AN AIR TRAFFIC MANAGEMENT (ATM) SYSTEM PERFORMANCE MODEL

Edward J. Modzelesky, Tejal Topiwala, Sarah E. Stalnaker, Michael A. Hermes,
The MITRE Corporation, McLean, VA

Abstract

In this paper, we discuss our approach in developing a system performance model of the Enhanced Traffic Management System (ETMS), which is a part of the Federal Aviation Administration's (FAA's) Air Traffic Management (ATM) system. The ETMS, developed and operated at the Volpe National Transportation Systems Center, is a complex system with a centralized processing architecture and a geographically wide data distribution system. The model we are presenting here is being developed using QASE, a Commercial-Off-The-Shelf (COTS), systems engineering tool. The model predicts the system performance of ETMS (e.g., server utilization, network utilization, and response time) if potential hardware, software, communications, and workload changes are implemented.

Keywords

Modeling, performance, system engineering, simulation, ETMS, ATM, ATC, QASE

Background

The role of traffic management is becoming more important in the FAA's overall strategy to cope with increasing air traffic. As a result, a number of functional and infrastructure enhancements are planned for the ETMS. Currently, decision-makers must, however, authorize changes to ETMS with little or no quantitative data on the impact these changes will have on system performance. For instance, when decisions to implement software changes are made, it is not clear if hardware and communications upgrades are needed to support the software changes. With a model, the impact of software changes can be estimated by modeling the software components; executing simulation runs; and monitoring server utilization, network utilization, and response time. Areas of risk or even potential failure may be discovered. Furthermore, the hardware and communications upgrades can be modeled to determine if they are sufficient and by how much. Finally, the software changes can be modeled with different implementation options. For instance, the software can be allocated to different processors or modeled to execute different data flows. The model can thus provide a useful, quantitative "what-if" analysis capability.

ETMS

The ETMS acquires airspace, flight intent and position, and Traffic Flow Management (TFM) data from across the U.S., Canada, England, and soon Mexico. The system integrates all of these data into a national picture of aviation activity and then redistributes this information to FAA facilities, National Airspace Users, and International Air Traffic Management facilities.

The ETMS is composed of two major subsystems: the ETMS Hub and the ETMS Field Site. The ETMS Hub, the system's main computing and communications complex, is located in Cambridge, MA. The ETMS Field Sites, each consisting of a File server for local and Hub data acquisition and processing and multiple Workstations for graphic display and human interface, are located at over 80 FAA facilities. The ETMS has a centralized architecture – the ETMS Hub acquires data, performs processing and forwards data for display at the Workstations. The ETMS Hub uses flight intent information to predict the number of aircraft that will arrive at airports, depart from airports, and occupy adapted volumes of airspace for 15-minute intervals over the next several hours. If the predicted traffic counts exceed adapted thresholds, an alert is generated. The prediction of traffic counts and the generation of alerts are a part of the Monitor Alert capability.

The ETMS Hub provides flight data, weather data, and Monitor Alert data to the Field Sites for display at the Workstations. Figure 1 provides a high level data flow for the ETMS.

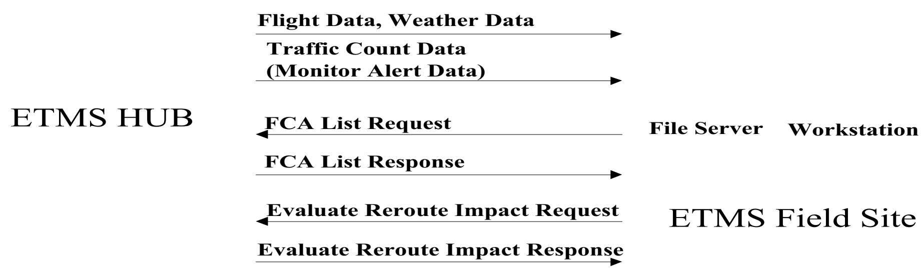

Figure 1. ETMS Data Flow

The ETMS Hub and the ETMS Field Sites are connected via a partially meshed, hierarchical FAA network that guarantees 256 kbps of bandwidth between the ETMS Hub and each of the ETMS Field Sites.

The FAA has decided to add two new capabilities – Flow Constrained Area (FCA) processing and Reroute processing - to the ETMS. The FCA capability allows the traffic manager to define an area around severe weather or congested area and to determine the flights that will traverse the defined area (via the FCA List Request). The Rerouting capability allows the traffic manager to develop a plan to reroute flights around severe weather and to determine the impact on traffic counts (via the Evaluate Reroute Impact Request) before actually rerouting any flights.

Approach

In developing the system performance model, we followed a spiral development cycle. We first concentrated on building the hardware components and then added the software components. We ran experiments and monitored system performance attributes of interest. As better information became available, we incorporated it into the model and again ran appropriate experiments, continuing to refine model components and experiments.

Model Design

Our first step was to determine a useful model abstraction level. Initially, we wanted to model the ETMS Hub and the multiple ETMS Field Sites, but we quickly realized that the simulation runs would take too long. Instead, we decided to just model the ETMS Hub, a single ETMS Field Site, and the Wide Area Network (WAN) connecting the two sites. This is acceptable since the data flow from the ETMS Hub to each ETMS Field Site is nearly identical. The ETMS Hub and the network, however, do handle loading from all of the ETMS Field Sites. In order to represent the workload imposed by the ETMS Field Sites that were not explicitly modeled, we injected workload directly into the ETMS Hub servers. In addition, the workload from external sources is also injected directly into the ETMS Hub servers.

In order to simplify the modeling of the software, we did not include the following processing in the model:

- Initialization Processing: Initialization processing of the system and individual processes is not modeled.
- Generation and maintenance of recovery data: For instance, the primary copy of a software unit transmits data to its backup copy. The generation, transmission, and maintenance of this recovery data is not modeled.
- Failure and recovery processing: If a software unit determines that its database is not current, it will recover data from another database. The generation and transmission of this recovery data is not modeled.

We also decided that the model should be developed in two major phases. In the first phase, the model should reflect ETMS prior to integrating FCA and Rerouting. In the second phase, the components necessary to model FCA and Rerouting capabilities should be added.

Our next step was to understand what information was needed in order to develop the hardware and software components of the model. The hardware components require a hardware architecture and the specification details for processors (type, speed, operating system), storage devices (type, transfer speed), and communication devices (type, protocol, packet size, transfer speed).

Figure 2 illustrates the hardware architecture of the model. The key model inputs are as follows:

- ETMS Hub Server -- HP-K570, 150 million instructions per second (MIPS), HP-UNIX
- Edge Router-- Cisco7206, 10MIPS, Cisco Integrated Operating System (IOS)
- Hub Site Router -- Cisco 8510, 100MIPS, IOS
- Autonomous System (AS) Border Router -- Cisco 7513, 50MIPS, IOS
- Network Node Router -- Integrated Digital Network Exchange (IDNX), 25MIPS, IOS
- Hub Communications Processor -- HP-C360, 90MIPS, HP-UNIX
- File Server, Workstation -- HP-C360, 90MIPS, HP-UNIX
- Wide Area Network link -- 256 kilo-bits per second (Kbps)

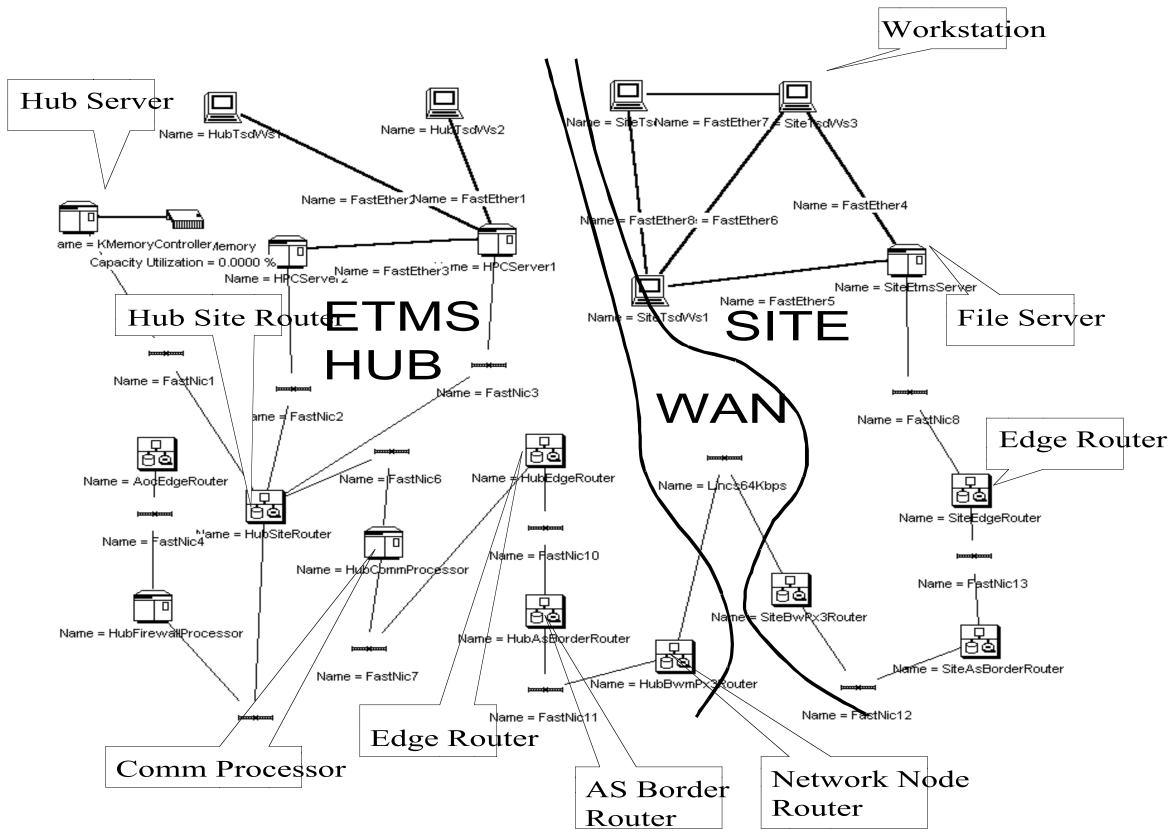

Figure 2. Hardware Architecture

The software components consist of an execution flow for each major processing thread and the workload for each execution flow. At a high level, an execution flow is a sequence of functions that must be performed on the stimuli (or workload) that enter the system. For example, in Figure 3, the execution flow contains Functions A, B, C, and D. These Functions are connected by Flow Connectors F1, F2, F3, and F4.

A Function specifies the processing that must be performed on stimuli. For instance, Function C specifies that for each stimuli, a database is accessed and a set of "x" instructions are executed. Note that an adapted number of instructions are executed for each database access. In addition, the input/output access times are modeled.

The Flow Connectors F1, F2, F3, and F4 specify the probability that the stimuli will traverse its path and size of data.

Note that in this example, Functions A, B, and C are allocated to processor 1 and Function D is allocated to processor 2. Each of the stimuli will incur a load on each processor and communication device between processor 1 and 2.

The key execution flows modeled for the ETMS include the following:

- Flight Data
- Monitor Alert
- Weather Data
- List Request
- Airport Demand List (ADL)
- Flow Constrained Area (FCA)
- Reroute

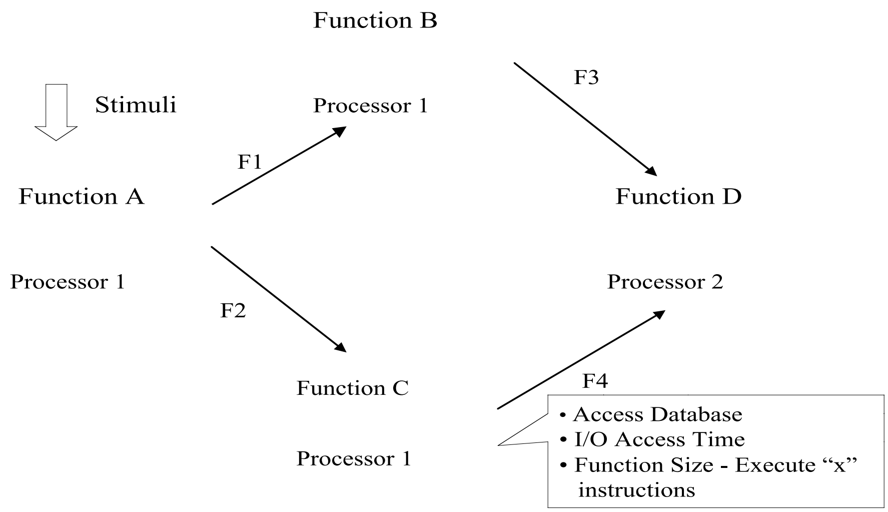

Figure 3. Execution Flow

Key Information Sources

The key information sources that we used to model the hardware components, software components, and new software components are identified. In many cases, we estimated the input.

Hardware Components

No formal architecture or hardware documents were available. We were, however, able to piece together the hardware architecture based on various informal documents, even though they were incomplete and inconsistent with each other. We were able to verify our findings with test experts at the FAA's William J. Hughes Technical Center (WJHTC) and with a communications expert at the Volpe National Transportation Systems Center. In addition, the test experts were able to provide the necessary details of the processors and communication devices.

Software Components

Execution Flows – We used the ETMS Software Design Document, the ETMS Functional Description Document, and an ETMS Version Design Document to develop the execution flows.

Workload – We were able to obtain and analyze archives of data received at an ETMS Field Site from the Hub Server. We used this data to estimate the workload into the Hub Server and into the WAN.

Function Size - Because we did not have access to the source code, we were not able to use a profiler tool to determine the number of instruction per input type. For the ETMS Hub Server, we were able to use a Hewlett-Packard (HP)-Openview client process called Measure Ware Agent (MWA) to collect process utilization data on 1- minute intervals as the system normally processed input. A process is a software unit that can contain processing for multiple input types. We used the collected MWA data along with our estimate of the workload into the Hub Server to estimate the number of instructions executed per input type.

For the File Server, we estimated the function sizes based on the design. For the Workstation, we used the processor speed and the time to bring up a display (measured using a stopwatch) to determine the number of instructions that are executed to bring up a display.

I/O Access Times – We obtained sample data from the WJHTC that indicated that disk access times are frequently 5-6 seconds. Our understanding is that it is the communications protocol used to access the databases (which are in virtual memory) and not the database search algorithms that take up much of 5-6 seconds. We will incorporate measured disk access times into the model when we have more complete data.

New Software Components

We used the FCA Design Document to develop the FCA Execution Flow. We derived a design for the Reroute capability and then developed the Reroute Execution Flow based on the design. We estimated the workload and function sizes for FCA and Reroute capabilities. We were able to obtain process CPU data for FCA and Reroute capabilities that have been implemented in a prototype and used this data to sanity check our model predictions.

Results

Figures 4 and 5 show predicted WAN Link utilization in going from generating and distributing traffic counts or Monitor Alerts once every 5 minutes to once every minute. The Monitor Alert (MA) files contain 50 KB of data. Note the reduction in spare WAN capacity. This analysis was a factor in our customer's decision to increase the frequency of Monitor Alert processing less aggressively, at least initially.

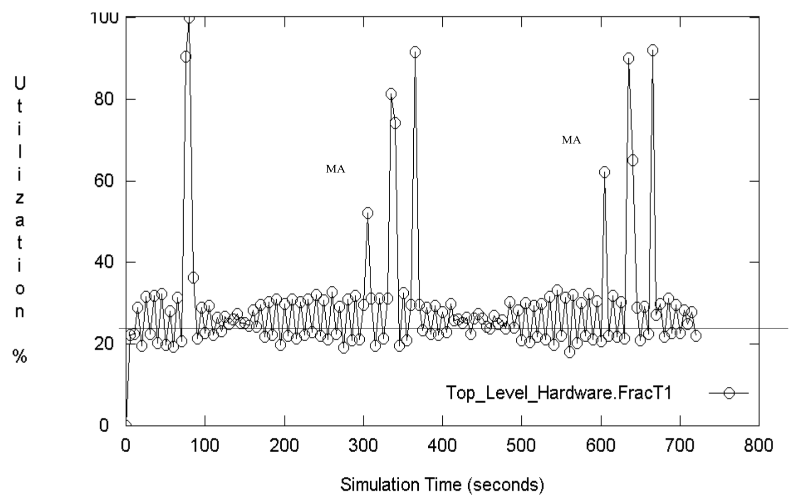

Figure 4. WAN Link Utilization with 5 Min Monitor Alert (MA)

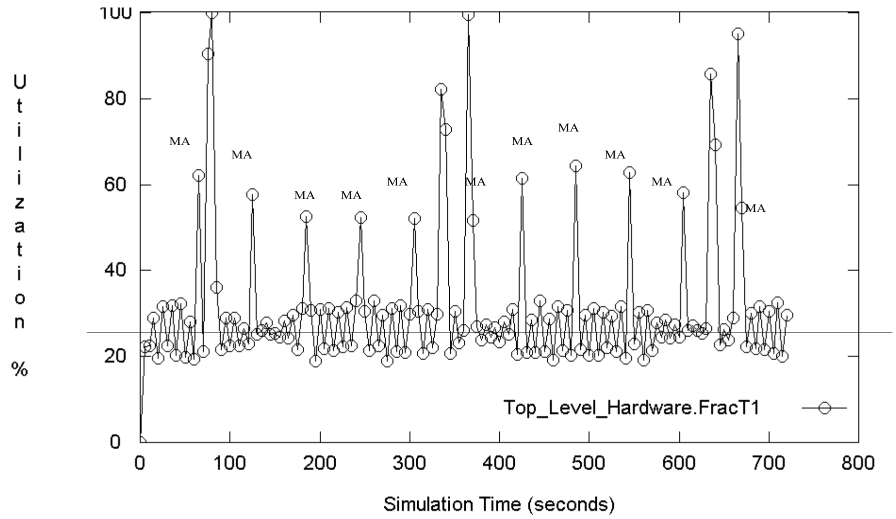

Figure 5. WAN Link Utilization with 1 Min Monitor Alert (MA)

Figures 6 and 7 show the predicted impact to Hub Server utilization to process FCA List Requests and Evaluate Reroute Impact Requests. In Figure 6, the Hub Server is modeled with a 120 MIPS machine. In Figure 7, the Hub Server is modeled with a 150 MIPS machine. Note the small but significant reduction in the Hub Server's peak utilization in Figure 7 versus Figure 6. The ETMS Hub Server was recently upgraded to 150 MIPS.

Figure 8 shows the end-to-end response time for executing the Evaluate Reroute Impact Request as a function of Disk Access time. In order to process the Evaluate Reroute Impact Request, the Hub Server needs to access the disk multiple times. Since we currently do not have good data for disk access times, we modeled a range for disk access times from 0 milliseconds to 50 milliseconds. If a disk access time of 50 milliseconds is used, it may take over 35 seconds for a traffic manager to see the results of an Evaluate Reroute Impact Request.

Conclusion

Our approach to modeling the ETMS has been to use a model abstraction level that allows us to avoid modeling unnecessary details but still provides us insight into system performance. We have followed a spiral development cycle, adding and refining model components, executing simulations, and analyzing results. We have used many sources to acquire the necessary data to develop the model components. A key benefit of the model is that it is a compilation of system information in one place.

Initial results indicate that the model can identify potential performance and capacity risks (or the lack thereof) with a selected implementation option. The customer has already noted the usefulness of the Monitor Alert analysis (see Figures 4 and 5). Further work is being planned to increase the model fidelity and to validate its predictions so that the model can more effectively support decision-makers.

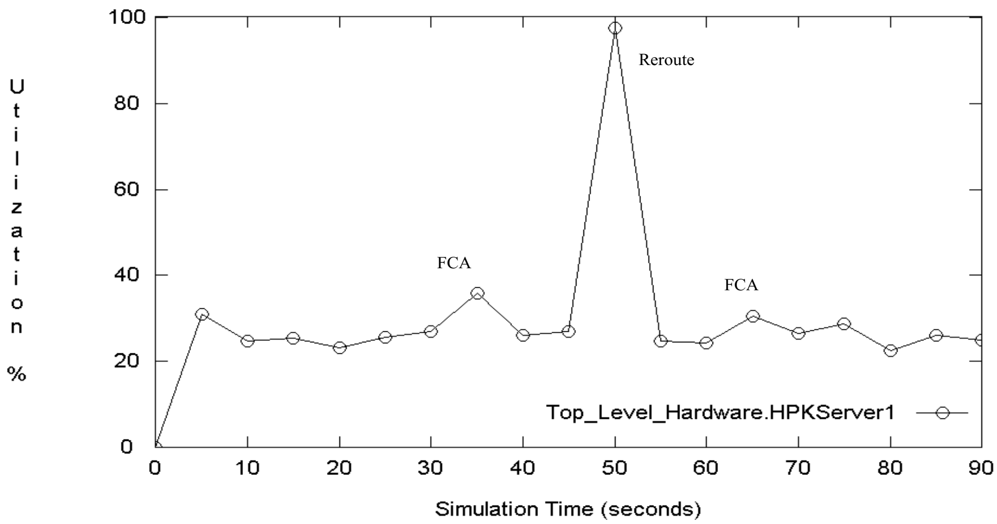

Figure 6. HUB Server Utilization (Hub Server—120 MIPS)

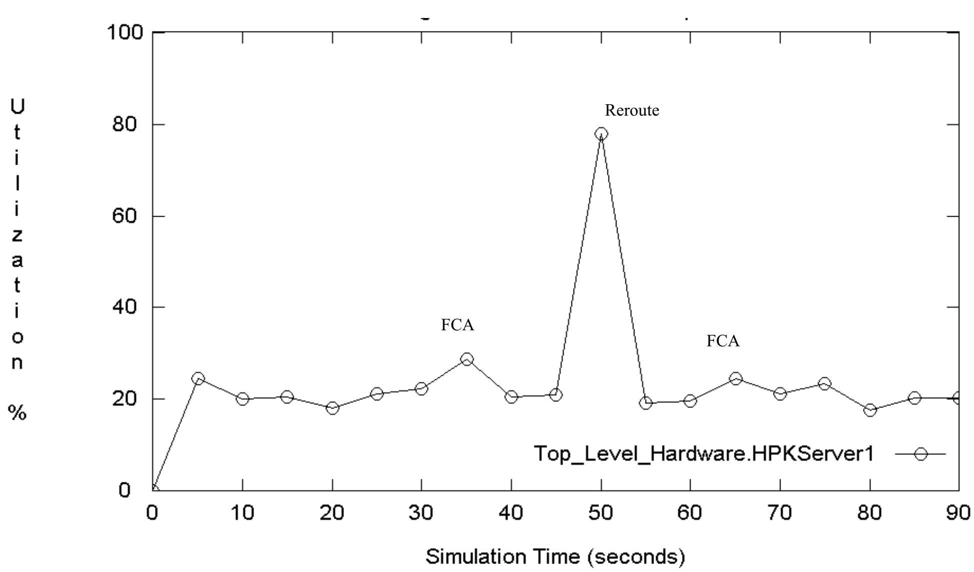

Figure 7. HUB Server Utilization (Hub Server—150 MIPS)

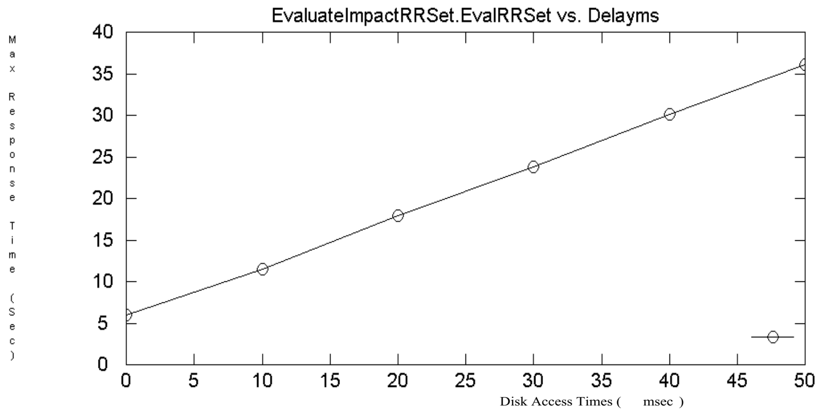

Figure 8. Response Time for Evaluate Reroute Impact