

## Visualizing Antenna Design Spaces

Kent Wittenburg, Tom Lanning, Darren Leigh, Kathy Ryall

TR2008-017 June 2008

### Abstract

This paper describes a long-term project exploring advanced visual interfaces for antenna design. MERL developed three successive prototypes that embodied an evolution towards larger scales and more concrete semantics for visualization of large sets of candidate designs and then winnowing them down. We experimented with multidimensional scaling and then collective line graphs before settling on linked scatterplots to visualize performance in a design of up to 10 million antennas at a time. In the end, the scatterplot solution was most successful at balancing intelligibility with visualization of the space as a whole. The design allows for adding more 1D or 2D linked feature visualizations if needed, and it smoothly transitions to other details on demand views for final tweaking.

*AVI 2008*

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.



# Visualizing Antenna Design Spaces

Kent Wittenburg, Tom Lanning, Darren Leigh, Kathy Ryall

Mitsubishi Electric Research Laboratories, Inc.

201 Broadway  
Cambridge, MA 02139

USA  
1-617-621-7500

wittenburg@merl.com, tom.lanning@att.net, leigh@merl.com, ryall@acm.org

## ABSTRACT

This paper describes a long-term project exploring advanced visual interfaces for antenna design. MERL developed three successive prototypes that embodied an evolution towards larger scales and more concrete semantics for visualization of large sets of candidate designs and then winnowing them down. We experimented with multidimensional scaling and then collective line graphs before settling on linked scatterplots to visualize performance in a design space of up to 10 million antennas at a time. In the end, the scatterplot solution was most successful at balancing intelligibility with visualization of the space as a whole. The design allows for adding more 1D or 2D linked feature visualizations if needed, and it smoothly transitions to other "details on demand" views for final tweaking.

## Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation]: User Interfaces - Graphical User Interfaces; I.3.6 [Computer Graphics]: Methodology and Techniques - Interaction Techniques; J2 [Computer Applications]: Physical Sciences and Engineering

## General Terms

Measurement, Design, Human Factors

## Keywords

Antenna Design, Information Visualization, Line Graphs, Human-Guided Search, Multivariate Visualization.

## 1. INTRODUCTION

This paper describes a long-term project exploring advanced visual interfaces for antenna design. Over five years ago, Mitsubishi Electric Research Laboratories (MERL) began a collaboration with Mitsubishi Electric Corp. Information Technology R&D

Center, one of two main corporate R&D centers in Japan, and Mitsubishi Electric Corp. Electronics Systems Division, a business unit that designs antenna systems for commercial and Japanese government applications. End products range from satellites to aircraft communication systems to automobile collision avoidance radar to RFID readers and cell phones. Ongoing to this day, the collaboration was initially triggered by MERL prototypes in the area of Human-Guided Search, where we and others before us had explored how humans can solve a large and complex optimization task through visualization and interaction with an evolving search space [1][3].

This paper will focus on a subset of the prototypes in a discussion of lessons learned as we sought to bring our expertise as visual interface researchers to the problem of antenna design. The individual systems have been described in brief in other publications, mostly in IEEE antenna engineering conferences [12][13][14][17]. Details regarding our contribution to antenna engineering may be found there. Our purpose here is to examine the overall project from the perspective of information visualization and human-computer interaction methods. Our conclusions will be based on qualitative evaluations gleaned from our collaborators and users who carry out the actual tasks involved in antenna design in Japan. The typical pattern of our work on this project was comprised of an annual fiscal year cycle in which initial plans were made in the April timeframe, a prototype was created and shown to our customers and collaborators in the late summer, and then based on feedback from this prototype, a final prototype was created and delivered to Japan by the end of the fiscal year in March. Our systems were intended for internal company use by antenna design engineers, and our most recent tool is in use today.

The project began with an experiment in designing classic Yagi-Uda wire antennas, and over time progressed to more sophisticated types of phased array antennas. (For background information, see, e.g., Wikipedia entries on these topics.) We believe the lessons we have learned about interactive visualization has application not only to antenna design but to multivariate design problems generally. Not surprisingly, we found that it is important to apply different visualization methods to different sized sets to visualize the performance of antenna designs under consideration [2]; the larger task is to select one or perhaps a few candidate designs from what is initially a very large (perhaps infinite) set, trying to maximize certain performance goals while minimizing costs. In the end, we found it most valuable to visualize the performance of very large sets of candidate antenna designs through a pixel-based technique--parallel linked scatterplots of domain spe-

cific performance metrics. Through brushing and querying within parallel scatterplots, users can prune the space down to a hundred or so. At that point a table view of numerical performance values is useful in which standard sorting operations are possible. For viewing a single antenna from the table, we provide a visualization of the physical antenna array elements themselves and the corresponding 2D radiation patterns. Individual designs can be cloned and certain physical parameters tweaked and visually compared to the original to arrive at a final design.

On the way to the current system design we experimented with other information visualization techniques that we will discuss here, beginning with dimensionality reduction employed in our first prototype inspired by Human-Guided Search. In our second prototype we focused on line graphs for visualization and (soft) querying. Our third prototype, the main topic of this paper, illustrates the use of scatterplots in groups of three, each of which reveals a face of a cube representing a 3D design space of high-value performance measures.

## 2. RELATED WORK

The largest body of previous work relating visualization to the antenna design process is characterized by tools that visualize the radiation pattern of a simulated antenna in 3 or perhaps 4 dimensions [5]. Commercial and open source software is available today that in some cases is also compatible with MatLab and with the open source program NEC2 used for simulating the performance of the most common types of antennas (<http://www.nec2.org>). Visualizing the radiation patterns of individual antennas is undoubtedly important for the final stages in a design process. It is also valuable as a first step when the design process begins with a known example of a successful antenna design and then adapts it to the requirements of a new application.

However, as pointed out in [15], an antenna design process that begins only with known designs is highly limiting. The effects of even small variations on a myriad of parameters can have unexpected results--positive or negative--on the performance of an antenna. There is a huge number of specification variations possible for even basic antennas, and thus finding the best design may require looking into unexpected territory that can be reached only by considering millions or billions of combinations. Not surprisingly, genetic algorithms have thus become an important thread of research in antenna design [7][15]. The art and science of utilizing genetic algorithms for antenna design depends on careful design of an objective function for optimization.

However, as with all fully automatic systems for optimization, expert users may be left with the sense that they would like to look beneath the hood. As articulated by Spence and his collaborators, complex engineering design problems require experience and human judgment [2][19][20]. They have proposed methods to be used in digital circuit design, among other application domains, that make use of dynamic histograms and the "Prosection Matrix," a scatterplot-based visualisation of performance targets against input parameters. Our research shared the goal of providing visualization tools to yield insight in the design process. A goal is to allow the human to make connections between input parameters and performance results in a design space that can be only partially sampled. Perhaps one difference in the domain of antenna design from those that Spence et al. explored is that the performance goals for antenna design are devilishly hard to specify while performance characteristics for common types of antennas

are relatively cheap and easy to simulate. The opposite seems to be true in the domain of digital circuit design [20].

## 3. METHODOLOGY

Conclusions in this paper are based on qualitative evaluations gathered from our collaborators and users of our system--antenna designers and engineers in Japan. We collected input for designs and feedback on prototypes in a number of ways, including attendance at an annual antenna meeting in Japan, e-mail exchanges throughout the year with our collaborators, face-to-face planning meetings both in Japan and the US, and interviews with other visitors to our lab with antenna domain expertise. The email exchanges were most often used for clarifications on points raised in person; visitor interviews usually confirmed feedback given at the annual meetings and during the face-to-face planning meetings. More detail on our annual meeting follows.

### 3.1 Annual Meetings

Each year one of the authors attended an annual antenna meeting in the late summer, presenting and demonstrating that year's prototype. The audience typically had 50-75 people, all antenna designers and engineers from different parts of our parent company's organization around Japan. After a technical presentation and brief demonstration of the prototype, the audience was free to ask questions about the approach and try out the prototype. Presentation and discussions were conducted mostly in English. As the prototypes evolved, the amount of interest and interaction with the audience grew. Once the third prototype was shown, there was great enthusiasm and requests to use the tool on the spot.

### 3.2 Observational Study

From the beginning of our collaboration, we had asked if we could meet directly with antenna designers, in particular to observe their current practice. While our direct collaborators were antenna engineers themselves, at this point in their careers, they were primarily scientists and managers and less involved with hands-on design.

After the presentation of our second prototype, we were eventually able to arrange a visit to observe a designer in practice. One of the authors spent an afternoon with an antenna designer in his workplace in Japan. We discovered that his process was manually intensive. He would use batch processing to generate a large number of antenna designs and then use COTS or open-source tools to review the results. There was a lot of manual editing of files and typing at the command line to launch multiple instances of the same application. The designer usually reviewed several candidate antenna designs in parallel, opening perhaps twelve graphs at a time in a four-by-three configuration on the screen. He would then pick two at a time and do a more in-depth side-by-side comparison that included looking at statistics in a spreadsheet.

As he explained, the work proceeded with iteratively re-running simulations in batch mode overnight, having tweaked the input parameters based on what he had seen. It became clear that our customers were functionally and numerically oriented. We realized that while advanced visual methods would be helpful to the designers in evaluating candidate antennae, numerical precision had to be maintained, and the tools needed to be compatible with their current practices.

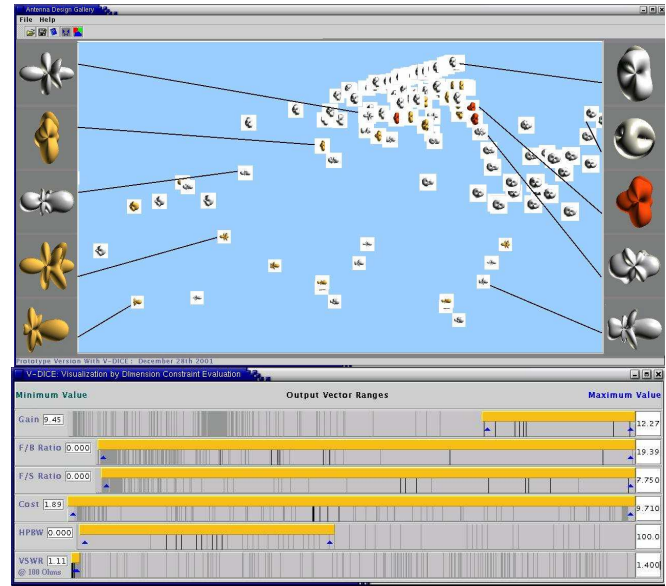
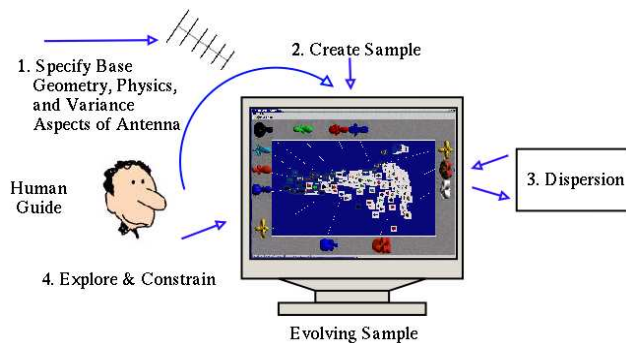


Figure 1: The design cycle shown at left was supported in our first prototype. On the right are the two primary visualization panes for the evolving candidate sample set: the top pane shows multidimensional scaling; the bottom, linear attribute widgets.

#### 4. FIRST PROTOTYPE: ABSTRACT VISUALIZATION OF AN EVOLVING SEARCH SPACE

Our first prototype [17] applied human-guided search [1] [3] to the process of exploring and repeatedly refining a search through a very large space of possible designs for Yagi-Uda antennas. Yagi-Uda antennas are a classic type of directional wire antenna such as one sees with conventional rooftop TV antennas. Among the design parameters are the number, length, and spacing of the elements in the antenna array.

A primary focus of this initial stage of our project was to explore the idea of designing antennas through achieving maximal dispersion in a sampling of a very large design space that could not be enumerated exhaustively. Maximal dispersion here refers to a property of a set such that the set members are maximally distant from one another in a multidimensional space. Such an approach had been applied earlier in a tool for graphics and animation design [16] as well as scheduling [1]. An overview of the design process that the system was designed to support is shown at the left of Figure 1, described in [17]. After the user set some initial ranges over which the design parameters could be varied, the system would simulate a large set of antennas and then select candidates for human inspection based on their maximal dispersion in a multidimensional evaluation space. The multidimensional space was defined through weighted vectors of performance values. As the user honed in on a subset of the design space through interaction with the visualization tool, the system would again try to achieve maximal difference among potential candidates in the chosen subspace. At any time a user could inspect an individual antenna design candidate by visualizing its radiation pattern with an open source 3D visualization module (shown in thumbnails in the margins of the right side of Figure 1) or by inspecting the actual numbers of its performance simulation.

There were two methods deployed for visualizing the collective set of candidate antenna designs. As shown in the right side of Figure 1 in the main area of the top pane, dimensionality reduction was used in a layout in which icons representing individual candidate antennas were projected onto a 2D plane. Their relative distance from one another was intended to reveal distance in the design space, i.e., the Euclidean distance between weighted multidimensional performance vectors. The layout algorithm, multidimensional scaling [11], attempted to find a positioning such that the distances on the 2D plane best correlated with the relative Euclidean distances in the performance space [16]. The overall goal was to reveal clusters of similar designs as well as outliers through this 2D layout.

The second method, shown in the bottom pane, visualized multidimensional performance attributes on widgets representing linear scales. Each of the parallel widgets revealed the distribution of a performance attribute across the entire set. The attribute values all fit a linear scale whose minimum and maximum corresponded to the actual performance range globally. The numbers were indicated in text boxes at the beginning and end of each attribute widget. Each candidate antenna was visible along each dimension as a vertical line. Selections could be controlled by sliders that would set minimum and maximum values. As subsets of candidates were selected by restricting the value range with one widget, the candidates would be visually highlighted in the other widgets and in the 2D layout above. The attribute widgets are a basic form of one-dimensional data visualization sliders [4] that incorporate brushing techniques [21].

What were the lessons learned from this first prototype? First, multidimensional scaling as a visualization technique was not well received by our users. The reduction of a multidimensional performance space to 2D did not help the designers understand what they were looking at. We were told that they couldn't make sense of this visualization, that it would be better to have a series of 2D views with axes whose semantics were clear. The parallel attribute

visualization sliders were received more positively since they indeed did have an understandable semantics.

Second, our notion of dispersion as a search mechanism got a lukewarm reception, probably because, again, its effects were not transparent to these antenna experts. We came to the conclusion that as the project moved to its next phase, it would be best to decouple the exploration of the design space from the visualization and filtering of the results. We would postpone our original goal of supporting an end-to-end human-guided search tool and instead focus on interactive visualization and winnowing of a large design set. We introduced new modularity in the system that allowed candidate sets to be generated independently by our clients or ourselves. The sets could then be loaded into a tool for visualization and filtering.

## 5. SECOND PROTOTYPE: VISUALIZATION AND FILTERING THROUGH QUERY LINES

The next phase of our project took on the problem of designing certain types of phased array antennas [6]. Another piece of feedback we had gotten from the first round of the project was that Yagi-Uda antenna design was not enough of a challenge for this group of antenna experts. As with Yagi-Uda antennas, phased array antennas are also directional. They are distinguished by varying the phases of the signals feeding the elements of the array such that a desired radiation pattern is achieved. Our approach was to utilize 2D line graphs both for visualizing the performance of candidate designs and for querying and filtering the candidates [13][18]. The most important of these graphs, examples of which can be seen in Figure 2 (b-e), are the radiation patterns. The x-axis is observation angles (degrees) and the y-axis is the array factor directivity. In general, designers are looking for a gain peak in the center with minimal energy in the off angles. At least for linear arrays, a 2D radiation graph is a good indicator of the primary performance design goal.

As mentioned, in this phase of the project we assumed that a set of candidate designs would be generated independently. For testing purposes we were able to generate an exhaustive set of possible candidates for phased array antennas with some simplifying assumptions. We assumed that the arrays were uniformly linear and looked at variants of phase-only synthesis. The generator took as input the number of elements in the array and a set of quantized values for phase and amplitude coefficients. It then exhaustively enumerated all possible combinations of excitation parameters to compute a set of candidate designs. We were able to load in on the order of 10,000 design variants at a time.

From an information visualization perspective, naively plotting the radiation patterns of all the generated designs in a large candidate set would result in an undifferentiable blob (Figure 2(b)). Line graphs were not a solution for visualising the space as a whole. However, our hypothesis was that it would be desirable for designers to explore the design space by filtering with queries that could be created directly with and on 2D line plots. Visual querying with 2D line graphs has been tackled before [8], but it is generally not straightforward to specify 2D constraints on line graphs. Approximate matching is even more of a requirement than with conventional Boolean queries since lines are highly unlikely to find exact matches. As with querying generally, result lists characterized only by hard matches do not reveal anything about the set of candidates that almost matched and very little about the space

of solutions as a whole. R. Spence has articulated the need for information visualization systems to reveal *sensitivity information* that can help guide users to explore parts of the design space that they hadn't previously considered [19].

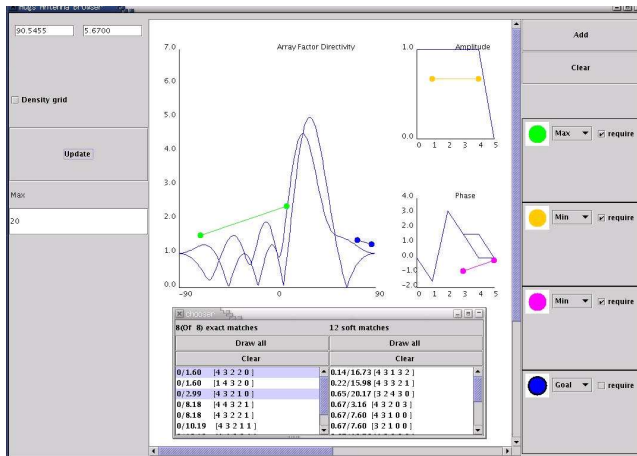
The main contribution of this work was in developing a set of approximate 2D graph-based query methods that could reveal sensitivity information. We will touch on only the main features of the system here. Figure 2(a) shows a screenshot of the overall system. Three types of linked performance graphs are shown in the main screen, the most important of which is the radiation pattern, but any of the graphs can be the basis of querying. In the embedded window at the bottom of Figure 2(a) is a set of results of a previous query in which the list on the left represents hard matches for the expressed constraints and the list on the right represents an ordering of soft matches, i.e., matches that are close to the constraints expressed by the query lines but do not fall strictly within them. Figure 2 (c-e) shows examples of different types of query specifications and their resulting matches. Figure 2(c) shows query lines that represent minimum and maximum constraints over the (x,y) plots shown and a set of "hard match" patterns that fall within those constraints. Figure 2(d) shows two soft matches, i.e., patterns that do not fully meet the specifications of the min./max query lines but are nevertheless close to the constraints. Figure 2(e) shows a different type of query line--a goal or preference. The contribution of a goal or preference query line to the results returned by the system is to sort the hard and soft matches on the basis of similarity to this 2D preference pattern. If no other constraints are given, a goal query line amounts to a query by example and all results are soft matches.

From one perspective, the second round of prototyping for antenna design represented an extension of the 1-D attribute visualization widgets of the first round to a 2D approach. Selecting and filtering with 1D attributes is straightforward. Selecting and filtering with 2D line patterns is a harder problem that required inventing these new approximate matching methods and interactions.


The second prototype drew more interest than the previous one -- both because of the more realistic problem domain and the more intuitive visualization techniques; there were many requests to try different queries using the prototype. While there was some indication that the query mechanism itself might be too complicated for engineers to use, the bigger question (and excitement) was whether the approach would scale to more complicated antennas and problems of larger size. Given the success with this small (constrained, yet realistic) problem specification, we were again directed to move onto a more challenging and larger problem task. It seemed clear to us that a different approach to visualizing larger sets globally was needed.

## 6. ROUND THREE: A PIXEL-BASED APPROACH WITH LINKED SCATTERPLOTS


In our next round [14], we were given a particular optimization problem to focus on, which was as follows. Our goal was to maximize performance of sparse linear array antennas with uniform excitation, i.e., the assumption of linear spacing of the phased array elements in our previous round was relaxed but the constraint of uniform signal excitation across the array elements was fixed. The number of array elements and their spacing were the variables to explore. With a supercomputer cluster, we were able




(a)



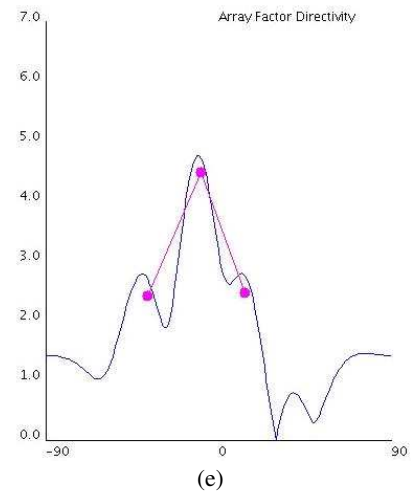
(b)



(c)



(d)



(e)

**Figure 2: The QueryLines System: (a) a snapshot of the overall system, (b) a large set of graphs, (c) min and max hard constraints and several matching results, (d) two soft matches that do not fall within the hard constraints, (e) a goal query and a result that is the closest match.**

to generate and simulate on the order of 1 billion variants of such antennas.

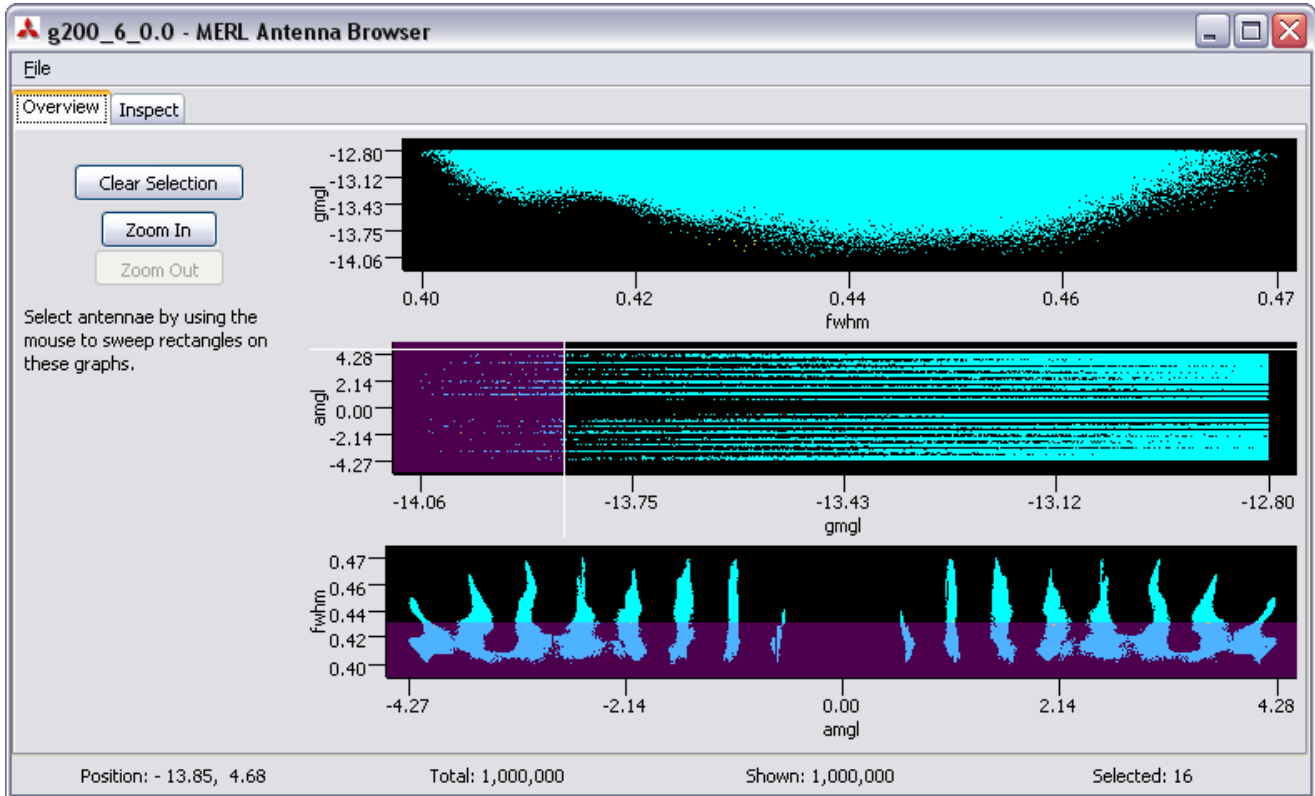
Our goal for the visualization tool was to handle up to 1 million candidates at a time and to develop interaction methods for exploring and filtering that would respond almost instantaneously. It should come as no surprise that our solution for visualization of such a large set of elements utilized pixel-based techniques [9]. Although the most important antenna performance design goal is best visualized as a line graph (or 3D plot) of the radiation pattern; neither of these techniques are appropriate for viewing in large sets. It was more efficient for machine computation and human interaction to use simple performance numbers that would characterize the radiation pattern indirectly. These were as follows:

- The width of the main lobe (full-width half maximum).
- The gain of the highest side lobe.
- The angle of the highest side lobe.

These performance measures are suitable for visualization with standard 1-D widgets, but we explored a variation in which these three dimensions defined the dimensions of a cube. We could then plot three faces of the cube as linked scatterplots; 1 million design variants are shown in Figure 3. The main interaction method is to sweep out selections of pixels in any of the scatterplots, which will be painted in all of the scatterplots. At any time a user can reduce the set (zoom in) by filtering to the current selection.

In order to meet the requirement for quick response time, we came up with a code design that utilized the resolution on the screen to organize the data. Each time there is a screen resize, a one-time process sorts the data into bins, one for each pixel. When rendering, the pixel is lit if it contains any data. The result is that a rendering of a scatterplot with one unit per pixel, can happen within a second on most standard desktop or laptop computers.

The striking striations visible in Figure 3 are an example of unexpected results that may be revealed by a visualization tool such as ours. The antenna experts we consulted are not sure why these



**Figure 3: The overview pane of the latest version of our visualization tool (shown an small window size for readability). Three linked scatterplots are views into a 3D space of performance measures. The current selections are updated in all views.**

patterns emerged, indicating a non-uniform distribution of the shape of the main lobe across the angles of the highest side lobes. All of the antenna variants in this particular example set had the same number of antenna elements in the array. But clearly the possible positioning of these elements left gaps in the distribution of main lobe energy with respect to highest side lobe angles.

An example based on observing how an antenna expert used the tool follows. The expert first swept out the lower region of the bottom scatterplot, the results of which are (subtly) visible in Figure 3. These antennas would have the narrowest main lobes (indicated on the y-axis), a measure of high directionality irrespective of the angle of the highest side lobe (indicated by the x-axis). The expert interactively played with the maximum setting on the y-axis in order that some selections appeared at the left side of the middle scatterplot, an area of sparse distribution. This area contained antennas whose highest side lobe has low gain, irrespective of its angle. Again, in general, designers are looking for high energy in the main lobe with minimal energy in the side lobes. Then the expert swept out the rectangle in the left area in this middle pane. This further constrained the selection set to those antennas with the desired properties. From a million antennas, the expert was able quickly to narrow down the set to a size of 16 or so, which he then looked at more closely in the Inspect pane, where line graphs of radiation patterns and the performance numbers themselves were visible.

A screen shot of the Inspect pane is shown in Figure 4. The upper part of the pane contains a table of the selected antenna

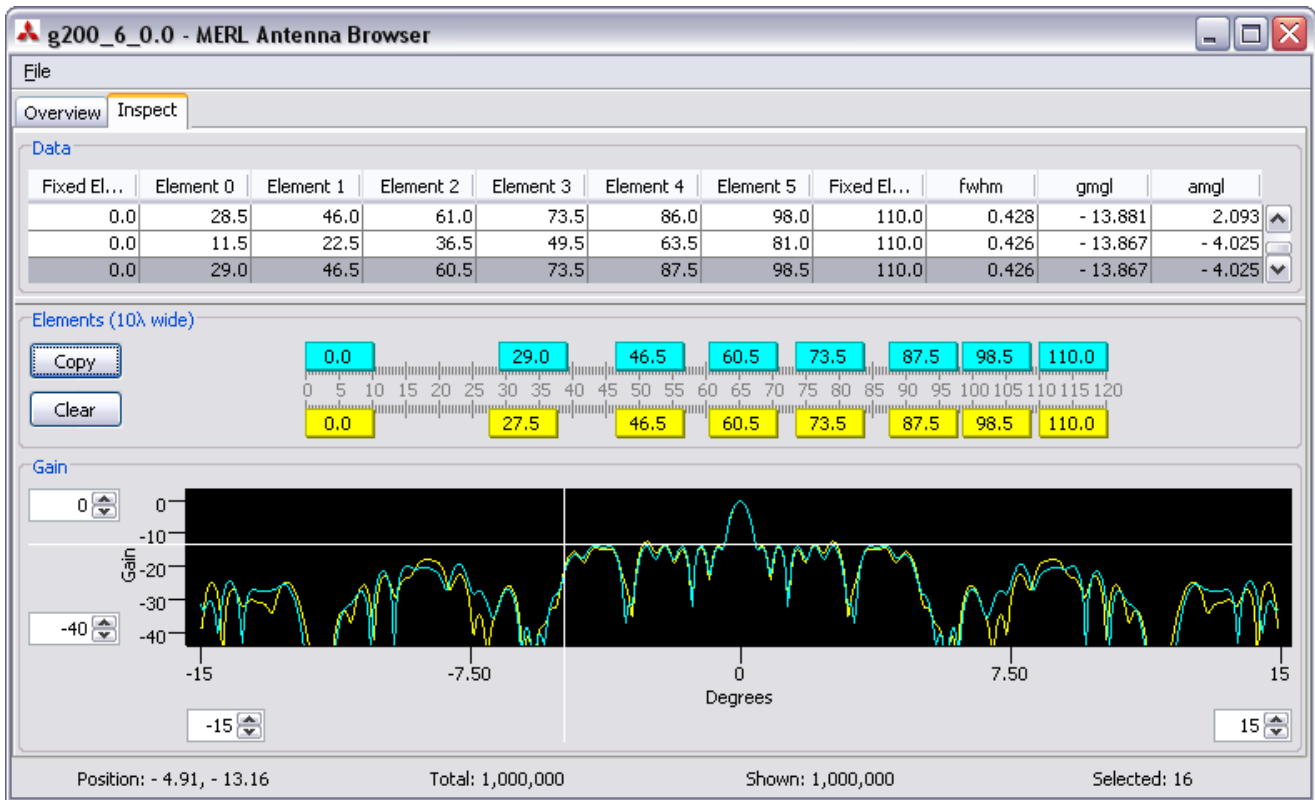
design candidates. The columns contain the numbers for the position of each included element as well as the three performance measures mentioned above. The table rows may be sorted on the basis of any of the columns in the usual way. Such a table method is useful and usable when there are no more than a few hundred design variants under consideration. A common interaction pattern we noted is for users to sort the antenna units along one column and then hold down an arrow key to traverse the list from top to bottom, causing a rapid serial visual presentation (RSVP) of the 2D gain pattern [19].

For viewing details of an individual design, a user may select a row in the table, an example of which is visible in the lower part of Figure 4. The graphic in the middle of the pane represents the physical position of the array elements and the line graph at the bottom is the radiation pattern. An individual antenna (shown in blue) may be copied (shown in yellow). The position of certain array elements in the copy may be interactively moved and the radiation pattern of the copy compared graphically to the original.

This visualization tool contributed to the finding published in [14] that it was possible to achieve essentially the optimal performance of uniformly spaced arrays with fewer elements (thus less cost) spaced non-uniformly in certain configurations. The tool can of course be used to explore other kinds of design issues as long as the data basically conforms to the patterns shown here.

With this third prototype we were pleasantly surprised to have several audience members at our annual design meeting ask if they could try the tool on the spot; at previous meetings they most





**Figure 4:** The inspect pane of the latest version of our antenna visualization tool (shown at small window size for readability). The table at top shows performance numbers for the selected antennas; the graphic in the middle shows individual elements and their positioning; the bottom shows a radiation pattern and its copy, which may be tweaked.

often made their request verbally and the presenter (one of the authors) would run a query in the prototype. There was also much discussion and side conversation in Japanese. This prototype was extended and deployed later that year and is in use today.

## 7. LIMITATIONS OF THIS STUDY

We see two main scientific limitations to our study. First, due to our clients' request we changed the problem specification (task) every year. As the antenna types shifted over the course of the study we were not easily able to compare prototypes head-to-head on the same set of antenna designs. Thus it was difficult to tease apart which changes in our design were due to the change in task and which were a result of improvements to the interaction and UI design. In a more controlled study we would have fixed the task, or perhaps re-run some of the later prototypes on the earlier data sets. Unfortunately, neither of these approaches was feasible logistically. We do believe that the visualizations used in all three prototypes would be appropriate for all antenna types with the exception of the antenna element visualization in prototype three, which would require modification for Yagi-Uda antennas.

A second limitation of our study is that our evaluation methodology was flawed. In a real-world setting it is difficult to do a comprehensive user study. In a more controlled setting, it is easier to design and evaluate systems and recruit participants. We faced some additional challenges due to geographical distance as well as differences in language and domain expertise (Antenna Engineers vs. Computer Scientists). In the end, we believe it is the latter that had the largest impact. It was sometimes difficult to get buy-in on our design process and requests for feedback. For example, convincing our colleagues that we would benefit from directly meet-

ing with and observing practicing antenna designers was a lengthy process. In a perfect world, we would increase the scientific rigor of our work by introducing surveys, increasing sample sizes, and retesting across the prototypes with a constant dataset.

## 8. CONCLUSION

In this paper, we have presented a longitudinal design study that resulted in what we believe to be a successful interactive visualization tool for antenna designers in a Japanese industrial setting. The primary visualization methods employed were parallel scatterplots, sortable tables, and 2d line graphs--not novel in themselves, but we believe novel in their application to this domain.

Dimensionality reduction, a popular technique among visualization researchers, was not successful in the eyes of our users. In order for a global visualization of a very large set to make sense, it is important that the semantics of the visualization be concrete. One can see from the example of use described in Section 6 that an expert would know how to filter a large space and understand the patterns if the data is presented with dimensions easily related to the task at hand. Dimensionality reduction, useful in many ways, may not be so useful if the first concern is concrete performance numbers.

Our experience with QueryLines showed that, although it is attractive to consider querying directly with line graphs representing radiation plots, there is complexity in specifying such queries as a set of constraints. Our conclusion was also that line graphs would not scale well in an overview mode. In retrospect, we should note that the methods in Line Graph Explorer [10], which we were not aware of at the time of this work, do provide some ability to scale

up to larger sets. However, in order to achieve the scale of, say, 10 million designs in a single view, each graph displayed using Line Graph Explorer would have to represent an aggregation on the order of 10,000 radiation patterns since one line of the display is needed per graph and there are order 1000 horizontal pixel rows available on desktop displays. We don't know whether the computational demands of such an approach would be able to offer suitably rapid response or how it might be received by antenna designers, but it may be worth a look.

## 9. FUTURE WORK

Since the development of the visualization tool described here, the project has returned to the problem of algorithms for enumerating the search space [12]. We have also extended the Inspect pane of the visualization tool to handle circular antenna arrays. In the future, we imagine that it may be useful to consider radiation plots in 3D rather than 2D as more complex types of antennas come within the scope of the tool. In such cases, we suppose that again we will need to come up with numerical functions that can be visualized with scatterplots and other types of easily understandable parallel widgets to handle large design spaces.

Also, one question we are left with is how our use of three 2D scatterplots to represent a 3D design space might compare to actually rendering the scatterplot itself in 3D. (See, e.g., [21].) We will have to leave the answer to that question to future research.

## 10. ACKNOWLEDGMENTS

A large number of people have contributed to the work described in this paper. They include, among our Japan colleagues, I. Chiba, Y. Hara, K. Hirata, Y. Konishi, S. Makino, H. Miyashita, and T. Sakura. On the MERL side, besides the authors, contributors include C. Lee, N. Lesh, J. Marks, and A. Quigley.

## 11. REFERENCES

- [1] Anderson, D., Anderson, E., Lesh, N.B., Marks, J.W., Mirtich, B., Ratajczak, D., and Ryall, K. 2000. Human-Guided Simple Search. In Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence (August 2000), 209-216.
- [2] Apperley, M., Spence, R., and Wittenburg, K. 2001. Selecting One from Many: The Development of a Scalable Visualization Tool. In Proceedings of IEEE Symposium on Human-Centric Computing Languages and Environments (Stresa, Italy, Sept. 5-7, 2001). HCC '01. IEEE Computer Society. 366-372.
- [3] Colgan, C., Spence, R., and Rankin, P. 1995. The cockpit metaphor. *Behavior and Information Technology* 14, 4, 251-263.
- [4] Eick, S. Data Visualization Sliders. 1994. In Proceedings of ACM Symposium on User Interface Software and Technology (Marina Del Rey, California, USA, Nov. 2-4, 1994). UIST '94. ACM Press, New York, NY, 119-120.
- [5] Chakrabarti, S., Wong, J. C., Gogineni, S. and Cho, S. 1990. Visualizing Radiation Patterns of Antennas. *IEEE Computer Graphics and Applications* (January 1990), 41-49.
- [6] Hansen, R. C. ed., 1985. *Microwave Scanning Antennas*. Peninsula Publishing, Los Altos, CA, 1985.
- [7] Haupt, R. L., and Werner, D. H. 2007. *Genetic Algorithms in Electromagnetics*. Wiley-IEEE Press, Hoboken, NJ, 2007.
- [8] Hochheiser, H., Shneiderman, B. 2004. Dynamic Query Tools for Time Series Data Sets, Timebox Widgets for Interactive Exploration. *Information Visualization* 3, 1 (March 2004), 1-18.
- [9] Keim, D. A. 2000. Designing Pixel-Oriented Visualization Techniques: Theory and Applications. *IEEE Transactions on Visualization and Computer Graphics* 6, 1 (January-March 2000), 1-20.
- [10] Kincaid, R., and Lam, H. 2006. Line Graph Explorer: Scalable Display of Line Graphs Using Focus+Context. In Proceedings of Advanced Visual Interfaces (Venice, Italy, May 23-26, 2006). AVI 06. ACM, 404-411.
- [11] Kruskal, J. B., and Wish, M. 1978. *Multidimensional Scaling*. Sage Publications, 1978.
- [12] Lee, C., Leigh, D., Ryall, K., Miyashita, H., and Hirata, K. 2006. Very Fast Subarray Position Calculation for Minimizing Sidelobes in Sparse Linear Phased Arrays. In Proceedings of European Conference on Antennas and Propagation (Nice, France, November 6-10, 2006). EuCAP 2006. ESA Publications Division, Noordwijk, The Netherlands.
- [13] Leigh, D., Lanning, T., Lesh, N., and Ryall, K. 2004. Exhaustive Generation and Visual Browsing for Radiation Patterns of Linear Array Antennas. In Proceedings of International Symposium on Antennas and Propagation (Sendai, Japan, August 17-24, 2004). ISAP.
- [14] Leigh, D., Ryall, K., Lanning, T., Lesh, N., Miyashita, H., Hirata, K., Hara, Y., and Sakura, T. 2005. Sidelobe Minimization of Uniformly-Excited Sparse Linear Arrays using Exhaustive Search and Visual Browsing. In Proceedings of IEEE Antennas and Propagation Society International Symposium (Washington, D.C., July 3-8, 2005), Vol. 1B, IEEE, 763-766.
- [15] D. S. Linden, D. S., and E. E. Altshuler, E. E. 1999. Evolving Wire Antennas Using Genetic Algorithms: A Review. In Proceedings of the First NASA/DoD Workshop on Evolvable Hardware (Pasadena, CA, USA, July 19-21, 1999). IEEE, 225-232.
- [16] Marks, J., et al. 1997. Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation. In Proceedings of International Conference on Computer Graphics and Interactive Techniques (Los Angeles, CA, USA, August 1997). SIGGRAPH '97. ACM, 389-400.
- [17] Quigley, A., Leigh, D. L., Lesh, N. B., Marks, J. W., Ryall, K., and Wittenburg, K. B. 2002. Semi-Automatic Antenna Design via Sampling and Visualization. In Proceedings of IEEE Antennas and Propagation Society International Symposium (San Antonio, TX, USA, June 16-21, 2002), Vol. 2, IEEE, 342-345.
- [18] Ryall, K., N. Lesh, N., T. Lanning, T., D. Leigh, D. H. Miyashita, H., S. Makino, S. 2005. QueryLines: Approximate Query for Visual Browsing. In ACM Conference on Human Factors in Computing Systems (Portland, OR, USA, April 2-7, 2005). CHI 2005. Extended Abstracts, ACM, 1765-1768.
- [19] Spence, R. 2007. *Information Visualization: Design for Interaction*, Second Edition, ACM Press, 2007.
- [20] Spence, R. 1999. The Facilitation of Insight for Analog Design. *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, 46, 5 (May 1999), 540-548.
- [21] Swayne, D. F., Cook, D., and Buja, A. 1998. XGobi: Interactive Dynamic Data Visualization in the X Window System. *Journal of Computational and Graphical Statistics*, 7, 1 (March 1998), 113-130.