Modeling & Simulation Capability for Consequence Management

Vic Baker
Advanced Systems Technologies
Mid-Atlantic Technology, Research & Innovation Center (MATRIC)
Morgantown, WV, USA
vic.baker@matricresearch.com

***Abstract*—The Modeling & Simulation Capability for Consequence Management prototype allows emergency planners to design and run dynamic, time aware, ArcGIS 10 based "what-if" gaming simulations depicting the impacts upon critical infrastructure and resource consumption during a large-scale mass evacuation from the National Capital Region to a host region. The capability enables emergency planners to run multiple iterations of a simulated event with different resource allocations including the ability to activate and deactivate shelters, place barriers, and enforce fuel rationing.**

The technical aspects of the capability rely on the ESRI ArcGIS API for Flex coupled with custom geoprocessing services housed on an ESRI ArcGIS Server.

The prototype system was developed for the FEMA Regional Catastrophic Preparedness Grant Program initially focusing on the eastern panhandle of West Virginia. The simulation logic is designed for extensibility to encompass all state-wide counties, neighboring states, and geographical regions/districts.

Keywords-simulation; geographic information systems; evacuation; national capital region

I. Introduction

The Modeling & Simulation (M&S) Capability for Resource Consumption and Consequence Management is a prototype system that was developed to assist with attack and disaster preparation, recovery and response planning for a mass evacuation event from the National Capital Region to a host State. The M&S capability enables emergency planners to design and run dynamic, "what-if" simulations depicting the impacts of a mass evacuation upon critical infrastructure and resource consumption.

The M&S prototype system simulates the impact of evacuees on the transportation infrastructure through congestion modeling and the corresponding consumption of fuel, water, first aid, and shelter availability. Users can activate and deactivate shelters, place barriers, and enforce fuel rationing using the simulation tool. Activation of resources (such as shelters and first aid) have a startup and per occupant cost which are deducted from a "warchest" which is defined by the user for their particular simulation experience. The type and inventory of supplies (i.e., cots, quantity 3,000) available in the warchest are defined and customized by the user during the creation of a simulation.

The simulation logic utilizes over 30 subject matter expert approved scenario parameters that are user-defined including the number of evacuees, percent of evacuees seeking medical treatment, average gallons of water used per shelter occupant, and which points of ingress to use.

Users can pause, play, fast forward, and rewind the "time-aware" simulation maps to visualize road congestion and simulated resource consumption over time. Additionally, users can generate temporal reports for specific resources such as hospital bed availability, or broader reports encompassing all resources and their depletion rates.

The M&S system consists of a presentation layer, geoprocessing services layer, and a data layer. The presentation layer is a web based user interface that enables users to sign into the system, create new simulations by entering scenario parameters, review prior simulation results, and launch simulations through an Adobe Flex application based on the ESRI ArcGIS for Flex API. The geoprocessing services layer consists of custom simulation logic services hosted on an ESRI ArcGIS Server 10. The simulation logic services were written in Python [6] and C# [7] and optimized for performance and scalability with a focus on algorithm efficiency, runtime efficiency, and utilization of parallel processing. The data layer consists of information for the four key resources (water, first aid, shelter, fuel) including location, inventory, capacity, and resource specific fields. This information was collected by combining assets from a variety of data sources including academic, state, and federal data agencies.

FEMA Regional Catastrophic Preparedness Grant Program

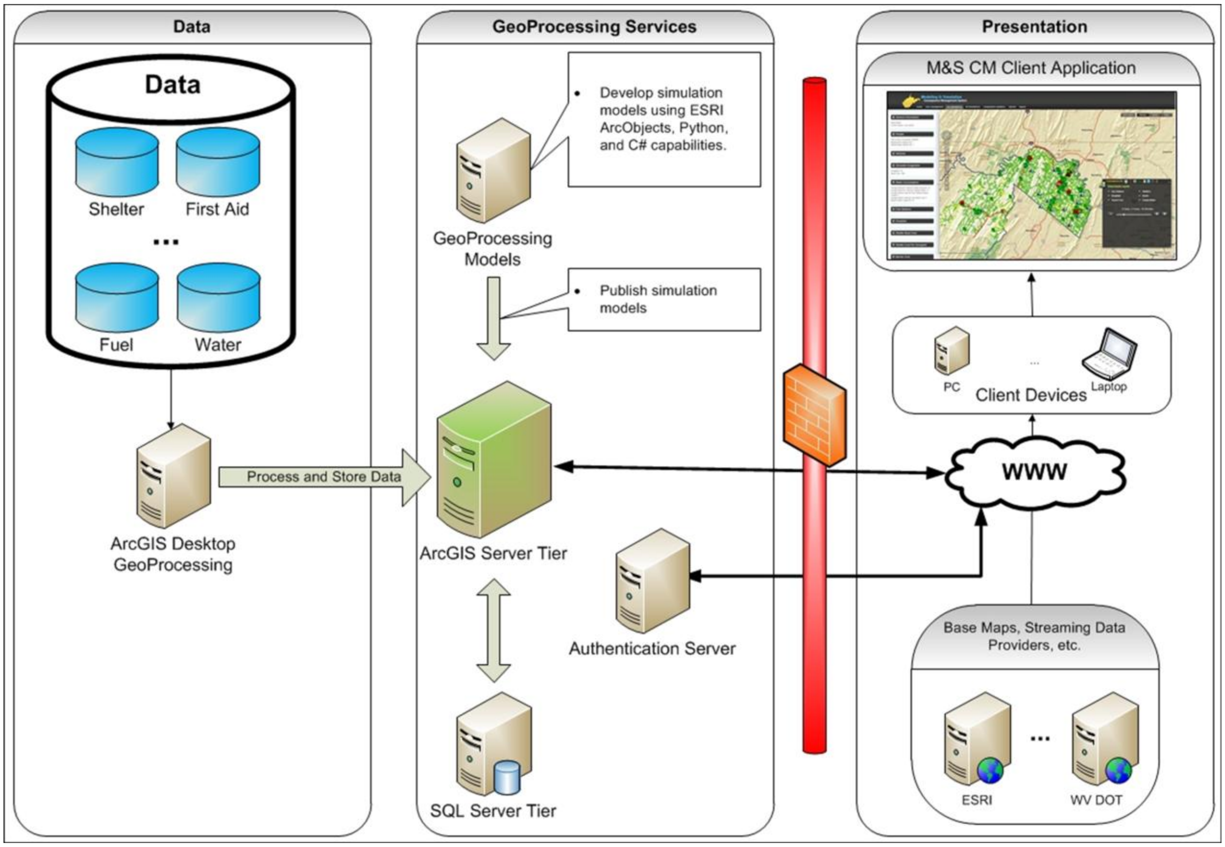

Figure 1. System Design

The prototype was developed for the FEMA Regional Catastrophic Preparedness Grant Program focusing on the four West Virginia Eastern panhandle counties of Jefferson, Berkley, Morgan and Hampshire [1]. The scope of the prototype project focused on data sets for the four selected counties; however, the simulation logic is designed for extensibility to encompass all state-wide counties, neighboring states and geographical regions/districts based on availability of resource data.

II. System Architecture

A. System Design

The M&S prototype system contains three components: Presentation, Geoprocessing Services, and Data (Fig. 1). The ArcGIS Desktop is used to create the maps and models that are published onto the ArcGIS Server. The ArcGIS Server communicates with a Microsoft SQL Server Database for data storage and retrieval. Geoprocessing Models on the ArcGIS Server utilize the stored data from the SQL Server Database in the simulation process.

The ArcGIS Server communicates with client devices through the web to display outputs and receive inputs. The authentication server is used to validate user credentials to grant access to the system. The ArcGIS Server can use outside third party data on the web that can be integrated into the User Interface displayed on the client devices.

B. Presentation Layer

The M&S capability provides emergency planners with the ability to explore the consequences of their decisions in regards to the allocation of available resources. Emergency planners can run simulation exercises and receive reports on resource consumption and availability and compare parameters and outputs of multiple simulations to compare strategies. The M&S capability provides a cost effective way to re-run scenarios and examine different plans or procedures for handling an event.

The M&S capability addresses the emergency planner's need to view data on available resources on hand, convey resource changes as a simulated scenario unfolds, enable changes to resource allocation and display the effect of these changes, and enable running of multiple iterations of the M&S simulation capability for the same scenario with different resource allocations.

The M&S user interface (Fig. 2, Fig. 3, and Fig. 4) is a web based portal with an embedded Adobe Flex component used to interact with the GIS information. The Flex component uses the ArcGIS API for Flex which takes full advantage of the mapping, geo-coding, and geo-processing capabilities of ESRI ArcGIS services. The ArcGIS API for Flex allows for the creation of internet applications that interact with ArcGIS Servers. Coupling the Adobe Flex framework with ESRI ArcGIS API enables advanced functionality for users to interact with the M&S application to perform tasks such as adding and removing data layers, modifying simulation parameters, and displaying GIS data such as raster grids, vector layers, or even streaming data, or simulation results.

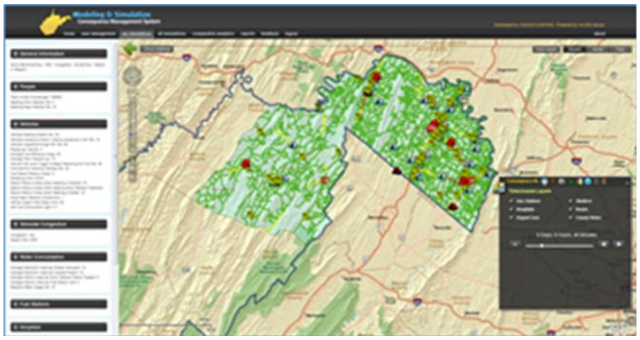

Figure 2. Modeling & Simulation Portal: Flex Viewer - User Interface

The M&S system is extensible to regions outside of the initial four eastern panhandle counties of West Virginia. To illustrate extensibility of the system the M&S team also utilized a nine county region surrounding the state's capital of Charleston, West Virginia to simulate a mass evacuation.

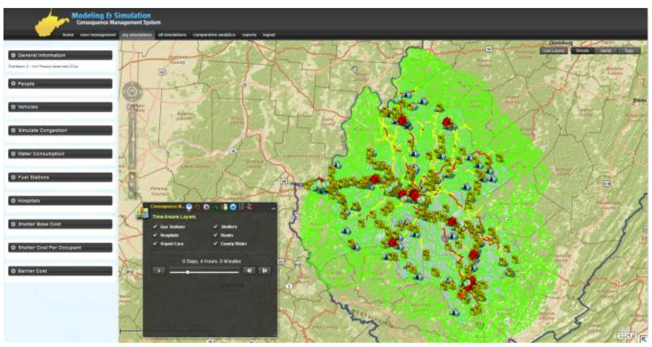

Figure 3. Consequence Management Simulation of Charleston, WV

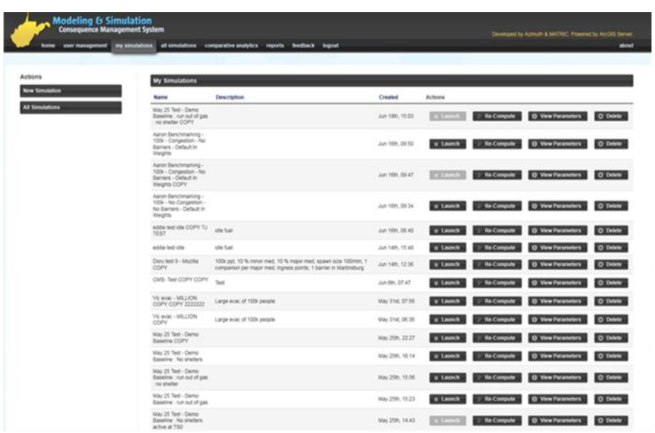

Figure 4. Modeling & Simulation Portal – "My Simulations"

C. *Consequence Management / Simulation Logic*

The M&S development team designed and implemented modular powerful, efficient, scalable event simulation and analytical geoprocessing models in Python [6] and C# [7] utilizing the ESRI ArcObjects API. The simulation models were published to an ArcGIS 10 Server where they are accessible by the web enabled user interface of the Portal via the embedded Adobe Flex plugin and custom tools built using the ArcGIS API for Flex.

A simulated event requires many geoprocessing models that work together and contain logical interdependencies of both the model's input and output parameters. Each geoprocessing model is a logical component responsible for a specific portion of the overall simulation. The system utilizes concepts from Object Oriented Programming to design the component models, taking into account interdependencies into the model design. Object Oriented Simulation design principles are ideal for breaking a large problem into manageable, testable modules and we fully utilized these approaches in the design and implementation of the event simulation system. For example, the routing geoprocessing model was designed to accept a logical set of inputs (barrier locations and route start and destination) and generate outputs (a polyline of the newly generated route). Thereafter, any simulation model that requires the routing functionality does not need to replicate the routing logic, but rather can call the existing routing model.

D. *Data*

The M&S system uses geographic information systems (GIS) data that was aggregated from many data sources. The main sources of data include: West Virginia Geographical Information System Technical Center, HSIP Gold, HSIP Freedom, TeleAtlas, NavTeq, National Shelter System (NSS), FEMA Headquarters, Department of Homeland Security, WV Infrastructure and Jobs Development Authority.

There is an ongoing effort from the development team to find the best available data, and most of the time it requires a large amount of effort, time, research and funding. Fragmented datasets are located in various government agencies and commercial organizations; frequently these data resources need to be merged, and new geodatabases and metadata were created. Google Earth Pro was used for verification purposes.

III. Design Rationale

The critical issues of design for this project are scalability, performance, and reliability. Performance decisions have been made by conducting various tests to determine the server deployment configuration setup. To determine the best scalability, performance, and reliability of the architecture, the development team consulted with ESRI on design strategies. Consultation from ESRI coupled with the project team's experience on similar projects has provided a firm basis for determining the server configuration setup. The development team used the following configuration for internal development and testing:

The development team utilized VMware's virtualization technology to setup three virtualized 64-bit servers:

Virtual Server #1 – ArcGIS Server: The ArcGIS Server was configured with Microsoft Windows Server 2008 R2, 16 gigabytes of RAM, 4-core Intel Xeon 5640 CPU @ 2.67 gigahertz, and a 4-core licensed version of ArcGIS 10 with Service Pack 1.

Virtual Server #2 - SQL Database Server: The Microsoft SQL Server was configured with Microsoft Windows Server 2008 R2, 16 gigabytes of RAM, 2-core Intel Xeon 5640 CPU @ 2.67 gigahertz, and Microsoft SQL Server 2008.

Virtual Server #3 – Portal Host Server: The Portal server was configured with Ubuntu version 11.04,2 gigabytes of RAM, 4-core Intel Xeon 5640 CPU @ 2.67 gigahertz, and Apache 2.2.17.

Throughout development of the M&S capability, MATRIC conducted system performance testing of the virtual server configurations and their distribution on physical hosts. The MATRIC server hardware consists of two HP blades each having dual 6-core Intel Xeon 5640 CPUs @ 2.67 gigahertz, 32 gigabytes of RAM, and a 16 terabyte storage area network. MATRIC conducted a variety of benchmarking tests to identify limitations and performance "best practices" opportunities of where the M&S development effort could push the boundaries of computational performance for this project.

The development team analyzed many simulation performance factors including: ESRI ArcGIS Server performance tuning, database design, virtualization optimization, runtime performance of software development languages, and parallel processing tuning.

Benchmarking tests and performance tuning tasks included runtime performance evaluation of implementation languages and performance optimization tuning for the virtualized ArcGIS Server and Microsoft SQL Server.

Performance optimization tuning for the virtualized ArcGIS Server and Microsoft SQL Server consisted of:

- Experimenting with the number of CPUs per virtual machine.

- Placing the virtual machines on the same physical host blade versus different physical host blades.

- ArcGIS Server parameters (number of concurrent asynchronous geoprocessing services per Server Object Container).

- Performance testing of synchronous versus asynchronous geoprocessing services.

- Virtualized core performance: tests were conducted evaluating performance a one-core virtualized ArcGIS server versus four-core virtualized ArcGIS server. The virtualized four-core ArcGIS Server had better performance than the virtualized one-core ArcGIS Server.

Runtime performance evaluation of implementation languages included:

- Testing of ArcGIS Modelbuilder [4] generated geoprocessing services versus Python [6] versus managed C# [7] versus unmanaged C++.

- Modelbuilder [4] GUI was helpful for prototyping of early M&S models.

- Python with ArcObjects had similar runtime performance to Modelbuilder generated models with the added benefits that Python enabled the development of more complex simulation logic and supported asynchronous geoprocessing.

- Managed C# with ArcObjects had over 300% better runtime performance over Python based geoprocessing services and consumed less process memory.

- Unmanaged C++ using ArcObjects via COM demonstrated better runtime performance than managed C# however the unmanaged C++ / COM required additional development overhead and it was decided not to pursue the unmanaged C++ ArcObjects approach due to the short development window.

Evaluation of data storage performance of file geodatabases versus ArcSDE included:

- File geodatabases had the best performance for loading GIS data into simulation logic; however file geodatabases have a read lock enabled while being accessed which limits them to a single concurrent user running a simulation.

- ArcSDE connection enables the accessing of GIS data by multiple concurrent simulations. Read times were slower using the ArcSDE (on average twenty seconds) compared to a file geodatabase (approximately six seconds), however for multiple users it was observed that the ArcSDE cached the data and load times decreased from twenty seconds for subsequent users.

A four-core virtualized ArcGIS Server configuration was chosen and tested to identify optimal performance on the virtualize server configuration. Benchmarking tests were conducted that compared how varying the number of SOC processes and the number of asynchronous geoprocessing services affect compute time for a re-routing geoprocessing service [3]. Table 1 below illustrates the results obtained from performance testing on the virtualized development servers.

Table 1. Asynchronous Geoprocessing Benchmakrs (virtualized)

Max SOC Processes	Max Concurrent Threads	Time (seconds)
1	1	585
1	3	463
1	6	474
1	999	462
6	1	597
6	3	280
6	6	267
6	999	270
999	1	596
999	3	289
999	6	266
999	999	275

Testing has confirmed that three threads at a time are ideal because additional threads have insignificant returns, causing unnecessary load on the server.

To determine anticipated simulation performance comparisons between the virtualized development environment and the nonvirtualized hosting environment, performance comparison metrics were calculated based on CPU benchmark scores obtained from CPU Benchmark [2] and virtualization performance conducted by Miller and Sakowicz [5]. Table 2 illustrates the CPU benchmarks of the physical CPUs and the anticpated 32% performance hit of running ArcGIS Server in a virtualized environment.

Table 2. Comparison of CPU Benchmarks and Virtualization Penalty

CPU	CPU Benchmark
MATRIC Nonvirtualized: Intel Xeon E5640 @ 2.67 gigahertz	5053
MATRIC Virtualized: 32% estimated performance decrease	3436
WVGISTC Nonvirtualized: Intel Xeon E5540 @ 2.53gigahertz	4298

Based on benchmark tests, it has been theorized that the WV GIS Tech Center's nonvirtualized servers may show increased performance over the virtual machines that MATRIC used for research, development, and testing. ESRI has determined 32% performance degradation when ArcGIS Servers are deployed in a virtualized environment [5].

The M&S system is hosted by the WV GIS Technical Center (WVGISTC) in a nonvirtualized configuration for the ArcGIS Server and Microsoft SQL Server. The WVGISTC has the following server specifications:

Physical Server #1 – ArcGIS Server: The ArcGIS Server was configured with Microsoft Windows Server 2008 R2, 12 gigabytes of RAM, 8-core Intel Xeon CPU @ 2.53 gigahertz, and a 4-core licensed version of ArcGIS 10 with Service Pack 1.

Physical Server #2 - SQL Database Server: The Microsoft SQL Server was configured with Microsoft Windows Server 2008 R2, 12 gigabytes of RAM, 8-core Intel Xeon CPU @ 2.53 gigahertz, and Microsoft SQL Server 2008.

Virtual Server #3 – Portal Host Server: The Portal server was configured with Ubuntu version 11.04, 2 gigabytes of RAM, 4-core Intel Xeon CPU @ 2.67 gigahertz, and Apache 2.2.17, MySQL 5.x, PHP5.

Table 3 illustrates the benchmarking results compared how varying the number of SOC processes and the number of asynchronous geoprocessing services affect compute time for a re-routing geoprocessing service used in the virtualized tests of Table 1.

Table 3. Asynchronous Geoprocessing Benchmarks (nonvirtualized)

Max SOC Processes	Max Concurrent Threads	Time (seconds)
1	1	368
1	3	314
1	6	316
1	999	314
6	1	358
6	3	194
6	6	176
6	999	172
999	1	361
999	3	198
999	6	174
999	999	172

Table 4 illustrates the comparison between virtualized and nonvirtualized server configurations for the simulation system comparing the “dev” virtualized configuration at MATRIC vs the “host” nonvirtualized configuration at the WVGISTC. The table shows execution time for major components of the simulation including data loading, computing, and publishing. The benchmark comparisons were each conducted three times with the average times shown. Furthermore, simulations were conducted with and without congestion modeling enabled. Congestion modeling is computed within the “Computing” module listed in the below table. The nonvirtualized “host” servers at the WVGISTC performed approximately 19% faster than the virtualized development “dev” servers at MATRIC even though the MATRIC hardware had better CPU benchmark scores.

Table 4. Execution Times of Virtualized vs Nonvirtualized Server Hosts

Time (seconds)	**Mean: Dev Virtual (No Conges-tion)**	**Mean: Host Nonvirtual (No Conges-tion)**	**Mean: Dev Virtual (Conges-tion)**	**Mean: Host Nonvirtual (Conges-tion)**
Data Loading	55.67	47.67	54.67	50.33
Computing	9.33	7.33	10.00	6.67
Publishing	16.00	10.33	19.00	10.33
Totals	81.00	65.33	83.67	67.33

IV. CONCLUSION

The Modeling & Simulation project provides a unique opportunity to blend simulation, GIS, and gaming fundamentals into a hybrid client/server system with performance requirements not typical for current state of the art GIS applications. The development team implemented custom geoprocessing tools on the ESRI ArcGIS 10 platform for congestion modeling, resource consumption modeling, and robust reporting and comparative analytics tools. Runtime optimization and tuning was important to the success of the project. The team conducted detailed performance assessments of simulation logic based on Python, C#, and C++ implementations and hosting benchmarks of virtualized vs nonvirtualized server configurations. Future work for this project will focus on an expanded geographic region encompassing the NCR, the potential to include additional resources beyond the initial four utilized for this study, and performance optimizations such as the utilization of server-side graphics processing units for GIS geoprocessing.

ACKNOWLEDGMENT

MATRIC would like to acknowledge our teaming partners Azimuth Incorporated, WV Division of Military Affairs and Public Safety (WVDMAPS), All Hazards Consortium, and the WV GIS Technical Center for their collaboration on these efforts.

REFERENCES

[1] Azimuth Incorporated, Modeling and Simulation (M&S) Capability for Consequence Management (CM) Project Proposal, 2010

[2] CPU Benchmarks, 2011, Retrieved from http://www.CPUBenchmark.net/high_end_cpus.html

[3] ESRI Thread Forum on: Use low isolation map services to support more 4X more concurrent users? (9.3.1). Retrieved from http://forums.arcgis.com/threads/18969-Use-low-isolation-map-services-to-support-more-4X-more-concurrent-users-%289.3.1%29

[4] ESRI, 2011. ArcGIS Spatial Analyst. Retrieved from http://www.esri.com/software/arcgis/extensions/spatialanalyst/graphical.html

[5] Miller, E. and Sakowicz A., 2010. ArcGIS Server Performance and Scalability— Performance Factors and Optimization. Presented at ESRI Developer Summit, Palm Springs, CA, March 22-25, 2010

[6] Python Programming Language, 2011. Retrieved from http://www.python.org/

[7] Visual C# Developer Center, 2011. Retrieved from http://msdn.microsoft.com/en-us/vcsharp