at it again later. We allow users to register a (free) account so that they can save their charts.

We provide a "management portal" for those users that administer the Spectrum Observatory measurement stations that perform the raw data collection. This allows station administrators to monitor the upload status, see statistics about the total amount of data collected, and request permissions.

3. Technical Deep Dive

The rest of this document is intended to be a completely describe all of the internal workings of the Microsoft Spectrum Observatory. Special notes have been added to indicate possible implications on the public-facing view of the portal.

3.1. Motivation

The electromagnetic frequency spectrum is a limited resource and its use is very tightly regulated, typically by governmental entities (for example, the Federal Communications Commission[10] in the United States, and the Belgian Institute for Postal Services and Telecommunications in Belgium[11]). It is argued that current spectrum allocations are not optimal, leaving many frequency bands unused for outdated technologies and not providing enough bandwidth to support emerging technologies.

To manage spectrum more effectively, and increase usage efficiency, regulators want to know how the different bands are being used and which are being left vacant. The Spectrum Observatory project aims to provide a common location for spectrum usage data to be collected, analyzed, and presented to those interested in understanding wireless spectrum utilization.

This document will describe the high-level architecture for the Microsoft Spectrum Observatory project which, when fully complete, will automatically collect frequency usage data, perform useful analyses over the data, and present the data in useful ways.

3.2. Architecture

The Spectrum Observatory project consists of three main components:

- Data collection
- Data analysis
- Data presentation

Figure 3 shows these components.

[10] http://www.fcc.gov/
[11] http://www.bipt.be/en/1/Home/Home/Home.aspx

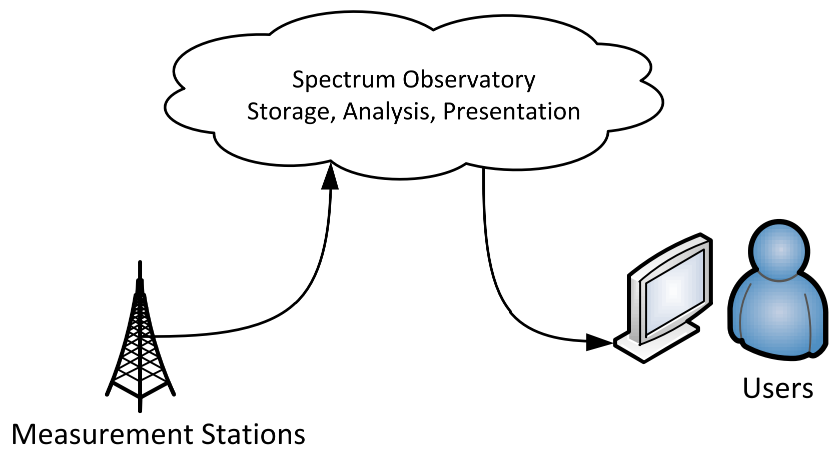

Figure 3 – High-level architecture diagram

3.2.1 Data Collection

The system is capable of collecting and storing measured frequency data from a number of measurement stations. These stations are geographically distributed throughout the world. Typically a measurement station consists of a computer-controlled receiver attached to a wideband antenna. The data collection process is discussed in more detail in Section 3.3.

3.2.2 Data Storage and Analysis

The system is capable of storing the data from the measurement devices in a central location so that further analysis is possible. Currently the analyses include: normalization to a fixed time/frequency grid and aggregation over various time ranges (hours, days, weeks, months). The storage formats as well as the individual analyses are discussed in more detail in Section **Error! Reference source not found.**.

3.2.3 Data Presentation

The analyzed data is presented to the user in ways that allow policymakers to view and interact with the collected data. The modes of presentation make it possible to view data across different times, different frequency bands, and different measurement locations. The main goal of the Microsoft Spectrum Observatory website is to provide insight and influence policy decisions related to frequency spectrum usage. More details about data presentation are presented in Section **Error! Reference source not found.**.

Now that we've covered the three main subcomponents of the Microsoft Spectrum Observatory at a very high level, let's rewind to the beginning and take a much deeper dive into each of the three areas.

3.3. Data Collection and Import

This section explains in detail the process by which data is collected, starting from the antenna.

3.3.1 Measurement Stations

Currently the Spectrum Observatory system only supports the collection of data from a single type of device, though more types of devices could be supported in the future. The type of supported device is called an RFEye[12] and is produced by CRFS[13]. A photo of the device is shown in Figure 4.

Figure 4 – A mounted RFEye measurement device. Antennas can be seen to the left and right.

For the most part, all four measurement stations are identical (they each have an RFeye and a computer), but there are slight differences. The next section will list the details about each of the four stations, and point out differences where relevant.

3.3.1.1 Redmond, WA, USA

3.3.1.1.1 Physical Location

The antenna and RFeye are mounted on the northeast corner of the roof of Building 99. The GPS location is 47.64251, -122.1416. See Figure 5.

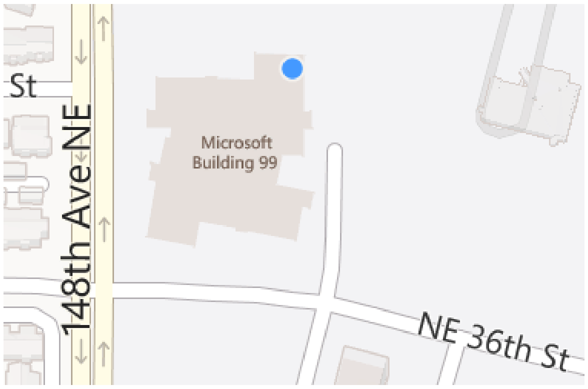

Figure 5 – Location of the Redmond measurement station on top of Building 99.

There is a large metal 15'-tall privacy screen mounted vertically all the way around the perimeter of the building (ostensibly to hide the heating/cooling units from view). A broadband antenna is mounted to

[12] http://www.crfs.com/technology/rfeye/
[13] http://www.crfs.com/

the top of this structure so that it protrudes above the surrounding metal wall. The RFeye enclosure is mounted on the pole just below the top of the privacy screen. There is also a plastic outdoor enclosure nearby which contains a power divider and a number of filters in-line between the antenna and the RFeye. A picture of the installation is shown in Figure 6, though this is an old photograph. The RFeye device has since been moved below the top of the wall, and the plastic enclosure cannot be seen in this picture.

Figure 6 – Mounting location of the Redmond RFeye

The server is mounted inside a small room at the top of the elevator. The power-over-Ethernet (PoE) injector device is also mounted next to the server.

3.3.1.1.2 History
This measurement station was originally installed in 2011.

3.3.1.1.3 Schematic
A schematic of the installation is shown in Figure 7.

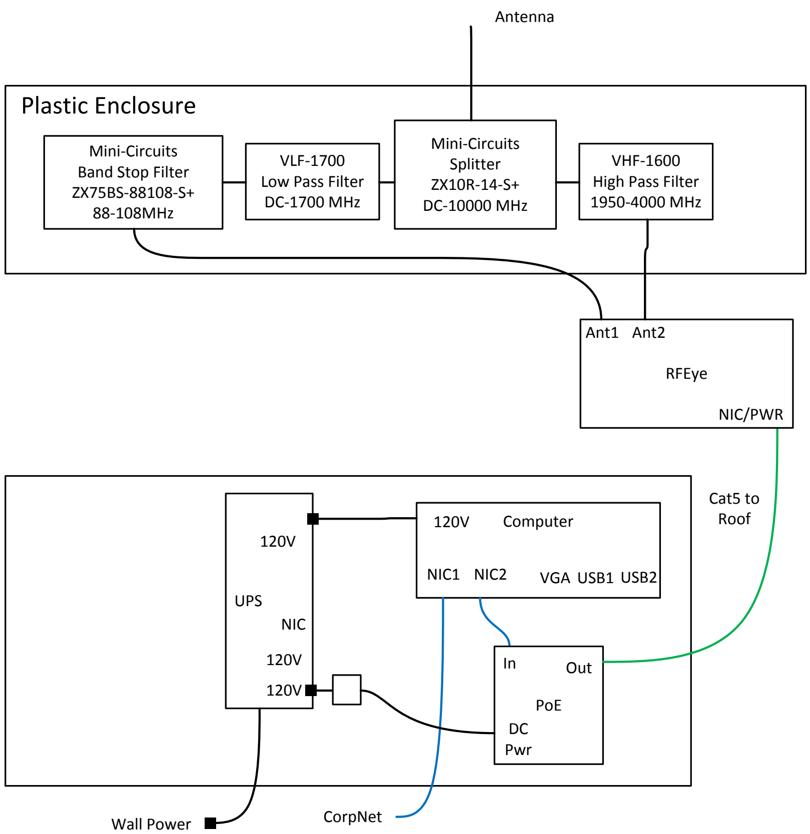

Figure 7 – Schematic of the Redmond measurement station

Note: The part numbers for the filters in the above schematic are approximate. I need to dig up my notes and update the diagram.

3.3.1.2 Seattle, WA, USA

3.3.1.2.1 Physical Location

The antenna and RFeye for the Seattle station are mounted on top of the building at 320 Westlake in Seattle.

The GPS location is 47.621571, -122.33808.

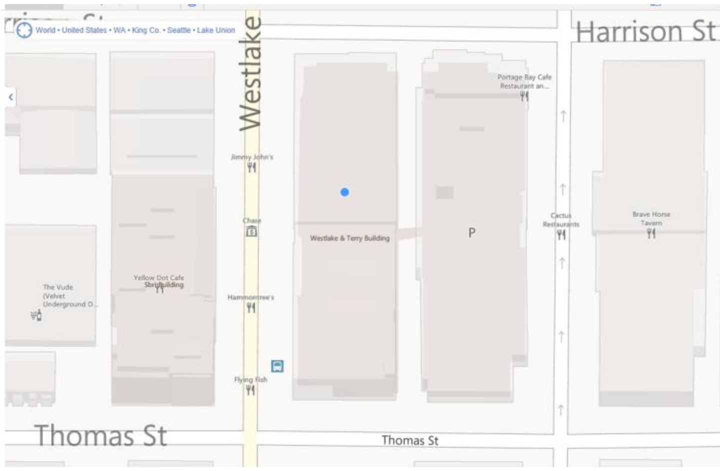

Figure 8 – Location of the Seattle measurement station at 320 Westlake

A picture of the antenna, RFeye, and plastic external enclosure is shown in Figure 9.

Figure 9 – Mounting location of the antenna for the Seattle measurement station

Note the black plastic enclosure directly beneath the RFeye which houses the power splitter and a few in-line filters. The power-over-Ethernet injector is mounted in a separate location in an enclosure towards the middle of the roof where the Ethernet cable enters the building. A picture of this enclosure is shown in **Error! Reference source not found.**.

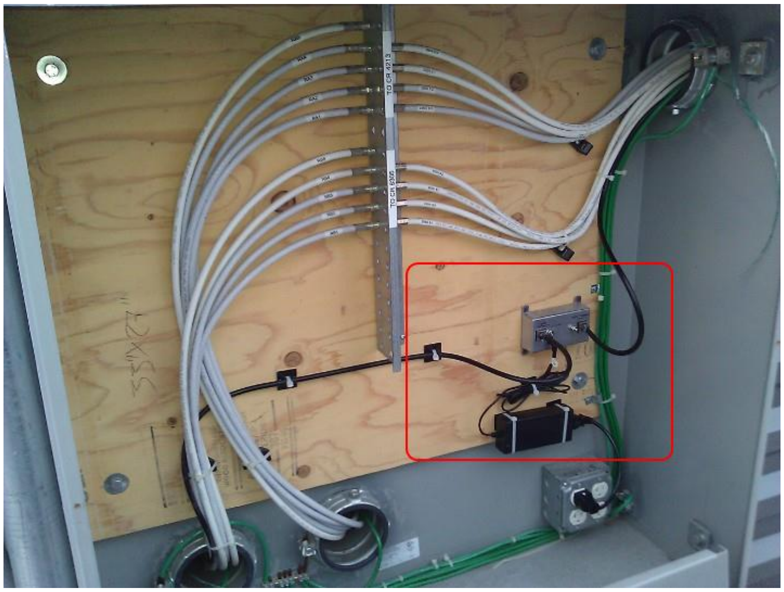

3.3.1.2.2 Schematic

A schematic of the Seattle measurement station installation is shown in Figure 10.

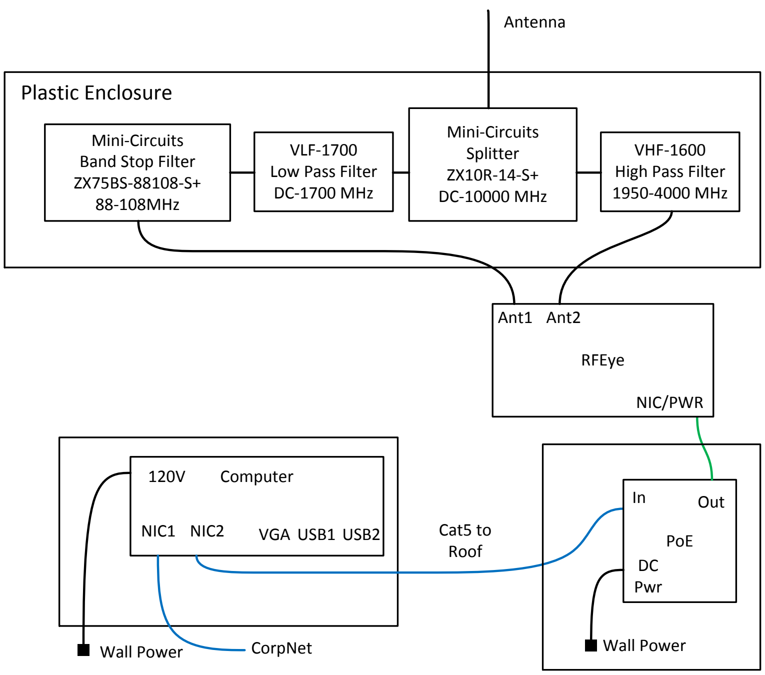

Figure 10 – Schematic of the Seattle measurement station

3.3.1.3 Washington DC, USA

3.3.1.3.1 Physical Location

The Washington DC measurement station is located on the top of a building that Microsoft leases at 901 K St NW, Washington D.C. The GPS location is 38.90288, -77.02565, as shown in Figure 11.

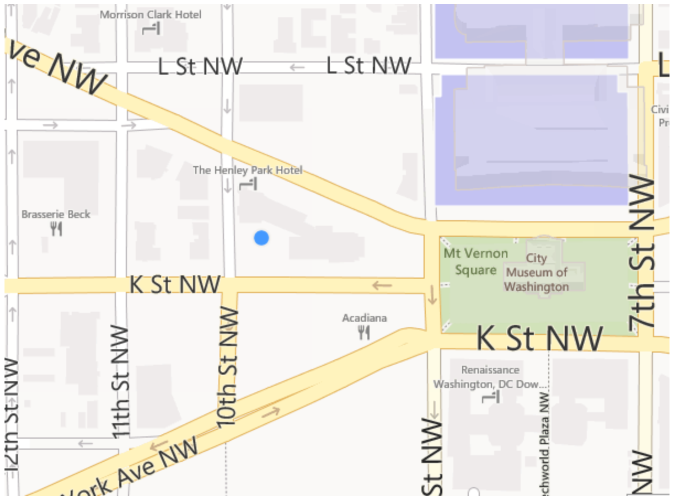

Figure 11 – Location (blue dot) of the Washington DC measurement station

The antenna and RFeye device are mounted on a support column on the roof. See Figure 12.

Figure 12 – Location of the RFeye at the Washington DC measurement station

3.3.1.4 Brussels, Belgium

3.3.1.4.1 Physical Location

The Brussels measurement station is located on the roof of the Microsoft Executive Briefing Center in Brussels. The GPS location is 50.838338, 4.389914. See Figure 13.

Figure 13 – GPS location (blue dot) of the Brussels measurement station

3.3.1.4.2 History

The device was installed on 2012-12-20 and began capturing data starting on 2012-12-21.

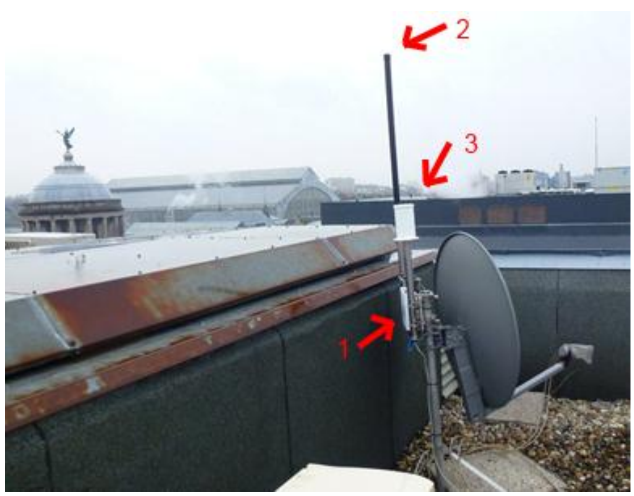

Figure 14 – Location of the RFeye (1) and antennas (2,3) at the Brussels measurement station

3.3.2 Generated Data Files

The generated binary files (.bin) are written using a proprietary data format specific to products manufactured by CRFS. The files consist of a binary header (typically < 200 bytes), followed by a stream of “data blocks”. Each data block has an integer field which specifies its “type”. “GPS timestamp data” and “power spectral density data” are the two most common types of data blocks. The data blocks containing power spectral density data have an integer field which specifies the “thread ID”, which corresponds to the section in the RFeye configuration file that specified a "scan". Each scan can specify the start frequency, stop frequency, and number of samples.

There are currently two versions of the RFeye binary file format – the second is an extension of the first and includes additional types of data blocks.

After the binary data files have been written to the measurement station server by the RFeye, this data is imported into the Spectrum Observatory so that it can be analyzed and presented to users.

The Microsoft Spectrum Observatory currently supports all major browsers: Internet Explorer, Chrome, Firefox and Safari, as well as Windows 8 devices using touch input with Internet Explorer 10 (e.g. the Surface RT or Samsung Series 7).

3.3.3 Data Visualisation

The measured spectrum data is essentially a four-dimensional data set that we want to visualize in useful ways. The data is structured such that there is a scalar quantity (power spectral density) that is measured over a f*requency* range at different *locations*, over *time.* One way to think about that is that there is a number (a power spectral density reading) at every point in the 3D space with time, frequency, and location axes, as shown in Figure 15.

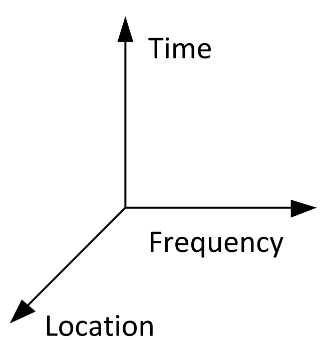

Figure 15 – The three dimensions of the measured spectrum data

Other scalar quantities can also be computed from this four-dimensional data set. For example, minimum, maximum, and average power spectral density values and occupancy percentages can be calculated from this source.

Because visualizing a four-dimensional data set is difficult for humans to grasp, it makes sense to allow users to choose to fix (or "pin") one or more of the independent variable axes. Each possibility will be explored in the next sections.

3.3.3.1 Fixed Location

If we fix only the location variable at a single measurement station, we are left with the visualization of power spectral density as a function of both time and frequency.

One way to visualize this would be with a spectrogram (sometimes called heatmap) – functionality that we currently have implemented. In this case, an image is created where the power spectral density value for a particular time (row) and frequency (column) is mapped to a color. An example is shown in Figure 16.

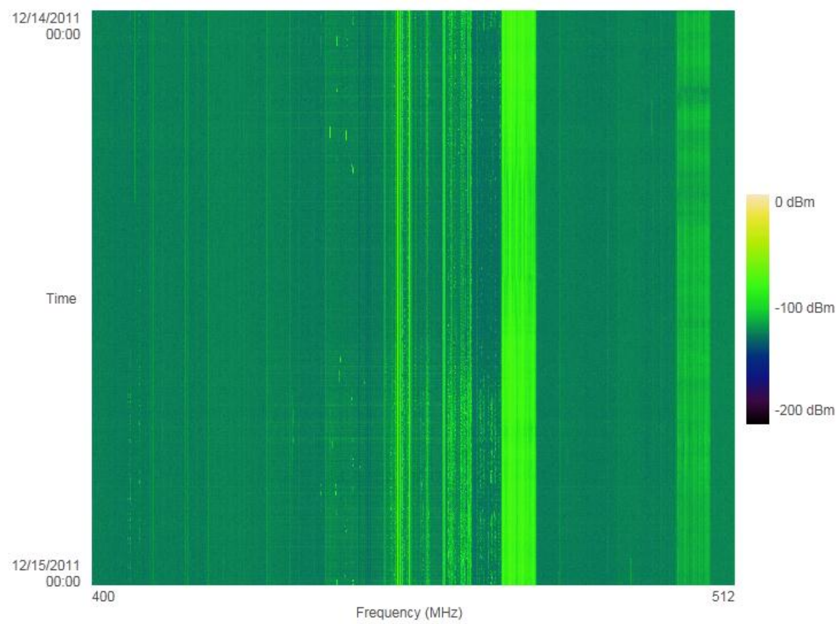

Figure 16 – An example spectrogram image.

Other ways to visualize a two-dimensional matrix of data exist, such as using 3-D surface plots – but are not implemented, at this stage.

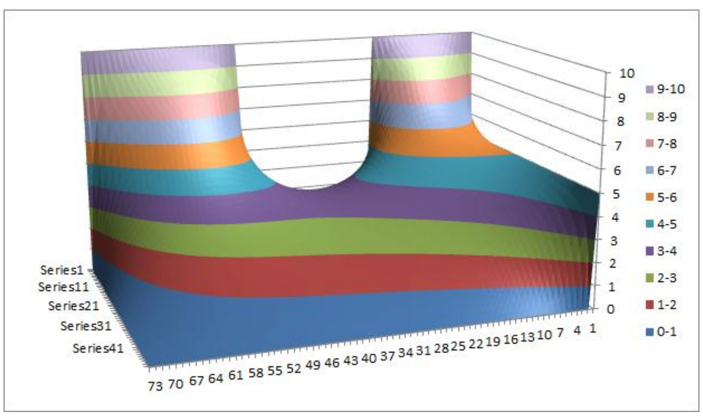

Figure 17 – An example of a 3D surface plot using Excel

3.3.3.2 Fixed Time

If we fix only the time variable at a specific instant, we are left with a visualization of power spectral density as a function of both frequency and location. This could be attempted in a few ways.

If the number of measurement stations is small (i.e. < 10), then it could be accomplished by using an X-Y chart of power spectral density as a function of frequency, with multiple series lines for each location. This type of view would be very useful and thus be a higher priority.

If the number of measurement stations is larger and information about the geographic "closeness" is important to be conveyed, a map could be used but would be much more difficult. It could be approximated using something like this:

- A map control is used to display the location information, but tilted at an angle to add perspective
- At each location, a vertical line is drawn to indicate an axis perpendicular to the plane containing the map.
- The PSD value is plotted as circles stacked on the line at z-axis values corresponding to the frequency value.
- The radius of the circle is proportional to the PSD value.

This type of view is likely of low priority due to the complexity of the implementation.

3.3.3.3 Fixed Frequency

If the frequency variable is fixed, this leaves a visualization of power spectral density as a function of both time and location. Similarly to the "Fixed Time" angle, this could be attempted in a few ways. Visualizations could be constructed in a manner similar to those described in the "Fixed Time" description, though with the frequency and time variable roles reversed.

3.3.3.4 Fixed Location and Time

If the location and time are fixed, this leaves a graph of power spectral density as a function of frequency.

This is already implemented by the Power Spectral Density graph page.

3.3.3.5 Fixed Location and Frequency

If the location and frequency are fixed, this leaves a graph of power spectral density as a function of time.

This is not currently implemented, but could be a very useful graph to show time-series information.

> *Note: There is a relationship between this graph and the "Fixed Location and Time" graph – they both correspond to rows and columns of data on the spectrogram ("Fixed Location"). One could easily conceive of a user interface where the spectrogram is shown and the user could easily choose a row (or column) and be presented with a graph of power spectral density against frequency (or time).*

3.3.3.6 Fixed Time and Frequency

If the time and frequency are fixed, this leaves a graph of power spectral density as a function of location.

This is not currently implemented but could be visualized easily using a map. For example, a circle with radius value corresponding to the power spectral density value could be plotted over each location.