8-1-2010

DTIME: Discrete Topological Imaging for Multipath Environments

Robert Ghrist
University of Pennsylvania, ghrist@seas.upenn.edu

H. Owen
HSOwen LLC, hank@hsowen.com

Michael Robinson
University of Pennsylvania, robim@math.upenn.edu

DTIME: Discrete Topological Imaging for Multipath Environments

Abstract

This report is presented to summarize work completed under a DARPA seedling project for the imaging of urban environments, using radio multipath measurements and topology extraction algorithms. This report provides an overview of the mathematical theory behind the work, as well as a description of the simulation and results that accompanies the theory.

Keywords

sensors, signals of opportunity, imaging, topology, multipath

Disciplines

Applied Mathematics | Geometry and Topology | Signal Processing

TECHNICAL REPORT

Title: "DTIME: Discrete Topological Imaging for Multipath Environments"

Grant: DARPA STO - HR0011-09-1-0050

Place of Performance: University of Pennsylvania

Period of Performance: 8 July 2009 – 8 April 2010

Investigators: Prof. Robert W. Ghrist
                          Mathematics & Electrical/Systems Engineering
                          University of Pennsylvania

                          Dr. Michael Robinson
                          Mathematics, University of Pennsylvania

                          Henry Owen, HSOwen LLC

OVERVIEW AND DELIVERABLES:

This report is presented to summarize work completed under a DARPA seedling project for the imaging of urban environments, using radio multipath measurements and topology extraction algorithms. This report provides an overview of the mathematical theory behind the work, as well as a description of the simulation and results that accompanies the theory.

Volume 1 of the grant application states:

The research for this seedling will be focused on building theoretical and computational methods for imaging in problems with signals of opportunity, with a particular emphasis on the multipath scenario. The deliverables will include methods, algorithms, and simulations directed toward these ends as detailed below. Effort will be directed towards two major objectives.

1. We will delineate the discretization of the domain and sample space as postulated in Conjecture 1. This also entails the development of algorithms to extract the topology of simulated multipath environments as a means of demonstrating our methodology.

2. We will extract the geometric information inherent in the channel impulse response. This will entail the development of "canonical" problems and the assessment of their computational load using simulated data. This will lead to an examination of the degrees of freedom in each canonical problem and compatibility conditions between them, with the goal of an efficient, obstruction theory-driven computational model for use with simulated data.

The deliverables of this proposal are theorems and algorithms to extract the topology and geometry of simulated multipath environments, specified as per the plan of work and accomplished as per the above timeline.

There were three components to the project results:

1. Topology extraction
2. Geometry extraction
3. Simulation

The novel tools introduced arise from Topology, specifically:

1. Transversality, generalizations of the Whitney embedding, and tubular neighborhood theorems;
2. Nerve constructions and Čech homology;
3. Cellular sheaf cohomology leading to an obstruction theory for imaging.

MATHEMATICAL THEORY

The mathematical theory associated with this seedling is discussed in two sections: one on topology extraction and a second on geometry extraction.

TOPOLOGY EXTRACTION

We considered the problems of localization, disambiguation, and mapping in a domain solved by multiple receivers utilizing signals-of-opportunity generated by transmitters located throughout the domain. A variety of different scenarios were addressed, including varying signal types (TOA, TDOA, DOA, etc.), incorporating mobile receivers, and discretizing the signal space as a model of coarse/uncertain signal processing.

Our strategy for executing topology extraction from a collection of signals began with a theorem on unique channel response. This fundamental result gave the critical bound on the number of (primary or reflected) signals necessary to disambiguate receiver location.

Theorem: *Let D be a compact manifold with corners. Fix a collection of transmitters in D whose signals are stably readable over subdomains which form a cover, **C**, of D. For generic signal placement, the set of points in D on which the TOA channel response is non-injective is of dimension 2 dim D – dep(C), where dep denotes the depth of the cover **C** – the minimal number of signals hearable in any given location.*

This result holds with or without unique signal ID. TOA, alone and without transmitter ID, suffices.

Corollary: *On a planar domain, five transmitter signals (primary or reflected) suffice to localize receivers uniquely.*

The results do not rely on TOA: TDOA, DOA, and FDOA are also covered, with appropriate changes to the signal profile mapping dimensions.

From this disambiguation result, topology extraction is produced by means of two ingredients:

1. **Cech homology**: to reconstruct the nerve of the transmitted coverage domains, and thus ascertain domain topology; and
2. **Clustering**: to differentiate connected components in the (potentially non-acyclic) cover.

GEOMETRY EXTRACTION

Theoretical progress on geometry extraction was accomplished in the limiting scenario of the urban canyon, where one approximates the underlying domain with an embedded, planar, metric graph.

The principal technical tool introduced was that of a flow sheaf: a novel class of constructible sheaves over metric graphs, which collate data about solutions to second-order wave-like differential operators.
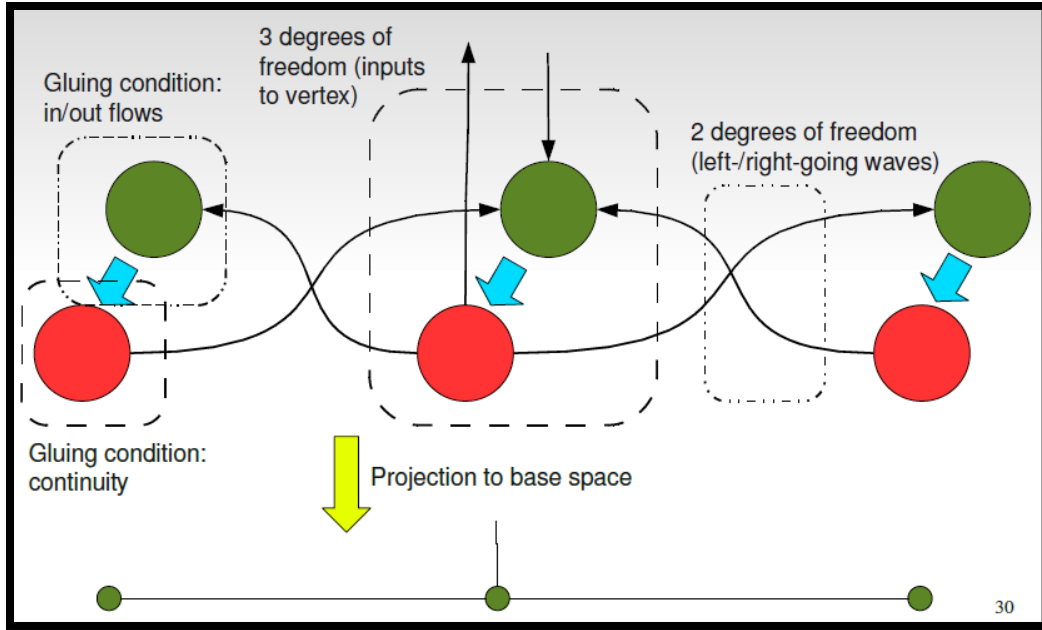
3

Figure 1: The flow sheaf construction.

Our results are encapsulated in the following general statement:

Theorem: *Given a metric graph* G *with flow sheaf* **F***, the cohomology* H*(G;**F**) *suffices to characterize the geometry of* G *(up to phase ambiguity).*

Implicit in this result is a computational algorithm for the incremental collapse of a flow sheaf, allowing for simplification methods. The algorithm for domain geometry is, via this method, linear in the size (number of vertices) of the underlying graph.

SIMULATION

A simulation was developed for this project to evaluate specific algorithms for topology extraction and geometry extraction of an urban environment, using realistic radio data. The goal of the overall simulation is to retrieve a map of building obstructions, building locations, and the associated propagation paths in an urban canyon (maze like) environment.

The topology extraction is largely implemented, while the geometry extraction is not yet implemented. The implementation of the radio data simulation as well as the topology extraction simulation, with initial results, is discussed here.

SIMULATION SCOPE

The simulation developed for this project consists of two main components. These include a radio propagation simulation, and an algorithm simulation to retrieve topological and geometric information from the radio measurements.

The radio propagation simulation focused on urban propagation of radio frequency (RF) signals, using operationally relevant modulations and geometries. Several goals were kept in mind as the simulation was developed:

4

- To provide realistic radio measurements for the urban canyon (maze like) environment, including first order effects imperfect measurements and noise;
- To use radio characteristics that are found either in commercial infrastructure today or are available (or soon will be available) operationally to the military;
- To capture the first order effects of urban multipath; and
- To provide a mechanism to allow for easy adaptation of transmitter and receiver distribution in a typical urban environment.

Table 1 provides an overview of several candidates considered for use in the radio simulation. These candidate radio protocols combine commercial infrastructure protocols and operational military tactical radios. The electronic warfare (EW) mode of Soldier Radio Waveform (SRW) was selected to be used in the radio simulation. The soundings retrieved from the SRW EW mode synchronization preamble form the basis for the radio simulation. These soundings are expected to be similar to those also retrieved by the commercial DTV (8VSB) field synchronization sequence. However, since SRW equipped radios are more likely to be used in a configuration that results in purely urban canyon propagation (that is, low height transmitter and low height receiver, both located within the urban canyon environment), it was chosen for this study.

An overview of salient issues involved in urban propagation is summarized well in ITU-R P.1441 [1] and in [2]. Of the various situations that can be encountered, Figure 2 and Figure 3 provide an overview of two basic scenarios. In the first scenario, a transmitter mounted above the level of many roofs is transmitting to a receiver at ground level. In the second case, both the transmitter and receiver are operating at ground level. The first case involves over-the-roof propagation, diffraction to the ground level in the vicinity of the receiver, and multiple reflections in the vicinity of the receiver. Additionally, occasional reflections from tall, sometimes distant, scatterers occur. In the second case, propagation is primarily (although not completely) confined to a maze-like propagation mechanism where a small percentage of the energy diffracts over rooftops.
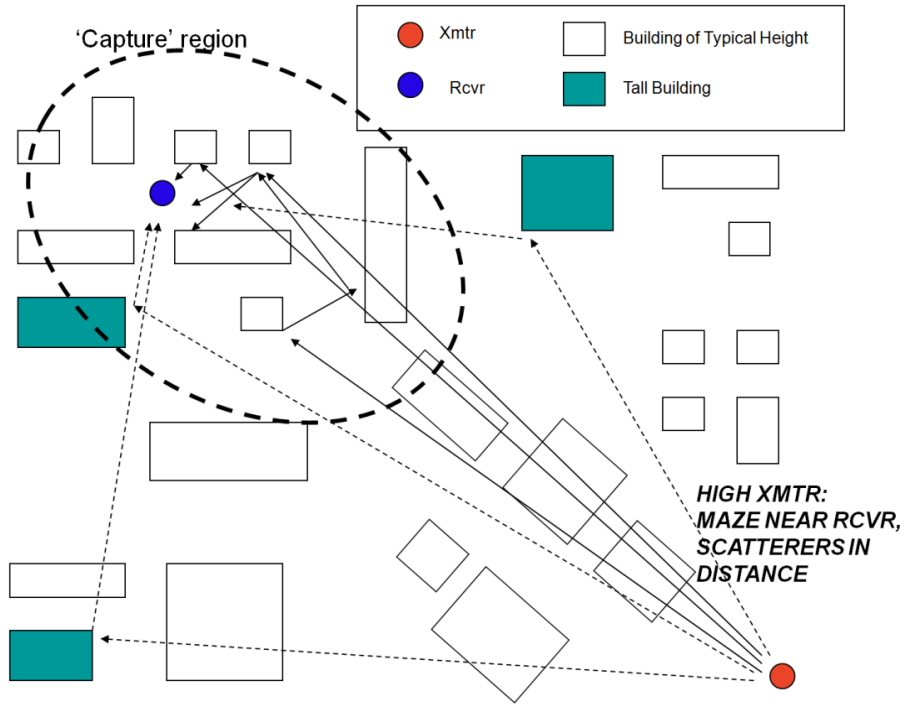
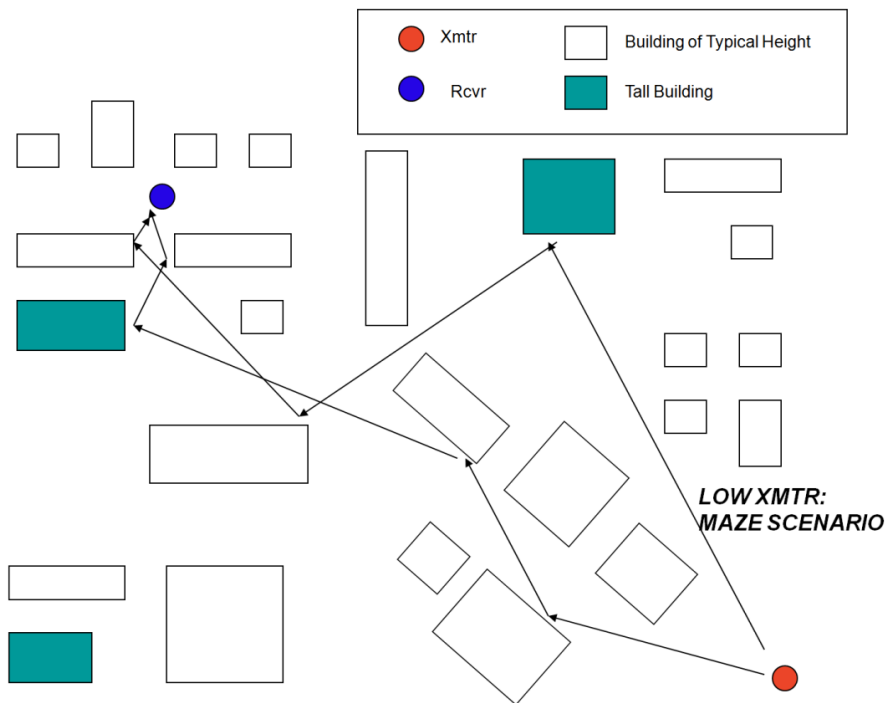**Figure 2: Overview of multipath conditions, elevated transmitter case.**

Figure 3: Overview of multipath conditions, ground level transmitter case.

6

For the purposes of this project, the focus is on the second scenario, identified in Figure 3, where maze-like propagation conditions are assumed to dominate. To simplify the model while still addressing propagation mechanisms accurately to a first order, the focus of the model is reflected energy, and diffraction is only considered in a very simplistic manner. It is expected that the techniques investigated under this project will still apply as diffraction effects and scattering effects are more accurately modeled.

Protocol Name	Support Provided for Channel Estimation	Likely Channel Estimation Approach	Modulation Characteristics Relevant to Channel Estimation	Notes
Soldier Radio Waveform (EW Mode)	Training sequence embedded.	Correlation estimator.	2048 chip acquisition sequence (variable length chips, from 2.4-32 Mcps). (circa 2003)	Recent military wireless protocol.
EPLRS	Training sequence embedded.	Correlation estimator.	320 chip preamble (0.6 uS each chip). Min possible update period is 2 mS (variable).	Older military wireless protocol. Freq: 420-450 MHz.
DTV: 8VSB (North America)	Training sequence embedded.	Correlation estimator or equalizer estimation.	700 symbol sequence (0.09 uS each symbol). Update period is 24.2 mS.	Freq: 54-88 MHz, 174-216 MHz, 470-806 MHz US (may change in US, varies around world).
DTV: DVB-T (Europe) *[Note 1]*	Pilot carriers embedded.	Frequency domain estimator.	Swept pilots that monitor every 3th carrier (2048 carrier mode: 4 kHz carrier spacing, 568 sweeping pilots at every 12th pilot, jumping by 3 carriers every 250 uS symbol).	Configurable modulation, with variable configurations of number of carriers, symbol length, guard interval, etc.

Table 1: Candidate wireless protocols considered for radio simulation.

Note 1: Has some similarities to WNW and WiMAX in terms of general use of OFDM.

RAYTRACE IMPLEMENTATION

Standard ray trace propagation modeling routines were used to build a ray trace simulation, using the MATLAB programming language.   Building data and received power measurement data, useful in validating the basic performance of the ray trace model used here, are publicly available and described in [3].   In developing the ray trace simulation, measured reflection characteristics

were taken from [4] and [5]. These references provided guidance on expected signal losses due to reflection from various material types relevant to urban buildings.

From the ray trace simulation, a set of idealized echo profiles is generated. Figure 4 provides an illustration of sample results from the ray trace simulation. In Figure 4, ray paths that originate from a transmitter (upper left) and eventually intersect a receiver (lower center) are illustrated in red in the left figure. Building outlines are shown as solid blue lines. The echo profile that results, considering the effects of free space propagation loss and attenuation due to reflection, is shown in the figure on the right. Four sets of echoes are visible: direct path echoes (1 uS delay), echoes that bounce once behind the receiver (1.3 uS delay), echoes that bounce behind the transmitter and then reach the receiver (1.9 uS delay), and echoes that bounce behind the transmitter, pass by the receiver, and then bounce off the wall behind the receiver before being received (2.2 uS delay).
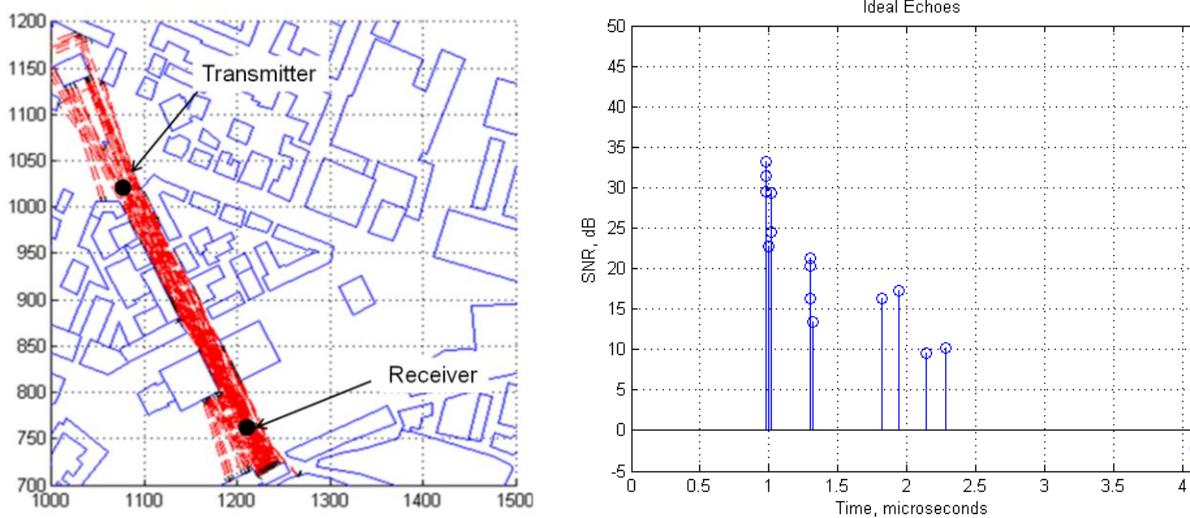

Figure 4: Ray trace example for one transmitter – receiver pair.

The simulated received power levels were compared to observed power levels provided in the study described in [3]. The simulated receiver path is shown in Figure 5 as a sequence of red points. The transmitter location is shown in the lower center part of the figure in black. Receivers are numbered from number 1 (upper left) to number 48 (upper right), travelling in a counter clockwise direction. The associated echo profiles for each receiver are shown in Figure 6. The x-axis represents the progression of receiver locations, as the receiver moves in a counter clockwise manner. The y-axis represents echo detections as a function of absolute time delay from transmission.
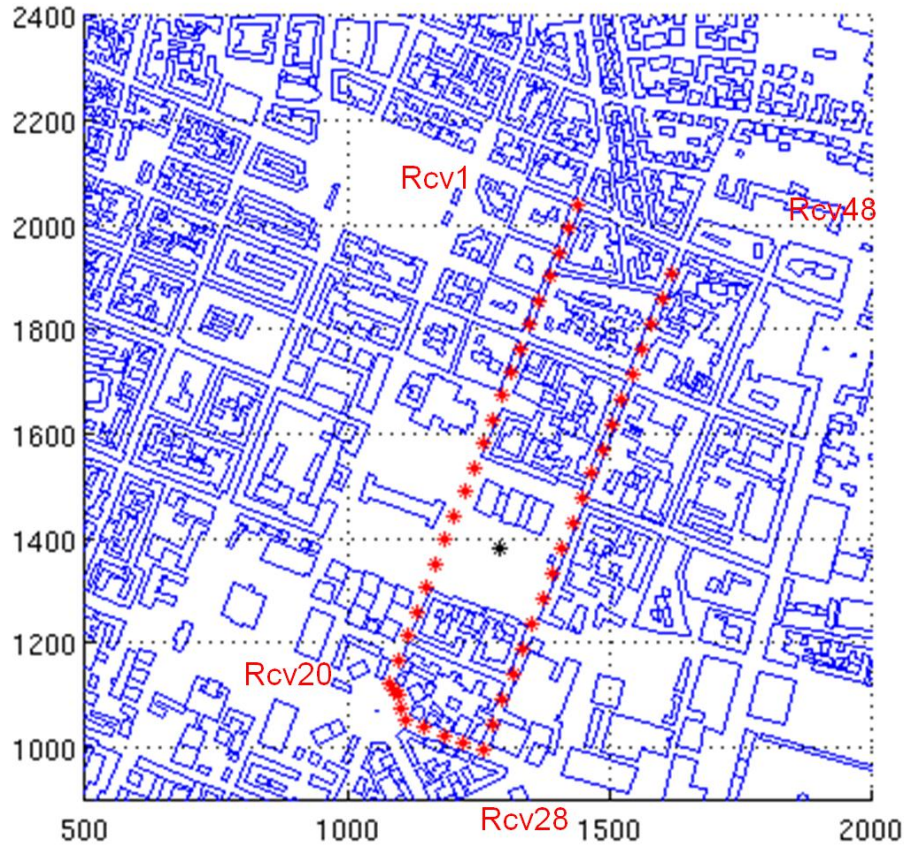
8

Figure 5: Transmitter location and receiver path for power level validation.

Simulated power levels are compared to actual observed power levels from COST 231 in Figure 7. The figure on the left shows simulated power levels (blue) as compared to actual power levels (red), by receiver number. In the center of the figure, there is a region where the simulation, discounting over roof top diffraction, shows an absence of received signal. The actual measurements, which include real roof top diffraction, fills in that region with a very low signal level. The figure on the right identifies the same comparison, plotting the received power levels as a function of range. Those points associated with roof top diffraction (absent in the simulation) are identified in the figure on the right. Overall, the two sets of measurements compare reasonably well, to the degree that is necessary for this study. The difference is that the simulation power level rolls off more slowly as a function of range. Although this could be partly due to imperfect reflection loss characterization, it is more likely due to overly optimistic path loss assumptions, where the simulation approximates free spaced path loss between reflections that are assumed to occur from perfectly flat, very large surfaces. As discussed in [1] the power roll of is expected to be greater, in a manner that is dependent on the details of the building characteristics in a given urban environment. So the results compare reasonably well, and the observed difference is expected when standard propagation models are consulted.
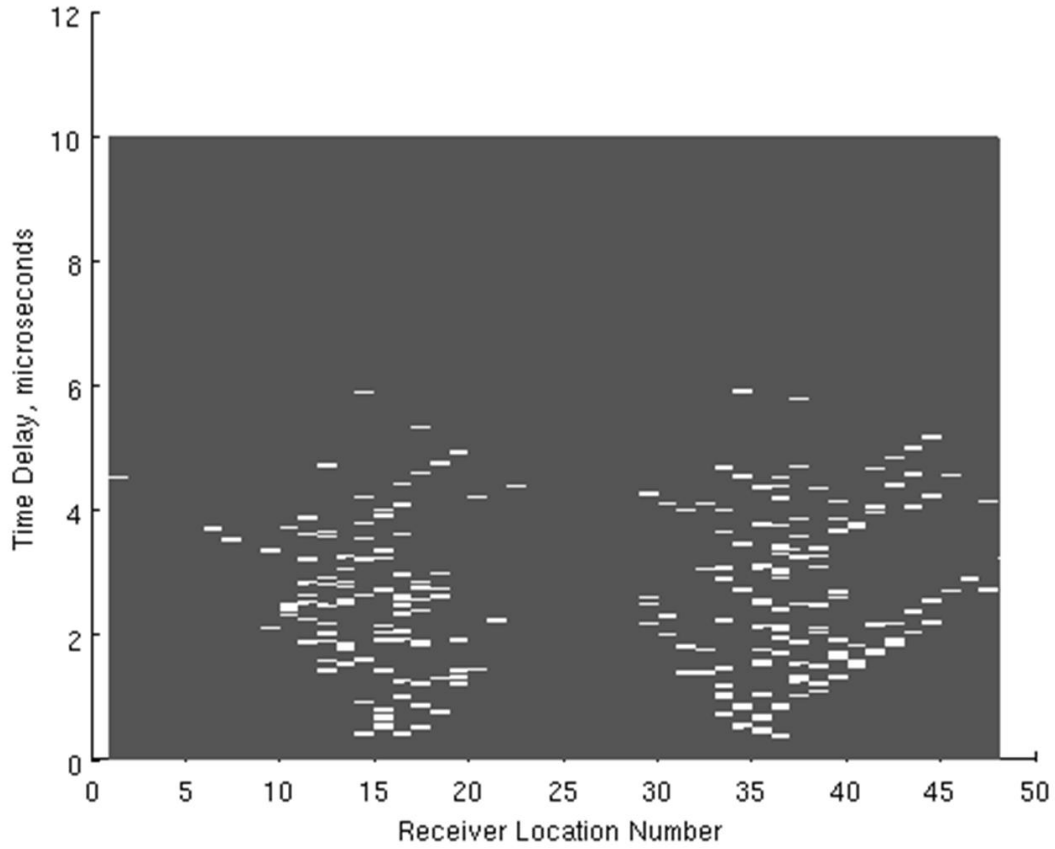
9

Figure 6: Echo profiles for the collection of receiver locations during power validation.
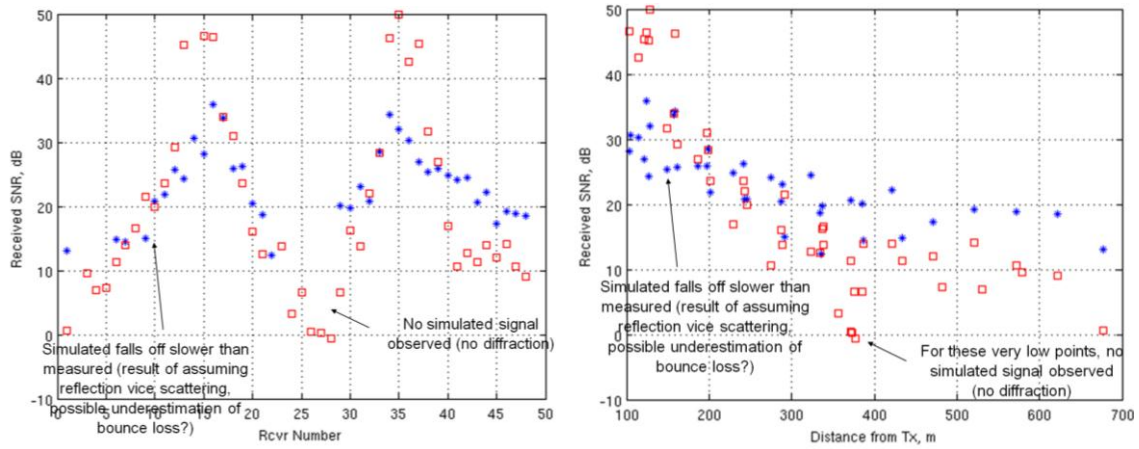

Figure 7: Received SNR by receiver and distance.

RADIO SIGNAL SIMULATION

The second portion of the ray trace program translates ideal echo profiles into realistic radio signals. Figure 8 provides an overview of the full process, from generating idealized echoes to creation of the realistic radio echoes to be used to drive the topology and geometry algorithms. The first two steps, generation of idealized echoes and realistic echo locations, including reflection losses and propagation losses, have been described. Next, the effects of signal modulation, reception, and noise are added. The third step in the figure shows the translation of the idealized echo amplitudes to distorted (realistic) echo amplitudes, and the fourth step shows the detection process by which echoes are declared. Depending on the level of distortion due to the modulation and noise characteristics, some smaller echoes may be lost, and others may be incorrectly inserted. The last figure shows two of the original echoes accurately represented, one that is missed, and another that is falsely inserted.
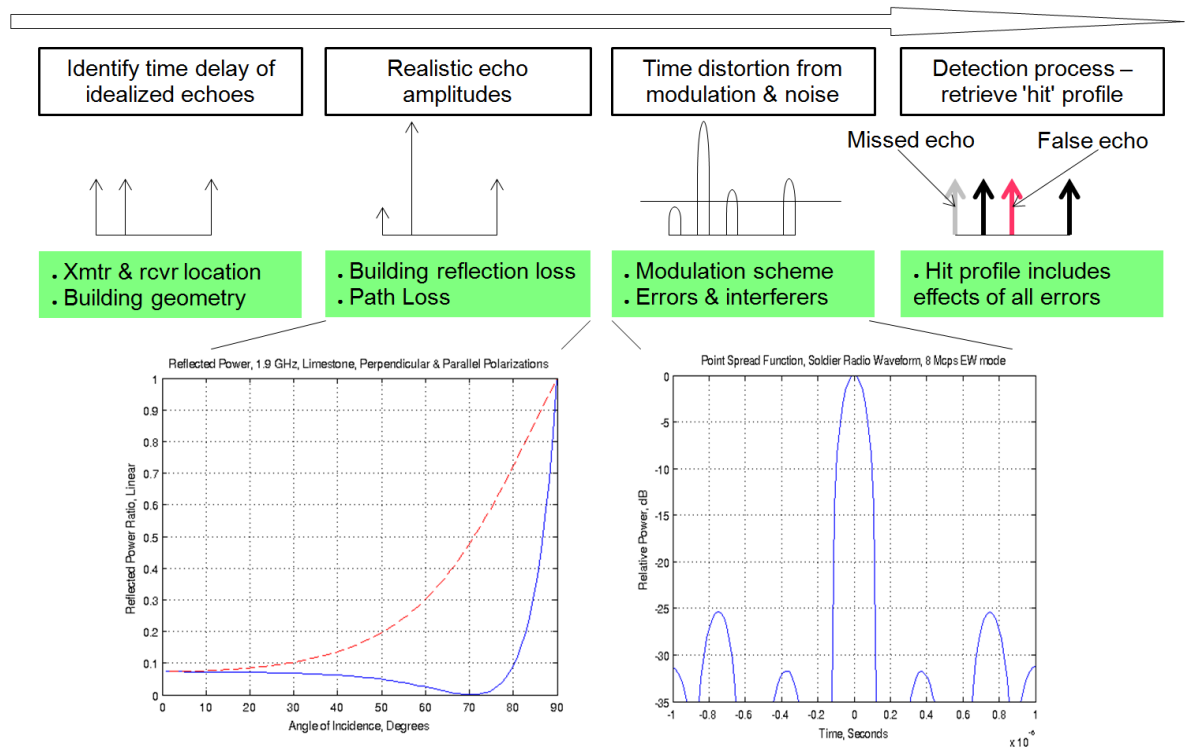

Figure 8: Overview of radio signal simulation process.

The basic modulation characteristics are captured via a point spread function that describes the modulation and detection characteristics of a preamble sequence in the radio packet. This point spread function consists of the convolution of a complex transmit sequence with a matched filter used in the preamble detection process, with the band limiting filtering associated with the MSK modulation scheme of the waveform. Figure 9 provides an overview of the point spread function associated with the Soldier Radio Waveform synchronization sequence, operating in an 8 MHz mode. This is very similar to operation at other bandwidths, with appropriate scaling applied as bandwidth is shifted upwards toward 16 MHz or 32 MHz EW operating modes. Figure 10 provides a close in view of the same point spread function, zoomed to the first few time sidelobes of the preamble sequence.
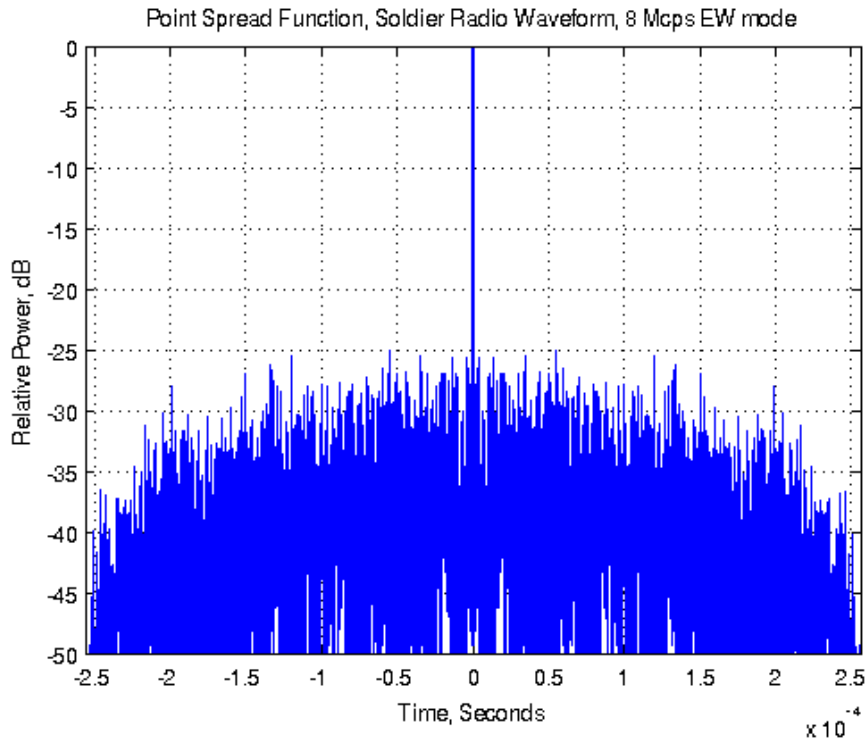
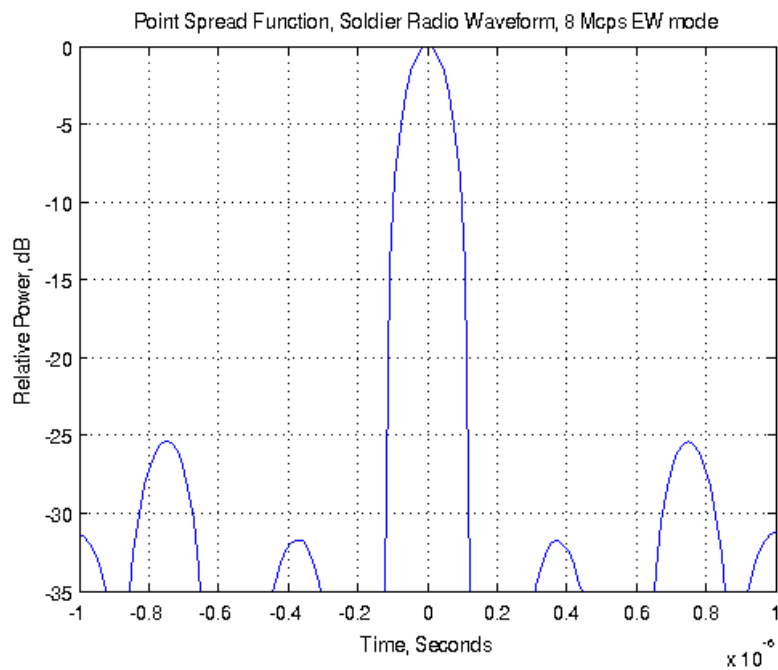**Figure 9: Example point spread function, 8 MHz band width.**

Figure 10: Example point spread function, 8 MHz band width, zoomed view.

Figure 11 and Figure 12 provide an overview of the progression through the detection process as an oversampled realistic signal is created, detection processing is conducted, and appropriate sampling (to a rate close to the inverse bandwidth of the modulated signal) is performed. Figure 12

13 provides a direct comparison of the idealized echo and the simulated, more realistic, echo. This comparison is provided for the scenario described in Figure 4. As expected, several of the echo groupings are accurately retained. Some shifts occur in the detailed signal samples for the first three echo groups, but the overall structure is still evident. However, the last echo group (associated with the ray path that bounces off walls behind both the transmitter and receiver) is lost due to SNR limits of the receiver. This level of fidelity is an important part of the simulation, to ensure that any algorithms subsequently evaluated have reasonably robust characteristics when it comes to realistic echo profiles.
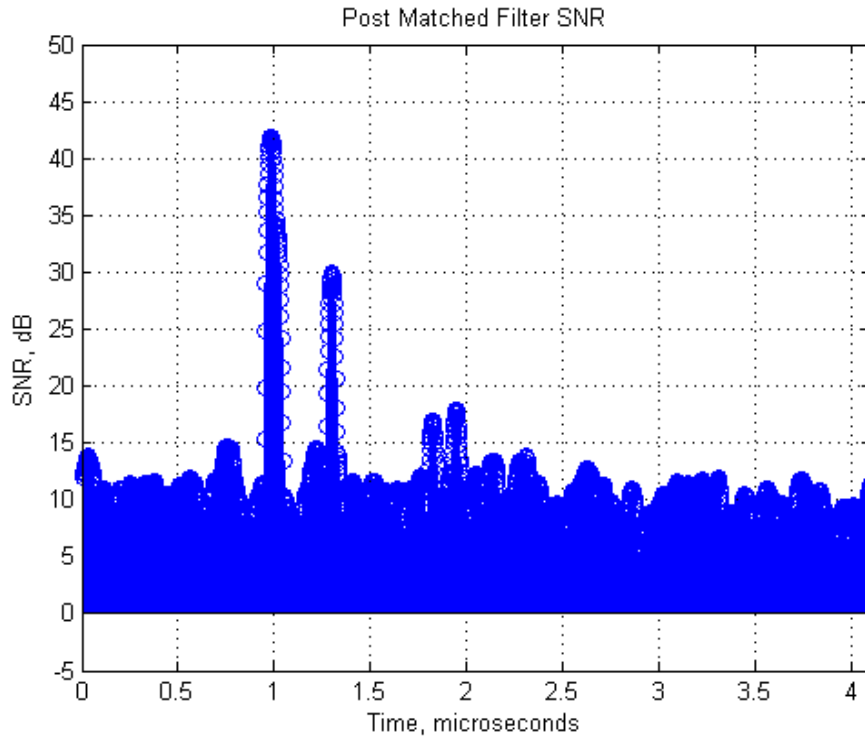

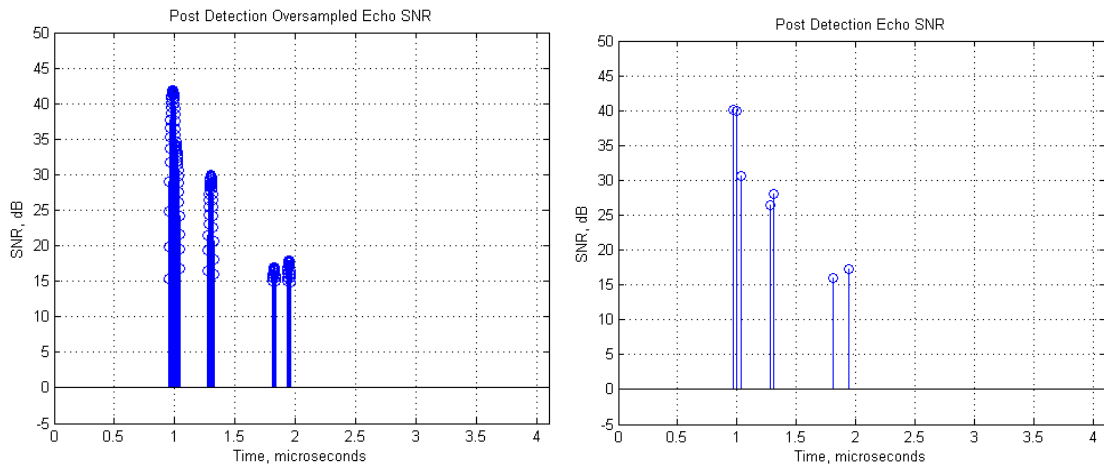**Figure 11: Post matched filter power profile from receiver simulation.**

Figure 12: Detected echoes - oversampled (left) and sampled at inverse bandwidth (right).
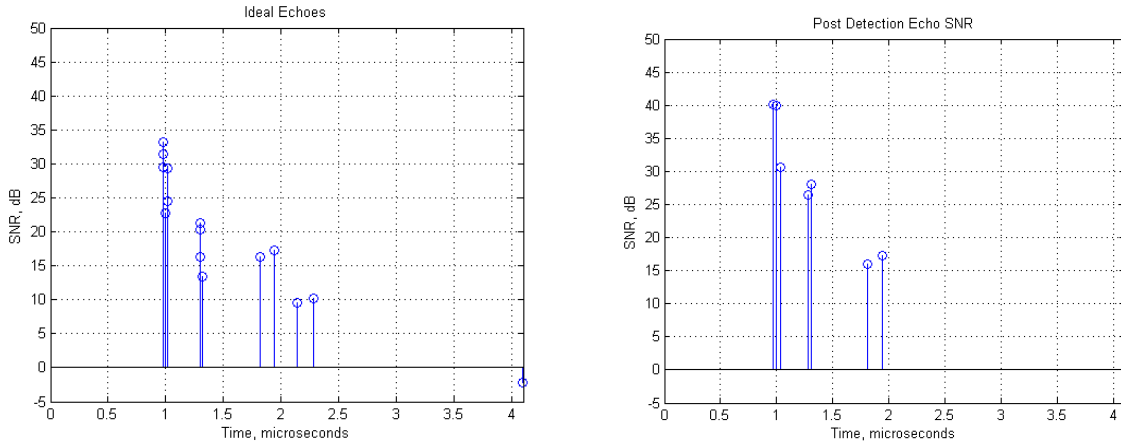
13

Figure 13: Ideal echoes (left) and simulated realistic echoes (right).

The echo profiles shown here are reasonably close to actual profiles that result in urban canyon environments, when the number of propagation paths of various lengths is limited. As rooftop diffraction paths become more significant (cases where building heights are lower, or transmitter or receiver heights are higher), then more reflections with varying delays, resulting in a more continuous profile, is expected.

SCENARIO DESCRIPTION

The simulation work was conducted using COST 231 building footprint information, from Munich, Germany. Transmitters and receivers were placed in one region of the city. Transmitters were placed randomly, and receivers were placed in a routine to simulate movement through the city streets. Figure 14 provides an overview of the transmitter (left side) and receiver (right side) locations. A total of 180 receivers were positioned to allow for a variety of receiver densities to be tested during simulation. The right hand side figure shows a subsampling of 3:1, so that 60 receivers are shown. This was a standard density used in many of the simulations.
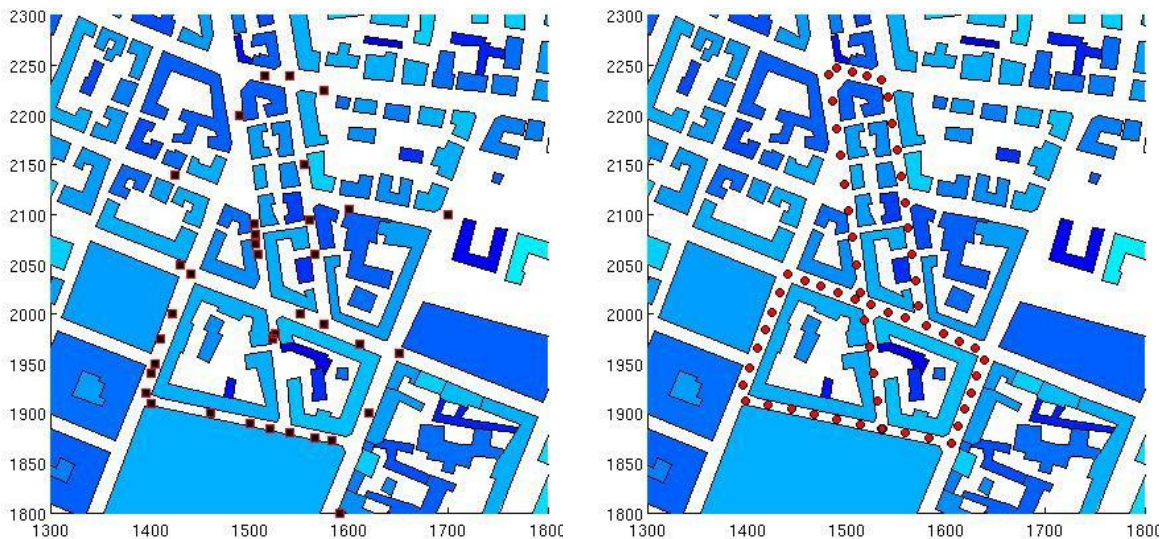

Figure 14: Transmitter (left) and receiver (right) locations associated with topological extraction.

14

All transmitters and receivers were positioned at ground level, with the intention to limit the impact of roof diffracion propagation paths.

TOPOLOGY EXTRACTION ALGORITHM SIMULATION

The second part of the simulation, the topology extraction portion, was also created using the MATLAB programming language. The goal of this part of the simulation is to extract key topological features of the complex that represents receiver visibility of transmitters, and therefore identify RF propagation obstructions in the vicinity of the receiver paths. It uses the echo profiles generated in the ray trace simulation as the basic input measurements, and it produces the 0th and 1st homology groups of the simplicial complex that represents transmitter coverage of the receivers in the urban environment.

An overview of the topology extraction algorithm simulation is provided in Figure 15. The echo data input to the simulation is shown as a black arrow to the left of the figure, and the loops (1-cycles) associated with building obstructions are shown in black on the right side of the figure. Intermediate data, also shown in black (left to right), include:

- Refined echo data after echo power SNR thresholding and receiver and transmitter subsampling,
- Coverage matrix identifying which transmitters cover which receivers,
- Signal space distance matrix between received signals (calculated for each receiver across all transmitters),
- A refined coverage matrix (following the cluster refinement algorithm),
- Nerve generated from signal space metrics, and
- 0^{th} and 1^{st} homology groups and associated boundary maps of the resulting complex.
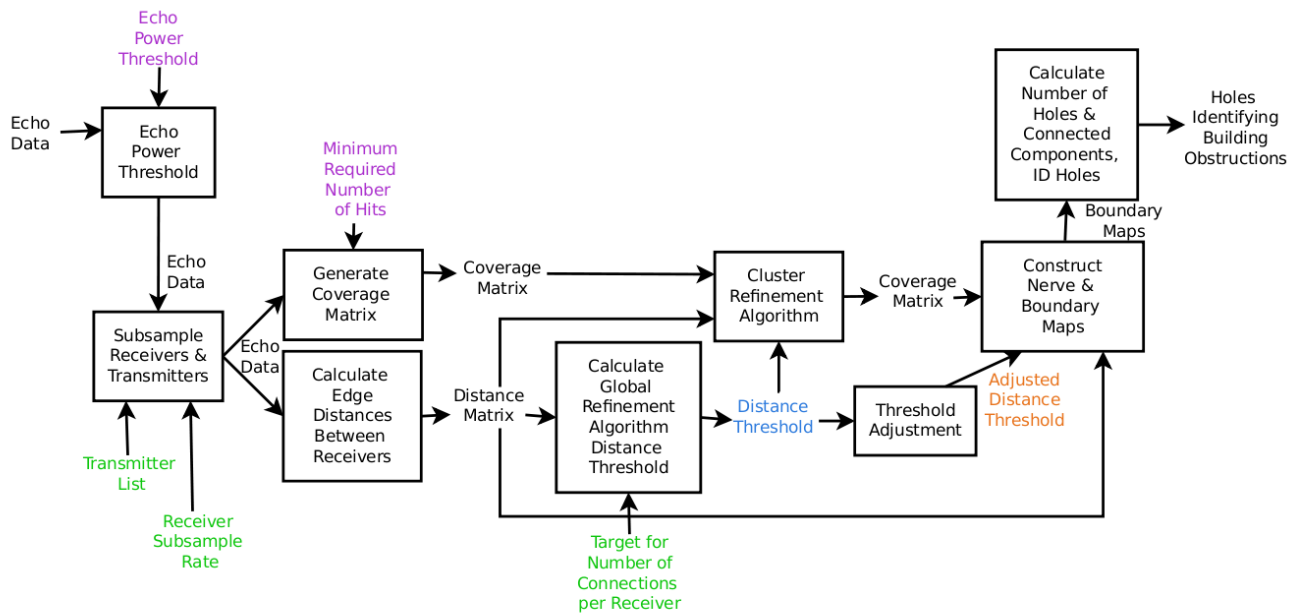

Figure 15: Process graph for topology detection.

Parameters are provided to conduct initial thresholding of the signal echo data. These thresholds are provided to help meet basic contractibility requirements for transmitter coverage of receivers, and to discard those echo profiles that are not feature rich and that may result in ambiguity in subsequent clustering algorithms. By allowing for an SNR threshold, coverage can be limited to contractible regions (examples shown later). By thresholding echo profiles using a 'minimum number of hits' metric, simple (e.g. single echo) profiles can be discarded, minimizing the chance that physical separated (and very simple) echo profiles may be grouped closely in signal space in subsequent clustering algorithms. Transmitter and receiver subsampling lists are also provided to allow for manipulation of the overall transmitter and receiver distributions.

Several signal space distance metrics are available for use in the simulation code. Three are discussed here. The first distance metric investigated was an L^2 distance metric operating on a vector of received power levels between pairs of receivers. A variation that was included is an L^2 metric operating on a vector of binary echo hits. In both cases, echo profiles are recorded as one value for sample intervals spaced at one half the inverse bandwidth of the transmitted radio signal (approximately one-half the resolution distance of the radio signal). The use of the L^2 norm on these types of vectors results in good performance when receiver spacing is fine enough to ensure that all echoes have not shifted by more than the resolution distance of the radio echoes. To address performance concerns as radio spacing becomes sparse, another metric was introduced which operates on the set of received echo detection times (and ignores information about echo power). This second metric also removes any concerns regarding relative calibration of radio receivers on signal power measurements. A more general p-norm metric was used on the distance between a set of echo delay times from one profile and the set of echo delay times of another profile from another receiver. This distance metric allows for sparse receiver spacing while still providing a good metric of similarity between echo profiles. It is the metric generally used in the simulation. In all cases, metrics were calculated either assuming perfect absolute time measurements, or alignment of pairs of echo profiles by lining up the first echo in each profile. The calculation of the signal space distance metric is reflected in Figure 15 in the box titled 'Calculate edge distance between receivers.'

Figure 16 provides an example coverage diagram, illustrating the result of the coverage refinement algorithm. On the left is a coverage diagram after SNR and 'number of hits' thresholding. A red square indicates that a transmitter (y-axis) covers a receiver (x-axis). The refinement algorithm evaluates distance metrics between each receiver, to determine whether receivers covered by pairs of transmitters should be further broken up. For receivers that have common transmitter coverage but have larger signal space distance metrics (that is, dissimilar echo profiles), additional transmitters are added to provide for new coverage regions, and receivers with similar echo profiles (over all transmitters) are grouped together in these new clusters. The result is a finer clustering of receivers, the addition of transmitters to represent these new cluster groups, and the addition of receivers to tie the new groups together (across transmitters). After the clustering algorithm is run, the overall cluster sizes are reduced, and the number of receivers and transmitters are increased.
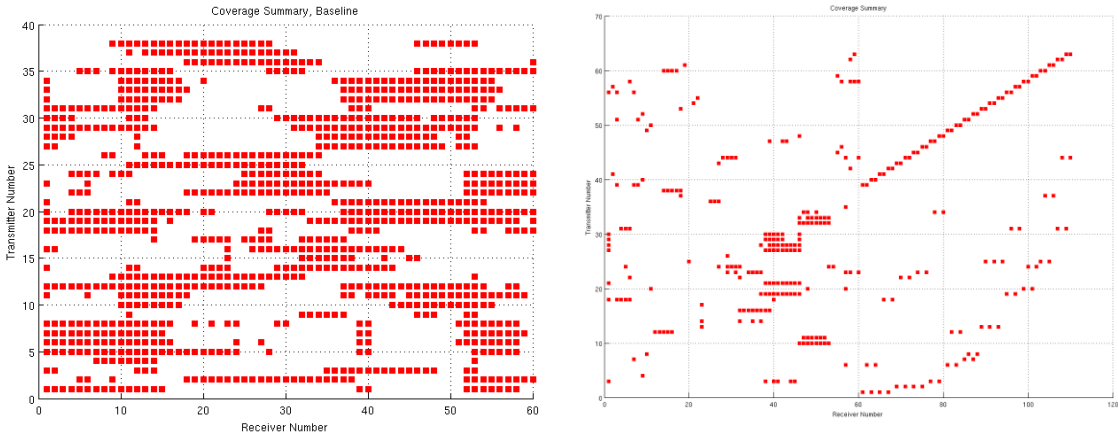
16

Figure 16: Coverage diagrams – without (left) and with (right) refinement.

A 'number of connections per receiver' target is provided to allow for meaningful control of the cluster refinement algorithm distance threshold using a view of average connectivity of vertices. This target parameter in turns generates a cluster refinement distance threshold, used to break up larger transmitter cluster groups into smaller, more closely clustered, groups when appropriate. The refinement distance threshold is shown in blue in Figure 15. A distance metric threshold is also provided in the last stage of the algorithm set, at the nerve construction step, to allow for the 'breaking' of edges that represent signal space distances between vertices that are too great to represent any meaningful connectivity. At this point, 1 and 2 simplices are identified from the refined coverage matrix, and a nerve is generated.

Figure 17 provides an example of the nerve construction results from the algorithm. The left side of the figure represents an overhead view of the nerve that results (building outlines are in blue) when the refinement algorithm is running. In this case, initial thresholding parameters are set high to remove significant signal, so that the propagation paths between buildings to the lower right are not maintained. The right side of the figure shows the result when the refinement algorithm is not running. Use of the refinement algorithm to more tightly cluster receiver groups results in a much more clear representation of certain building obstructions, whereas the decision not to use the refinement algorithm results in 1 and 2 simplices in the nerve that physically cross over some elements of building obstructions. Figure 18 provides a second example of a nerve, from two different aspect angles.
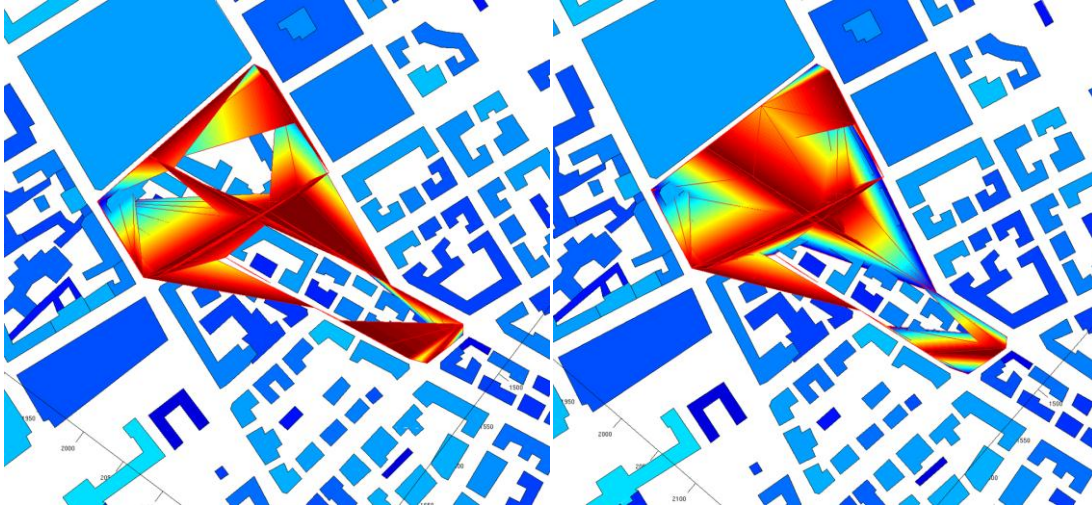
17

Figure 17: Example nerve results with refinement (left) and without (right).

Another example of a nerve using the refinement algorithm, showing two perspectives, is provided in Figure 18.
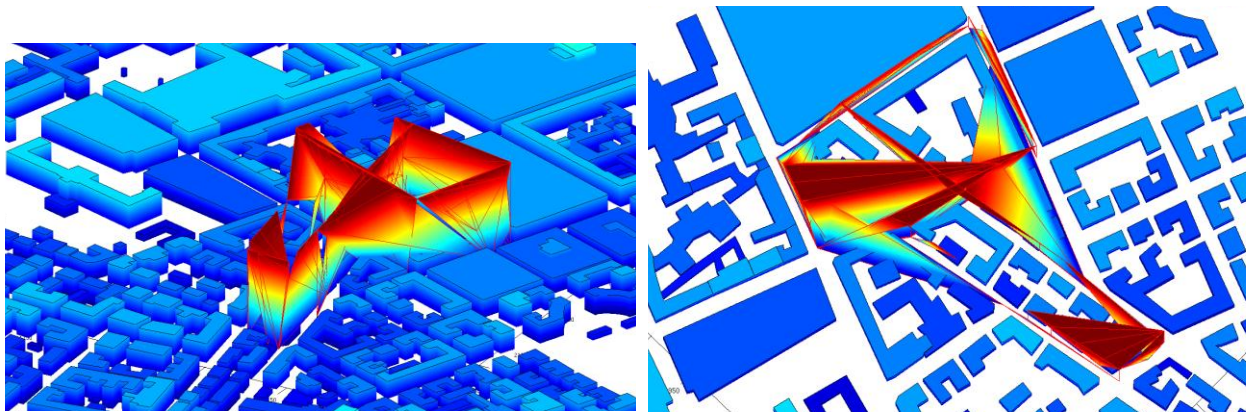

Figure 18: Second example of refined nerve results, with algorithm threshold adjustments.

As is expected, the characteristics of the coverage regions of each transmitter play a significant role in the effectiveness of the homology groups to identify meaningful structure in the RF propagation paths. Several cases are of particular interest, and some specific examples are shown in Figure 19 and Figure 20. These examples help illustrate the role of various algorithm elements described previously in Figure 15. In figure 19, two examples are provided for non-contractible receiver groups. The case on the left illustrates a scenario that might result if initial thresholding (via SNR and 'number of hits' thresholds) is not completed. A single transmitter is shown as a black dot, and receivers that are covered by the transmitter are shown in red. In this case, a large building (or building block) obstruction is circled by a receiver set covered by a single transmitter. Increased thresholding (in the future, likely adapted at a local level) is required to manage this situation and maintain a cover that is contractible. On the right, the case results where receiver coverage at the extent of the coverage region is not continuous, due to non-ideal signal fall-off and realistic detection limits. In this particular case, the intermittent coverage is caused by echo profile variations as the path lengths of different echoes vary along the receiver path. With these echo profile variations comes some variation in the number of

18

echoes observed in a profile – some echoes merge together for short periods of time as the receiver moves and relative path lengths of echoes periodically line up. As the number of observed echoes drops below the 'minimum number of hits' that is required for initial thresholding, this type of intermittent coverage results. As a result of this non-contractible cover, it is likely in this case that extra holes may be inserted in the resulting nerve when it is constructed.
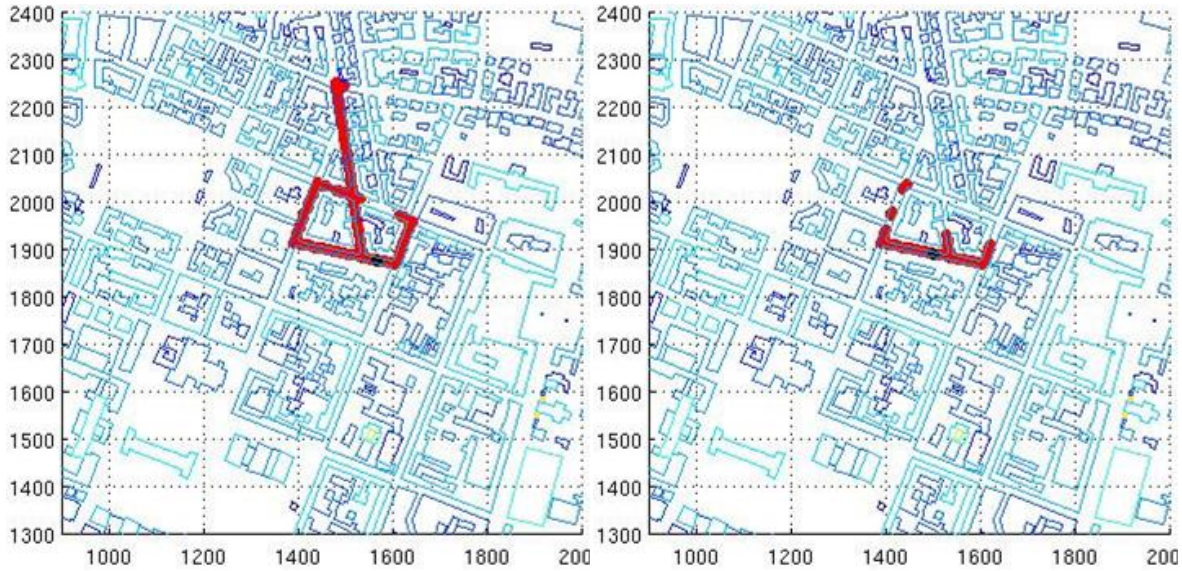

Figure 19: Receiver coverage examples – non-contractible cases, due to imperfect thresholding.

Figure 20 shows two other cases of interest. On the left hand side, a transmitter is located far from the receiver vertices (black circle to the lower left of the figure). As a result, the transmitter illuminates two separate sections of the receivers (the middle section is obstructed by buildings). This is a non-contractible cover that can be controlled by eliminating poorly placed transmitters located at long distances from the receiver sets. On the right is a case where a receiver located in the center of the figure illuminates a large number of receivers. Some receivers are located down different streets, and others are illuminated by RF energy passing between buildings (the receivers to the upper right of the figure). In this case, there are significant differences in the signal space distance metrics between receivers and the cluster refinement algorithm plays a significant role by further clustering the receivers in new 'transmitter groups', separating this relatively large cover region into smaller, more locally similar coverage regions.
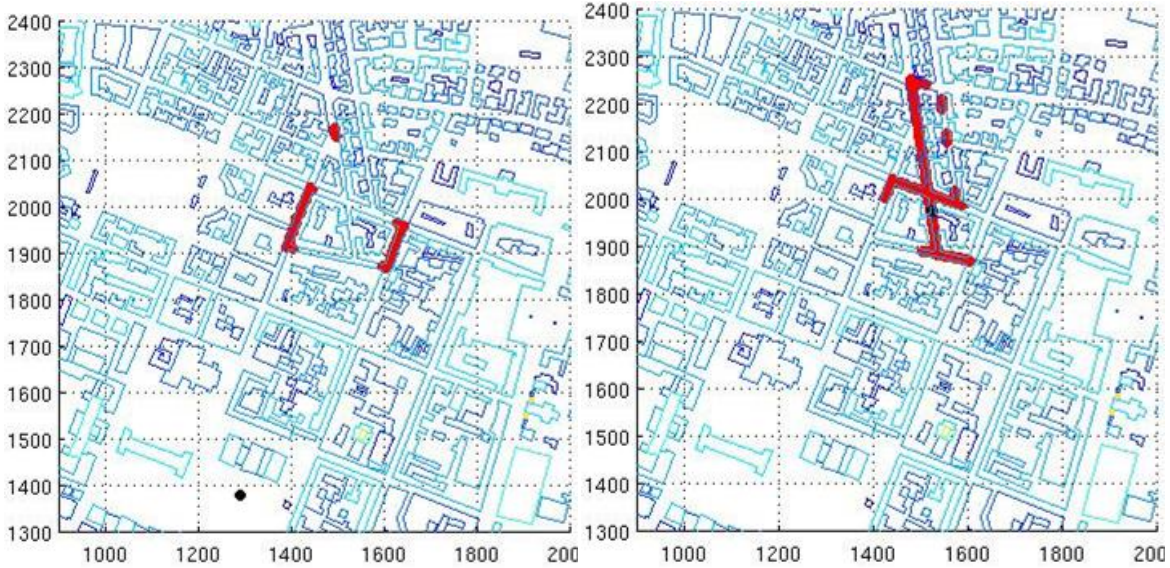
19

Figure 20: Receiver coverage examples – distance transmitter (left) and case requiring refinement (right).

Figure 22 provides an overview of the location of building obstructions that are expected to be identified by the homology extraction algorithm set. Black circles represent buildings that obstruct RF propagation. The larger circles are expected to show up readily, while the smaller circles represent more detailed features of the RF propagation environment.

Figure 21: Overview of Expected Propagation Holes

Figure 22 shows the holes that are retrieved using two sets of threshold parameters. The figure to the left illustrates the case where both initial thresholding and the adjusted distance threshold used in the final nerve construction are relatively low. The figure to the right illustrates the case where initial threshold parameters and later distance thresholds applied directly before nerve construction are lowered. In the first case, much of the detail of individual buildings in the upper portion of the covered region is recovered, as are the more coarse building obstructions to the

20

south. In the second case, much of the detail in the north is lost, and even one major edge to the south east is lost.
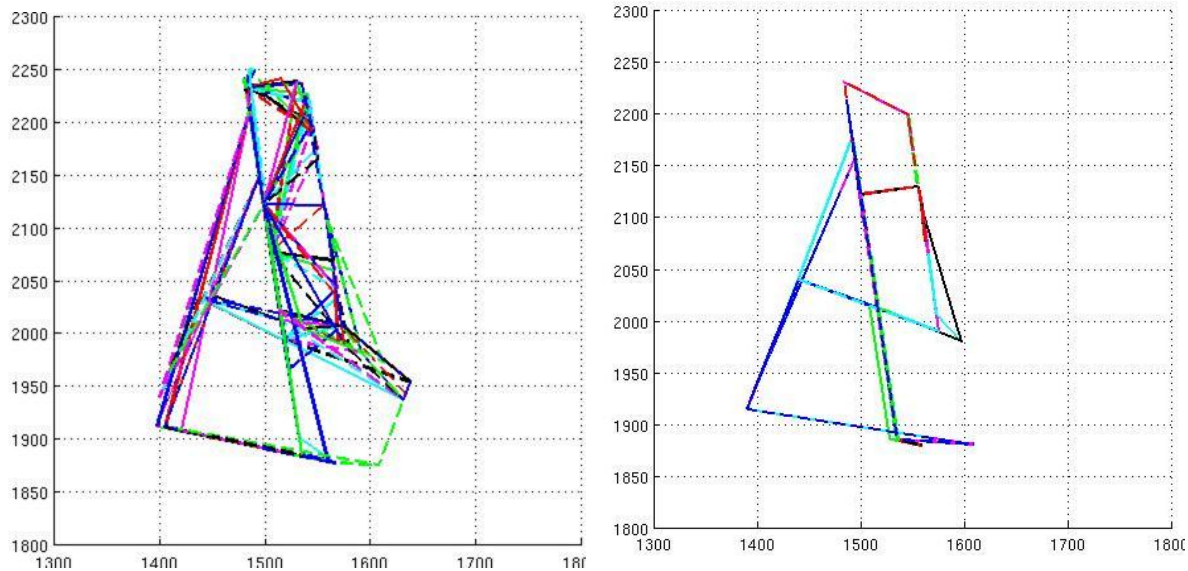

Figure 22: Recovered holes – varying thresholds.

At this point, a detailed study of threshold adjustments versus measurements of holes and connected components (0^{th} and 1^{st} Betti numbers) has not yet been completed. These initial results have some good qualities, in that detailed characteristics of the buildings appear to be recoverable with coarse radio measurements. With absolutely no previous knowledge of transmitter or receiver location and relatively simplistic thresholding rules, key features of the building structure (and associated obstructions) are recovered. However, the number of holes ($B1$) appears to be too large when fine detail is recovered, resulting in the detection of additional obstructions that were not present in the building structure used in the simulation.   Of major concern is the nature of the coverage regions being used to generate the nerve. It is expected that there still exist coverage regions that are not contractible, as illustrated in Figure 19. These will likely insert additional holes where no obstruction actually exists. Additionally, the general rules for local threshold adaptations have not been adequately optimized, so that hole density may vary significantly across an area (e.g. lower right region and upper region of left plot of Figure 22).

Work continues to improve upon conditioning of the coverage regions for each transmitter, and on local adaptation of thresholds used in the topology extraction algorithms.

Once these topology extraction results can be improved upon, the geometric extraction code development can begin.

CONCLUSIONS

At the close of this seedling, we note:

1. The techniques developed, being mathematical in foundation, are more widely applicable than anticipated. The unique channel response theorem applies to sensor modalities far beyond the TOA originally assumed.

2. Sheaf theory provided an unexpectedly useful context for the imaging problems in this project. We anticipate huge possibilities for growth in this class of sheaf-theoretic tools.
3. Our initial assumptions about contractible coverage regions were too optimistic: most realistic data sets seem to violate this assumption. To compensate for this, we learned to resort to a topological preprocessing step (via clustering). It appears that this methodology of preprocessing for topology control is more broadly applicable.

REFERENCES

[1] Recommendation ITU-R P.1411-1, *Propagation data and prediction methods for the planning of short-range outdoor radiocommunication systems and radio local area networks in the frequency range 300 MHz to 100 GHz,* International Telecommunications Union, 2001.

[2] J. Fuhl, A. F. Molisch, E. Bonek, "Unified channel model for mobile radio systems with smart antennas", IEE Proc. - Radar, Sonar navigation, Vol. 145, No. 1, Feb. 1998, pp. 32-41.

[3] European Commission, Directorate General XIII - Information Society: Telecommunications, Markets, Technologies, *COST Action 231: Digital mobile radio towards future generation systems, Final report*, 1999.

[4] O. Landron, M. J. Feuerstein, T. S. Rappaport, "A Comparison of Theoretical and Empirical Reflection Coefficients for Typical Exterior Wall Surfaces in a Mobile Radio Environment", IEEE Transactions on Antennas and Propagation, Vol. 44, No. 3, March 1996, pp. 341-351.

[5] B. Davis et al, "Complex Permittivity of Planar Building Materials Measured with an Ultra-Wideband Free-Field Antenna Measurement System", Journal of Research of the National Institute of Standards and Technology, Vol. 112, No. 1, Jan.-Feb. 2007, pp. 67-73.