

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

**SPECTRAL ANALYSIS AND SONIFICATION OF
SIMULATION DATA GENERATED IN A FREQUENCY
DOMAIN EXPERIMENT**

by

Hsin-Fu Wu

September 2002

Thesis Co-Advisors:

Susan M. Sanchez
Russell D. Shilling
Paul J. Sanchez

Second Reader:

**This thesis is done in cooperation with the MOVES Institute.
Approved for public release; distribution is unlimited.**

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY		2. REPORT DATE September 2002	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE: Spectral Analysis and Sonification of Simulation Data Generated in a Frequency Domain Experiment			5. FUNDING NUMBERS	
6. AUTHOR(S) Hsin-Fu Wu				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) In this thesis, we evaluate the frequency domain approach for data farming and assess the possibility of analyzing complex data sets using data sonification. Data farming applies agent-based models and simulation, computing power, and data analysis and visualization technologies to help answer complex questions in military operations. Sonification is the use of data to generate sound for analysis. We apply a frequency domain experiment (FDE) to a combat simulation and analyze the output data set using spectral analysis. We compare the results from our FDE with those obtained using another experimental design on the same combat scenario. Our results confirm and complement the earlier findings. We then develop an auditory display that uses data sonification to represent the simulation output data set with sound. We consider the simulation results from the FDE as a waveshaping function and generate sounds using sonification software. We characterize the sonified data by their noise, signal, and volume. Qualitatively, the sonified data match the corresponding spectra from the FDE. Therefore, we demonstrate the feasibility of representing simulation data from the FDE with our sonification. Finally, we offer suggestions for future development of a multimodal display that can be used for analyzing complex data sets.				
14. SUBJECT TERMS Simulation output analysis, frequency domain experiment, sonification, auditory display, design of experiment, data farming.			15. NUMBER OF PAGES 112	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited.

**SPECTRAL ANALYSIS AND SONIFICATION OF SIMULATION DATA
GENERATED IN A FREQUENCY DOMAIN EXPERIMENT**

Hsin-Fu Wu
Lieutenant, United States Navy
B.S., University of Kansas, 1995

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2002**

Author: Hsin-Fu Wu

Approved by: Susan M. Sanchez
Thesis Co-Advisor

Russell D. Shilling
Thesis Co-Advisor

Paul J. Sanchez
Second Reader

James N. Eagle
Chairman, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

In this thesis, we evaluate the frequency domain approach for *data farming* and assess the possibility of analyzing complex data sets using data *sonification*. Data farming applies agent-based models and simulation, computing power, and data analysis and visualization technologies to help answer complex questions in military operations. Sonification is the use of data to generate sound for analysis. We apply a *frequency domain experiment* (FDE) to a combat simulation and analyze the output data set using spectral analysis. We compare the results from our FDE with those obtained using another experimental design on the same combat scenario. Our results confirm and complement the earlier findings. We then develop an *auditory display* that uses data sonification to represent the simulation output data set with sound. We consider the simulation results from the FDE as a waveshaping function and generate sounds using sonification software. We characterize the sonified data by their noise, signal, and volume. Qualitatively, the sonified data match the corresponding spectra from the FDE. Therefore, we demonstrate the feasibility of representing simulation data from the FDE with our sonification. Finally, we offer suggestions for future development of a multimodal display that can be used for analyzing complex data sets.

THIS PAGE INTENTIONALLY LEFT BLANK

THESIS DISCLAIMER

The reader is cautioned that the computer programs developed in this research may not have been exercised for all cases of interest. While every effort has been made, within the time available, to ensure that the programs are free of computational and logic errors, they cannot be considered validated. Any application of these programs without additional verification is at the risk of the user.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION.....	1
II.	FREQUENCY DOMAIN APPROACH TO DATA FARMING.....	3
A.	INTRODUCTION AND MOTIVATION.....	3
1.	Data Farming	3
2.	Simulation Output Analysis.....	4
3.	Frequency Domain Approach.....	6
B.	DESIGN OF FREQUENCY DOMAIN EXPERIMENTS FOR TERMINATING SIMULATIONS.....	8
1.	Selection of Simulation and Scenario.....	8
2.	Selection of Input Factors	8
3.	Selection of Driving Frequencies	9
4.	Selection of Output Responses.....	11
5.	Determination of Indicator Frequencies.....	12
6.	Spectral Analysis of the Output Responses	12
7.	Analysis of Results	16
III.	THE FREQUENCY DOMAIN EXPERIMENT.....	19
A.	APPLICATION OF THE FREQUENCY DOMAIN APPROACH	19
1.	Simulation and Scenario: MANA Peace Enforcement Scenario ..	19
2.	Input Factors	25
3.	Driving Frequencies.....	26
4.	Output Responses.....	29
5.	Indicator Frequencies.....	31
6.	Spectral Analysis of the Output Responses	33
7.	Analysis of Results	37
B.	STATISTICAL ANALYSIS OF RESULTS.....	38
C.	COMPARISON OF RESULTS.....	41
D.	CONCLUSIONS AND RECOMMENDATIONS.....	42
1.	Conclusions.....	42
2.	Recommendations	43
IV.	DEVELOPING AN AUDITORY DISPLAY FOR FDE USING DATA SONIFICATION.....	45
A.	SOUND.....	45
B.	AUDITORY DISPLAYS AND MULTIMODAL DISPLAYS.....	47
1.	Introduction.....	47
2.	General Design Principles	49
C.	SONIFICATION.....	50
1.	Purpose.....	50
2.	Attributes of Sound Synthesis.....	51
D.	EXAMPLES OF AUDITORY DISPLAYS USING SONIFICATION.....	53
1.	Classic Examples.....	53

2.	Innovative Examples.....	54
E.	AN AUDITORY DISPLAY OF DATA SONIFICATION USING JASS	57
1.	Java Audio Synthesis system (JASS).....	57
2.	Sonification Procedure	58
3.	Results	60
4.	Discussion of Results.....	62
5.	Other Concepts for Sonification	62
F.	CONCLUSIONS	62
G.	RECOMMENDATIONS FOR FUTURE RESEARCH.....	64
V.	SUMMARY OF RESULTS	67
	APPENDIX A: MANA SCENARIO INFORMATION	69
	APPENDIX B: RESPONSE SPECTRA OF THE FREQUENCY DOMAIN EXPERIMENT	75
	APPENDIX C. JASS CODE OF SONIFICATION PROGRAM.....	83
	LIST OF REFERENCES	87
	INITIAL DISTRIBUTION LIST	89

LIST OF FIGURES

Figure 1.	An example of frequency domain approach.	7
Figure 2.	A comparison between sample rates and the consequent aliasing due to undersampling.	10
Figure 3.	Effect of Nyquist frequency on aliasing.	11
Figure 4.	Layout of the MANA peace enforcement scenario.	20
Figure 5.	Schematic diagram of data farming using the Frequency Domain Approach.	22
Figure 6.	Pair-wise comparisons of five-factor variations in FDE.	28
Figure 7.	Pair-wise comparisons of five-factor variations in LH design of experiment.	28
Figure 8.	Combined spectrum of the number of Blue Agents killed.	35
Figure 9.	Combined spectrum of the number of Red Agents killed.	35
Figure 10.	Combined spectrum of the Fractional Exchange Ratio (FER).	36
Figure 11.	Combined spectrum of the Exchange Ratio (ER).	36
Figure 12.	Combined SNR of the number of Blue Agents killed.	39
Figure 13.	Combined SNR of the number of Red Agents killed.	39
Figure 14.	Combined SNR of FER (Blue/Red).	40
Figure 15.	Combined SNR of ER (Red/“Blue + 1”).	40
Figure 16.	Spectrum of the number of Blue Agents killed in Batch 1.	76
Figure 17.	Spectrum of the number of Red Agents killed in Batch 1.	76
Figure 18.	Spectrum of the FER in Batch 1.	77
Figure 19.	Spectrum of the ER in Batch 1.	77
Figure 20.	Spectrum of the number of Blue Agents killed in Batch 2.	78
Figure 21.	Spectrum of the number of Red Agents killed in Batch 2.	78
Figure 22.	Spectrum of the FER in Batch 2.	79
Figure 23.	Spectrum of the ER in Batch 2.	79
Figure 24.	Spectrum of the number of Blue Agents killed in Batch 3.	80
Figure 25.	Spectrum of the number of Red Agents killed in Batch 3.	80
Figure 26.	Spectrum of the FER in Batch 3.	81
Figure 27.	Spectrum of the ER in Batch 3.	81

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	List of Driving Factors and Frequency Assignments.	26
Table 2.	Assignment of Factors.	29
Table 3.	List of Indicator Frequencies for Each Run of the Experiment.	32

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

The author wishes to acknowledge God who inspired this endeavor and faithfully and patiently carried it on to completion. The author also wishes to recognize the people God used to motivate, encourage, and support this endeavor both in person and in prayer; I thank God for all of you.

Specifically, the author recognizes his family and friends for their love, patience and understanding throughout the thesis process. Mom and Dad, thank you so much for taking care of things at home and keeping me well fed. MAJ Jake Adams, USMC, thank you for your personal example and for holding me accountable to the task at hand. To all the friends who sustained me in prayer, thank you for your faithfulness and true friendship in Christ.

The author wishes to thank the faculty and staff in the Operations Research Department at the Naval Postgraduate School for the incidental help they provided throughout this thesis research. Particularly, the author recognizes Dr. Nita Miller, Dr. Lyn Whitaker, and Dr. Sam Buttrey; thank you all for taking the time to help me and encourage me.

The author also wishes to thank the Project Albert team at MITRE in Woodbridge, Virginia, for sponsoring this thesis research. Thank you for encouraging me to think freely, and inspiring me with Einstein's immortal words: "If we knew what it was we were doing, it would not be called research, would it?"

The author also thanks LTC Tom Cioppa, PhD., for providing the reference scenario from his doctoral dissertation.

Finally, the authors wishes to recognize his wonderful team of thesis advisors and second reader: Dr. Susan Sanchez, thank you for being the "good cop;" Dr. Russ Shilling, thank you for being the "bad cop;" both of you got the right information out of me to write this thesis, be it gently or sternly! And to Dr. Paul Sanchez, who has gone above and beyond the call of the second reader, thank you for your enthusiasm and support throughout this thesis.

God bless you all.

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

We have two key objectives for this thesis research:

1. Evaluate the frequency domain approach as a *data farming* technique.
2. Assess the possibility of analyzing complex data sets using data *sonification*.

We seek to accomplish these objectives in the context of the overall data farming environment. *Data farming* is a “meta-technique” that exploits advancements in three core disciplines: 1) Agent-based models and simulations; 2) Computing power; and 3) Data visualization. Data farming is the application of these disciplines to help answer complex questions in military operations [Brandstein and Horne, 1998]. Basically, data farming is similar to real agricultural farming in that, just as we grow crops and raise livestock to feed our bodies, we grow data and analyze the results to answer our questions.

We achieve the first key objective by developing a *frequency domain experiment* (FDE) appropriate for use with terminating simulations. Just like agricultural farming, we begin data farming by planting “genetically-engineered seeds” of data. We then sow our seeds in the data landscape of a peace enforcement scenario using a terminating combat simulation. We allow the data to “grow” from the simulation, and then we “reap” the data using spreadsheets and examine the yield using spectral analysis. By applying spectral analysis to the output data sets from the FDE, we separate “the wheat from the chaff” in the data sets. In the frequency domain, the input factors that contribute significantly to the output response parameters of the simulation show up as significant spectral power peaks in the response frequency spectra. We compare the significant factors from our FDE with those obtained using an alternative experimental design on the same combat simulation scenario. Our results confirm and complement the earlier findings. Both share common significant terms, and some interactions not identified using the other experimental design are significant in our FDE. Moreover, and perhaps most importantly, the results of the spectral analysis pass the “common sense test.” The significant factors in the response spectra are attributes of the combat setting that intuitively would have substantial effects on the output responses we consider. Based on

the success of our FDE, we propose some suggestions for further investigation of applying FDE to data farming.

We accomplish our second key objective by developing an *auditory display* (AD) using a data sonification technique. An AD is a display that represents information using sound. Three reasons lead us to consider the use of an AD for harvesting the data from our FDE. The first reason is the similarity between the spectral analysis of FDE and the spectral analysis of acoustic signals; we want to exploit the advantages of spectral analysis with respect to the decomposition of higher-order components of acoustic signals, which are analogous to higher-order terms and interactions in our model of the FDE. Another reason for considering the use of an AD to harvest the data from our FDE is the difficulty of visualizing data sets with high dimensionality. Visually representing data sets with high dimensionality can be difficult because human beings are limited in visual perception to three dimensions in space. Finally, we seek to exploit the natural “robustness” of auditory acuity to minimize the tendency to overfit the data set visually. The mantra of data collection—“Garbage in, garbage out”—cannot be overemphasized, as we are all familiar with the tendency to forget the quality of a data set we are analyzing and perform analysis on the data set to a precision not commensurate with its quality. Therefore, we seek to develop an auditory display to provide the decision-maker an adequate answer that literally “sounds good” in a shorter amount of time than performing an unintentionally more rigorous examination of the data set by visualization.

Sonification is the “use of data to control a sound generator for the purpose of monitoring and analysis of the data” [Kramer, 1994]. We apply the following procedure to use the output data sets from our FDE to generate sounds for analysis:

1. Serialize the response data sets into data streams.
2. Perform 0th-order mapping of data by mapping each element of a response data stream directly to the amplitude of the response waveshape of the data stream.
3. Specify the sample rate and upload a response data stream into an audio buffer using an open-source software development kit called Java Audio Synthesis System (JASS).
4. Use JASS to store the buffer and stream the data in the buffer to the sound card at the specified sampling rate.

5. The sound card synthesizes sounds based on the variations of the data stream in the audio buffer.
6. Repeat the sonification for the remaining data streams from our FDE.

When we hear the sounds of the sonified data streams, we can characterize at least three aspects of the sounds: noise, signal, and volume. We are able to distinguish the data streams by these three attributes of the sounds from the sonification. As we qualitatively compare the sounds of the data streams with the corresponding visual spectra, we conclude that our sonification of the data streams produces sounds that match the response spectra. Therefore, we believe that we demonstrate the feasibility of representing simulation data from the FDE with our sonification scheme.

Furthermore, based on our results, we assert two implications of our sonification with respect to data analysis. First of all, data analysis using our sonification may reduce the number of simulation runs required for data collection, while enabling the analyst to inject more complexity in the response by simultaneously varying more factors in the FDE. When we examine an “orchestrated” selection of observations over the entire data space, we will be able to see, and hear, a more representative rendering of the chaotic behavior and/or the hidden periodicities induced by our FDE. Secondly, data analysis by our sonification may be performed quicker than visualization. By listening to one entire output data stream we can qualitatively differentiate between data streams within a few seconds. Thus, each observation contributes to the analysis, and the overall sound is a “symphonic” representation of the data space.

We are very encouraged by our attempt in integrating simulation output analysis and human factors. We believe there is significant value in further research to develop an auditory display using sonification that will benefit data farming in the frequency domain. We embarked on our research having in mind the ultimate goal of a virtual environment for the analysis of complex data sets. We imagine that someday an immersive environment, created through a multimodal display, would enable the operations analyst to use more than just visual and auditory perceptions in order to improve understanding of the complexity of military operations. Through this research effort we believe we have advanced one step closer toward this goal, and strongly recommend continued research and development to make this goal a reality.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

This thesis implements an interdisciplinary approach to operations research. We seek to integrate traditional operations research applications (e.g., modeling and simulation and data analysis) with human factors. Human Factors is a field of research that is increasingly being explored for application in operations research problems. The goal of this integration is a platform-independent program that will assist operations analysts in performing effective and efficient factor screening of complex multidimensional data sets. This thesis proposes an alternate and possibly more efficient method of factor screening of complex multidimensional data sets by representing the data using an auditory display technique known as *sonification*. Moreover, it may also instigate further applications of human factors in operations research, particularly in the area of *data farming*.

Chapter II begins with an introduction of a concept used for data exploration called *data farming*. Next, it introduces the principles of simulation output analysis. An introduction to the principles of data farming in the frequency domain follows, along with a discussion on the procedure used to set up the *frequency domain experiment* (FDE) conducted in this thesis.

In Chapter III, we present the FDE that we conducted using a scenario in a combat simulation. The procedure established in the previous chapter is applied to the FDE in order to assess factors affecting the simulation outcomes of a particular peace enforcement scenario. The results of the FDE are presented at the end of the chapter with a discussion that compares them with those obtained from an alternative data exploration method.

Chapter IV describes the development of an auditory display using data *sonification* to represent data sets from the FDE. It begins with an introduction to sound, auditory displays, and sonification. Examples of auditory and multimodal displays are presented next in order to demonstrate existing applications of auditory and multimodal displays. We then discuss the development of an auditory display that applies a sonification technique to the data sets of the FDE. Finally, we present our

recommendations for future research on sonification of complex data sets and development of an auditory display for data analysis.

Chapter V summarizes the results from the FDE and the development of an auditory display using sonification. We briefly discuss the future potential of integrating the two parts of this thesis in the context of the ultimate goal of this thesis research.

II. FREQUENCY DOMAIN APPROACH TO DATA FARMING

A. INTRODUCTION AND MOTIVATION

1. Data Farming

Stochastic computer simulations can produce large data sets with high dimensionality in both the response surfaces generated from the output as well as the number of input factors and Measures of Performance (MOPs). Analysis of the data sets thus may be challenging because of the potential for numerous relationships and interactions between simulation parameters, as well as the random component of the output. *Data farming* is a “meta-technique” that exploits advancements in three core disciplines: 1) Agent-based models and simulations; 2) Computing power; and 3) Data visualization. Data farming is the application of these disciplines to help answer complex questions in military operations in a process called Operational Synthesis as a part of Project Albert. Project Albert is a research program sponsored by the Marine Corps Combat Development Command (MCCDC) [Brandstein and Horne, 1998].

As the name implies, data farming is similar to agricultural farming. Data farming utilizes the following four principal processes:

- *Fertilize* the minds of military professionals and other experts with ideas on how to capture the important aspects of conflict that have not been well-captured in the past, such as morale, leadership, timing, intuition, adaptability, etc.
- *Cultivate* ideas from these professionals concerning what might be important in a given situation.
- *Plant* these ideas in models to the degree made possible by the model in use and run the model over a landscape of possibilities for variables of interest.
- *Harvest* the data output from the model using innovative techniques for understanding scientific data.

We do not want to call the actions just described “steps” because they are all intertwined into the inquiry process of the scientific method that allows us to grow in our understanding. But, just as you do not grow crops or raise livestock in a vacuum, the growth resulting from data farming has a larger purpose. The reason for data farming is to feed our desire for

answers to questions. We can grow an overwhelming amount of data, so we continually re-focus on the question at hand and grow data which promises to add to our understanding. [Brandstein and Horne, 1998]

The propensity to produce large multi-dimensional data sets is inherent in the need for data farming. Care must be taken in generating data, because the time required to examine all potential factor level combinations is astronomically large [Lucas et al. 2002]. Capturing the essence of the data set in order to answer our questions is difficult because it may be difficult for analysts to understand and interpret the relationships between the numerous parameters of a simulation. Furthermore, because Project Albert uses agent-based simulations to model military operations, the simulation parameters themselves represent aspects of military operations that are often difficult to conceptualize. Thus, the high dimensionality of data sets and the obscure meaning of the parameters compound the difficulty of any attempt at analysis. Therefore, a key objective of this thesis is to evaluate the frequency domain approach as a means of planting and harvesting data *efficiently* in order to better help the analysts and decision makers answer complex and difficult questions about military operations and/or other complex operations.

2. Simulation Output Analysis

As “data farmers,” we plant data by running simulations of models that are *distillations* of the military operations from which our questions arise. A distillation is a simulation of a model that captures the essence of the questions we seek to answer. Because they are relatively simpler than the detailed models on which many complex military simulations are based, distillations require less computing power to run and can be quickly replicated [Brandstein and Horne, 1998]. Furthermore, we develop strategies to plant the data efficiently. We first develop a design of experiment we would use to plant our data that we expect would produce the responses we seek. We then harvest the output data produced from the experiment for analysis. We develop regression models to relate the input parameters, which we call *factors*, and output parameters, which we call *responses*, as part of the output analysis. We use these regression models, which are “meta-models” of the distillations, to estimate and predict responses of the distillations. Finally, we analyze the results of the experiment using a variety of analysis methods.

In order to apply these analysis methods, we first categorize simulation models into terminating and non-terminating simulations. In a terminating simulation, the simulation runs until it satisfies a particular condition or a set of conditions, and then terminates. For example, in a combat simulation, the simulation can terminate when all friendly or enemy forces are destroyed, or the simulation can terminate after a user-specified number of time steps. In a non-terminating simulation, the condition at which to terminate the simulation is ambiguous [Law and Kelton, 2000]. A simulation using an M/M/1 queuing model is an example of a non-terminating simulation.

We also must consider the *experimental unit* of a simulation. This affects how analysts perform the statistical analysis of the data. An experimental unit is a set of data from which one observation of the statistical sample can be collected. In a terminating simulation, we consider one experimental unit as one run of the simulation. When a terminating simulation is run n times, we have n experimental units toward our statistical sample. However, since data may be collected from a non-terminating simulation at specified points in time during the simulation, more experimental units may be gathered from one run of a non-terminating simulation than the single experimental unit obtained from a terminating simulation. For example, suppose we sample a non-terminating simulation at m points in time. Furthermore, suppose we also replicate the simulation n times. We would now gather $m \times n$ experimental units, even though we only replicate the simulation n times.

Whether the data are statistically independent or interdependent is another concern in simulation output analysis. Since it is easier to analyze independent data than dependent data, we might attempt to design our simulation experiment such that the simulation runs will produce independent data. Nevertheless, various methods of analysis are still available to enable the analysts to extract meaning from the data even if independence cannot be achieved.

Therefore, another key objective of this thesis is to propose an experimental design approach for a terminating simulation that minimizes the number of experimental units necessary for output analysis. Moreover, we would like the data generated from the proposed design of experiment to have the benefits of independence for the analysis.

3. Frequency Domain Approach

When harvesting data, we want to separate “the wheat from the chaff” in the data set—we analyze the sensitivity of output responses to the input factors of the distillation. We propose using the frequency domain approach to simulation output sensitivity analysis for data farming. Schruben and Cogliano [1981] first introduced the frequency domain approach to simulation sensitivity analysis and applied it to a simulation whose input factors could be varied during the run. In this approach, the factors are oscillated at specified frequencies, called *driving frequencies*, throughout the run. At specified intervals of time steps during a simulation run, the simulation samples the responses and collects them in a data set. Once the simulation has terminated and all data from the simulation have been collected, the analyst then applies spectral analysis to each response, in turn, in order to decompose the variations in each of the responses into a spectrum of frequencies. In the response frequency spectrum, factors that oscillated in the simulation show their relative contribution to the response by the magnitude of their spectral power peaks. Frequencies that are multiples of driving frequencies and the sums and differences of driving frequencies, are called *indicator frequencies*. Spectral power peaks at the indicator frequencies that have significant contribution to the response, as displayed in the spectrum, correspond to the contribution of the oscillated factors, their higher order terms, and their interactions with one another. The frequency domain approach is an appealing sensitivity analysis technique because many experimental units can be collected from one run of the simulation experiment. Moreover, the analyst can simultaneously assess the contribution of all factors and their interactions that are included in the regression model using the frequency spectrum. Therefore, the frequency domain approach may be an efficient way for the data farmers to plant and harvest data.

Figure 1 is a simple example of the frequency domain approach. The linear function represents a non-terminating simulation with one input parameter and one output response. As the input parameter continuously oscillates over the range of interest at the driving frequency, the oscillations induce corresponding continuous oscillations in the response. We can then apply spectral analysis to determine the spectral power peaks of the response at the parameter driving frequency and assess the sensitivity of the response due to the input parameter.




Figure 1. An example of frequency domain approach [Schruben and Cogliano, 1981].

This thesis proposes to plant data in the data landscape by applying a frequency domain approach in a design of experiment similar to Schruben and Cogliano [1981] and Sanchez and Buss [1987]. Sanchez and Buss proposed a model for frequency domain experiments (FDEs) that provides a technique for factor screening of simulations. In FDEs, input factors are oscillated at assigned driving frequencies. Then, by careful selection of the uniquely determined driving frequencies, the effects of the input parameters and their interactions on the response can be identified at the indicator frequencies. Whereas the simulation in Schruben and Cogliano allows factors to vary during the simulation run, we use a terminating simulation where the factors cannot be varied during the run. We will explain our design in detail in the following section.

B. DESIGN OF FREQUENCY DOMAIN EXPERIMENTS FOR TERMINATING SIMULATIONS

In order to assess the utility and appropriateness of applying the frequency domain approach to data farming, we base the design of our experiment on the suggestions of Schruben and Cogliano [1981] by using the following procedure:

1. Selection of Simulation and Scenario

As mentioned previously, when applying the frequency domain approach, we must consider the type of simulation used with respect to its terminating characteristic. We design our experiment and assess the proper experimental units based on this characteristic of the simulation. Furthermore, we must also determine whether the data generated in each experimental unit of the simulation are independent of each other. Obtaining independence of data may reduce the complexity of data analysis methods.

Another characteristic of the simulation that must be considered is the ability to vary input factors and collect output responses while the simulation is running. Schruben and Cogliano [1981] require this characteristic to be designed into the simulation. However, most simulations do not have this characteristic, and we as data farmers might not have any participation in how the simulation that we are analyzing is designed.

We must also consider the scenario for the simulation. As a feasibility study of the frequency domain approach, we pick a scenario that has been analyzed using other data farming methods. Thus we are able to compare our results with the results of existing analysis as measures of performance.

2. Selection of Input Factors

Similar to agricultural farming, it is essential for farmers to understand which types of crops to plant, given the soil and weather conditions at the farm, so that the land may yield abundant crops. We data farmers must also know and understand what input factors we should plant for our simulation in order to harvest data that would enable us to answer the questions we are asking. However, since the simulation of a complex system may have many input factors, we can easily become overwhelmed by the choices of the input factors available for planting the data landscape. We try a combination of factors that might have significant effects on the output responses. We select these initial factors based on intuition about the scenario and the simulation. We also perform test runs to

confirm our intuition in a smaller scale, as in planting the data in a small portion of data landscape to evaluate the amount and quality of data we might harvest when we plant the entire data landscape. Thus, we judiciously select input parameters based on intuition, prior experience and experimentation.

3. Selection of Driving Frequencies

In the frequency domain, we decompose the signal we seek to analyze into its component frequencies by applying spectral analysis to the signal. The component frequencies are then displayed in a frequency spectrum. We use angular frequency expressed in radians per observation of data for our analysis. Thus, one cycle of oscillations per observation equals 2π radians of oscillations per observation. Furthermore, it is sufficient to display only frequencies ranging from $[0, \pi]$ in the spectrum because of the phenomenon called the *Nyquist frequency*, which establishes the upper bound of the spectrum. The Nyquist frequency is the highest frequency that can be composed by two consecutive observations of data sampled at equal intervals. This phenomenon can be explained by the following example: Suppose we sample from our data set in equal intervals. We can only conclude with certainty that at most one half cycle, i.e., π radians, of oscillations occurs in between any two observations because there is no way for us to tell how many more oscillations have occurred between the observations unless we sample in between the observations. Hence, the highest frequency, i.e., oscillations per observation, is one half cycle, or π radians per cycle.

A phenomenon related to the Nyquist frequency is frequency *aliasing*. This phenomenon can be explained using the same example above. Suppose the signal in the example now has a frequency of 3 cycles per observation. Since the highest frequency that can be resolved in the spectrum is one-half cycle per observation, when the signal is decomposed by spectral analysis, the indicator frequency of the signal would be “folded” back below the Nyquist frequency and “aliased” by the *zero frequency*, i.e., zero cycle per observation. Thus, the actual frequency of the signal, 3 cycles per observation, has an alias at zero cycle per observation in the frequency spectrum.

Figures 2 and 3 illustrate the relationship between the Nyquist frequency criterion and aliasing.

Figure 2 shows the differences between an adequately sampled sinusoidal signal with an undersampled sinusoidal signal at the same frequency. The circles on the signals indicate the sample values. The two signals oscillate at the same frequency, but the signal on top is sampled more frequently at equal intervals than the signal on the bottom. Hence there are fewer signal oscillations possible between two consecutive samples. When we apply spectral analysis to the bottom signal, the undersampling causes ambiguity in determining the frequency of the signal. More signal oscillations than can be resolved are not sampled between two consecutive samples. Hence, an “alias” signal shows up in the frequency spectrum. It has a frequency lower than the frequency of the actual signal, and it also fits the sample intervals of the actual signal.




Figure 2. A comparison between sample rates and the consequent aliasing due to undersampling [National Instruments Corporation, 2000].

Figure 3 illustrates how the Nyquist frequency causes aliasing. The figure displays a continuous frequency spectrum that spans frequencies above and below the Nyquist frequency. We define *sampling frequency* (f_s) as the number of signal cycles sampled per unit time. In the figure, f_s is 100 cycles per second, or 100 Hertz (Hz). The Nyquist frequency criterion dictates that the Nyquist frequency ($f_s/2$) of the signal at this sampling frequency is 50 Hz because two consecutive samples compose at most one-half of a cycle of the signal; therefore, the Nyquist frequency is half the sampling frequency. Thus, all signals with frequencies above $f_s/2$ of 50 Hz are aliased back below 50 Hz. Four

signals in the figure, F1 through F4, are on different locations on the spectrum. F1 (25 Hz) is well within $f_s/2$ and is not aliased. F2, F3, and F4 are all aliased back below $f_s/2$, to 30, 40, and 10 Hz, respectively.




Figure 3. Effect of Nyquist frequency on aliasing [National Instruments Corporation, 2000].

Therefore, when selecting driving frequencies of input factors, we must consider the Nyquist frequency criterion and prevent aliasing from masking indicator frequencies in the response spectrum. In order to accomplish this objective, we use software developed by Paul Sanchez [Sanchez et al., 2002] to select the frequency assignments of the factors. The program implements the algorithm of Jacobson et al. [1991]: it considers the number of factors varied and assigns driving frequencies that prevent aliasing of frequencies.

4. Selection of Output Responses

We select the appropriate Measure of Effectiveness (MOE) from the set of output responses from the simulation. However, we first categorize the response into Measures of Performance (MOPs) and Measures of Effectiveness (MOEs). We define MOP as a quantitative parameter that provides indication of one aspect of system performance. We define MOE as a numerical means of assessing the overall system performance with respect to an objective set by the decision maker. We may select the MOE from the set of responses. We may also consider an MOP to be the MOE. Nevertheless, we are not limited by the set of responses. We can seek to combine responses and/or MOPs to form an aggregate MOE, such as a ratio of two similar responses, if appropriate. The seminal textbook on operations analysis, “Naval Operations Analysis” by Wagner, Sanders, and Mylander [1999], provides the following guidance for selecting the appropriate MOE:

- a. It must be quantitative.
- b. It must be measurable or estimable from data and other information available to the analyst.
- c. A significant increase {decrease} in MOE value must correspond to a significant improvement {worsening} in achieving the decision-maker's objective.
- d. It must reflect both the benefits and the penalties of a given course of action.

Therefore, we select an output response that is based on the guidance above and that we believe might be affected by the oscillations of the input factors when we plant the data landscape using the frequency domain approach.

5. Determination of Indicator Frequencies

Recall that one benefit of the frequency domain approach is the convenience of evaluating higher-order and interaction terms in the regression model. In the frequency domain, the frequency spectrum displays all frequencies contributing to the variations in the response. Furthermore, the indicator frequencies for higher-order effects of the oscillated factors on response show up at the multiples of the driving frequencies to which the first-order main effects are assigned. For example, suppose the factor X is assigned a driving frequency ω_1 . The first-order effect of X on the response frequency spectrum has an indicator frequency of ω_1 . The quadratic effect of X on the response has an indicator frequency of $2\omega_1$. Similarly, the n^{th} -order effect of X on the response has an indicator frequency of $n\omega_1$ in the response frequency spectrum. For interaction terms, the indicator frequencies are the sums and differences of the driving frequencies. For example, suppose there is a second-order interaction effect on the response from two factors, X_1 and X_2 , where X_1 and X_2 are assigned driving frequencies of ω_1 and ω_2 , respectively. The second-order interaction term in the response, i.e., $\beta X_1 X_2$, for some constant β , has two indicator frequencies: one at the sum and the other at the difference of the driving frequencies of X_1 and X_2 , i.e., $\omega_1 + \omega_2$ and $\omega_1 - \omega_2$, respectively. This result also emphasizes the importance of judicious assigning driving frequencies so as to prevent frequency aliasing.

6. Spectral Analysis of the Output Responses

By this stage of the data planting process, we have selected the factors to investigate and have assigned driving frequencies to these factors. We have also

determined the indicator frequencies based on these driving frequencies that represent all possible terms in our regression model. The driving frequencies are assigned such that they all remain within the Nyquist frequency and prevent aliasing of factors at the same indicator frequency. We then plant the data by running and replicating the simulation scenario and collecting the data set of the responses. Now we harvest the data by applying spectral analysis to the response data set in order to obtain the response spectrum. We interpret and analyze the response spectrum to seek answers to our questions

Fourier spectral analysis is the analysis of the frequency spectrum resulting from the approximation of a function using Fourier series. Although spectral analysis can be performed using other function sets as a basis, for the remainder of this thesis, we will use the term *spectrum* to refer to a Fourier spectrum. The Fourier series approximation of a function consists of two orthogonal components, which are sine and cosine functions. We summarize the derivation of spectral analysis in the following paragraphs based on Chatfield [1996].

Consider the model,

$$X_t = \mu + \alpha \cos \omega t + \beta \sin \omega t + Z_t,$$

where Z_t is a random process, parameters μ , α , and β are to be estimated, and t is the index of observations; $t = 1, \dots, N$. We can represent this model in the matrix form,

$$\mathbf{E}(\mathbf{y}) = \mathbf{A}\boldsymbol{\theta},$$

where

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}, \quad \mathbf{A} = \begin{bmatrix} 1 & \cos \omega & \sin \omega \\ 1 & \cos 2\omega & \sin 2\omega \\ \vdots & \vdots & \vdots \\ 1 & \cos N\omega & \sin N\omega \end{bmatrix}, \quad \boldsymbol{\theta} = \begin{bmatrix} \mu \\ \alpha \\ \beta \end{bmatrix},$$

for some angular frequency ω . This matrix representation is a general linear model of the original model. Therefore, we apply regression to the general linear model. The least squares estimate of $\boldsymbol{\theta}$ is thus: $\hat{\boldsymbol{\theta}} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}$. $\hat{\boldsymbol{\theta}}$ is called the *Fourier transform* of $\mathbf{y}$.

Furthermore, the matrix $(\mathbf{A}^T \mathbf{A})^{-1}$ becomes a diagonal matrix for $\omega_i = 2\pi \frac{i}{N}$, for $i = 1, \dots, \frac{N}{2}$.

Moving from the Fourier transform to the spectrum involves squaring the estimated coefficients for the sine and cosine terms, and summing them by frequency. It

can then be demonstrated that $\frac{1}{N} \sum_{i=1}^N (\mathbf{y}_i - \bar{\mathbf{y}})^2 = \sum_{i=1}^{\frac{N}{2}} (\alpha_i^2 + \beta_i^2)$, i.e., the Fourier spectrum partitions the variance. Under mild assumptions [Chatfield, 1996], the estimated spectral coefficients have a Chi-square distribution.

Transforming the original model into its Fourier series representation enables us to use all observations of the data set to fit the entire data set. Hence, the error term in the model is unnecessary and therefore omitted. Furthermore, the coefficients of this Fourier series representation at a given frequency ω are now the least squares estimates of the original model. In essence, Fourier analysis partitions the variability in the data by the contribution of each frequency in the spectrum to the overall variability in the data.

The Wiener-Khintchine theorem relates the Fourier spectrum of the model to the Fourier transform of the *autocovariance* function of the observations in the data set. Autocovariance is a measure of the covariance of a sequence of observations with each other. Because these observations are in a sequence, even if there is strong correlation between consecutive observations, it is intuitive that the contribution of variance from one observation in the sequence to another observation in the sequence reduces as the observations become further and further apart in the sequence for a stationary process. A technique called *windowing* is thus developed to weigh the contribution of the autocovariance values of all observations in the data set. Windowing applies the weighting factors to a specified number of observations (M) that is less than the number of observations in the entire data set (N). M is also referred to as the *window size*. One principle of windowing is to select M such that as $M \rightarrow \infty$ and $N \rightarrow \infty$, the ratio $\frac{M}{N} \rightarrow 0$.

One way to accomplish this is to select M to be proportional to $\sqrt{N}$.

A more thorough explanation of Fourier analysis can be found in Chapter 7 of Chatfield [1996].

The frequency spectrum resulting from the spectral analysis of the data shows all possible frequencies within $[0, \pi]$ as discussed in Section 3 above, which includes both indicator and non-indicator frequencies. Nevertheless, both types of frequencies are used for our analysis. The *spectral power* of a frequency is its contribution to the estimate of variance. We compare the spectral power of the indicator frequencies to the spectral power of the non-indicator frequencies in the frequency spectrum. We attribute the spectral power at non-indicator frequencies to variability of the response due to random noise. On the other hand, because the magnitude of the spectral power at a particular indicator frequency is an estimator of the contribution of the indicator frequency to the variance of the response, we consider the spectral power as the contribution of the term in the regression model corresponding to the indicator frequency. The spectral power thus is analogous to the regression coefficient of the term in the model. Therefore, if an indicator frequency has a high spectral power in the response spectrum, the corresponding term in the regression model is also significant in its contribution to the response.

For this thesis, we use software designed by Paul Sanchez [Sanchez et al., 2002] to perform spectral analysis of the output responses. The program is written in Java, and thus we can use command-line arguments in the command shell of any computer to specify the input parameters and the input data file, as well as the output file for the resulting spectrum. The program requires the following inputs: the number of frequencies into which the response is to be partitioned, the window size, the type of windowing, and the number of observations in the input data set. The program then estimates the spectrum of the observations and produces a response spectrum. The program automatically adds one more partition to the user-defined number of frequencies to partition the spectrum. This additional partition accounts for the *zero frequency*. The zero frequency in the response spectrum corresponds to the constant term in the regression model. Thus, the spectral power at the zero frequency signifies the contribution of the constant term in the regression model to the response.

7. Analysis of Results

Once we have harvested the data using spectral analysis to develop the response spectrum, we can now “reap” the harvest for answers to our questions. We perform a first-pass inspection of the results. We can immediately notice which indicator frequencies literally stand out from the normal noise levels. We can associate these relatively significant indicator frequencies to the effects they represent and make “quick and dirty” inferences about how the response is affected by the oscillated factors. We can also infer that the oscillated factors whose indicator frequencies do not show significant differences from noise probably do not contribute significantly to the response.

As mentioned in the previous section, the spectral power of each frequency in the spectrum is an estimator of the variance of the response at that frequency. Thus, under the null hypothesis that there is no factor effect, the spectral power of the response spectrum has a Chi-square distribution. We determine the degrees of freedom of the Chi-square distribution based on the type of windowing used to produce the spectrum, the window size, and the size of the data set, which is the sample size. We then pool together the spectral power values of the non-indicator frequencies by summing the values and then determining the average spectral power of the non-indicator frequencies. We also determine the degrees of freedom of this pooled noise. We do not include spectral power at the zero frequency in the pool because it represents only the constant term of the regression model. We determine the *signal-to-noise ratio* (SNR) of the spectrum by dividing the spectral values at each of the indicator frequencies by the average spectral power of the noise. The quotient has an F distribution. We apply the F-test for variability such that the null hypothesis indicates that there is no difference between the variability contributed by the indicator frequencies on the response and the variability contributed by noise. Since we are simultaneously comparing the variability of all indicator frequencies, we determine a Bonferroni level of significance for simultaneous comparisons from a level of significance for a single comparison. From the results of the F-test, we then can determine the significant factors in the regression model because the indicator frequencies that have significant SNRs correspond to the significant terms in the regression model.

All of the above is only possible because the experimental units are independent, by construction. It follows from the Wiener-Khintchine theorem that the spectrum of independent observations is flat. Thus, under the null hypothesis that there is no factor effect, the heights of the indicator and non-indicator frequencies have the same expected value, and their ratio has expected value 1.

THIS PAGE INTENTIONALLY LEFT BLANK

III. THE FREQUENCY DOMAIN EXPERIMENT

A. APPLICATION OF THE FREQUENCY DOMAIN APPROACH

We apply the frequency domain approach discussed in Chapter 2 to a combat simulation used as one of the data farming tools in Project Albert. We then analyze the harvest of data to gain insight into the significant factors and compare our results with an existing study.

1. Simulation and Scenario: MANA Peace Enforcement Scenario

MANA (Map Aware Non-uniform Automata) simulation was developed by Roger Stephen and Michael Lauren for the New Zealand Army and Defense Force. It has a graphical user-interface for specifying initial conditions and trigger states of the agents, as well as displaying the simulation run. A MANA simulation run will terminate after the user-specified number of time steps has elapsed. The input parameters are set at the beginning of each run and cannot be varied while the simulation run is in progress. Nevertheless, MANA offers the user the ability to specify levels of input parameters easily from a formatted input file. This ability to submit input levels in batches makes designing the planting data using FDEs easy.

We use a scenario in MANA that models a peace enforcement mission for our data landscape. LTC Tom Cioppa developed the scenario for his doctoral dissertation. Below is a description of the scenario:

The devised scenario is a challenging one since the Blue force is subjected to a series of encounters with the Red force and an original non-hostile force (Yellow) turns hostile as the scenario progresses. Blue's mission is to clear area of operation (AO) Cobra ... within the next two hours in order to facilitate United Nations (UN) food distribution and military convoy operations. Blue uses a light infantry platoon composed of three nine-man rifle squads and a platoon headquarters (HQ) of seven soldiers containing two machine gun teams. Their movement scheme is one squad up and two squads back with the platoon HQ following the lead squad (2nd squad). The 1st squad task is to follow and support 2nd squad with the purpose of clearing AO Cobra. Their follow-on task is to clear AO Python for subsequent UN food distribution and military convoy operations. The 2nd squad task is to conduct a movement to contact with the purpose of clearing AO Cobra. Their follow-on task is to clear AO Cobra for subsequent UN food distribution and military convoy

operations. The 3rd squad task is to follow and support 2nd squad with the purpose of clearing AO Cobra. Their follow-on task is to clear AO Boa (a small urban area with four building structures) for subsequent UN food distribution and military convoy operations. After 2nd squad clears AO Cobra, the platoon HQ moves to AO Boa to provide supporting fires for 3rd squad.

Red has a five-member element located in the vicinity of AO Cobra and two two-member elements patrolling along the movement routes of Blue squads 1 and 2. Additionally, Red has a two-member element in vicinity AO Boa. An originally non-hostile Yellow three-member element is initially in Blue's starting location. After discovering no safe water in vicinity AO Rattler, Yellow becomes hostile against Blue, seeks small arms from vicinity AO Boa, and moves to vicinity AO Python. The overall scenario is deemed doctrinally correct and plausible by the U.S. Army Infantry Simulation Center at Fort Benning, Georgia... [Cioppa, 2002]

Appendix A contains the full scenario description. Figure 4, best viewed in color, is the layout of the scenario in MANA.




Figure 4. Layout of the MANA peace enforcement scenario [Cioppa, 2002].

Since MANA is a terminating simulation, we modify the basic approach of frequency domain analysis for a non-terminating simulation to accommodate this characteristic. Recall that the experimental unit for a terminating simulation is one run of the simulation, and each run is an observation in our sample. We specify the time step at which a simulation run terminates to 400, based on consultations with LTC Cioppa. We determine the driving frequency assignments for the five factors we choose to oscillate, as discussed in Sub-section 3. We also consider one oscillation of a factor as varying from its maximum value down to its minimum value and back to its maximum. Since the lowest driving frequency assignment for this experiment results in one oscillation in eighty-one experimental units, we vary the levels of the five factors in eighty-one equal increments. Thus, in one set of eighty-one ordered experimental units we would complete at least one oscillation for each factor. In data farming terminology, we plant a “row” of eighty-one “genetically engineered seeds”.


Next, we arbitrarily select to replicate the set of eighty-one experimental units five hundred times. We consider this one batch of experimental units. Essentially, we produce five hundred sets of eighty-one experimental units and line them up end-to-end to produce an indexed series of digitized oscillations for spectral analysis. We note that even though the experimental units are serialized, they are independent from each other. The independence property of the experimental units is advantageous for our statistical analysis, which is discussed in Section B. In data farming terminology, we plant three batches of seeds. Each batch has five hundred rows of seeds, and each row has eighty-one seeds. Figure 5 is a schematic diagram of data farming using the frequency domain approach for our FDE.

Data Farming Using the Frequency Domain Approach

Step 1. Genetically engineer seeds of data.

Factor	Range of Values		Driving Frequencies (cycles per observation)		
	Min	Max	Batch 1	Batch 2	Batch 3
U	72	200	1/81	29/81	10/81
F	-64	64	4/81	1/81	17/81
G	-64	64	10/81	4/81	29/81
P	-64	64	17/81	10/81	1/81
V	-64	64	29/81	17/81	4/81


The gene pool of data seeds.




Comparison of the genetics of one row of seeds in Batch 1.

Figure 5. Schematic diagram of data farming using the Frequency Domain Approach.

Step 2. Sow seeds in rows.



Step 3. Let crops grow.




Schematic diagram of data farming using the Frequency Domain Approach (cont'd).


Step 4. Reap the harvest.

Serialize each batch by lining rows end to end in order.



$\Rightarrow N = IJ \Rightarrow \bar{x}_{nk} = n^{\text{th}} \text{ seed in batch } k, n = 1 \dots n.$



Step 5. Examine the yield.



Spectral Analysis



Spectrum of a response parameter of interest.

Schematic diagram of data farming using the Frequency Domain Approach (cont'd).

We ran our design of experiment using the “Gilgamesh” cluster at the MITRE Corporation in Woodbridge, VA. The Gilgamesh cluster consists of 15 nodes of Windows NT Workstations. 12 of the Workstations had PIII 550 CPUs with 64MB RAM. The other 3 nodes had PII 450 CPUs and 64MB RAM.

Each batch of runs has five hundred sets of the eighty-one levels from our frequency domain experiment. Thus, each batch contains 40,500 experimental units that were planted using the Gilgamesh cluster. The first batch of runs took about 9 hours to complete. The second and third batches each took about 11 hours each to complete. Hence, the total number of experimental units we used in our experimental design was 121,500.

2. Input Factors

LTC Cioppa applied a near orthogonal Latin Hypercube (LH) experimental design to the peace enforcement scenario and examined twenty-two factors with 129 levels [Cioppa, 2002]. For our experiment, we choose the five most influential factors that affect the outcome of the scenario, based on consultations with LTC Cioppa. We leave the remaining factors at their nominal values in all runs of the scenario. Below are the five factors we choose to oscillate in our frequency domain experiment (FDE). Refer to Appendix A for a description of all twenty-two input parameters considered in the scenario.

F. Blue Squad 1 in contact personality element w1 - controls the propensity to move towards agents of same allegiance, i.e., this factor represents the unit cohesiveness of Squad 1 when it encounters enemy Red agents.

G. Blue Squad 2 in contact personality element w1 - controls the propensity to move towards agents of same allegiance, i.e., this factor represents the unit cohesiveness of Squad 2 when it encounters enemy Red agents.

P. Blue Squad 3 injured personality element w1 – controls the propensity to move towards agents of same allegiance, i.e., this factor represents the unit cohesiveness of Squad 3 when any members of the squad are injured.

U. Blue movement range for all squads - controls the movement speed of Blue agents.

V. Red personality element w8 – controls the propensity to move towards enemies (Blue) in situational awareness map which are of threat level 1, i.e., the aggressiveness of Red agents to pursue a perceived threat.

Factors F, G, P, and V take on values ranging from -64 to 64. Factor U takes on values ranging from 72 to 200.

3. Driving Frequencies

In order to construct our meta-model of the scenario, we assume a second-order regression model with all interaction terms and determine the driving and indicator frequencies of the five factors and their interactions by applying the frequency assignment program called *Design*, developed by Paul Sanchez. Table 1 is the output of the program showing the driving frequencies for five factors.

Table 1. List of Driving Factors and Frequency Assignments.

Factor	Assigned Frequency		
	Run 1	Run 2	Run 3
1	1 / 81	29 / 81	10 / 81
2	4 / 81	1 / 81	17 / 81
3	10 / 81	4 / 81	29 / 81
4	17 / 81	10 / 81	1 / 81
5	29 / 81	17 / 81	4 / 81

The driving frequency assignments are listed in fractions, with the numerator as the number of cycles oscillated and the denominator as the number of observations over which the oscillations occur. Thus, for example, the driving frequency assigned to Factor 1 in Run 1 is 1 cycle in 81 observations. Similarly, the driving frequency assigned to Factor 1 in Run 2 is 29 cycles in 81 observations, and so on. The number of observations over which the oscillations occur is the same for all driving frequencies. The *Design* program determines that the spectrum must be partitioned into 81 discrete frequencies in order to prevent frequency aliasing of the indicator frequencies while insuring that all driving frequencies remain within the Nyquist frequency, i.e., one-half cycle per observation. Because the spectrum is thus partitioned, we consequently partition the values of the five factors into 81 discrete settings, with the settings at the beginning of each oscillation assigned to their respective maximum, as mentioned previously. Hence, rather than having continuous oscillations, our factors oscillate discretely, with the lowest

driving frequency oscillating one complete cycle through the range of possible values in 81 runs of the MANA distillation.

Note that there are three frequency assignments for each factor in three runs. We consider each run with respect to the frequency assignment to be one batch of five hundred “rows” of eighty-one “genetically engineered seeds” that we plant using the MANA distillation. In other words, the frequency assignments remain the same for all factors for each batch of data planted. Because we have three frequency assignment schemes for the factors, we plant three batches of data in the data landscape. The different frequency assignments enable the detection of possible frequency dependence of responses on the oscillated factors.

Figure 6 shows the pair-wise variations of five factors in an FDE. “F1” is factor 1 and so on. Each factor is assigned a driving frequency using the `Design` program. F1 has the lowest driving frequency, and F5 has the highest driving frequency. Note the patterns of variations between pairs of factors with high driving frequencies. For a pair of parameters, when the difference between the driving frequencies is proportionately large, the FDE tends to sample points at the edges of the parameter levels, e.g., the pattern of F1 and F5. When the difference between the driving frequencies is proportionately small, the FDE forms interesting patterns in the parameter levels, e.g., the pattern of F3 and F4. Despite the presence of patterns, all designs are mutually orthogonal.

(As an aside, we compare the patterns of variations between pairs of factors in an FDE with those in an LH design of experiment. Figure 7 shows the pair-wise comparison between five factors in LH design of experiment. We see an interesting and obvious difference: the FDE generates more patterns than the LH. FDE designs are also denser at the edges, sparser to the center, while still spanning the entire two-dimensional space of each pair-wise comparison.)




Figure 6. Pair-wise comparisons of five-factor variations in FDE.




Figure 7. Pair-wise comparisons of five-factor variations in LH design of experiment.

Table 2 shows how the five factors in MANA we oscillate for the frequency domain experiment correspond to the five factors whose driving frequencies are determined by the `Design` program developed by Sanchez [Sanchez et al., 2002]:

Table 2. Assignment of Factors.

Factor in <code>Design</code>	Factor in Peace Enforcement Scenario
1	U
2	F
3	G
4	P
5	V

Hence, Factor 1 represents factor U in the peace enforcement scenario, and so on.

4. Output Responses

When performing batching runs in MANA, the distillation stores output parameters for the batch in a spreadsheet file such that the output parameters may be manipulated for data analysis. We select two Measures of Performance (MOPs) from the set of available response parameters: The number of red agents killed and the number of blue agents killed in each simulation run. Individually, each MOP represents one aspect of the outcome of the course of action (COA). For example, an increase in the number of blue agents killed for one COA over another reflects only the penalty of the COA and provides no indication of the benefits for choosing the COA. Similarly, the converse applies to the number of red agents killed. It only reflects the benefits of a COA and does not inform the decision maker of the penalties incurred by choosing the COA. Hence, we seek to determine an MOE based on a combination of the two MOPs.

Cioppa uses Exchange Ratio (ER) as an MOE for his dissertation. ER is defined as the ratio of the number of red agents killed to the number of blue agents killed [Cioppa, 2002]. In our search for the appropriate MOE, however, we also consider three other ratios between the number of red agents killed and the number blue agents killed as

MOEs. One is the ratio of the number of blue agents killed to the number red agents killed. We call this ratio the Fractional Exchange Ratio (FER). FER is the reciprocal of ER. Another ratio is the percent of blue agents killed to the percent of red agents killed. Lastly, we consider the ratio of the percent of red agents killed to the percent of blue agents killed.

Recall the guidance for selecting the appropriate MOE in Section II.4. ER seems the intuitive choice for an MOE. It is quantitative because it is the ratio of the number of red agents killed to the number of blue agents killed for each experimental unit in our experiment. An increase in ER, i.e., an increase in the number of red agents killed compared to the number of blue agents killed, may correspond to an improvement in the blue agents performance in accomplishing the objectives of the scenario. Furthermore, the number of casualties incurred by a course of action reflects the benefits and penalties of the selected course of action. However, we recognize one particular significant limitation to ER as an MOE. In the MANA scenario, there are a total of fourteen red agents and thirty-four blue agents. Suppose blue agents choose a COA so effective that they do not incur any casualties. The ER for this course of action becomes indeterminate because it is the quotient of the number of red agents killed divided by the number of blue agents killed, which is zero. Hence, even though the blue agents may select the best COA, we cannot draw conclusions from the ER because it is indeterminate. In fact, we discover that there actually are runs where there are no blue agents killed. (We replicated a few of these runs by setting the same random number seeds and factor levels. We discovered that not all blue agents reached their objectives and thus were not exposed to enemy fire when the runs terminated after 400 time steps.) Therefore, we reduce the volatility of ER by adding one to the number of blue agents killed to all runs. We recognize that this treatment solves the indeterminate ER problem at the expense of a slight shift in distribution of the MOE.

In general, the FER may have the volatility of division by zero if no red agent is killed in any of the outcomes. After inspecting the data for such occurrences, however, we determine that there is no such occurrence in our result; in all runs at least one red agent is killed. Thus, in this peace enforcement scenario, the FER is more stable than the ER because it does not have the indeterminate volatility of the ER for our experimental

design. Nevertheless, the FER does not completely satisfy the four stated criteria for an appropriate MOE. Its behavior is inversely proportional to the measure improvement of a COA. An increase in FER, i.e., an increase in the ratio of the number of blue agents killed to the number of red agents killed, indicates that the blue agents have chosen an inferior COA since more of them are killed. Therefore, the FER is a stable ratio for our experimental design but does not truly satisfy the criteria for an appropriate MOE.

Alternatively, we also consider the ratio of the percentage of the number of red agents killed out of the initial red force strength to the percentage of the number of blue agents killed out of the initial blue force strength as an MOE. We consider the reciprocal of this ratio as an MOE as well. However, these two candidates have the same problem in volatility and representation as both the ER and FER, respectively. We first consider these as candidates because percentages have the benefit of standardizing the proportions of agents killed out of the initial force strength. Comparing proportions of force attrition may be beneficial because it is more representative of the actual scenario than the raw attrition values. For instance, the initial red force strength is five agents and the initial blue force strength is twenty agents. Suppose the blue force chooses a COA that results in five blue agents killed and five red agents killed. If our MOE were either ER or FER, it would show that the COA is not very effective; there is a one-to-one exchange in attrition. However, if we compare the percentages of attrition, we would find that the COA might have some merit: twenty percent blue loss to one hundred percent red loss! Nevertheless, because these ratios of percentages may still suffer the same problems as the ER and the FER, they are not considered any more favorably than the ER and the FER for the appropriate MOE.

Because the focus of our thesis is in the feasibility of applying the frequency domain approach to data farming, we simply choose for analysis the two MOPs, (the number of red agents killed and the number of blue agents killed) and the two straight attrition ratios (ER and FER).

5. Indicator Frequencies

The `Design` program not only provides information regarding driving frequency assignments in three simulation runs, but also provides a list of indicator frequencies and their corresponding terms for all three batches of MANA distillation runs. We use this

list of indicator frequencies to match the resulting response spectrum to the corresponding terms in the regression model. Table 3 below lists these indicator frequencies and their corresponding terms.

Table 3. List of Indicator Frequencies for Each Run of the Experiment.

Indicator Frequency			Factors		
Fractional	Decimal	Run1	Run2	Run3	
-----	-----	-----	-----	-----	
1 / 81	(0.012346)	1:0	2:0	4:0	
2 / 81	(0.024691)	1:1	2:2	4:4	
3 / 81	(0.037037)	2:1	3:2	5:4	
4 / 81	(0.049383)	2:0	3:0	5:0	
5 / 81	(0.061728)	2:1	3:2	5:4	
6 / 81	(0.074074)	3:2	4:3	5:1	
7 / 81	(0.086420)	4:3	5:4	2:1	
8 / 81	(0.098765)	2:2	3:3	5:5	
9 / 81	(0.111111)	3:1	4:2	4:1	
10 / 81	(0.123457)	3:0	4:0	1:0	
11 / 81	(0.135802)	3:1	4:2	4:1	
12 / 81	(0.148148)	5:4	5:1	3:2	
13 / 81	(0.160494)	4:2	5:3	5:2	
14 / 81	(0.172840)	3:2	4:3	5:1	
16 / 81	(0.197531)	4:1	5:2	4:2	
17 / 81	(0.209877)	4:0	5:0	2:0	
18 / 81	(0.222222)	4:1	5:2	4:2	
19 / 81	(0.234568)	5:3	4:1	3:1	
20 / 81	(0.246914)	3:3	4:4	1:1	
21 / 81	(0.259259)	4:2	5:3	5:2	
23 / 81	(0.283951)	5:5	1:1	3:3	
25 / 81	(0.308642)	5:2	3:1	5:3	
27 / 81	(0.333333)	4:3	5:4	2:1	
28 / 81	(0.345679)	5:1	2:1	4:3	
29 / 81	(0.358025)	5:0	1:0	3:0	
30 / 81	(0.370370)	5:1	2:1	4:3	
33 / 81	(0.407407)	5:2	3:1	5:3	
34 / 81	(0.419753)	4:4	5:5	2:2	
35 / 81	(0.432099)	5:4	5:1	3:2	
39 / 81	(0.481481)	5:3	4:1	3:1	

The fractional and numerical representations of the indicator frequencies are listed. Recall that the unit for frequency is cycles per observation. The numerical representation of the frequency is in angular form and the unit for frequency is radians per observation, e.g., 1 cycle per observation is 2π radians per observation. The pair of numbers in each of the columns is a second-order representation of the factors. Thus, for example, the lowest indicator frequency of Run 1 is 1 cycle per 81 observations, or 0.012346 radians per observation. This indicator frequency represents Factor 1, which is

the factor U in the MANA scenario. Similarly, the driving frequency of Factor 2, the factor F in the MANA scenario, has an indicator frequency at 4 cycles per 81 observations in Run 1, and so forth.

The second indicator frequency listed is 2 cycles per 81 observations, or 0.024691 radians per observation. Recall that, in the frequency domain, products of the same factor appear in the response spectrum at multiples of the driving frequency. Therefore, this indicator frequency represents the quadratic effect of Factor 1 in Run 1. The third indicator frequency listed is 3 cycles per 81 observations, or 0.037037 radians per observation. This indicator frequency represents the first-order interaction between Factor 1 and Factor 2 in Run 1. Interaction terms of the factors occur at the sums and differences of the driving frequencies of the two factors in the frequency domain. Hence, this indicator frequency is the difference between the driving frequencies of Factor 2, which is 4 cycles per 81 observations, and Factor 1, which is 1 cycle per 81 observations. Moreover, the next indicator frequency, 5 cycles per 81 observations, also represents the interaction term between Factor 2 and Factor 1 because it is the sum of the two driving frequencies. Notice that the indicator frequency for the quadratic effect of Factor 5 in Run 1, which is 23 cycles per 81 observations, is not double the driving frequency of Factor 5 in Run 1, which is 29 cycles per 81 observations. This is because the indicator frequency for the quadratic effect is aliased back to the spectrum that is within the Nyquist frequency. Doubling the driving frequency of Factor 5 in Run 1 would result in 58 cycles per 81 observations. Note that 58 cycles per 81 observations is 23 cycles per 81 observations from 1 cycle per observation, i.e., 81 cycles per 81 observations. Frequency aliasing “folds” the difference back within the Nyquist frequency and thus the doubling of the indicator frequency for Factor 5 in Run 1 now appears at 23 cycles per 81 observations.

6. Spectral Analysis of the Output Responses

After data are planted in the data landscape using MANA, we harvest the data by collecting the results and analyzing the results in the frequency domain. We first tabulate the number of blue agents killed and the number of red agents killed from each batch using Microsoft® Excel. We also determine the ER and FER by calculating the ratios of

these MOPs. We then process the MOPs, ERs and FERs for all three batches through the spectral analysis program, *Fourier*, designed by Paul Sanchez.

We use the *Fourier* programs to perform Fourier analysis on the two MOPs, the ERs and the FERs for all three batches. For each spectrum, we specify the program to partition the spectrum into 81 discrete frequency bins. We choose a window size of 10,000 observations ($M = 10,000$) and the default windowing method, which is a truncation window. We specify the program to use all 40,500 observations ($N = 40,500$) in each batch to determine the Fourier spectra. Finally, we manipulate the spectra and represent them visually using Microsoft® Excel. Thus, we harvest four spectra from each batch of simulation runs: one for the number of blue agents killed; one for the number of red agents killed; one for the FER; and finally one for the ER with the addition of one blue agent to each observation to prevent division by zero. Appendix B contains all of the individual spectra from the harvest.

We present the response spectra containing all three batches by factors in the following figures. Because each factor has a different frequency in each batch, we sort the frequencies by correlating the indicator frequencies in each batch with the associated terms in our regression model. We omit the non-indicator frequencies in these figures because we assume they collectively become the noise term in our model. Furthermore, because the spectrum is a partition of eighty-one discrete frequency bins, we present the spectra as stacked bar graphs rather than continuous linear graphs.

In the following figures, Figures 8 through 11, the main and quadratic effects are grouped to the left, while the interaction terms are grouped to the right. The quadratic effects are identified by the number “2” after the name of the factors. We separate the terms into our model in these two groups based on their degrees of freedom. Recall that each interaction term has two indicator frequencies. The indicator frequencies of each interaction term are at the sum and difference of the driving frequencies of the factors in the interaction term. Therefore, the interaction terms have twice the degrees of freedom of the main and quadratic terms.




Figure 8. Combined spectrum of the number of Blue Agents killed.




Figure 9. Combined spectrum of the number of Red Agents killed.




Figure 10. Combined spectrum of the Fractional Exchange Ratio (FER).




Figure 11. Combined spectrum of the Exchange Ratio (ER).

7. Analysis of Results

A visual inspection indicates that some terms dominate the spectra. For example, factors U , U^2 , F , and V most often have the highest spectral power values. Because the spectra are indications of the variability of factors on the response, we can qualitatively infer that these factors contribute most to the variability in the responses. Furthermore, we observe that the interaction term, GP , also contributes noticeably to the variability in the responses. We interpret these “quick and dirty” observations of our responses in the context of the scenario. Because the spectra are representations of the variance, we cannot determine whether each term affects the response positively or negatively without further analysis. Nonetheless, the following are the preliminary assessment of the results from this qualitative inspection.

The movement speed of the blue agents (factor U) significantly affects the outcome of the scenario. This concurs with intuition. Having the advantage in movement speed over the enemy means that the warrior can outmaneuver the enemy and position for attack before the enemy has an opportunity to attack. On the other hand, moving at higher speeds means that the warrior will likely run into more enemy contacts, and thus will have greater exposure to enemy fire.

Unit cohesion in the heat of battle also significantly affects the outcome of the scenario. This behavior is demonstrated by factor F , the propensity of squad 1 to move toward agents of the same allegiance when it is in contact with the enemy. This observation agrees with conventional intuition that “there is strength in numbers.” When engaging in battle, the outcome favors the side with the numerical superiority. We observe similar behavior for the protection of the injured, as represented by the term GP . This term is the interaction of the propensity for squad 2 to move toward agents of the same allegiance when in contact with the enemy (factor G), and the propensity for squad 3 to move toward agents of the same allegiance when injured (factor P). However, massing on the enemy enables the enemy to concentrate firepower from its own positions. Therefore, unit cohesion can affect losses on both sides.

Additionally, enemy aggressiveness (factor V) significantly affects the outcome of the scenario. This observation is also intuitive. The more aggressive the enemy, the

more likely the engagement will occur and thus affect the number of casualties on both sides.

Based on observations from this first-pass qualitative inspection, we can thus invest more computational time and plant more data to investigate interesting regions of the data space. Moreover, as we screen the factors, we can fit a regression model using only the significant terms from our results in order to determine whether the terms have positive or negative effects on the response. Most importantly, these observations show that the harvest of data planted using the frequency domain approach passes the “common sense test” by producing results that agree with intuition.

B. STATISTICAL ANALYSIS OF RESULTS

We pool the spectral power values corresponding to the same terms in all three batches for each of the four responses. We also pool the spectral power values at all non-indicator frequencies, except the zero frequency, for all three batches. Because the spectral power values are unbiased estimators of the variance of the response under the null hypothesis that there is no factor effect, we assume that they have Chi-square distributions with degrees of freedom equal to the quotient of the number of observations in the sample (N) divided by the window size (M) of the spectral analysis. We calculate the Signal-to-noise ratio (SNR) for the responses, which are the ratios of the quotients of spectral power values of the regression terms divided by the associated pooled degrees of freedom to the spectral power values of the noise divided by the associated pooled degrees of freedom. These SNRs are F-statistics. We then perform a simultaneous test of the variance attributable to every one of the terms in the regression model for each of the four responses using a F-test at ($\alpha = 0.01$) level of significance. Figures 12 through 15 display the resulting SNRs. The horizontal lines in the figures indicate the F-test statistic for the different degrees of freedom. Recall that the main effects and quadratic effects have the same degrees of freedom, while the interaction effects have twice the degrees of freedom as the main and quadratic effects because there are two indicator frequencies for each interaction term.




Figure 12. Combined SNR of the number of Blue Agents killed.




Figure 13. Combined SNR of the number of Red Agents killed.




Figure 14. Combined SNR of FER (Blue/Red).




Figure 15. Combined SNR of ER (Red/“Blue + 1”).

Figure 12 shows the combined spectral ratio of the number of blue agents killed for all three batches. We see that the overwhelmingly dominant factor is the aggressiveness of the red agents (Factor V). This agrees with our qualitative assessment of the behavior that the more aggressive are the red agents, the more likely is the number of blue agents killed to increase.

Figure 13 is the combined spectral ratio of the number of red agents killed for all three batches. We see that nearly all of the terms in our regression model are significant. The most dominant term is Factor F, the propensity of Blue Squad 1 to move toward agents of the same allegiance when in contact with the enemy. This is one of the factors that represent unit cohesiveness. The next dominant terms are the movement speeds of the blue agents, Factor U, which also contributes a significant quadratic term, and the unit cohesiveness of the blue agents, Factors G and P. The quadratic terms of both F and G also significantly affect the number of red agents killed, but are not dominant. It is interesting to note that all interaction terms are significant with respect to the number of red agents killed. This result agrees with the qualitative assessment that speed, unit cohesion, and enemy aggressiveness all affect the number of red agents killed.

Figures 14 and 15 show the combined spectral ratios of the FER and ER, respectively, for all three batches. These ratios of the MOPs enhance the dominant terms and diminish the remaining terms that are also significant. Note that the spectral ratios of these two MOEs are nearly identical. This similarity is somewhat reasonable because the ratios are reciprocals of each other.

C. COMPARISON OF RESULTS

From these spectral ratios, we note that all first order effects of the factors oscillated are significant: F, G, P, U, and V. We compare these factors from our frequency domain approach with Cioppa's [2002] regression model for ER based on his near orthogonal LH design:

$$\begin{aligned} ER = & 1.890 + (1.928 \times 10^{-7})U^2 + (.000457)B + (.000736)E + (.00237)F + \\ & (.00568)G + (.000826)P - (.00898)U - (.00327)V - (4.866 \times 10^{-6})BU - \\ & (3.021 \times 10^{-5})GU - (2.688 \times 10^{-5})FV + (1.378 \times 10^{-5})IJ + (2.225 \times 10^{-6})BN. \end{aligned}$$

Qualitatively, we see that the terms we oscillate agree with the terms in Cioppa's model, as we should expect because we used this model to select the factors to oscillate for our experiment. If any of these five factors we oscillate were insignificant, it might be a cause for concern about using the frequency domain approach as a data farming technique. Conversely, it might mean that further work was needed to determine why the results differed, and the relative strengths and weaknesses of the different approaches. Furthermore, we note that the significant interactions in our approach include the two terms in Cioppa's model of factors we choose to oscillate, namely interactions GU and FV. The remaining interactions in Cioppa's model do not show in our result because they include factors that are not considered in our experiment. Our results also show that in addition to the quadratic effect of Factor U, Factors G and V also contribute significant quadratic effects to our model of ER.

D. CONCLUSIONS AND RECOMMENDATIONS

1. Conclusions

Using our spectral ratios, we compare results of data farming using the frequency domain approach with an existing regression model of the scenario. Based on our comparison, we conclude that frequency domain approach is a feasible technique for data farming. For our experiment, we select five significant factors from a peace enforcement scenario using the MANA distillation. We then apply the frequency domain approach in designing an experiment to verify that the same five factors are also significant in the frequency domain. The results qualitatively agree with the regression model from which we select the five factors. Furthermore, we show that harvesting the data using the frequency domain approach provides a useful visual display for simultaneous comparison of factors and interactions that we seek to evaluate. Because the spectral ratios of the terms to the overall noise term represent the variability of the response, the magnitude of the spectral ratio for each term indicates the contribution of the term to the variability of the response. We screen the factors by applying F-tests to simultaneous comparisons of the SNRs of the terms. Hence, factor screening can be efficiently performed in the frequency domain. Therefore, we conclude that the frequency domain approach is not

only a feasible method for data farming, but also a useful technique for factor screening that is easy to generate.

2. Recommendations

In achieving the objective of determining the feasibility of the frequency domain approach to data farming, we recommend the following issues for further research.

a. “How many observations are enough?”

For our example, we arbitrarily determined the number of replications for each batch of experiments, because we wanted to make sure that we have plenty of observations to obtain sufficient statistical significance for our experiment. Now that we have demonstrated the feasibility of the frequency domain approach, we recommend evaluating the number of observations that are sufficient to achieve the same model identification results. This assessment will help determine the efficiency of the frequency domain approach in terms of computing power requirements.

b. “What are the signs of the regression coefficients?”

Because variance is the square of deviation, the spectral ratio indicates the relative magnitude of the coefficients with each other, but not the signs of the regression coefficients associated with the terms. For factor screening, it is sufficient to determine the relative contribution to variability of the term to the response. However, a regression model is necessary to fit the data in order to determine whether term affects the response positively or negatively. Therefore, a regression analysis for the model should be performed in order to determine the signs of the regression coefficients associated with the factors that have been screened using the frequency domain approach.

c. “What about other factors?”

We only oscillated five factors for our experiment. Our comparisons with the original model from which the five factors were selected are limited, because the original model considers twenty-two factors. Therefore, we recommend selecting the same twenty-two factors for oscillation, planting data in MANA using the frequency domain approach, and comparing the results with the existing regression model. Because the spectrum is continuous, there is no limit to the number of indicator frequencies.

Therefore, the frequency domain approach can potentially enable factor screening of all factors simultaneously. However, assigning driving frequencies to prevent frequency aliasing of the indicator frequencies may become more difficult. For example, increasing the number of factors to oscillate from five to six increases the number of discrete frequency bins in the spectrum from 81 to 119. Oscillating twenty-two factors simultaneously increases the number of discrete frequency bins to 2,367. The increase in the number of frequency bins means an increase in the number of observations required for one set of simulation runs. Therefore, the number of factors to oscillate directly affects the number of observations required for the frequency domain approach.

d. “What about higher-order terms?”

Similar to increasing the number of factors for comparison, assuming a higher-order model increases the number of frequency bins in the spectrum. Therefore, unless the complexity of the response cannot be sufficiently modeled by second-order models, we recommend simply assuming a second-order regression model.

IV. DEVELOPING AN AUDITORY DISPLAY FOR FDE USING DATA SONIFICATION

In order to develop an auditory display for FDE using data sonification, we first discuss the attributes of sound, the principles of *auditory display*, and the process of *sonification*.

A. SOUND

We define the following attributes of sound for our thesis research:

The *frequency* of sound is the number of cycles at which sound propagates per second. We measure frequency in the Hertz (Hz); one Hertz is one cycle per second. The frequency spectrum for human hearing ranges from 20 to 20,000 Hz. Note that acoustic frequency is synonymous with the frequency we define for our FDE. Perceptually, we consider the frequency attribute of the sound as *pitch*. When a high-frequency sound reaches the human listener, we describe the sound as having a high pitch.

The *intensity* of the sound is the magnitude of energy in the propagation of the sound per unit area. Sound intensity is usually measured logarithmically in units of decibels (dB) as the ratio of the energy per unit area of the sound to a reference energy level per unit area. Logarithmic measurements are used because sound intensity can vary over a large range of values. Note that this is analogous to the spectral ratios resulting from our FDE, if we present the spectral ratios logarithmically. Perceptually, we consider sound intensity as *loudness*.

The *complexity* of a sound is the most difficult attribute of the sound to define. Generally, we associate the complexity of a sound with its waveform, i.e., the shape of the wave, as well as the harmonics inherent in the sound, i.e., the number of multiples of the fundamental frequencies in the sound. Fundamental frequencies are similar to notes on the musical scale. Fundamental frequencies are analogous to driving frequencies in our FDE, and the harmonics are analogous to indicator frequencies that are multiples of the driving frequencies. Perceptually, we consider the complexity of a sound as *timbre*.

For example, the timbre of a violin is different than the timbre of a flute, even when the violinist and the flutist play the same note.

Sound also has temporal and spatial attributes. For our thesis, we consider one temporal attribute—the duration of the sound, i.e., how long we should generate the sound. Spatially, we consider the location of the sound relative to the listener. We use the polar reference coordinate system to describe the location of the sound by its elevation, azimuth and radial distance from the reference location.

Detectability of sound signals varies with frequency, intensity, and duration. Because sound is always present in the natural environment, we define detectability as the ability to detect an audio signal embedded in background noise. Detectability of a signal in noise depends on the sound intensity, frequency and the duration of the signal and the background. We introduce a measure for distinguishing two similar sounds in the following paragraphs. With respect to duration of a sound, however, there are some neurological limitations for auditory perception that establish a minimum duration for the human listener to perceive the signal—a sound signal should last at least 500 milliseconds in duration for the listener to perceive the signal [Sanders and McCormick, 1993].

In order to differentiate one level of a sound attribute from another, we define *just noticeable difference* (JND) as “the smallest change or difference along a stimulus dimension (e.g., intensity or frequency) that can just be detected 50 percent of the time by people” [Sanders and McCormick, 1993]. For example, JND in sound frequency is the minimum difference in frequency between two sounds that have the same intensity and timbre for the human ear to distinguish the two sounds as different 50 percent of the time. Similarly, JND in sound intensity is the minimum difference in intensity between two sounds that have the same frequency and timbre for the human ear to distinguish the two sounds as different 50 percent of the time. Experiments using pure tones, i.e., sounds generated from pure sinusoidal oscillations, show that for two pure tones having the same frequency, the JND in sound intensity between the two tones is smallest when the tones have high intensity. Sound intensity also affects the JND in frequency between two pure tones. The JND in frequency is smallest between two pure tones at low frequency having

the same intensity. Between two low-frequency pure tones, the JND in frequency is smallest if the tones have high intensity [Sanders and McCormick, 1993].

Spatially, a sound source directly in front of a human listener must displace about one arc-degree laterally for the listener to notice a difference in location. However, the listener would not be able to accurately perceive changes in location unless the sound source is displaced as much as 15 arc-degrees laterally from the front of the listener. Furthermore, spatial acuity of sound varies with orientation from the listener and dimension. A sound source directly to the side of a human listener must be displaced 10 arc-degrees before the listener notices the change in location. Changes in distance of a sound source from the listener are generally difficult for the listener to distinguish [Shilling and Shinn-Cunningham, 2002].

B. AUDITORY DISPLAYS AND MULTIMODAL DISPLAYS

1. Introduction

An *auditory display* (AD) is a display that represents information using sound. Familiar examples of auditory displays like a doorbell and a telephone announce to those within hearing range of visitors at the door and on the phone, respectively. Complex auditory displays that enhance data analyses and complement data visualization are emerging. We will present more of familiar and complex examples of auditory displays in Section D. For now, we present some benefits and limitations of AD.

One benefit of an AD is that it can complement a visual display. When information is presented using an AD for monitoring and warning purposes in conjunction with visual displays, the AD enables the user to freely and simultaneously perform other tasks that require visual focus. For example, auditory displays used in cockpits of aircraft for audible warnings and indications of the flight environment reduce pilot workload and enhance situation awareness. Similarly, adding sound to the visual picture engages and enhances the interest of the user, if properly designed. The strongest evidence for the enhancement of visual perception with auditory is the appeal of a good movie with good sound effects [Shilling and Shinn-Cunningham, 2002].

Unlike visual perception, which has a limited field of view, our “field of sound” is omnidirectional and continuous. We stop seeing, temporarily, when we close our eyes, but we cannot stop hearing at any time unless we somehow cover or plug our ears. Hence, another potential benefit of AD is spatial presentation of information around the user. Spatial auditory display (a.k.a. spatial audio, 3-D audio, surround sound, etc.) is an emerging field of research made possible by the advancement in computer technology. Experimental spatial auditory displays representing threat, navigational, and targeting information in the cockpit of an AH-64A attack helicopter simulator show promising results for development of spatial auditory displays [Shilling et al., 2000].

The lack of orthogonality of sound attributes is a major limitation of auditory display for data representation. Changing one attribute of sound may affect other attributes of the same sound. The aforementioned differences in JNDs with frequency and intensity of sound are examples of this limitation. Therefore, the lack of orthogonality of sound attributes can make representing data using sound difficult.

A more thorough discussion of the benefits and limitations of AD can be found in Kramer [1994].

Three reasons led us to consider the use of an auditory display for harvesting the data from our frequency domain experiment. First, there is commonality between spectral analysis (as applied in our FDE) and *acoustic signal analysis*; in fact, they are identical. When we applied the spectral analysis to our data set, we decomposed the oscillations of the response into their component frequencies. Similarly, acoustic signal analysis decomposes signals, i.e., data, in the acoustic range of the frequency spectrum and analyzes the component frequencies. The similarities between the two applications of the same analysis technique are familiar personally to the author because of his experiences with sonar.

Another reason for considering the use of an auditory display to harvest the data from our FDE is the difficulty of visualizing data sets with high dimensionality. For our FDE, we examined five factors out of the many available factors in MANA from which we may choose and relate them to two responses. Suppose we assign each of the five factors to a dimension and examine each response as a function of those five factors.

Representing the function visually is impossible because human beings are limited in visual perception to three dimensions in space even though the function has six dimensions. Granted, there are advanced data visualization computer programs and techniques that can present many dimensions of a complex data set. Nevertheless, these displays still cannot adequately convey orthogonality between dimensions beyond much more than three dimensions. Furthermore, visualization of too many parameters can saturate the visual perception. Hence, in a sense, representing multidimensional data using a visual display has similar difficulties as using an auditory display. For example, the spectral ratios that we harvest from our data are visual representations of responses with respect to all five factors and their interactions. Nevertheless, the spectral ratios are only summaries of the relationships between the factors and the responses and are not representations of the data space per se. Therefore, we also want to consider alternative methods of displaying data besides visualization.

Finally, we seek to exploit the natural “robustness” of auditory acuity to minimize the tendency to overfit the data set visually. The mantra of data collection—“Garbage in, garbage out”—cannot be overemphasized. However, the average analyst tends to forget the quality of the data with respect to accuracy and variability, and attempt to analyze the data, e.g., perform regression analysis, to a precision that is not commensurate with the quality of the data. Hence, we tend to “read more into” the data than we should when we analyze the data set. For operations analysts, overfitting data costs the operational decision-maker time and resources while waiting for the analysis results. Therefore, we seek to develop an auditory display to provide the decision-maker an adequate answer that literally “sounds good” in a shorter amount of time than performing an unintentionally more rigorous examination of the data set by visualization.

One of the goals of this thesis is to assess the feasibility of using an auditory display for data analysis. Hence, we develop and describe an auditory display prototype by using data sonification techniques.

2. General Design Principles

Just as there are sound principles for graphical representations of data, e.g., Tufte [1983], there are some principles for developing effective auditory displays, e.g., Kramer and Smith et al. [both in Kramer, 1994]. These principles can be distilled into one

fundamental principle: “Represent the information in a way that is understandable.” With respect to AD, this means mapping parameters of the data set to different attributes of sound in an effective way, with consideration of the benefits and difficulties of auditory displays. Some data are more easily represented with sound than others. For example, data from the frequency domain such as acoustic signals, seismic data, and FDEs, can be directly represented using an AD. On the other hand, data that do not have natural relationships to sound require some level of abstraction and subjective judgment in order to map parameters from one data set to attributes of sound. For example, Bly sonified categorical data such as the famous Fisher’s iris data set for her dissertation using pitch, volume, duration and waveshape [Bly, 1982].

Nevertheless, designing an effective auditory display is not a trivial task. Some principles for data sonification are intuitive, but difficult to implement. Therefore, this thesis attempts to design and experiment with an auditory display in order to facilitate the process of factor screening in FDEs with multidimensional data sets for data farming. The goal of this attempt is an auditory display that will effectively communicate the FDEs in data farming efficiently.

C. SONIFICATION

1. Purpose

Sonification is the “use of data to control a sound generator for the purpose of monitoring and analysis of the data” [Kramer, 1994]. We define *sound generator* as a means of mapping dimensions in the data set to attributes of sound, and subsequently rendering the sound for the human listener. The purpose of sonification is to increase a user’s ability to perceive and analyze a greater number of data dimensions. Whereas our visual perception limits the number of dimensions graphics display can represent, we attempt to present more dimensions of the data set to the user at one time by sonification.

An often-cited paper on data sonification is the doctoral dissertation of Sara Bly [1982]. Bly compared the effectiveness of sonification and visualization of data. She encoded three types of data—multivariate, logarithmic, and time-varying data, into sound. Bly experimented with human participants to compare the rates of correct

identifications using three treatments: sound-only representation, graphics-only representation, and a combined sound-and-graphics. The results of Bly's experiments demonstrated significant differences in the identification rates among the treatments. The combined sound-and-graphics representation of the data had the highest percentage of correct identifications, followed by sound-only representations, then by graphics-only representations. We note that Bly used a computer battle simulation data as a test data set:

Professor Sam Parry of the Naval Postgraduate School... suggested a time-varying application. Computer battlefield simulations which run from start to finish without human interaction provide information about the state of the battle at each time step. To an analyst interested in the results of the simulated battle, this information is often an overwhelming collection of statistics. Nevertheless, it is important to note the battle characteristics which yield various results. Thus the information at each time step encoded into sound results in a song for each battle. Listening to the songs provides a quick view of the battle in progress and draws attention to critical points during the battle [Bly, 1982].

Although the MANA distillation collects data from each simulation run in time-step increments, we do not examine the progress of each run by time-steps, but rather the summary results from the simulation runs because we perform a large number of runs. However, MANA does use limited sound effects to indicate the occurrence of certain key events in the simulation run with sound, though the collection of these sounds in each simulation run does not even come close to resembling a battle song. Nevertheless, the above suggestion for sonification of combat simulations from over twenty years ago is still valid. To date there is no combat simulation that integrates sonification as part of the output analysis technique. Based on the success of Bly's and subsequent experiments, sonification of combat simulation data may be a more useful technique than graphical visualization of data for analyzing the complexity and multidimensional nature of combat.

2. Attributes of Sound Synthesis

A sound generator has two general components: the data processing component and the hardware component. After determining the mapping of the dimensions of data to the attributes of sound, the sonification designer interfaces with the data processing component of the sound generator and uses it to perform the mapping. Once data are

mapped to sound attributes, the sound generator produces the sound using the hardware component for the user to hear and thus analyze the data. An example of a sound generator is a computer algorithm that maps data dimensions to attributes of sound and then renders the sound for output through the speakers of the computer. This type of computer algorithm is a *sound synthesis algorithm*.

We consider the following attributes of sound synthesis as the basis of our sonification technique:

We define *sampling* as the process of digitizing the analog oscillations of a sound in equal time intervals. The soundboard in a computer receives its signal to generate sound from an input audio *data stream*, i.e., an ordered series of data elements representing digitized samples of oscillations of the sound to be generated. The rate at which the soundboard samples the data stream is the *sampling rate*. The user can specify the sampling rate; however, the current nominal sampling rate of a personal computer is 44,100 samples per second. In order to maintain the specified sampling rate, the data stream is stored in and retrieved from a “First In, First Out” (FIFO) *audio buffer*. We define the *buffer size* as the number of data elements that can be stored in the audio buffer. Each time the soundboard samples a number of data elements equaling the buffer size, the soundboard completes sampling one cycle of oscillations in the data stream. The output frequency of the sound from the data stream depends on the sampling rate and buffer size of the data stream. The output frequency is proportional to the sampling rate and inversely proportional to the buffer size:

$$\text{Output frequency (Hz; cycles per second)} = \frac{\text{sampling rate (samples per second)}}{\text{buffer size (samples per cycle)}}.$$

Furthermore, audio data streams may be mapped to the parameters of sound in the following order:

In a 0th-order mapping, the data stream itself is listened to as a stream of digital audio samples.

In a 1st-order mapping, the data stream controls a parameter or parameters of a synthesis model (e.g., the data controls the amplitude of an oscillator).

In a 2nd-order mapping, the data stream controls the parameters of a synthesis model that controls the parameters of another synthesis model

(e.g., the data controls the amplitude of an oscillator that, in turn, controls the frequency deviation of another oscillator). [Scaletti; in Kramer, 1994].

Another sound synthesis technique is *waveshaping*. Waveshaping is the transformation of component frequencies into complex frequencies. The transformation is the result of applying a *transfer function* to the elements in the data stream. “Each element in a data stream g could be interpreted as an argument to a function f , where only the output of f is actually heard (or used to control the parameter of another sound)” [Scaletti; in Kramer, 1994]. An example of waveshaping transformation is Taylor approximations of a trigonometric function, where the input function $g(t)$ is a sinusoid, e.g., a cosine function, and the transfer function $f(x)$ is a polynomial, and thus $f(x)$ becomes the n^{th} -order Taylor approximation of $g(t)$. We note that this is very similar to our process of spectral analysis of MANA simulation data.

D. EXAMPLES OF AUDITORY DISPLAYS USING SONIFICATION

1. Classic Examples

An example of a common auditory display is the household smoke detector. A smoke detector senses the amount of particulates in the room due to smoke and triggers an audible alarm to warn occupants in the room of fire. Another example of an effective auditory display is the Geiger-Mueller radiation detector, commonly called the Geiger counter. The Geiger counter detects ionizing radiation particles and displays the amount of radiation it senses visually and sonically. Each radiation particle reaching the detector causes an ionizing event in the detector, which the Geiger counter converts into a voltage deflection in its detection circuitry. The voltage deflection is converted into rate of detection, in counts per unit time. The radiation level measured in rates is displayed using a mechanical or digital meter visually. In addition, the voltage deflection also causes a “click” to sound from the speaker or headset of the Geiger counter. The number of “clicks” in an interval of time thus directly represents the radiation level detected. It is a well-known fact that the auditory representation of radiation level in the Geiger counter is more sensitive and responsive to changes in the radiation level than the visual display.

The two examples above are general applications of auditory display. The smoke detector is a pure auditory display in that information is only presented with the sound of

the smoke alarm. However, the Geiger counter presents information not only with an auditory display, indicating the radiation level with the amount and rapidity of audible clicks, but also with a visual display to provide a meter reading of the radiation level. Hence, the Geiger counter can also be classified as a *multimodal display* in that it presents information using more than one sensory modality. The Geiger counter also represents an example of a simple sonification model, while the smoke detector is merely an auditory display.

A classic example of a multimodal display is sonar. Sonar (Sound Navigation and Ranging) onboard a ship or submarine receives sounds in the ocean via hydrophones. A hydrophone is basically a microphone designed for use in water. A hydrophone transduces acoustical energy of sound into electrical energy for signal processing. Sources of sounds in the ocean include marine animals, a.k.a. biologics, ships and submarines, as well as seismic activities. As these sounds propagate in the ocean, the array of hydrophones of a sonar system senses the sounds and transduces the sounds into electrical signals. The sonar signal processor analyzes the signals and converts the signals into visual and auditory displays. The visual display portion of sonar presents sounds as pixels on a video screen. The visual display is a display of sound duration with respect to time, with the most recent pixels of sounds appearing along the top of the display by the directions from which they are received. The sonar operator can also wear a headset and listen to a sound appearing from one direction on the visual display by “steering” the signal processor to sonify the sound at that direction. As the sonar operator hears the sound, he evaluates the type of sound aurally. If the sound is from a source of interest, e.g., a ship, he can analyze the sound and classify the ship. Hence, sonar is a multimodal display that represents acoustic information both visually and auditorily for the sonar operator to monitor and analyze.

2. Innovative Examples

For the inaugural conference of the International Conference on Auditory Display (ICAD) in 1992, Bly [in Kramer, 1994] solicited several auditory displays of two multivariate data sets using data sonification. One data set was relevant for discriminatory tasks, i.e., determining the similarities and differences between data sets. Another data set was a multidimensional time-varying data set relevant for pattern

recognition. Each data set had six variables, i.e., six dimensions. Three different displays using different sound mapping techniques were applied to the first data set. The most successful display was the one that mapped the sum of squares of the six dimensions to pitch. Bly's conclusion enforces the fundamental design principle of AD: "Serious consideration must be given as to which factors will make the process of data exploration, especially sonification, most effective" [Bly; in Kramer, 1994].

Fitch and Kramer [in Kramer, 1994] experimented with an auditory display that represented eight time-varying physiological variables of a computer-simulated patient:

1. Body temperature,
2. Heart rate,
3. Blood pressure,
4. Blood carbon dioxide level,
5. Respiratory rate,
6. Atrio-ventricular dissociation,
7. Fibrillation, and
8. Pupillary reflex.

The first five variables varied continuously with time while the last three had binary states. Fitch and Kramer used sounds that mapped naturally to the variables for their sonification: The heart rate was sonified with a sound resembling the beating of a heart and the respiratory rate was sonified with a sound resembling a person breathing. They then applied modifications to attributes of these two sounds to signify variations in the remaining variables. For example, they varied the pitch of the heart sound with variations in blood pressure. The design of their experiment was similar to Bly's dissertation research [Bly, 1982]. Participants for the experiment were given three treatments in random order: 1) auditory display only; 2) visual display only; and 3) combined auditory visual display. The visual display used in the experiment was similar to the nominal visual display of the physiological data by used in the medical community. Participants were asked to respond with the proper remedial actions when abnormal indications of the variables manifested in the treatment display. The results showed that the participants responded to indication of abnormalities faster when using the auditory display than the other two treatments. This example reinforces the assertion that effective AD enables the user to perceive information better than a visual display of the same information.

Blattner et al. sought to enhance a two-dimensional static graphics displays, i.e., maps, with sonification [Blattner et al.; in Kramer, 1994]. First, they examined the structure of sound and organized it similar to the linguistic hierarchy. The basic parameters of sound such as frequency and volume belong in the lexical level of sound. The next level of sound is the syntactic level. In this level, *earcons*, the auditory equivalent of icons, are formed by *motives*. Motives are short sequences of tones created by manipulating the lexical parameters of sound. The semantic level of sound is the highest level in the structure of sound. This is the level in which a combination of earcons represents an expression that can be interpreted and understood by the user. This structuring of sound is similar to object-oriented programming in computer languages, where hereditary relationships exist between objects representing program elements. Blattner et al. used these so-called “earcons” to represent traditionally visual cartographic data on a digitized two-dimensional map displayed using a computer. As the user pointed to a location on the map using a mouse, earcons would sound to present information about the location that could not be displayed on the map.

Barras and Zehner [2000] developed a Responsive Workbench that allows the user to interact with the multidimensional data from well-logs. A well-log is the recording of a geological attribute, such as neutron density or radiation level, along the path of a hole in the ground drilled for geological survey. The Responsive Workbench sonifies well-logs of different geological attributes of a drill hole by representative audible clicks, similar to a conventional Geiger counter, to indicate levels of the parameters with respect to depth along the well-logs. The user accesses the data with the probe of a “Virtual Geiger Counter” on a three-dimensional visual display of the well-logs. The user selects the well-log of an attribute to analyze with the Virtual Geiger and points the probe at the region of interest on the well-log. The Virtual Geiger then produces audible clicks representing the value of the geological attribute in a well-log at the region for the user to hear and analyze. The Virtual Geiger also allows the user to hear well-logs of several attributes at the same time. The simultaneous sonification of several well-logs enable the user to evaluate the relations between the well-logs. This bimodal display is a popular emerging technique in the field of oil and gas exploration.

Hermann, Meinicke, and Ritter [2000] applied Principal Curve Sonification (PCS) to multidimensional data sets in order to evaluate the structure of the data set acoustically. The principal curve of a data set with continuous parameters is the projection of the principal components of the data set onto one dimension. PCS uses a *model-based sonification* scheme to sonify a data element based on its relationship from its projection onto the principal curve. The model sonifies the data element in a way that is intuitive for the user to understand. PCS represents each data element with a tick sound; again, similar to a Geiger counter. The distance from the projection on the principal curve is proportional to the volume of the tick. The tick is spatially located relative to the reference orientation of the user along the principal curve. Any additional feature of the data element, e.g., its class label, is represented by the frequency of the tick. The auditory display of PCS presents the user time-variant auditory scenes of the data as the user proceeds along the principal curve. The user can thus assess the structure of the multivariate distribution of the data set.

We present examples of both pure auditory displays and multimodal displays because of our fledgling concept of a virtual environment for the analysis of complex data sets: An immersive environment such that the analyst can use more than just visual and auditory perceptions to extract information from the simultaneous display of many dimensions of a complex data set. This concept presumes that such a multimodal display improves perception and situation awareness by engaging more senses for perception. However, because we can't get there from here, yet, we want to first evaluate the feasibility of sonifying data from our FDE for analysis.

E. AN AUDITORY DISPLAY OF DATA SONIFICATION USING JASS

1. Java Audio Synthesis system (JASS)

Java Audio Synthesis System (JASS) is an open-source sound synthesis software developed by van den Doel and Pai [2001]. We reviewed other sound synthesis software, e.g., Csound, before choosing JASS for our sonification. We had two reasons for choosing JASS. First, JASS is written in Java and benefits from object-oriented programming and platform independence. The other sound synthesis programs are

written in programming languages other than Java, such as C++, or other specialized languages. Second, JASS is a cost-free, open-source program.

JASS is designed to produce model-based sound effects for simulations. JASS uses three core abstract classes called *unit generators* (UGs): `In`, `Out`, and `InOut`, and two interfaces, `Sink` and `Source`, as building blocks for sound synthesis. The `Source` interface contains methods for maintaining an audio buffer. The `Sink` interface contains methods for maintaining and storing `Source` objects. The unit generator `In` implements the `Sink` interface. Thus, an extension of the `In` abstract class enables retrieval of audio buffers. The unit generator `Out` implements the `Source` interface. Thus, an extension of the `Out` abstract class enables the production audio buffers. The unit generator `InOut` implements both `Sink` and `Source` interfaces. Thus, an extension of the `InOut` abstract class can produce and retrieve audio buffers.

The *engine* package of JASS contains the above unit generators and interfaces. JASS has two other packages. The *generator* package contains extensions of the abstract classes in the engine package for basic sound processing and synthesis. The *render* package also contains extensions of the abstract classes in the engine package, but these extensions are used to interface with JavaSound Application Programming Interface (API) in order to produce the desired sonification using the sound hardware in the computer. Classes in the render package also perform basic utility functions such as formatting audio data and designing simple graphical user interface (GUI).

JASS also provides some examples of sonification to simulate sounds and sound effects on the Internet: <http://www.cs.ubc.ca/~kvdoel/jass/>.

2. Sonification Procedure

We sonified the output data sets of the FDE in order to perform factor screening and analysis. The response data sets from our FDE were most similar to the sets of values of transfer functions as mentioned previously. In essence, the MANA distillation is the waveshaping transfer function, and it produces the response for sound synthesis.

For our sonification, we considered the simulation as the waveshaping function and attempted to synthesize sound from the batches of simulation output data. We performed the following six steps in order to sonify our data:

1. Serialize the response data sets into data streams.
2. Perform 0th-order mapping of data by mapping each element of a response data stream directly to the amplitude of the response waveshape of the data stream.
3. Specify the sampling rate and upload a response data stream into an audio buffer using JASS.
4. Use JASS to store the buffer and stream the data in the buffer to the sound card at the specified sampling rate.
5. The sound card synthesizes sounds based on the variations of the data stream in the audio buffer.
6. Repeat the sonification for the remaining data streams from our FDE. We sonify the data streams of the MOPs and MOEs from each batch of simulation run. Hence, we have twelve sonified data streams.

We performed Steps 1 and 2 using Microsoft® Excel. We created two Java classes using JASS to perform Steps 3 through 6. The `DataStreamSonfication` class is the user interface that allows the user to specify the sample rate, buffer size, and file name from which to read the data stream. It calls the `DataStreamBuffer` class, which reads the data stream from a file into the audio buffer and computes the buffer for sound synthesis. We then sonified all MOE and MOP data streams from the FDE.

Recall that the output frequency of the signal is the quotient of the sampling rate divided by the buffer size. Furthermore, recall that each data stream in the audio buffer is a batch of response parameters from our FDE. Because each batch of data has 500 rows of seeds, we actually have 500 cycles in each batch. Thus, the output frequency is the quotient of the sampling rate divided by the sample size, multiplied by the number of cycles in the buffer. Hence, if we sample a buffer contain one batch of data stream that has 40,500 samples at a sampling rate of 40,500 samples per second, the actual output frequency is 500 Hz. This output frequency corresponds to the lowest driving frequency in our FDE, which is one cycle per row of data. Hence all other driving frequencies and indicator frequencies are multiples of the unitary driving frequency. We did not alter the volume of the output manually because the amplitude of the response controls the output. Before we began sonification, we adjusted the volume of the speakers at our PC to an audible level and refrained from any manual adjustments until we completed our sonification unless the sound was too loud or too soft for comfort.

In addition to sonifying the data streams, we also created reference data streams that were digitized oscillations at the five driving frequencies. For example, the sinusoid

of the lowest driving frequency, 1 cycle per 81 samples, was digitized into 81 samples per cycle to represent a row of data from the FDE. The values of the digitized oscillations were real numbers between 1 and -1. We then replicated this digitized sinusoid to fill the data stream with five hundred rows of the same sinusoid. We also created a noise data stream using the random number generator in Microsoft® Excel and generated uniform random variates. Before we sounded the actual data streams we listened to sonifications of these data streams to verify our sonification. For example, the sonified reference data stream of the lowest driving frequency (1 cycles per 81 samples) in our FDE resulted in a 500 Hz pure tone. The sonified noise data stream resulted in white noise.

3. Results

When we heard the sounds of the sonified data streams, we were able to characterize at least three aspects of the sound: noise, signal, and volume. The noise in the sound indicated the random component of the response. The signal represented the response. The intensity of the sound indicated the amplitude of the signal, i.e., the strength of the response. The timbre of the signal indicated the complexity of the response, i.e., the number of indicator frequencies that significantly affect the response. Finally, a comparison between the presence of noise and signal indicated the relative intensity of these attributes in the sound. Note this is similar to performing a Signal-to-Noise Ratio (SNR) comparison real-time by listening to the sound. In order to determine the relative levels of each of these attributes, we listened to the data streams of the same MOP or MOE from all three batches.

The following is a description of a representative sonified data set that we heard. Wav files for the data streams are available from the author upon request.

Data streams of blue agents killed: The presence of white noise in all three sonified data streams sounded about the same. The white noise sounded like high-pressure air diffusing into the atmosphere. However, the presence of signals in each sonified data stream sounded different from the others. In the data stream of Batch 1, the signal had a dominant 500 Hz tonal component with some distortions and light buzzes. In Batch 2, the dominant signal was at a noticeably higher pitch that sounded hollow. There also were various tonal components at very high frequencies. In Batch 3, the

dominant signal had the highest pitch of the three sonified data streams. These differences made sense because we rearranged the driving frequency assignments for each of the three data streams. The volume of all three sonified data streams sounded about the same. However, the noise was “in front of,” i.e., masking to some extent, the signals.

Data streams of red agents killed: The white noise sounded similar to that in the data streams of red agents killed. In Batch 1, various tonal components could be heard, with the dominant tonal sounding similar to the 500-Hz pure tone. The dominant tonal component had a different timbre than the pure tone; it had more distortions and buzzes. In Batch 2, the dominant tonal in the signal was slightly higher than the one in Batch 1, and buzzes were more evident. In Batch 3, the dominant tonal had the highest pitch of the three data streams. The volumes were all about the same, but the signals were definitely in front of the noise.

Data streams of FERs: The white noise in these sonified data streams sounded soft and grainy. In Batch 1, the dominant tonal sounded like the 500-Hz pure tone, but it had other higher-pitch tonal components that were noticeable. In Batch 2, the dominant had a higher pitch than Batch 1, with marginally noticeable higher-pitch tonal components. In Batch 3, the dominant tonal had the highest pitch of the three data streams. Additionally, we noticed a rhythmic click in the sound of the data stream for Batch 3. The click occurred at the end of the buffer before the buffer was played back. The volumes were all noticeably lower than the data streams of the MOPs. The data stream of Batch 2 sounded a little louder than the rest, but Batch 3 had a noticeably lower volume than the other two data streams. The presence of noise and signal were about the same in all three data streams.

Data streams of ERs: The white noise was very grainy and crackly in these data streams, like the tearing of a piece of sandpaper. In Batch 1, again the dominant tonal sounded like the 500-Hz pure tonal. In Batch 2, the dominant tonal had a higher pitch. In Batch 3, the dominant tonal had the highest pitch of the three data streams. The volumes were lower than those of the MOPs and about the same as the FERs. The

graininess of the noise made the noise sounded in front of the signal in all three data streams more than data streams of the other parameters.

4. Discussion of Results

We can explain the similarity between the volumes by the sensitivity of human hearing to logarithmic change in volume. Recall that the total number of blue agents in the scenario is 34, and the total number of red agents is 14. The logarithmic variations in the number of respective agents killed are small. Hence, the differences in volumes between data streams of the MOPs were not very noticeable; they sounded roughly the same. Examining the visual spectra associated with the parameters from all three batches, we see that the noise patterns agree with the visual spectra. The noises for the MOPs account for more variability than the noises of the MOEs. Hence, the noises of the MOPs are not only louder, but also more saturating than the noises of the MOEs.

5. Other Concepts for Sonification

In addition to the methods employed in this thesis, there are many other methods available for the sonification of the simulation data set and subsequent manipulations and analysis of the auditory display. For example, one concept is an auditory display that supplements the visualization of the response parameters. We chose dimensions of the data set to display visually. We also chose from the remaining dimensions of the data set those we wish to sonify. We supplement the visual display with data sonification using slider bars to sonify regions of interest to our analysis. This is similar to the PCS example mentioned previously.

F. CONCLUSIONS

We seek to answer two questions from our development of sonification and auditory display:

Question 1. *How does this sonification display improve data analysis?*

When we compared the qualitative characterization of the sonified data streams to the visual spectra of the four response parameters, we saw—and heard—agreements between the visual display and the sonification of the data sets. Therefore, we believe that we have proven the feasibility of representing simulation data from the FDE with our

sonification. Furthermore, we note that each data stream contains the response from all observations from each batch, i.e., each data stream has 40,500 observations of one MOP or MOE from the simulation output data set. Recall that our sampling rate is 40,500 samples per second. Hence, in one second, we can hear the entire batch of MOPs. Furthermore, we are able to differentiate between data streams with respect to the three sonic attributes after listening to each data stream for just a few seconds. Therefore, we believe that data sonification may have the potential of becoming an efficient qualitative analysis technique of complex data sets that saves time in computational processes and data analysis.

Question 2. *What can be obtained by this sonification and this auditory display that you can't obtain with visualization?*

Based on our results, we assert two implications of our sonification with respect to data analysis. First of all, in addition to the possibility of efficiently sampling the data space using the frequency domain approach, data analysis using our sonification may reduce the number of simulation runs required for data collection while enabling the analyst to inject more complexity in the response by simultaneously varying more factors in the frequency domain experiment. When we examine an “orchestrated” selection of observations over the entire data space, the multimodal representation imparts a more representative rendering of the chaotic behavior and/or the hidden periodicities induced by our frequency domain experiment. Secondly, data analysis by our sonification may be performed more quickly than visualization. We hear the entire set of 40,500 observations in one second when we set the sampling rate to 40,500 samples per second. Based on our results, we can qualitatively differentiate between data streams within a few seconds. Thus, each observation contributes to the analysis, and the overall sound is a “symphonic” representation of the data space.

We also assert that our sonification method provides the first step towards a robust auditory display that will enable different users to arrive at the same conclusions. Recall that the output frequency of our sonification is determined by the sampling rate, the sample size, and any inherent cycles in the sample. Our sonification allows the user to specify the sample rate to sample the buffer for a given buffer size, within the

limitations of the computer and with consideration of the Nyquist criterion and aliasing. Therefore, the user may determine the frequency at which to analyze the data stream that best suits his or her hearing acuity. Because the entire data stream is sonified at the same proportion, theoretically all other attributes of the sound should remain the same.

G. RECOMMENDATIONS FOR FUTURE RESEARCH

Based on what we learned, we offer the following recommendations:

First, a user interface is needed to permit usability of the sonification. Currently, we perform the sonification using command-line arguments in DOS, and thus the current process is definitely not user-friendly. We suggest a graphic user interface (GUI) shell for the sonification. We believe the GUI should have at least the following functions:

1. File utility functions that allow the user to administrate the files of data streams.
2. A visual display that incorporates the spectral analysis portion of the FDE and displays the spectra from the analysis.
3. Sonification functions that enable the user to select data streams to sonify and analyze.

In particular, we strongly recommend including filter functions, e.g., notch filters, to permit the user to filter out the noise and analyze the signal, as well as other signal analysis techniques to decompose the signal into component frequencies for the user to correlate with the respective terms in the regression analysis of the response data set. Because the lowest meaningful frequency is the output frequency, a high-pass filter may be useful to minimize low frequency noise. Moreover, because the indicator frequencies are discrete, band pass filters may also be useful in filtering out noise at non-indicator frequencies above the output frequency. Finally, notch filters may be useful for listening to particular indicator frequencies.

Furthermore, we propose the following general guidelines for the design of a human participant experiment to validate our claim that this auditory display can improve the data analysis of multidimensional data sets. The experiment tasks participants to perform factor screening of a nonlinear model, e.g., a second-order model with interaction terms like our meta-model from FDE. The design of the experiment would be similar to Bly [1982] and Fitch and Kramer [in Kramer, 1994]:

1. Select participants with experiences in data analysis.
2. Train and apply three treatments to participants: 1) a visual-only display; 2) a combination of visual and auditory display; and 3) a “beta” version of this auditory-only display.
3. Task the participant to determine the factors in the model that contribute significantly to the response of the model using the treatment displays.
4. Measure the amount of time for the participant to complete the factor screening and the percentage of correct and incorrect identification. In addition, survey the participants for background information and for feedbacks about the treatment displays, as well as personal preferences of the treatment displays.

We recognize that it may require additional research to develop a bimodal display for such a comparison; thus it may be practical to compare a visual display with the beta version of this display. We recommend using typical data visualization and analysis programs such as S-Plus 2000 as the visual-only display.

We also recommend efforts to spatialize data streams using headphone-based spatialization techniques so that a user can analyze multiple parameters simultaneously. These techniques allow sounds to be presented in 3-D with complete externalization around the user’s head [Shilling & Shinn-Cunningham, 2002]. These techniques allow the user to hear and recognize multiple data streams simultaneously.

Because data sonification is still an emerging field of application, there are no established standards for designing sonification schemes—only intuition, art, and past examples of sonification techniques to emulate. A good resource is the International Community for Auditory Display (ICAD), formerly the International Conference on Auditory Display. The ICAD website, <http://www.icad.org>, contains papers and conference proceedings relevant to the diverse applications of auditory displays and data sonification.

Finally, based on our results and the examples of other sonification efforts, we believe that using sonification to harvest data in data farming has significant potential for success. Therefore, we strongly recommend future research to explore the possibilities of data farming with sonification.

THIS PAGE INTENTIONALLY LEFT BLANK

V. SUMMARY OF RESULTS

For this thesis, we attempted to apply an interdisciplinary approach to operations research. First, we examined the feasibility of data farming in the frequency domain and conducted FDEs using a peace enforcement scenario in MANA. By considering the simulation as a waveshaping function, we then attempted to develop an auditory display by sonifying data streams of four measures from the output data set using a direct-mapping technique that maps values of the measures to the amplitudes of the wave shape.

With respect to data farming in the frequency domain, we have achieved our key objectives of evaluating the frequency domain approach as a means of planting and harvesting data efficiently. The results from our FDE confirm the regression model from which we selected our factors. Furthermore, the resulting visual spectra from the FDE are useful for simultaneous comparison of factors and interactions that we seek to evaluate. Therefore, we conclude that the frequency domain approach is not only a feasible method for data farming, but also a useful technique for factor screening that is easy to generate. From our results, we believe that the frequency domain approach to simulation output analysis will help operations analysts and decision makers answer complex and difficult questions about military operations and/or other complex operations.

With respect to the purposes of developing an auditory display using sonification, we have developed a simple auditory display using data sonification that can be used for factor screening of multidimensional data sets for data farming. We arranged response parameters from our FDE in data streams and sonified the streams by performing direct mapping of response to amplitude of output oscillations. The resulting sounds contained noise and signals that agree with the visual spectra from harvesting our data in the frequency domain. Even though we did not conduct an experiment to validate our goals for creating a data sonification display, our informal results indicate that it is feasible to use an auditory display for data analysis in data farming environment.

We are very encouraged by our attempt in integrating simulation output analysis and human factors. We believe there is significant value in further research to develop an

auditory display using sonification that will benefit data farming in the frequency domain. One potential application for our display is in the training and development of analytical judgments of complex data sets. Entry-level data analysts can generate a data set using an FDE and sonify the resulting data streams using the display. The analysts can then use the display to explore the response data set and understand how different parameters contribute to the variability of the response parameters both visually and auditorily. In addition, we suggest a very interesting and worthwhile improvement to the display that renders simultaneous representation of several response parameters using spatial audio. We conjecture that this improvement may allow analysts using the display to appreciate the contributions of factors to responses from an overall perspective, thus gaining insight into the complexity of the responses.

We embarked on our research having in mind the ultimate goal of a virtual environment for the analysis of complex data sets. We imagine that someday an immersive environment created through a multimodal display will enable the operations analyst to use more than just visual and auditory perceptions in order to improve understanding of the complexity of military operations. Through this research effort we believe we have advanced one step closer toward this goal, and strongly recommend continued research and development to make this goal a reality.

APPENDIX A: MANA SCENARIO INFORMATION

Mana Scenario File
Nov 30 2000

The below summarizes which 22 factors will be examined, the overview of the mission, a definition of peace enforcement, and the rules of engagement which Blue forces would receive for conducting the operation.

INITIAL 22 FACTORS IDENTIFIED FOR EXTENSIVE EXAMINATION

- A. Blue Platoon HQ move precision - amount of randomness in blue movement
- B. Blue Squad 1 move precision - amount of randomness in blue movement
- C. Blue Squad 2 move precision - amount of randomness in blue movement
- D. Blue Squad 3 move precision - amount of randomness in blue movement
- E. Blue Platoon HQ in contact personality element w1 - controls propensity to move towards agents of same allegiance
- F. Blue Squad 1 in contact personality element w1 - controls propensity to move towards agents of same allegiance
- G. Blue Squad 2 in contact personality element w1 - controls propensity to move towards agents of same allegiance
- H. Blue Squad 3 in contact personality element w1 - controls propensity to move towards agents of same allegiance
- I. Blue Platoon HQ in contact personality element w2 - controls propensity to move towards agents of enemy allegiance
- J. Blue Squad 1 in contact personality element w2 - controls propensity to move towards agents of enemy allegiance
- K. Blue Squad 2 in contact personality element w2 - controls propensity to move towards agents of enemy allegiance
- L. Blue Squad 3 in contact personality element w2 - controls propensity to move towards agents of enemy allegiance
- M. Blue Platoon HQ injured personality element w1 - controls propensity to move towards agents of same allegiance
- N. Blue Squad 1 injured personality element w1 - controls propensity to move towards agents of same allegiance
- O. Blue Squad 2 injured personality element w1 - controls propensity to move towards agents of same allegiance
- P. Blue Squad 3 injured personality element w1 - controls propensity to move towards agents of same allegiance
- Q. Blue Platoon HQ injured personality element w2 - controls propensity to move towards agents of enemy allegiance
- R. Blue Squad 1 injured personality element w2 - controls propensity to move towards agents of enemy allegiance
- S. Blue Squad 2 injured personality element w2 - controls propensity to move towards agents of enemy allegiance
- T. Blue Squad 3 injured personality element w2 - controls propensity to move towards agents of enemy allegiance

U. Blue movement range for all squads - controls movement speed of agents

V. Red personality element w8 - controls propensity to move towards enemies (Blue) in situational awareness map which are of threat level 1

Notes:

Factors A-D will have settings of 1-513 in increments of 4 which will correspond to 129 levels

Factors E-T and V will have settings of -64 to 64 in increments of 1 which will correspond to 129 levels

Factor U will have settings of 72 to 200 in increments of 1 which will correspond to 129 levels

Firepower and sensor ranges of all allegiances will be equal to amplify personalities - furthermore a high firepower range in essence has blue destroying red right from the simulation start

Red and Blue will have same stealth settings, but Yellow will have increased stealth which represents that although we initially knew them to be of the same allegiance as Blue it is difficult to ascertain they have switched allegiances

MISSION

Blue Mission: Destroy red element of 5-7 soldiers, who are equipped with small arms, located in vicinity of area of operation (AO) Cobra within the next two hours in order to facilitate UN food distribution and military convoy operations.

Scheme of Maneuver: Blue uses a light infantry platoon composed of three nine-man rifle squads and a platoon HQ of seven soldiers containing two machine gun teams. Their movement scheme is one squad up and two squads back with platoon HQ following the lead squad (2nd squad). 1st squad task is to follow and support 2nd squad with purpose of destroying red element. Follow-on task is to secure area of operation Python for subsequent UN food distribution and military convoy operations. 2nd squad task to conduct movement to contact with purpose of destroying red element. Follow-on task is to secure area of operation Cobra for subsequent UN food distribution and military convoy operations. 3rd squad task is to follow and support 2nd squad with purpose of destroying red element. Follow-on task is to secure area of operation Boa (a small urban area with four building structures) for subsequent UN food distribution and military convoy operations. After 2nd squad secures area of operation Cobra, Platoon HQ moves to area of operation Boa to provide supporting fires for 3rd squad. Red has 5 member element located vicinity Cobra. Red also has two 2 member elements patrolling along movement routes of blue squads 1 and 2. Red has 2 member element in vicinity Boa. A non-hostile (and Blue allegiance) Yellow 3 member element is initially in Blue's starting location. After discovering no safe water in vicinity Rattler, Yellow becomes hostile against Blue, seeks small arms from vicinity Boa, and moves to vicinity Python.

PEACE ENFORCEMENT (From FM 100-23)

Peace Enforcement is the application of military force or the threat of its use, normally pursuant to international authorization, to compel compliance with generally accepted resolutions or sanctions. The purpose of Peace Enforcement is to maintain or restore peace and support diplomatic efforts to reach a long-term political settlement.

RULES OF ENGAGEMENT FOR SCENARIO

1. (U) Situation. Basic OPLAN/OPORD.

2. (U) Mission. Basic OPLAN/OPORD.

3. (U) Execution.

(U) Concept of the Operation.

(U) If you are operating as a unit, squad, or other formation, follow the orders of your leaders.

(U) Nothing in these rules negates your inherent right to use reasonable force to defend yourself against dangerous personal attack.

(U) These rules of self-protection and rules of engagement are not intended to infringe upon your right of self defense. These rules are intended to prevent indiscriminate use of force or other violations of law or regulation.

(U) Commanders will instruct their personnel on their mission. This includes the importance of proper conduct and regard for the local population and the need to respect private property and public facilities. The Posse Comitatus Act does not apply in an overseas area. Expect that all missions will have the inherent task of force security and protection.

(U) ROE cards will be distributed to each deploying soldier (see below).

(U) Rules of Self-Protection for all Soldiers.

(U) US forces will protect themselves from threats of death or serious bodily harm. Deadly force may be used to defend your life, the life of another US soldier, or the life of persons in areas under US control. You are authorized to use deadly force in self-defense when--

(U) You are fired upon.

(U) Armed elements, mobs, and/or rioters threaten human life.

(U) There is a clear demonstration of hostile intent in your presence.

(U) Hostile intent of opposing forces can be determined by unit leaders or individual soldiers if their leaders are not present. Hostile intent is the threat of imminent use of force against US forces or other persons in those areas under the control of US forces. Factors you may consider include--

(U) Weapons: Are they present? What types?

(U) Size of the opposing force.

(U) If weapons are present, the manner in which they are displayed; that is, are they being aimed? Are the weapons part of a firing position?

(U) How did the opposing force respond to the US forces?

(U) How does the force act toward unarmed civilians?

(U) Other aggressive actions.

(U) You may detain persons threatening or using force which would cause death, serious bodily harm, or interference with mission accomplishment. You may detain persons who commit criminal acts in areas under US control. Detainees should be given to military police as soon as possible for evacuation to central collection points.

(U) Rules of Engagement. The relief property, foodstuffs, medical supplies, building materials, and other end items belong to the relief agencies distributing the supplies until they are actually distributed to the populace. Your mission includes safe transit of these materials to the populace.

(U) Deadly force may be used only when--

(a) (U) Fired upon.

(b) (U) Clear evidence of hostile intent exists (see above for factors to consider to determine hostile intent).

(c) (U) Armed elements, mobs, and/or rioters threaten human life, sensitive equipment and aircraft, and open and free passage of relief supplies.

(U) In situations where deadly force is not appropriate, use the minimum force necessary to accomplish the mission.

(U) Patrols are authorized to provide relief supplies, US forces, and other persons in those areas under the control of US forces. Patrols may use deadly force if fired upon or if they encounter opposing forces which evidence a hostile intent. Nondeadly force or a show of force should be used if the security of US forces is not compromised by doing so. A graduated show of force includes--

(a) (U) An order to disband or disperse.

(b) (U) Show of force/threat of force by US forces that is greater than the force threatened by the opposing force.

(c) (U) Warning shots aimed to prevent harm to either innocent civilians or the opposing force.

(d) (U) Other means of nondeadly force.

If this show of force does not cause the opposing force to abandon its hostile intent, consider if deadly force is appropriate.

(U) Use of barbed wire fences is authorized.

(U) Unattended means of force (for example, mines, booby traps, trip guns) are not authorized.

(U) If US forces are attacked or threatened by unarmed hostile elements, mobs, and /or rioters, US forces will use the minimum amount of force reasonably necessary to overcome the threat. A graduated response to unarmed hostile elements may be used. Such a response can include--

(a) (U) Verbal warnings to demonstrators in their native language.

(b) (U) Shows of force, including the use of riot control formations.

(c) (U) Warning shots fired over the heads of the hostile elements.

(d) (U) Other reasonable uses of force, to include deadly force when the element demonstrates a hostile intent, which are necessary and proportional to the threat.

(U) All weapons systems may be employed throughout the area of operations unless otherwise prohibited. The use of weapons systems must be appropriate and proportional, considering the threat.

(U) US forces will not endanger or exploit the property of the local population without their explicit approval. Use of civilian property usually be compensated by contract or other form of payment. Property that has been used for the purpose of hindering our mission will be confiscated. Weapons may be confiscated and demilitarized if they are used to interfere with the mission of US forces.

(U) Operations will not be conducted outside of the landmass, airspace, and territorial seas of Somalia. However, any US force conducting a search and rescue mission shall use force as necessary and intrude into the landmass, airspace, or territorial sea of any country necessary to recover friendly forces.

(U) Crew-served weapons are considered a threat to US forces and the relief effort whether or not the crew demonstrates hostile intent. Commanders are authorized to use all necessary force to confiscate and demilitarize crew-served weapons in their area of operations.

(a) (U) If an armed individual or weapons crew demonstrates hostile intentions, they may be engaged with deadly force.

(b) (U) If an armed individual or weapons crew commits criminal acts but does not demonstrate hostile intentions, US forces will use the minimum amount of necessary force to detain them.

(c) (U) Crew-served weapons are any weapon system that requires more than one individual to operate. Crew-served weapons include, but are not limited to tanks, artillery pieces, antiaircraft guns, mortars, and machine guns.

(U) Within those areas under the control of US forces, armed individuals may be considered a threat to US forces and the relief effort, whether or not the individuals demonstrate hostile intent. Commanders are authorized to use all necessary force to disarm and demilitarize groups or individuals in those areas under the control of US forces. Absent a hostile or criminal act, individuals and associated vehicles will be released after any weapons are removed/demilitarized.

(U) Use of riot control agents (RCAs). Use of RCAs requires the approval of CJTF. When authorized, RCAs may be used for purposes including, but not limited to--

(1) (U) Riot control in the division area of operations, including the dispersal of civilians who obstruct roadways or otherwise impede distribution operations after lesser means have failed to result in dispersal.

(2) (U) Riot control in detainee holding areas or camps in and around material distribution or storage areas.

(3) (U) Protection of convoys from civil disturbances, terrorists, or paramilitary groups.

(U) Detention of Personnel. Personnel who interfere with the accomplishment of the mission or who use or threaten deadly force against US forces, US or relief material distribution sites, or convoys may be detained. Persons who commit criminal acts in areas under the control of US forces may likewise be detained.

(1) (U) Detained personnel will be treated with respect and dignity.

(2) (U) Detained personnel will be evacuated to a designated location for turnover to military police.

(3) (U) Troops should understand that any use of the feet in detaining, handling or searching Somali civilians is one of the most insulting forms of provocation.

4. (U) Service Support. Basic OPLAN/OPORD.

5. (U) Command and Signal. Basic OPLAN/OPORD.

ROE Card

Nothing in these rules of engagement limits your right to take appropriate action to defend yourself and your unit.

1. You have the right to use force to defend yourself against attacks or threats of attack.

2. Hostile fire may be returned effectively and promptly to stop a hostile act.

3. When US forces are attacked by unarmed hostile elements, mobs, and/or rioters, US forces should use the minimum force necessary under the circumstances and proportional to the threat.

4. You may not seize the property of others to accomplish your mission.

5. Detention of civilians is authorized for security reasons or in self-defense.

Remember

The United States is not at war.

Treat all persons with dignity and respect.

Use minimum force to carry out the mission.

Always be prepared to act in self-defense.

APPENDIX B: RESPONSE SPECTRA OF THE FREQUENCY DOMAIN EXPERIMENT

Figures 16 through 27 are spectra of MOEs and MOPs from each batch of MANA distillation runs. All spectra have window size of 10,000 ($M = 10,000$) and sample size of 40,500 observations ($N = 40,500$). The figures are also color-coded: The bars for the number of Blue Agents killed are in blue; the number of Red Agents killed are in red; the FER in green; and the ER in purple.




Figure 16. Spectrum of the number of Blue Agents killed in Batch 1.




Figure 17. Spectrum of the number of Red Agents killed in Batch 1.




Figure 18. Spectrum of the FER in Batch 1.




Figure 19. Spectrum of the ER in Batch 1.




Figure 20. Spectrum of the number of Blue Agents killed in Batch 2.




Figure 21. Spectrum of the number of Red Agents killed in Batch 2.




Figure 22. Spectrum of the FER in Batch 2.




Figure 23. Spectrum of the ER in Batch 2.




Figure 24. Spectrum of the number of Blue Agents killed in Batch 3.




Figure 25. Spectrum of the number of Red Agents killed in Batch 3.




Figure 26. Spectrum of the FER in Batch 3.




Figure 27. Spectrum of the ER in Batch 3.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX C. JASS CODE OF SONIFICATION PROGRAM

The following Java source codes are created for sonification of data streams from FDE using JASS [Van den Doel and Pai, 2001]. JASS requires the sonification designer to extend the abstract classes and implement one of the inherited methods, `computeBuffer()`, to perform basic sound synthesis. We created `DataStreamBuffer`, which extends the `Out` abstract class, to read the data stream into the audio buffer and implement the `computeBuffer()` method. We also created a main class, `DataStreamSonification`, to sonify a data stream from our FDE data sets using a `SourcePlayer` object in JASS.

```

/*
 * DataStreamBuffer.java
 *
 */

/**
 *
 * @author Hsin-Fu Wu, LT, USN
 */
package test;

import java.io.*;
import jass.engine.*;
public class DataStreamBuffer extends Out {
    private File file1;
    private FileReader reader1;
    private BufferedReader in1;

    /** Creates a new instance of DataStreamBuffer */
    public DataStreamBuffer(int bufferSize) {
        super(bufferSize);
    }

    public DataStreamBuffer(int bufferSize, String file1) {
        super(bufferSize);
        this.file1 = new File(file1);
        try {
            if (!this.file1.exists()) {
                throw new RuntimeException("No such file: " +
this.file1.getName());
            }
            reader1 = new FileReader(this.file1);
            in1 = new BufferedReader(reader1);
        } catch (Exception e) {}
    }

    /** Compute the next buffer and store in member float[] buf.
     * This is the core processing method which will be implemented
     * for each generator.
     */
    protected void computeBuffer() {
        try {
            for (int i = 0; i < getBufferSize(); i++) {
                float a = Float.parseFloat(in1.readLine());
                buf[i] = a;
            }
        } catch (Exception e){}
    }
}

```

```

/*
 * DataStreamSonification.java
 * This class synthesizes sound from a numerical data stream supplied
by the user.
 * The inputs are sampling rate, buffer size, and the file name of the
data stream.
 * The output frequency = sampling rate (samples per second) / buffer
size (samples per cycle)
 * If the data stream has inherent cycles, then the actual output
frequency equals the number
 * of cycles times the output frequency.
 * The buffer is computed using the MyOutReadFromBuffer object.
 *
 */

/**
 *
 * @author Hsin-Fu Wu, LT USN
 */
package test;
import jass.render.*;
import jass.engine.*;
import jass.generators.*;
import java.awt.*;
import java.applet.*;

public class DataStreamSonification extends Applet{

    /**
     * @param args the command line arguments
     *
     */
    public static void main(String[] args) {
        float srate = Float.parseFloat(args[0]);
        int bufferSize = Integer.parseInt(args[1]);
        DataStreamBuffer streamer = new DataStreamBuffer(bufferSize,
args[2]);
        try {
            new SourcePlayer(bufferSize,0,srate, streamer).start();
        } catch(Exception e) {}
    }

}

```

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

Barrass, Stephen, and Zehner, Björn, “Responsive Sonification of Well-logs,” paper presented at the International Conference on Auditory Display, Atlanta, Georgia, April 2000.

Brandstein, Alfred G. and Horne, Gary E., “Data Farming: A Meta-technique for Research in the 21st Century,” in F.G. Hoffman and Gary Horne, eds., *Maneuver Warfare Science*, United States Marine Corps Combat Development Command, 1998.

Bly, Sara A., *Sound and Computer Information Presentation*, PhD dissertation, University of California, Davis, Davis, California, 1 March 1982.

Chatfield, C., *The Analysis of Time Series: An Introduction*, Fifth Edition, Chapman & Hall, 1996.

Cioppa, T. M., *Efficient Nearly Orthogonal and Space-Filling Experimental Designs for High-Dimensional Complex Models*, Doctoral Dissertation, Naval Postgraduate School, Monterey, California, June 2002.

Hermann, T., Meinicke, P., and Ritter, H., “Principal Curve Sonification,” paper presented at the International Conference on Auditory Display, Atlanta, Georgia, April 2000.

Jacobson, S., Buss, A., and Scruben, L., “Driving Frequency Selection for Frequency Domain Simulation Experiments,” *Operations Research*, v. 39 n. 6, November-December 1991.

Kramer, Gregory, ed., *Auditory Display: Sonification, Audification, and Auditory Interfaces*, Santa Fe Institute Studies in the Sciences of Complexity, Vol. XVIII, Addison-Wesley, 1994.

Law, Averill M. and Kelton, W. David, *Simulation Modeling and Analysis*, 3rd ed., McGraw-Hill Series in Industrial Engineering and Management Science, 2000.

Lucas, T. W., Sanchez, S. M., Brown, L. P., and Vinyard, W. C., “Better Designs for High-Dimensional Explorations of Distillations,” in G. Horne and S. Johnson, eds., *Maneuver Warfare Science*, United States Marine Corps Combat Development Command, 2002.

National Instruments Corporation, Application Note 041, *The Fundamentals of FFT-Based Signal Analysis and Measurement*, by Michael Cerna and Audrey F. Harvey, July 2000.

Sanchez, S. M., Moeeni, F., and Sanchez, P. J., "So Many Factors, So Little Time..." forthcoming in *Technology and Operations Review*, 2002.

Sanchez, Paul J. and Buss, Arnold H., "A Model for Frequency Domain Experiments," paper presented at the Winter Simulation Conference, Atlanta, Georgia, December 1987.

Sanders, Mark S. and McCormick, Ernest J., *Human Factors in Engineering and Design*, 7th ed., McGraw-Hill, 1993.

Schruben, Lee W. and Cogliano, Vincent James, "Simulation Sensitivity Analysis: A Frequency Domain Approach," paper presented at the Winter Simulation Conference, Atlanta, Georgia, December 1981.

Shilling, Russell D., Letowski, Tomasz, and Storms, Russell, "Spatial Auditory Displays For Use Within Attack Rotary Wing Aircraft," paper presented at the International Conference on Auditory Display, Atlanta, Georgia, April 2000.

Shilling, Russell D. and Shinn-Cunningham, Barbara G., "Virtual Auditory Displays," to appear in *Handbook of Virtual Environment Technology*, K. Stanney, ed., Lawrence Earlbaum, Assoc., Inc., in press, 2002.

Tufte, Edward R., *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, Connecticut, 1983.

Van den Doel, Kees, and Pai, Dinesh K., "JASS: Java Audio Synthesis System for Programmers," paper presented at the International Conference on Auditory Display, Helsinki, Finland, August 2001.

Wagner, Daniel H., Mylander, W. Charles, and Sanders, Thomas J., *Naval Operations Analysis*, 3rd ed., Naval Institute Press, 1999.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Susan M. Sanchez
Naval Postgraduate School
Monterey, California
4. Russell D. Shilling
Naval Postgraduate School
Monterey, California
5. Paul J. Sanchez
Naval Postgraduate School
Monterey, California
6. Thomas W. Lucas
Naval Postgraduate School
Monterey, California
7. Thomas M. Cioppa
TRADOC Analysis Center
Monterey, California
8. James N. Adams
United States Marine Corps
Quantico, Virginia
9. Alfred G. Brandstein
Marine Corps Warfighting Lab
Quantico, Virginia
10. Dr. Gary E. Horne
Marine Corps Warfighting Lab
Quantico, Virginia

11. Mary L. MacDonald
Dumfries, Virginia
12. Hsin-Fu Wu
Naval Postgraduate School
Monterey, California