Interference Likelihood Mapping with Case Studies

Paul Alves, Carmen Wong, Matthew Clampitt, Eric Davis and Eunju Kwak
NovAtel Inc.

BIOGRAPHY

Paul Alves received a Ph.D. from the Department of Geomatics Engineering at the University of Calgary in 2006. He has worked for Leica Geosystems, Hemisphere GNSS, and NovAtel focusing on positioning algorithms research and development. He is currently a principal research engineer in the Applied Research Team at NovAtel.

Carmen Wong is a Geomatics Engineer working at NovAtel Inc. She tests receiver firmware and has served as the verification prime in the Position Algorithms team for many software releases. She received her B.Sc in Geomatics Engineering with Biomedical Specialization from the University of Calgary in 2008.

Matthew Clampitt graduated in 2014 with a B.Sc in Geomatics Engineering from the University of Calgary and is now a developer in the Positioning Algorithms group at NovAtel Inc.

Eric Davis has an undergraduate degree from the University of Calgary, with majors in both astrophysics and physics. He also earned an MSc in physics at the University of Calgary. His thesis involved working with very low frequency (VLF) waves and the effect on signals travelling through the ionosphere and atmosphere. He joined NovAtel Inc. in 2016.

Eunju Kwak received her Ph.D. from the department of Geomatics Engineering, University of Calgary in 2013. Her research interests have been Photogrammetry and integration of other sensors. She is currently working at NovAtel Inc. as a Geomatics Engineer.

ABSTRACT

Interference is a growing concern among GNSS users, particularly in parts of the world where radio frequency (RF) transmission is not strictly regulated. Intentional interference and jamming is cheap and relatively easy to obtain in the form of Personal Privacy Devices (PPD). These devices can sometime cause unintended interference and jamming to important infrastructure like an airport (Mitch, et. al. 2013). This paper describes a method for creating an interference map using NovAtel OEM7 Interference Tool Kit (ITK). ITK can be used to measure the power of a received interferer. When data is collected for an area around a static and continuously operating interference source then it can be used to map out the interference over the affected area. This paper describes a method for mapping the interference and, using a model of power loss over distance, create a map of the likelihood of the position of the interferer. This paper includes simulated results and three case studies with live (real-data) interference sources from India, Canada, and Japan.

INTRODUCTION

NovAtel introduced the ITK as a new feature in 2016. This feature is capable of detecting and eliminating interference. For more information about the implementation details of the ITK see Gao and Kennedy (2016).

The ITK's interference detection provides a list of interference sources, which includes an estimate of the frequency, bandwidth, and power of the measured interference. It also provides the power levels across the entire frequency band of the front-end. Either of these can be used as measurements of the received interference power levels. When the power levels for a given frequency are combined from multiple locations, they can be used to estimate the power and location of the interference source. The received power levels can also be combined to estimate the interference power as a function of location. The performance degradation experienced by one receiver at a given interference level can be extrapolated to other receivers at the estimated interference levels.

INTERFERENCE DETECTION

ITK is a collection of few very useful tools. These include the ability to visualize the power received across the input frequencies (front-end) bands. This can be used to quickly and easily identify any irregularities in the spectrum. These irregularities could be caused by internal interference, which is interference between electrical components introduced through hardware integration or installation. It can also be caused by external interference, such has a PPD or nearby radio.

There is also a detection feature that will identify potential interference and provide a list of interference power, frequency, and bandwidth. This is to make it easier for integrators to automate responses to potential interference without the need to scan the spectrum themselves. Figure 1 shows the received signal power and interference detection threshold for the GPS L1 frequency band. In this case there is no interference detected.

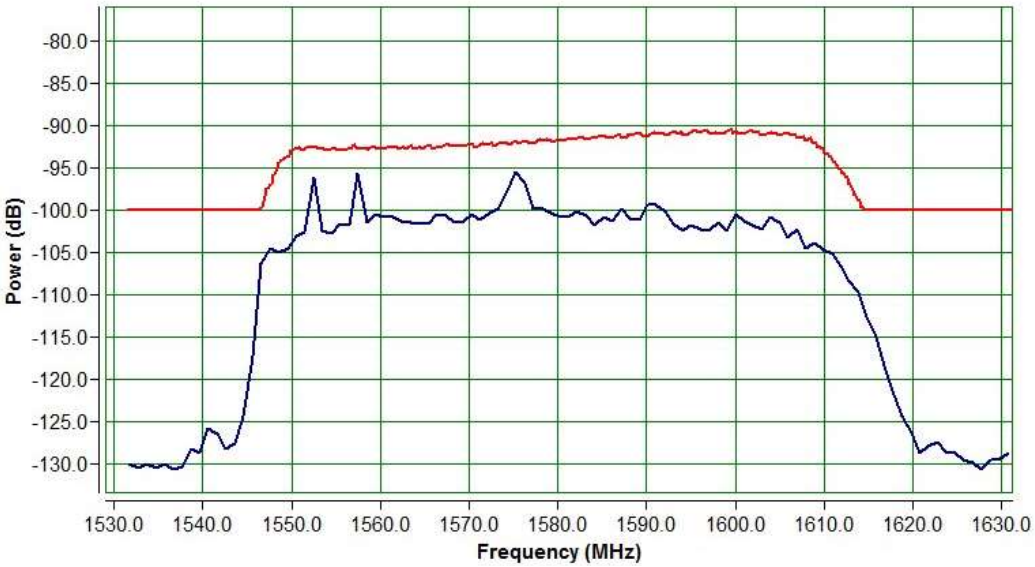

Figure 1: Received signal power (blue) and interference detection threshold (red) for L1

The detection threshold is adjustable, however, if it is set too high, it can cause interference to be undetected or if it is set too low, it can cause false detection. For this example, a fairly low value was chosen because we were willing to manually identify the interference source and ignore any false detection.

ITK also includes tools to mitigate interference to limit or eliminate the impact of the interference. This includes a High Dynamic Range (HDR) mode, which is effective in reducing the impact of interference. If this is not sufficient, then notch or low-pass filters can also be applied to completely cut out parts of the spectrum to neutralize the impact of interference or jamming.

FREE-SPACE LOSS

The mapping algorithm, which will be discussed later, requires a model of the power loss as a function of distance (d) to the transmitter. The attenuation experienced as an electromagnetic wave propagates is well known (Faruque 2015). As the wave spreads from the transmission source, the power is lost at the following rate:

$$L_p(dB) = 10\log\left[\left(\frac{4\pi d}{\lambda}\right)^2\right] \quad (1)$$

where, $L_p(dB)$ is the power lose in dB, d is the distance in m, and λ is the wavelength in m. This equation can be expanded into a function of frequency (f, in Hz) and distance (d, in m). Changing the units in this equation changes the constants.

$$L_p(dB) = 20\log(d) + 20\log(f) - 147.55 \quad (2)$$

For example, if the transmitter is broadcasting at 1.237 GHz then Equation (2) gives

$$L_p(dB) = 20\log(d) + 121.85 - 147.55$$
$$= 20\log(d) - 25.70 \quad (3)$$

This ideal power loss is significantly increased by physical obstructions that are common, such as vehicles, buildings, trees, the terrain type, etc. Different materials can have significantly different impacts on the power loss (Crowford 2002).

Khalifeh et al. (2015) uses a precomputed power map and map matching to do indoor positioning. This uses the expected received power to position a receiver. The same algorithm that is used to position the receiver could also be used to position the transmitter.

Figure 2 shows the received power as a function of distance that was observed for the Calgary test. There is a large variability in the power, which is likely due to the natural obstructions.

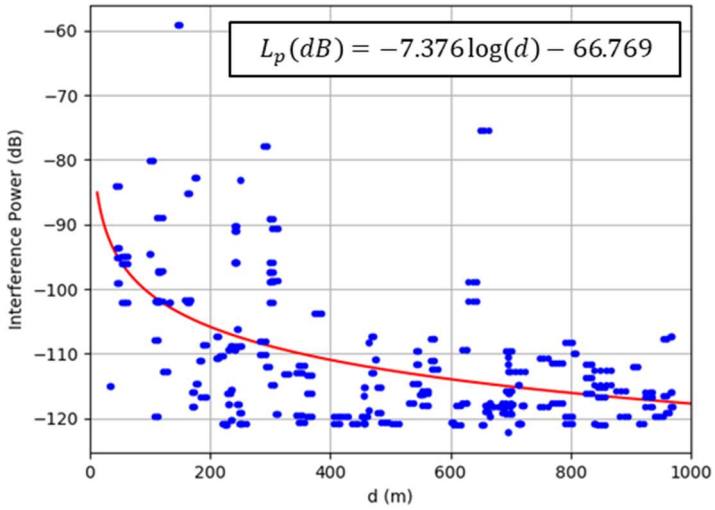

Figure 2: Received power as a function of distance from the transmitter

The equation for the line of best fit of this data is significantly different than Equation (3). This is likely due to the obstructions and limited number of data points. Due to problems with inaccuracies with this data fit, any further power calculations will use Equation (2).

MAPPING THE INTERFERENCE IMPACT

Using a single observation of the received interference power, a profile of the transmit power as a function of location can be created using the power decay curve shown in Figure 2. If we assume that the transmitter is at a given position and use the decay curve through the observed power, then we can estimate the transmit power at that location. When we do this for multiple locations, then a power profile is created. This process is shown in Figure 3. When these dots are connected continuously, then we get the power profile shown in Figure 4.

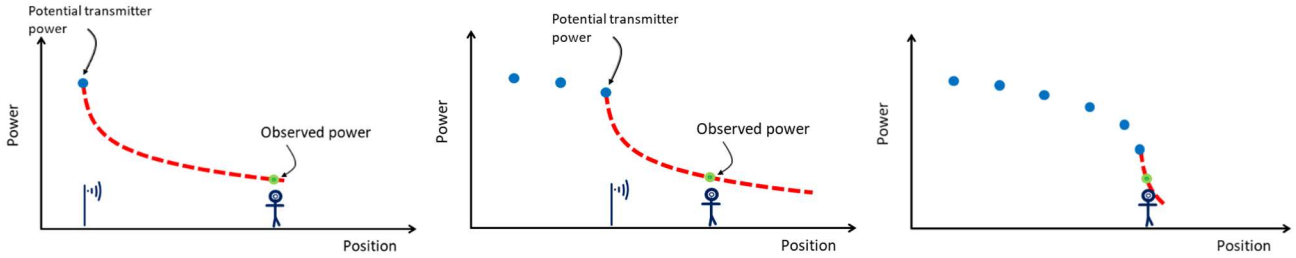

Figure 3: Profile of the transmit power as a function of location to the single observed power

Figure 4 shows an example of the estimated power of a transmitter as a function of location using the free-space loss shown in Equation (2). The green dot shows the location and power of the received interference. This power could have been received by a lower power transmitter that is relatively close to the receiving antenna or could be a stronger transmitter that is farther away. A single transmitter at any location could be responsible for the received power depending on the power of the transmitter.

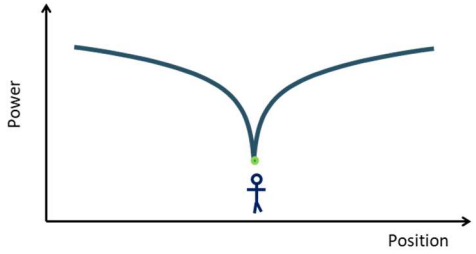

Figure 4: Estimated power of the transmitter (as a function of distance) based on the measured, observed power

When additional measurement points are added at different locations the estimated power of the transmitter for each individual observation can be combined together. Figure 5 shows an example of two observed power levels at two locations. The transmit power that would be required to receive the observed power is shown by the blue solid and green dotted lines. At some locations it will not be possible for a single transmitter to be received at both locations at their measured transmit power. For example, a single transmitter on the right side of the plot cannot be responsible for the received power for both locations. This allows the model to reduce the estimate power level in some locations because there could not be a transmitter with higher power without influencing the other measurement point. Figure 6 shows as example the resulting estimate power of the transmitter using this method.

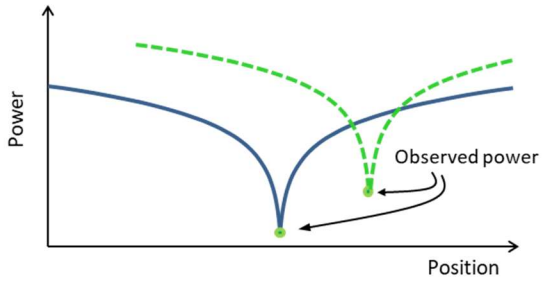

Figure 5: Estimated power of the transmitter (as a function of distance) for two separate measurements of received power at different locations

There are many methods that could be used to estimate the transmit power as a function of location. A simple method is to take the minimum estimated power. This model is very simple and quick to compute. Instead of calculating the minimum, we could instead estimate the transmit power by using a weighted average of the propagated power from each of the observation points. The weighting factor could be a function of distance so that nearby power observations are favored more heavily than observations that are farther away.

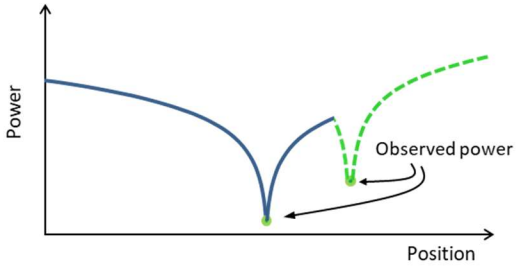

Figure 6: Estimated power of the transmitter (as a function of distance) for two separate measurements of received power at different locations using the minimum power method

The treatment of zero-observation points is also very helpful in this analysis. A zero-observation point is a location where no interference was detected. These measurements are just as useful as detected interference. If looking for the source of the interference then these points give vital information about the transmitter's direction. If a measurement to the left of the green dots (in Figure 6) shows that there is no interference detected, then we can infer that the interference source must be to the right.

Creating this type of power profile can be useful for pre-analysis. If we assume that none of the measurement locations can observe the interference, then the received interference must be equal to or less than the noise floor. If we assume that the received interference is at the noise floor then we can use this profile map to identify the power of any hidden, undetectable transmitters in a region. There may be interference that is broadcasting under the noise floor that would be undetectable at that power and distance. For example, if we want to monitor an area for interference around critical infrastructure, an airport for example, then we can deploy a network of ITK receivers. If no interference is detected, it is still possible for interference to be present if the power level of the transmitter is low enough that it does not reach any of the receivers above the noise floor. This analysis can be used to estimate the minimum detectable interference across the area. If one wanted to ensure that a minimum interference power would be immediately detected then this could be used to determine the receiver network spacing and locations to ensure the minimum detectable interference was detected.

Figure 7 shows an example of measurement points from the India case study. If no interference is detected anywhere at the measurement points then, this figure shows the estimated power of a potentially undetectable interference source. Lighter colors indicate a higher power of the possible undetectable interferer. Notice how it is possible to miss a weak interferer that is close or a high-powered interference source that is farther away. This also illustrates how much information we can gather from zero-observation points where interference could not be detected.

If one wanted to monitor a region to detect interference of a certain level, this method could be used to determine the path or spacing of receivers. With some history added into the model so that the uncertainty increased over time, a single receiver or a fleet of receivers could plan out their routes to monitor for interference.

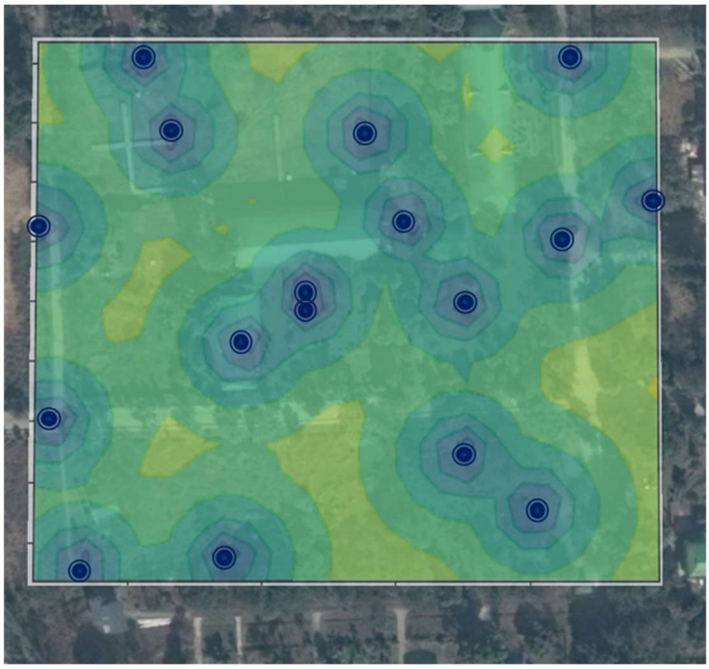

Figure 7: Locations and power of possibly hidden interference sources that would be undetectable by observation points, shown as blue dots (Map data: Google, DigitalGlobe)

The estimated interference source power can be used to determine the impact of the interference and give an estimate of the location of the interferer. A single static interferer will be assumed when estimating the location of the interferer using a goodness-of-fit model. A grid is created over the interference area. For each point in the grid, the attenuation (power loss) model is used to calculate the residual between the minimum transmit power and all power measurement points. If the residuals are low for all the observed power locations, then this is the most likely location of the interference transmitter.

Figure 8 shows an example of this goodness-of-fit test. The red dot shows the location of a potential transmitter location under test. Using the distance attenuation model, the predicted received power for each of the measurement points is calculated. The difference between the expected received power and the actual received power is an indication that this is not the correct transmitter location. The root mean squared (RMS) of the fit error for all the observed points gives a likelihood that the transmitter is at this location.

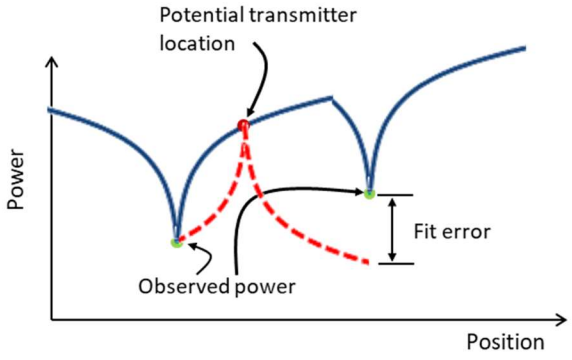

Figure 8: Example of the goodness of fit for potential transmitter location and power

SIMULATED RESULTS

Using the goodness-of-fit method, we can generate reasonable visualizations of the interference effect. Figure 9 to Figure 12 show maps produced from simulated interference. In these cases, the expected power attenuation model matches perfectly with the data because it is a simulation.

Figure 9 shows the result from an interference source on the east side. The yellow line shows a "roller-coaster" plot of the interference power. The height of the line shows the relative received power. Notice that it increases as we approach the interference and decreases as the path moves away from the interference. A combination of the roller-coaster plot and the map give a quick visualization of the impact and location of the interference. There is a slight ambiguity between the east and west side of the road because the transmitter is close to the road. The goodness of fit works very well in this case to identify the location of the interference source.

Figure 9: Interference map from a simulation where the interference is on the east side (Map data: Google)

Figure 10 shows the simulated interference on the west side. Due to the geometry of the roads, the model fits relatively well for a transmitter on both the north and south sides of the east-west road.

Figure 10: Interference map from a simulation where the interference is on the west side (Map data: Google)

Figure 11 shows the interference map when the simulated interference is in the north. This is the best fit of any of the simulations with very little ambiguity between either side of the road.

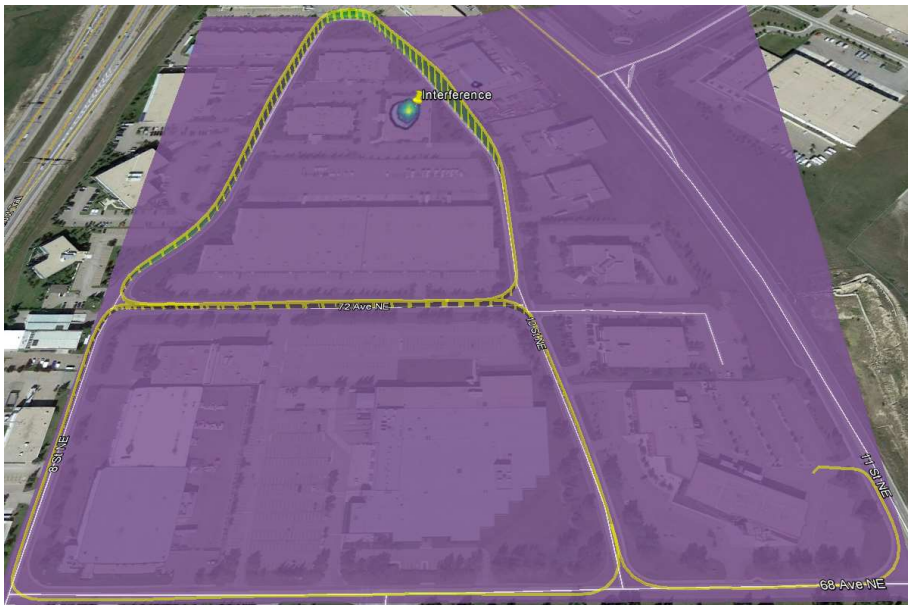

Figure 11: Interference map from a simulation where the interference is on the north side (Map data: Google)

Figure 12 shows a case where two interference sources are simulated. In this case the model breaks down because it assumes that there is only a single interference source. The model clearly has difficulties determining the location of the interference. Even with the model accuracy issues, this could still be used as a visualization of the interference that would be easier to interpret than looking at numbers from a table.

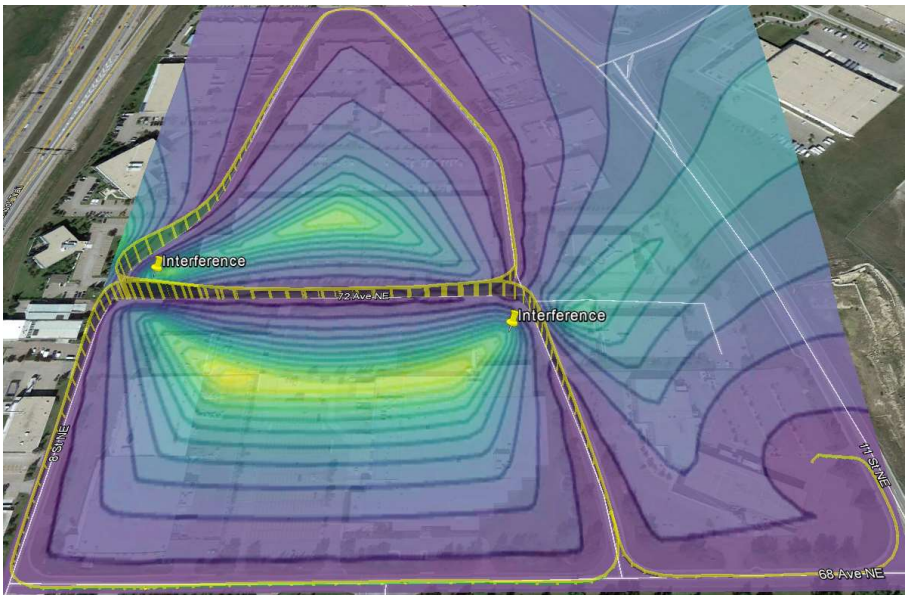

Figure 12: Interference map from a simulation with 2 interference sources (Map data: Google)

INDIA DATASET

This dataset was the initial motivation for this work. A customer reported intermittent tracking problems with a newly installed receiver. The receiver would stop tracking for a few hours every evening. Customer service visited the site to investigate. Due to the intermittent nature of the problem, interference was suspected. A OEM 729 receiver was walked around the affected antenna in an attempt to find the source of the interference and also to prove to the customer that interference was in fact the cause of the tracking problems. Figure 13 shows the measurements that were collected. The numbers shown are the received interference power at each location. It is possible to approximate the location of the interference and the impacted area by looking closely at the measurements but it takes some close examination and interpretation.

Figure 13: Received interference power measured when searching for interference in India

The source of the interference was identified using this approach. It was found to be a weather station (Figure 14) that performs a nightly upload of all the data that it collected throughout the day. This weather station broadcasts at 1580MHz, which was jamming L1. The customer was able to move the interfering antenna to another site. They could have alternatively used the interference toolkit to apply a notch filter, which would have also mitigated the interference's impact but it is better to remove the source of the interference if possible.

Figure 14: Photo of the weather station causing the interference in the India case study

Using the data points collected, an interference map can be generated using the method described. This map is shown in Figure 15. The lighter color indicates a higher likelihood that the interference transmitter is at that location. The location of the transmitter is also shown in the figure. The likelihood map is very close to the actual location of the transmitter. It gives a very quick and easy to interpret visualization as opposed to the individual measurement points.

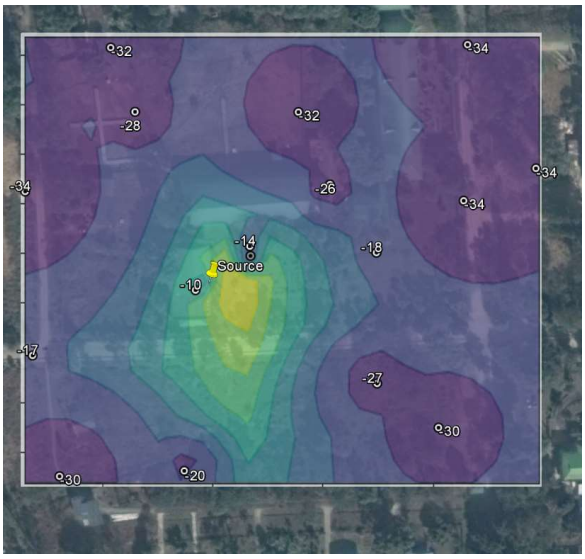

Figure 15: Interference map for the India case study (Map data: Google, DigitalGlobe)

CALGARY DATASET

We were made aware of a potential unintentional interference device and took it to Cross Iron Mills mall, north of Calgary, Canada, to investigate. Figure 16 shows a map of the area. The path shown in blue was driven to characterize the interference. Data was collected using an OEM729 receiver with the ITK feature. There are two buildings nearby the interference transmitter. A smaller building is to the north and a large building is south of the interference source. These buildings block and shield the receiver from the interference when it is between the interference and the receiver.

The interference device was a video transmitter intended for transmitting video from a drone to a monitor. Figure 17 shows the transmitter. This transmitter broadcasts at 1.2GHz at 800mW. It was purchased online with no warnings as the potential impacts it may have on other systems or devices. As recreational drones (and electronics) become more popular, it is possible that more unintentional jammers and interference sources like this one become commonplace. There is no continuous monitoring and enforcement for short range and short duration unintentional jammers like this.

Although many commercial grade receivers, like the ones common in cell phone and GPS watches, were unaffected because they currently only operate at L1, the box that the device came in also appears to also have a 1.5GHz model that is capable of broadcasting at 2W. At 2W at 1.5GHz, GPS L1 would be significantly jammed. This emphasizes the need for interference detection and mitigation. There is nothing stopping recreational hobbyists from accidently jamming a significant number of users and services.

Figure 18 shows the roller-coaster plot of the interference that was observed during the test. The height of the yellow bars indicates the received power for the L2 interference. The power is generally higher closer to the interference source. The power decreases generally as a function of distance however there is a lot of deviation. Physical obstructions also cause significant decreases in received power. For example, on the north end of the small building, shown on the right side of Figure 18, the observed interference power drops to almost zero despite being relatively close to the interference source. The large variations in power throughout the southern loop may be due to partial obstructions from parked cars or outcrops of the building. These physical obstructions cause larger decreases in received power than simply moving the antennas away from each other.

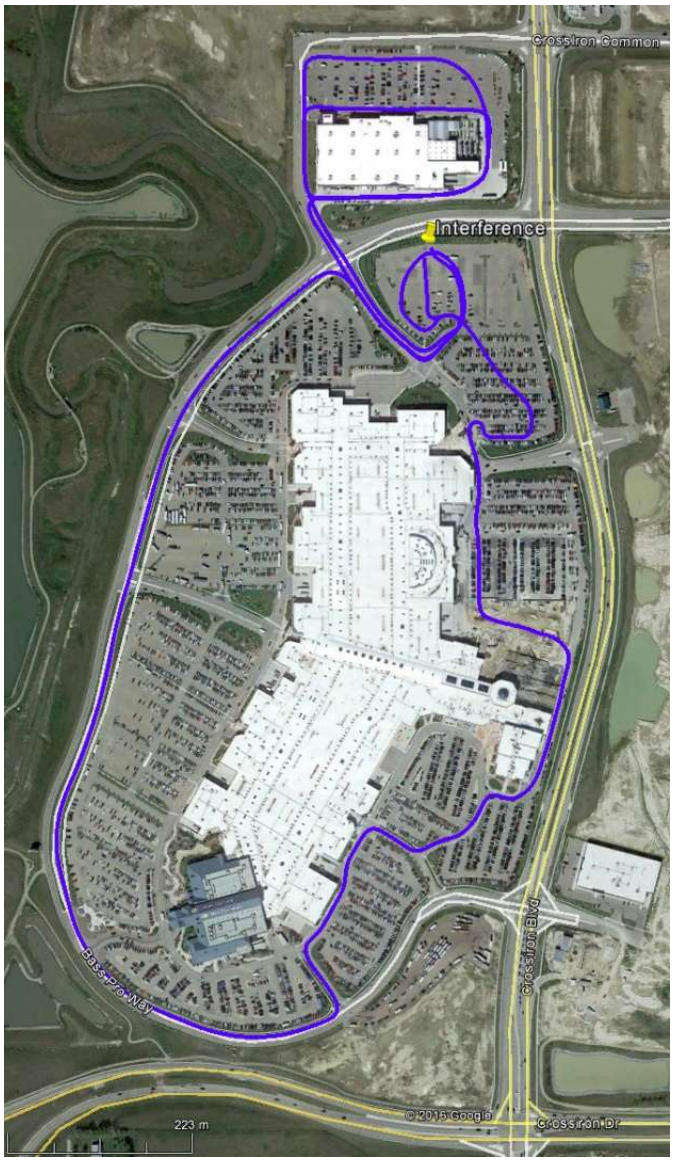

Figure 16: Map of the test area showing the location of the interference source.

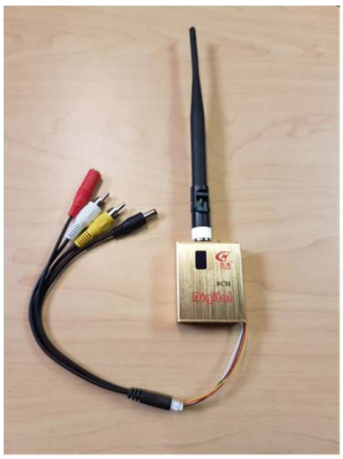

Figure 17: Video transmitter for a drone used as an L2 interference source

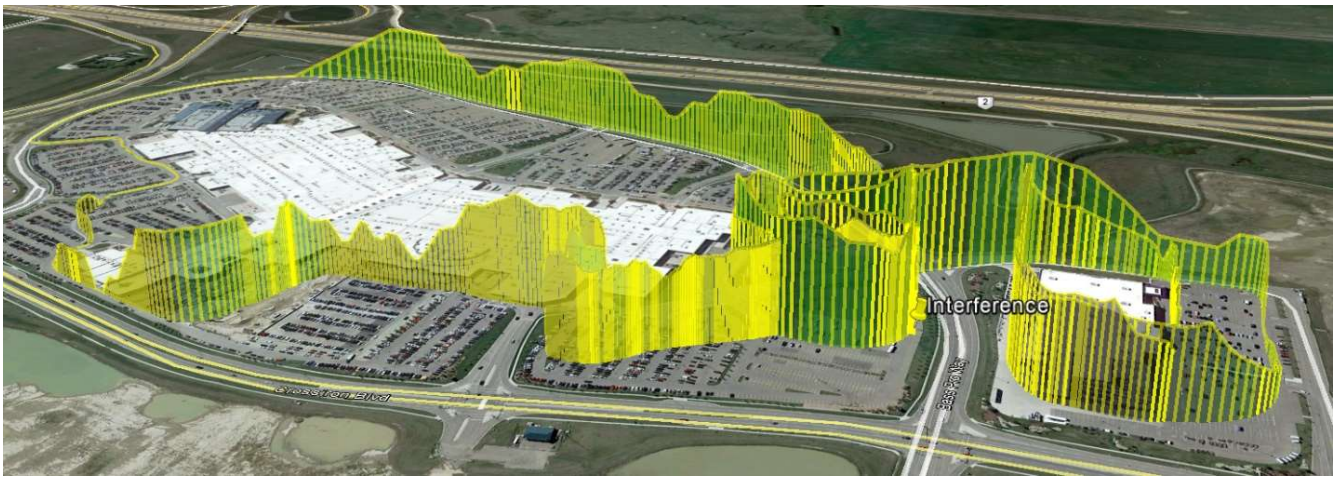

Figure 18: Observed power of the interference source (yellow) over the test course (Map data: Google, Landsat / Copernicus, DigitalGlobe)

Since the interference was only broadcasting on L2, a position is still available through the other GNSS frequencies. The GPS receiver has difficulty tracking GPS L2 signals because of the interference. Figure 19 shows the number of GPS L2 signals tracked. As the receiver approaches the interference source, it becomes more and more difficult to track the L2 signals. As the receiver moves away from the interference, or behind a physical obstruction (like a building), the impact of the interference decreases and the signals are reacquired.

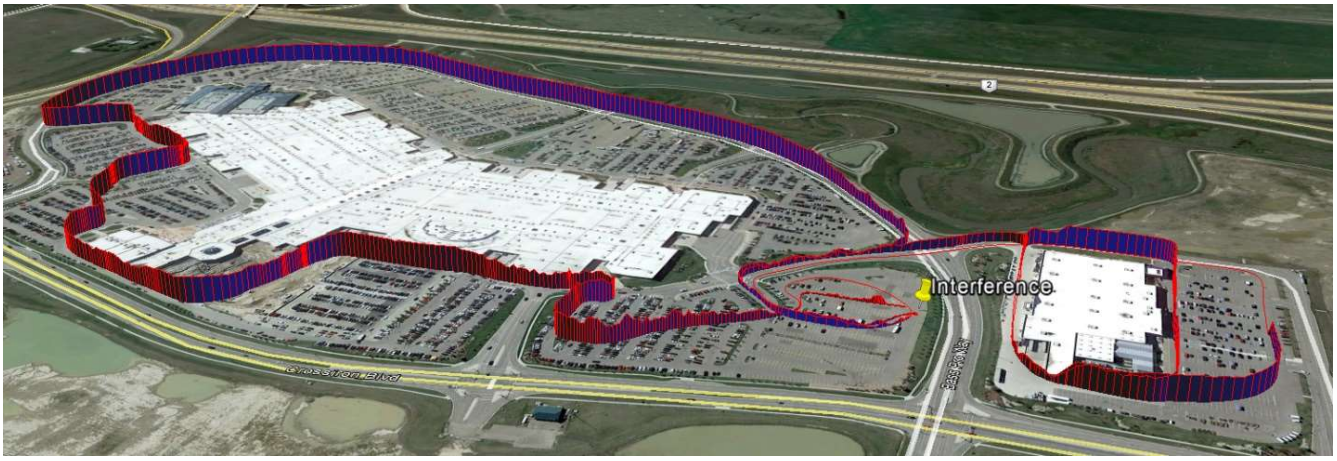

Figure 19: Number of L2 satellites tracked (red) over the test course (Map data: Google, Landsat / Copernicus, DigitalGlobe)

This shows how a simple device can inadvertently be harmful. Anyone could have purchased this device to simply transmit video from their recreational drone. Since this device only broadcasts on L2, the GPS of the drone and many nearby devices would have been unaffected, while almost completely jamming and disrupting any dual frequency receivers nearby.

The interference power at each location is used to estimate the power of the interference as a function of location. This can be visualized in a heat map of the interference power. The power attenuation as a function of distance from the transmitter is essential in calculating the location of the source using this goodness-of-fit method. This function can also be used to estimate the power of undetectable, hidden transmitters that are too far away and too weak to be detectable.

Figure 20 shows the interference goodness-of-fit map from the real data test. The map shows the correct trend but the peak of the map does not include the actual location of the interference transmitter. This is due to inaccuracies in the power attenuation model. For example, a significant shift to the south is due to the rapid decrease in power when moving behind the north building.

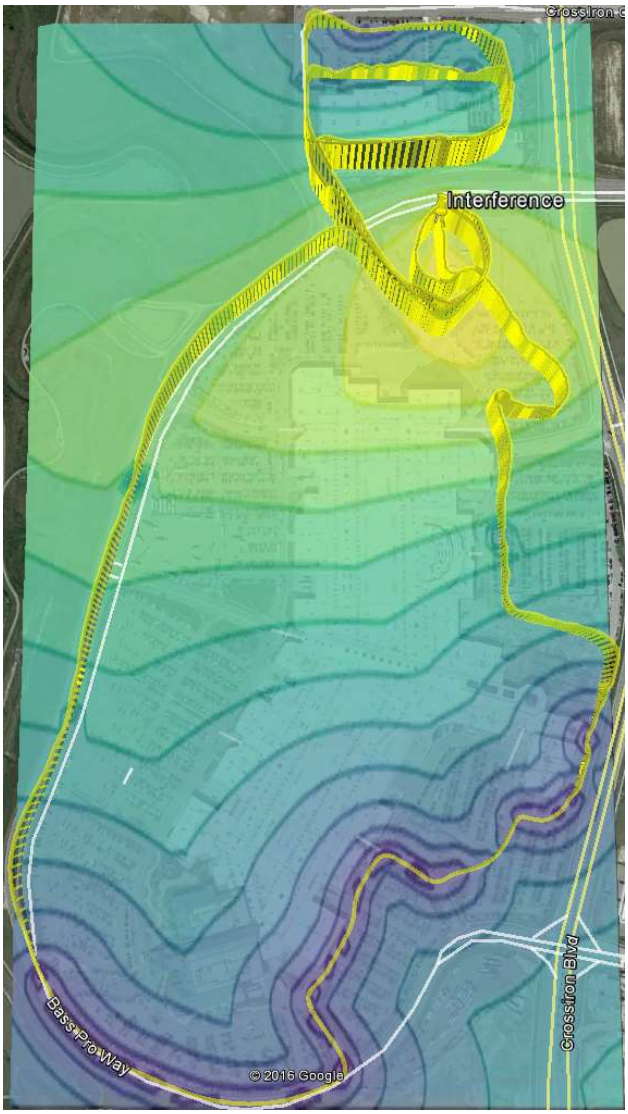

Figure 20: Interference map from the real-data test

Figure 21 shows the interference map when only the southern dataset is considered. Without the impact of the northern building, the map is more accurate. This is because the attenuation model does not account for obstructions. The performance of this kind of model could be significantly improved with a model that includes the topography and buildings, similar to the one proposed by Khalifeh et al. (2015).

Despite the inaccuracy of the map to precisely locate the interference source, these simple model maps give a nice visualization of the interference. Figure 22 shows the number of L2 satellites tracked on top of the interference map. Perhaps a more accurate map could be produced by using the tracking performance along with the power model described here.

Figure 21: Interference map from the real-data test using only the south dataset

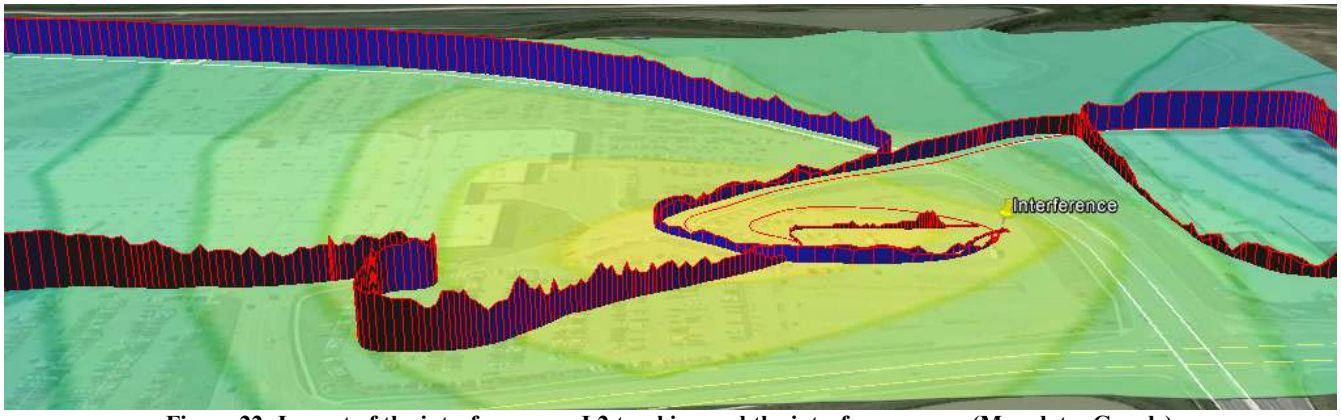

Figure 22: Impact of the interference on L2 tracking and the interference map (Map data: Google)

TOKYO REAL DATA RESULTS

We received a report of interference in Tokyo, Japan and took a receiver to investigate. Figure 23 shows the maximum received power throughout the dataset. The interference around 1570.69MHz is obvious and easily to identify in the figure.

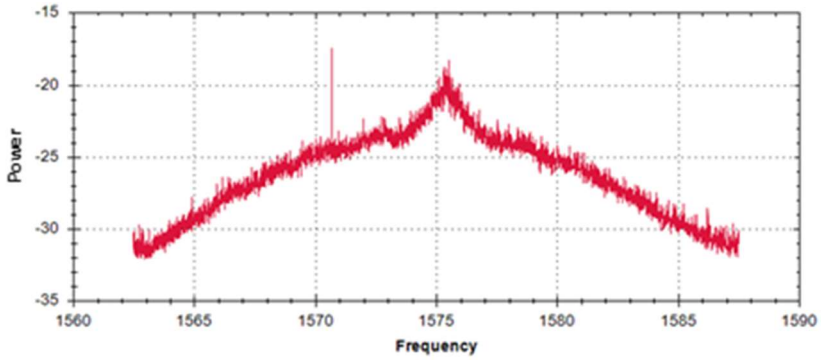

Figure 23 Spectrum power level for the Tokyo dataset

Figure 24 shows the observed power of the interference source when walking around the building. There is a peak in the received power when moving to one side of the building, while the observed power is relatively constant over the other three sides of the building. This strongly suggests that the interference source is along the one side of the building. This figure also shows the estimated goodness-of-fit interference map produced using the algorithm described earlier. The source of the interference could not be conclusively determined however we believe that the source was emanating from one of the vehicles in the parking lot.

This real example illustrates how useful this visualization of the observed power is in understanding the nature of the interference, identifying the source and localizing its effect. The interference in this case did not cause a noticeable change in the number of satellites or signals tracked.

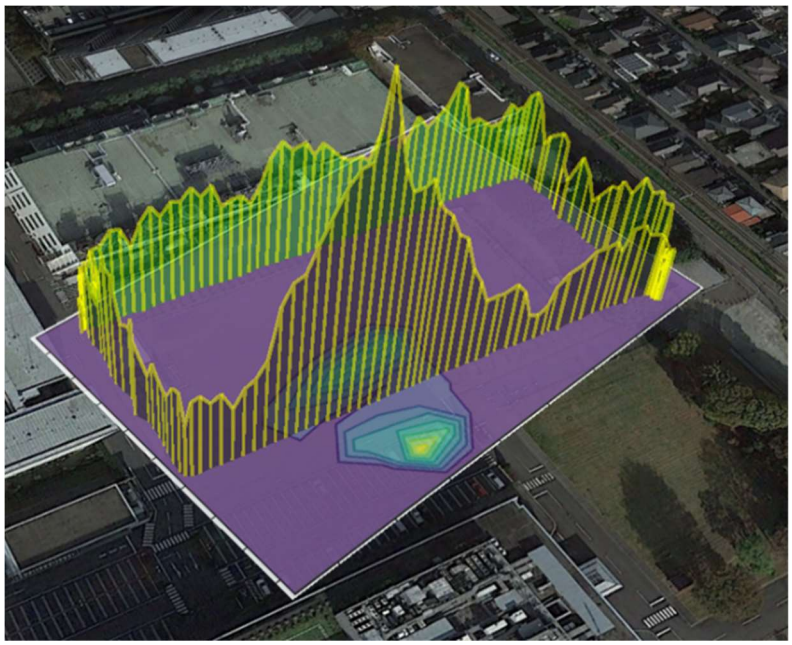

Figure 24: Observed power of the interference source (yellow) for the Tokyo dataset (Map data: Google, Zenrin)

CONCLUSIONS

This paper shows a creative and useful application of NovAtel's Interference Tool Kit that was recently released as a feature on their OEM7 line of receivers. The ITK can be used to create maps that show the estimated location of the interferer as well as the impact of the interference on other users. This was demonstrated first using simulated datasets where the agreement between the simulated and actual loss of power models made for overly optimistic results. Three case studies are also shown:

The original motivation for this work was a customer service case in India. The second is a case in Calgary where unintentional interference was being caused by a drone video transmitter. The third dataset was a similar example from Tokyo, where, unfortunately, the true interference source could not be conclusively identified.

The three interference case studies show the importance for interference detection and mitigation because intentional and unintentional interference sources are easy to obtain and are not easily monitored or restricted. In one of these cases, a device that was naively purchased online as a UAV video transmitter ended up jamming GPS L2 in an area roughly 2000m^2. With interference mitigation, it is possible to continue to work and operate in these environments without interruption or significant impact.

FUTURE WORK

There are many possibilities that are created by the Interference Took Kit. The accuracy of the maps can be improved by using the topology and building locations or combining it with tracking or positioning degradation observed while under the influence of the interference.

A difference in power from multiple receivers could be used to calculate a direction to the source of the interference. If included with a directional antenna, it can also be used to determine the direction of the interference by manually rotating the antenna and looking for changes in power. This can be done through looking for peak power or pointing the antenna null at the interferer and watching for the drop in power.

Using a slightly different model and multiple receivers distributed over a region, ITK could be used to track moving interferers as they travel through a network.

ACKNOWLEDGEMENTS

The authors would like to thank Bryan Leedham and Saravanan Karuppasamy for sharing their customer stories with us and providing us with the data for the case studies.

REFERENCES

Crawford, James (2002) “Propagation Losses Through Common Building Materials: 2.4 GHz vs 5 GHz, Reflection and Transmission Losses Through Common Building Materials”, *Technical Report E10589*, Magis Networks, Inc.

Faruque, Saleh (2015) “Radio Frequency Propagation Made Easy”, *SpringerBriefs in Electrical and Computer Engineering*, ISBN 978-3-319-11394-4

Gao, Feng, and Sandra Kennedy (2016) “Demonstrated Interference Detection and Mitigation with a Multi-frequency High Precision Receiver” *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, Portland, Oregon, September 2016, pp. 159-170.

Khalifeh, Joe J., Kassas, Zaher M., Saab, Samer S. (2015) "Indoor Localization Based on Floor Plans and Power Maps: Non-Line of Sight to Virtual Line of Sight," *Proceedings of the 28th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2015)*, Tampa, Florida, September 2015, pp. 2291-2300.

Mitch, R.H., Psiaki, M.L., Powell, S.P., O'Hanlon, B.W. (2013) "Signal Acquisition and Tracking of Chirp-Style GPS Jammers," *Proceedings of the 26th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2013)*, Nashville, TN, September 2013, pp. 2893-2909.