

Airport Simulations Using Distributed Computational Resources

William J. McDermott and Dr. David A. Maluf
NASA Ames Research Center

Yuri Gawdiak
NASA Headquarters

Peter B. Tran
QSS Group, Inc.

To increase air transportation throughput without raising accident rates requires a simulation environment of the highest achievable degree of fidelity. Achieving this national airspace-wide simulation environment demands a revolutionary technology leap in information management. In 2001, using simulation software running over a distributed network of computers, NASA researchers developed a working prototype of a virtual airspace. The Virtual National Airspace Simulation prototype utilized available large-scale resources from a computational grid to attack the problem of improving air transportation safety. It modeled daily flight operations at Atlanta's Hartsfield International Airport by integrating measured operational and simulation data on up to 2,000 flights. By identifying precursors to component failure, this nationwide simulation environment supports the development of strategies for improving aviation safety.

A goal of NASA and the Federal Aviation Administration (FAA) is to develop technologies that contribute to a significant reduction in aviation accidents in the next five to 10 years [1]. The U.S. national airspace is a large and complex system that is defined by the interaction of a large number of entities. In an air transportation system as geographically large as that of the United States, one of the most difficult challenges is transforming massive quantities of highly complex ever-changing data on airport terminal operations into knowledge. The Virtual National Airspace Simulation (VNAS) represented in Figure 1 illustrates this.

The challenge is the 10,000 to 15,000 aircraft in the air at any one time, represented by objects inside the globe in Figure 1. Organizational and aircraft component simulators and a network of computers – represented by models, boxes, and connecting lines in Figure 1 – make up a VNAS. The VNAS application is a suite of components or middleware that expands the scope of airport/air traffic simulation to include whole aircraft modeling [2]. It does this by automating the assignment of large-scale simulation workloads to computers distributed across the Internet, providing researchers in the air transportation community with low-impact access to additional computational resources.

The Simulation Concept

U.S. air carriers account for approximately 50,000 takeoffs and landings each day. To perform an airspace-wide simulation of six sub-assemblies per aircraft on a single day would require 300,000 simulations, generating gigabytes of data in the process. The numbers would be much larger if human and organizational models, such as pilots, air traffic con-

trollers, and airline dispatch were added. Researchers might want to run simulations with the large amount of air traffic monitoring data that are available. To

“Technological progress that enables the creation of large-scale computational grids now makes it possible to transparently assign large-scale scientific workloads to a network of computers distributed across the country.”


process this much data they would have to expend significant computational resources. Frequently they do not have

sufficient computing power for this level of processing.

Technological progress that enables the creation of large-scale computational grids now makes it possible to transparently assign large-scale scientific workloads to a network of computers distributed across the country. The VNAS is one such concept that mirrors that of SETI@Home, which accomplishes its computational workload by automatically selecting available computer resources connected to the Internet. It then uses these spare central processing unit (CPU) cycles in a search of radio telescope data.

In a similar fashion, the VNAS locates available resources from a network of distributed computers and assigns large-scale scientific workloads to them. Tapping into unused cycles allows researchers who are studying complex operational environments access to significantly more computational power than the resources normally available to the individual effort or group. They can then use these resources to drive multi-discipline simulations with the huge amount of air transportation

Figure 1: VNAS Structure



Altitude (ft)	High Speed RPM	Compressor Temp. In (° F)	Compressor Temp. Out (° F)	Thrust (lb)
2,000	8,055	326	855	50,879
2,000	8,057	326	856	51,008
2,100	8,058	326	856	51,082
2,100	8,062	327	857	51,377
2,200	8,063	327	858	51,482
2,300	8,065	327	858	51,571
2,400	8,065	328	858	51,612
2,600	8,063	327	858	51,456
2,700	8,064	327	858	51,468
2,900	8,065	327	858	51,346

Table 1: *Simulation Data Produced by Engine Simulator*

monitoring data that are available.

In 2001, NASA researchers developed a VNAS prototype to demonstrate the concept of complex and computationally intensive integration using distributed computer resources. With this demonstration, NASA researchers also illustrated the technical feasibility of integrating architecture, concepts, and technology covered in this article with a network of computers.

Prototype

The VNAS must provide the underlying infrastructure to support whole aircraft modeling. To achieve this, systems that simulate the operation of key aircraft sub-assemblies such as jet engines are used to provide risk assessments and performance parameters of incoming and outgoing flights from high-density airports.

To produce risk assessment, systems have to process large amounts of input data and return output data to requestors who may be located across the country. A typical request may involve processing flight operations data about a particular airport for multiple days. After simulation and performance parameters have been generated, output data in graphical or tabular form is sent to the user. The operational data, software simulators, and the user interface may be on different computers.

The initial prototype of a national airspace simulation runs on computers making up a VNAS grid. Each night operational data from up to 2,000 flights can be batch processed by three aircraft component simulators running on these machines. When a user makes a request

for simulation services, the VNAS system dispatches the work and manages distribution of input and output data. On remote computers, the Numerical Propulsion System Simulation (NPSS), an engine simulation developed by the NASA Glenn Research Center [3]; Wing Sim, a wing simulation application developed by researchers at the NASA Ames Research Center [4]; and a Boeing 737 landing gear simulation process flight data. During this process, they produce parameters reflecting performance of key aircraft sub-assemblies.

System Elements

Traditionally, taking advantage of distributed large-scale computational resources requires detailed, specific information about computers. Researchers have to know the names of hosts and manually submit and track individual jobs. The VNAS changes the user's view from numerous manually intensive steps to a single request.

The VNAS prototype consists of two major parts:

- Intelligent Information Management Node, controlling batch execution of distributed simulations and data integration.
- VNAS grid computers running NPSS, Wing Sim, and 737 landing gear applications at three distributed sites across the United States.

Domain experts plug in models and simulations that are integral parts of the prototype, incorporating them into a model of flight operations at Atlanta's Hartsfield International Airport. In a distributed environment such as the VNAS, however, data, machines, and

people are located in different places. Research teams that are developing and maintaining simulators used in the prototype are located in California, Ohio, and Virginia. A whole aircraft simulation comprised of multidisciplinary simulations requires bringing together applications and data that are developed by these different teams. This is achieved by regulating simulations and managing data from a central site, the Intelligent Information Management Node.

Intelligent Information Management Node

The NASA Ames Research Center has expertise in information technology and provides the Intelligent Information Management Node. It is a central executive, providing overall supervision of automated batch processing. It distributes large amounts of data between central and remote sites. It also schedules preprocessing steps on a local host and execution of engine, wing, and landing gear simulations on grid computers.

Simulation Applications and VNAS Grid Computers

Engine Simulation: Engine simulation is provided by the NASA Glenn NPSS. The generic turbofan engine model consists of components of the inlet-compressor-turbine chain and flight characteristics. The input data required by NPSS are altitude, temperature, barometric pressure, and Mach number. Inlet and output temperatures and pressures of compressor and turbine are examples of computed simulation parameters. Component revolution is an example of a category of risk assessment measurement that could be used to improve engine safety.

The Intelligent Information Management Node initiates an engine simulation by calling a client on a grid computer. The client then executes a script, passing parameters to the turbofan simulation program and starting the simulation. A flight's radar tracks are processed individually, producing an output record for each radar track. Sample engine parameters for a departing flight generated by NPSS are provided in Table 1.

Wing Simulation: NASA Ames' Wing Sim provides a simple wing model for one aircraft type that consists of flight conditions and manufacturer's technical specifications. The input data required by Wing Sim are Mach number and altitude. Computed simulation parameters are lift, drag, and side force coefficients, and pitching, rolling, and

yawing moment coefficients. The wing simulator has not been validated.

Clients on grid computers invoke a wing simulation through a shell script that sends several parameters to the Wing Sim executable. A flight's radar tracks are processed individually, producing one output record for each radar track. Data in these records include basic aerodynamic coefficients and forces and moments. Structural stress is an illustration of risk assessment measurement that could improve the safety of wings.

Graphs of sample lift coefficients generated by Wing Sim for arrivals are illustrated in Figure 2. The solid line represents a lift coefficient's baseline curve for simulated arrivals at Hartsfield International Airport. The broken lines represent lift coefficients for two simulated arrivals.

Landing Gear Simulation: The NASA Langley Research Center has expertise in airframe systems. A Boeing 737 simulator jointly developed by NASA and Boeing is used to simulate landing gear for arrivals [5]. This simulator contains the capability to auto-land a flight, given runway and aircraft locations and other flight conditions.

This is used as the landing gear model. The input data required by the simulator are aircraft latitude, longitude, speed, weight, and altitude, along with runway threshold altitude, latitude, and longitude. Computed simulation parameters are summed forces and moments for nose and left and right landing-gear subassemblies. Environmental exposure and structural stress on landing-gear components are examples of risk assessment measurements that could improve safety.

VNAS Grid Computers: The VNAS' concept of a computational grid involves interconnecting computers so that they appear to be a single, large virtual computer. To coordinate the various parts of the grid environment, the VNAS grid uses shared authentication, authorization, and auditing systems and directory services provided by Globus [6]. Globus services allow users to submit jobs that run on the spare cycles of computers operating within this network of machines.

For example, in the VNAS prototype each flight simulation may take from one to two minutes to generate engine simulation parameters. On a single grid machine, it will require approximately two days to run engine simulations for 2,000 flights. The same processing can be accomplished overnight by spreading




Figure 2: Simulated Arrival Lift Coefficients Versus Altitude

out the workload throughout several grid machines. Grid security services and interconnections between machines support the splitting up of large workloads such as a day's worth of engine simulations onto several grid computers. Grid services are also used to transfer input and output data between the Intelligent Information Management Node and other grid machines.

Lessons Learned

Lessons learned from the VNAS prototype can be divided into two categories:

- Establishing network connections between sites distributed across the country and grid security.
- Transforming measured operational data into a form that is compatible with simulators.

In distributed computing, network connections and grid security are important technical elements. Today, setting up network connections between sites requires person-to-person communication. Direct negotiations with networking officials at each site are required to work out solutions that are consistent with local policies and meet technical requirements.

Other organizations such as NASA's Information Power Grid are working on improving this technical area [7]. As automated means of assigning secure connections are developed, less effort will be required to expand the VNAS prototypes.

Grid security is another area of cutting edge research in which organizations such as the Globus project are making improvements [8]. As mentioned

earlier, Globus' security services provide the VNAS prototype with a mechanism for interconnecting remote supercomputers. It eliminates burdensome logons and replaces manually intensive steps with a single transaction.

Another problem is developing an interoperable interface to input data. Radar tracks and weather data are not in units or formats that are compatible with the simulation models. Separate executables are required to calculate and convert values such as latitude, longitude, temperature, and Mach number to satisfy the simulators' requirements.

System Operations

The Intelligent Information Management Node provides overall supervision across automated batch processing of flight data from up to 2,000 daily flights. FAA and other government data collection systems provide measured operational data to the VNAS prototype. The central executive transforms and integrates weather and radar track data. As the next step in the process, the Intelligent Information Management Node disseminates transformed data and invokes remote clients at simulation sites. When driven by this data, sub-assembly simulators produce simulation parameters that reflect components' operation.

The input data is processed by simulation software running at remote sites. For each flight, the appropriate simulator produces simulation data that reflect the state and performance of a particular aircraft subassembly such as an engine or a wing. When simulations

from multidisciplinary perspectives have finished, the Intelligent Information Management Node manages the retrieval and integration of simulation and monitoring data. When the integration process is complete, the data is stored at the central site for analysis.

The Intelligent Information Management Node also handles interactive simulation requests made through a Graphical User Interface (GUI). Monitoring and simulation data can be viewed via a visualization tool provided by the GUI. As an alternative, data can be downloaded to a user's computer where it can be trended and analyzed for risk exposure.

What Can Be Done With Whole Aircraft, Multidisciplinary Simulations?

The VNAS prototypes will support improvements in the design process and reductions in operational costs. Airport-wide aircraft simulations using distributed resources can be applied in three areas: engineering design process, airline operations, and airport/air traffic practices and procedures.

Engineering Design Process

Multidisciplinary airport simulations can be inserted at the front end of the engineering design process. Simulations of interdependent components can provide airframe manufacturers with risk exposure measurements for new designs or modification of existing ones. This early understanding of a new design or change can be fed back into the process, reducing overall design time.

Airline Operations

Whole aircraft simulations can be used by airlines to assess the impact of current or proposed policies and procedures on operations and maintenance. Insight from subassembly simulations of engines can lead to condition-based maintenance and more optimal scheduling of repairs. For example, use of reverse thrusters while landing an aircraft directly influences incidences of foreign object ingestion, a primary cause of engine wear. Feedback from airlines' maintenance staff indicates that they are especially interested in engine simulations that include foreign object ingestion.

Airport/Air Traffic Practices and Procedures

Airport flight simulations that involve

whole aircraft models, including noise and emissions provide a more complete picture of operational changes before they are made. Multiple simulations across a range of conditions can provide insight into possible safety and cost issues. This knowledge can then be used in making decisions on changes to airport and air traffic policy and procedures.

Future Research

Two important future research areas are safety and homeland security. A continuous real-time, national airspace, system-risk exposure evaluation is crucial to measuring security risk and evaluating new concepts for improvement of safety and security.

"Simulations of interdependent components can provide airframe manufacturers with risk exposure measurements for new designs or modification of existing ones."

The NASA research team will continue to work with leaders in industry and government who are actively engaged in the process of developing standards for computational grids.

Conclusion

For the national airspace to increase throughput of air transportation while retaining a low accident rate requires a nationwide simulation environment that has the highest achievable degree of fidelity. Technological progress that is enabling creation of large-scale computational grids now makes it possible to bring together data and applications and assign large-scale scientific workloads to a network of computers distributed across the country.

The VNAS is a software platform that is optimized for running computing applications in a distributed manner. It has an infrastructure that consists of a distributed network of computers connected through secure, high-speed links. Collaborative computing allows a suite of tools to direct the workload to the

best available network and hardware topology to meet demands. By dynamically distributing the processing load for simulation to computers with spare capacity, it provides access to significantly more computational power than the resources normally available to the individual effort.

With this approach, research teams that do not have enough in-house capacity have access to large-scale computing power for tasks such as airport simulations that include whole aircraft modeling. Large-scale airport simulations can improve the engineering design process, reduce operations costs, and provide a basis for making decisions on airport/air traffic policy and procedural changes. ♦

Collaboration

The authors would like to hear from people who are interested in sharing ideas on the concepts, architecture, and technology discussed in this article.

References

1. Goldin, Daniel S. Aeronautics and Space Transportation Technology: Three Pillars for Success. Washington, D.C.: NASA's Aerospace Technology Enterprise, Mar. 1977 <www.aero-space.nasa.gov>.
2. NASA. "Information Technology Base Program, Intelligent System Controls and Operations Project." Data Sharing FY00-FY01 Task Plan, 29 May 2001: 15-20.
3. Follen, G., A. Evans, C. Naiman, I. Lopez. Numerical Propulsion System Simulation. 98-3113. Washington, D.C.: American Institute of Aeronautics and Astronautics, July 1998.
4. Bardina, J., W. McDermott, et al. "Integrated Airplane Health Management System." First Joint Army Navy NASA Air Force Modeling and Simulation Subcommittee, Nov. 2000: 4-6.
5. Wallace, Lane E. Airborne Trailblazer: Two Decades With NASA Langley's 737 Flying Laboratories. Langley Research Center: NASA: Chap. 6:3 <http://oea.larc.nasa.gov/trailblazer/SP-4216/chapter6/ch6-3.html>.
6. Foster, I., and C. Kesselman. "Globus: A Metacomputing Infrastructure Toolkit." International Journal of Supercomputing Applications 11.2 (1997): 115-128.
7. NASA's Information Power Grid <www.ipg.nasa.gov>.
8. The Globus Project "Computational Grids." <www.globus.org>.

About the Authors



William J. McDermott is product manager and a software engineer in the Computational Sciences Division at NASA's Ames Research Center.

He is responsible for product management of a software tool that performs information integration. Previously at NASA, McDermott managed a project for three years as project lead on an L1 milestone for the Information Technology Base Level Data Sharing Project. He has also developed software for use in a number of research projects in the aerospace domain. McDermott has a master's degree from California State University, San Francisco.

NASA Ames Research Center
IC, Mail Stop 269-4
Moffett Field, CA 94035-1000
Phone: (650) 604-5650
Fax: (650) 604-4036
E-mail: william.j.mcdermott@nasa.gov



David A. Maluf, Ph.D., is a senior scientist, the Engineering Information Management element manager, and chief information officer for Engineering for Complex Systems at NASA Ames Research Center. Maluf manages from \$2 million to \$4 million annually and supports national-scale projects for the Federal Aviation Administration. Previously, he taught system engineering and control at McGill University and information management and databases at Stanford University. He has published more than 50 articles in industry journals and conference proceedings. Maluf has a master's degree and a doctorate in electrical engineering from McGill University and postdoctoral work from Stanford University.

NASA Ames Research Center
IC, Mail Stop 269-4
Moffett Field, CA 94035-1000
Phone: (650) 604-0611
Fax: (650) 604-4036
E-mail: maluf@email.arc.nasa.gov



Yuri Gawdiak is program manager of Engineering for Complex Systems Program at NASA headquarters in Washington, D.C. Gawdiak is responsible for strategic planning and implementing NASA's new initiative in Engineering for Complex Systems, and also for briefings to the NASA administrator, the Office of Management and Budget, and other government organizations. Previously he was principal investigator for the personal satellite assistant and principal investigator for the Wireless Network Experiment on STS-76/Mir 21 at NASA's Ames Research Center in Moffett Field, Calif. Gawdiak has a Bachelor of Science in information systems from Carnegie Mellon University.

He is responsible for strategic planning and implementing NASA's new initiative in Engineering for Complex Systems, and also for briefings to the NASA administrator, the Office of Management and Budget, and other government organizations. Previously he was principal investigator for the personal satellite assistant and principal investigator for the Wireless Network Experiment on STS-76/Mir 21 at NASA's Ames Research Center in Moffett Field, Calif. Gawdiak has a Bachelor of Science in information systems from Carnegie Mellon University.

NASA Headquarters
Code R
Room 6K17
Washington, DC 20546
E-mail: ygawdiaki@hq.nasa.gov



Peter B. Tran is a research and development computer scientist at QSS Group, Inc., and is a member of the Aerospace ExtraNet Intelligent Information Integration Group in the Computational Sciences Division at NASA Ames Research Center in Moffett Field, Calif. Previously he worked as a consultant, architect, and software engineer at several companies, including BEA Systems, XUMA, Computer Sciences Corporation, and Recom Technologies, emphasizing on distributed computing systems, Web-based technologies, and information management. Tran has a degree in electrical engineering from the University of California.

NASA Ames Research Center
IC, Mail Stop 269-4
Moffett Field, CA 94035-1000
Phone: (650) 604-4531
Fax: (650) 604-4036
E-mail: pbtran@mail.arc.nasa.gov

COMING EVENTS

June 29-July 3
INCOSE 2003
Washington, D.C.
www.incose.org/symp2003

July 14-17
JAWS S3 Symposium
Monterey, CA
www.usasymposium.com

July 15-18
Practical Implementation of Software Configuration Management
San Diego, CA
www.stittraining.com

August 18-21
2nd Annual CAISR Summit
Danvers, MA
www.paulrevereafa.horizons.com

September 8-12
International Conference on Practical Testing Technique
Minneapolis, MN
www.psqtconference.com

September 14-19
International Function Point Users Group Annual Conference
Scottsdale, AZ
www.ifpug.org

September 22-25
AUTOTESTCON 2003
Anaheim, CA
www.autotestcon.com

October 15-18
Diversity in Computing Conference
Houston, TX
www.ncsa.uiuc.edu

April 19-22, 2004
Software Technology Conference 2004



Salt Lake City, UT
www.stc-online.org