Scheduling the Operation Desert Storm Airlift: An Advanced Automated Scheduling Support System

MICHAEL R. HILLIARD
Center for Transportation Analysis, ORNL
PO Box 2008, Oak Ridge, Tennessee 37831-6270

RAJENDRA S. SOLANKI
Center for Transportation Analysis, ORNL

CHENG LIU
Center for Transportation Analysis, ORNL

INGRID K. BUSCH
Center for Transportation Analysis, ORNL

GLEN HARRISON
Center for Transportation Analysis, ORNL

RONALD D. KRAEMER
Hazardous Waste Remedial Actions Program
PO Box 2003, Oak Ridge, Tennessee 37831-7606

A typical airlift mission carrying troops and cargo to the Persian Gulf required a three-day round-trip, visited seven or more different airfields, burned almost one million pounds of fuel, and cost $280,000. During Operation Desert Storm, the Military Airlift Command (MAC) averaged more than 100 such missions daily as it managed the largest airlift in history. By August 7, 1991, more than 25,000 missions had moved more than 966,000 passengers and 774,000 tons of cargo to and from the Persian Gulf region. Each mission required scheduling aircraft, crew, and mission support resources to maximize the on-time delivery of cargo and passengers. To meet this challenge, MAC worked with the Oak Ridge National Laboratory to develop and deploy the Airlift Deployment Analysis System (ADANS). Within three months, ADANS provided a set of decision support tools to manage information on cargo and passengers to be moved and the available resources, as well as tools to schedule missions, to analyze the schedule, and to distribute the schedule to MAC's worldwide command and control system.

At 3:00 AM eastern standard time (EST) on Saturday morning, a crew consisting of an aircraft commander, a copilot, two flight engineers, and a load-

0091-2102/92/2201/0131$01 25

TRANSPORTATION—AIR
DECISION ANALYSIS—APPLICATIONS
MILITARY

INTERFACES **22:** 1 January–February 1992 (pp. 131–146)

master from the 436th Military Airlift Wing begins the day at Dover Air Force Base (AFB), Delaware, by filing flight plans and performing preflight checks on a C-5B Galaxy. At 6:15 AM EST, the aircraft departs for a two-hour flight to Fort Campbell, Kentucky, where the loadmaster supervises the loading of 73 servicemen and 60 tons of cargo. After refueling, the Dover crew flies the aircraft to Westover AFB, Massachusetts, where another crew is waiting to take the aircraft to an enroute airfield in Europe. Once in Europe, the Westover crew begins a rest period, and the aircraft is serviced and refueled with 190,000 pounds of fuel. An augmented crew consisting of three pilots, two engineers, and two loadmasters who are temporarily stationed at the European airfield, flies the aircraft to Riyadh, Saudi Arabia and arrives at 11:15 AM EST on Sunday. After unloading the cargo, the aircraft is refueled, and the same crew flies the C-5B back to Europe. The crew members arrive at the European enroute airfield 18 hours after they leave, sleeping in shifts in the bunks behind the flight deck. Another Dover crew then flies the aircraft home to Dover AFB, arriving at 9:30 AM EST on Monday. Between Saturday and Monday, the aircraft is serviced by seven different ground crews, accumulates 34 hours of flying time under the command of four different crews, and burns almost one million pounds of fuel at a cost of approximately $280,000.

During the height of Desert Storm, the Military Airlift Command (MAC) initiated such a scenario 100 times each day on average as it managed the largest airlift in history. As of August 7, 1991, more than 25,000 airlift missions had moved over 966,000 passengers and 774,000 tons of cargo to and from the Persian Gulf region. To schedule the airlift, MAC combined the capabilities of its in-house operations research and systems management staff with a new airlift scheduling tool, the Airlift Deployment Analysis System (ADANS). Since October 1990, ADANS, which is being developed by Oak Ridge National Laboratory (ORNL), has been used to schedule airlift missions for Operations Desert Shield and Desert Storm, for refugee relief, and for disaster response.

The Airlift Problem

MAC is a major command of the US Air Force and a component of the US Transportation Command (USTRANSCOM), with the primary mission of providing air transportation for personnel and cargo for all US military services worldwide. To accomplish this, MAC operates a fleet of more than 500 cargo aircraft, contracts with civilian carriers for cargo and passenger movements, and schedules the use of the Strategic Air Command's KC-10s when those aircraft are configured for cargo rather than for air refueling. In both peacetime and wartime, MAC schedules missions for these aircraft through a worldwide network of military and civilian airfields, each with its own capabilities and limitations. MAC's responsibility is to keep its crews, aircraft, and support resources throughout the world prepared for deployment at a moment's notice while concurrently operating the day-to-day peacetime airlift system.

MAC is headquartered at Scott AFB, Illinois (HQ MAC). HQ MAC delegates the responsibility for managing the execution of individual missions to the 21st and 22nd Air Forces (AFs). The 21st AF, located at McGuire AFB in New Jersey, has jurisdiction over missions scheduled in the region from the Mississippi River east to the eastern coast of Africa. The 22nd AF, located at Travis AFB in California, manages missions from the Mississippi River west to the eastern coast of Africa. The 834th Airlift Division (ALD), located at Hickam AFB in Hawaii, manages intra-theater airlift missions within the Pacific region. The 322nd ALD, located at Ramstein AFB in Germany, manages intra-theater missions within the Atlantic region.

MAC's challenge is to schedule aircraft, crew, and mission support resources to maximize the on-time delivery of cargo and passengers. The movement of passengers and cargo is constrained both temporally and spatially. Each military unit (or movement requirement) determines the earliest date it is available to load, and the commander of the operation specifies a delivery-time window constrained by earliest and latest arrival dates (Figure 1). A movement requirement also specifies the origin and destination of the movement and the amount and type of cargo and passengers to be moved. MAC's airlift scheduling specialists (flow planners) are responsible for deciding which resources to use to move the requirements. The resources include aircraft, airfield support, and crews. The use of aircraft is constrained by the number of aircraft available, the capacity and capabilities of each type of aircraft, and maintenance schedules. Each airfield is limited in the number and type of aircraft it can accept and what types of activities (onloading cargo, refueling, resting crews, and so forth) can be accomplished for each type of aircraft. An airfield is also limited by how much cargo and how many passengers it can accommodate loading and unloading each day and which hours of the day it is open for arrivals and departures. There are limited numbers of crews available at major airfields for each type of aircraft, and a crew can fly no longer than a duty day of designated length. All these constraints must be considered when building a schedule.

The Development of ADANS

ADANS development began at ORNL in 1987 under an interagency agreement be-

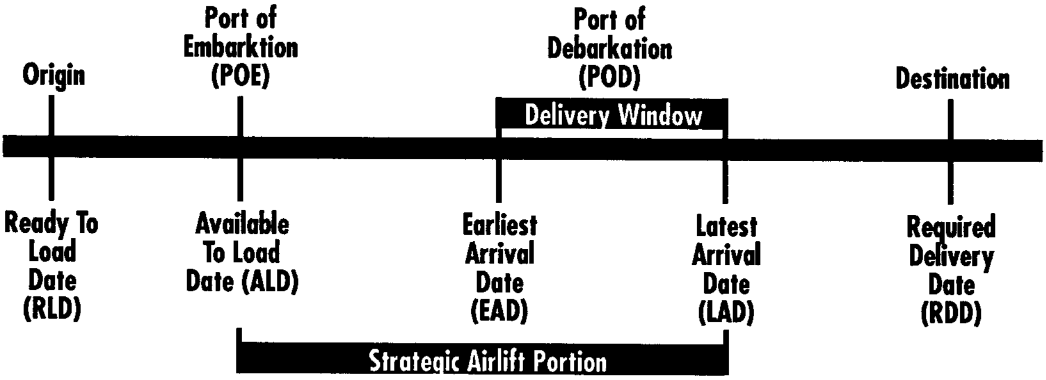

Figure 1: A unit's availability and the delivery window specified by the commander of the operation constrain the timeframe of the strategic airlift portion of a deployment.

tween MAC and the Department of Energy. The product of this collaboration has included both basic and applied airlift scheduling and analysis research at several universities and corporations [Hane, Jarvis, and Ratliff forthcoming; Hooker and Natraj 1991; Rappaport et al. 1991; Ray 1990; Solanki and Southworth 1991]. Although some of the peacetime scheduling tools were in place when Iraq invaded Kuwait on August 2, 1990, ADANS wartime execution tools were still on the drawing board or in the early stages of development. These tools were not scheduled to be developed until the summer of 1991. Developing and applying the software to a real-world deployment in a time-critical environment was an intensive endeavor (Figure 2).

MAC's part of the war effort (moving troops, equipment, and supplies to the Persian Gulf) began when President Bush ordered the deployment in early August. For each mission, such as the one described in the introduction, MAC must create a detailed itinerary, including the type of aircraft to be used, the operator assigned to fly the mission, arrival and departure times at each station, the activities at each station (onloading cargo, offloading cargo, changing crews, refueling), how much of each military unit will be transported by the mission, and any special instructions necessary for completing the mission. As the first missions departed carrying the equipment and personnel who would unload and service the aircraft in the Persian Gulf, flow planners and operations research analysts developed several concepts of operations for the deployment [Roherkasse and Hughes 1990]. A concept of operations is the overall plan for the airlift movement. It involves setting up the structure of the network over which the aircraft will fly and the aircraft, crews, mission support resources, and permissions at each airfield on the network.

The automated systems in place at MAC to support planning what was quickly becoming an enormous airlift operation were soon found to be inadequate. Due to their slow speed and lack of flexibility for incorporating real-world scheduling constraints, the existing systems did not provide sufficient automated support to meet the airlift scheduling challenge effectively. The flow planners had little help from any automated system; they were limited to plan-

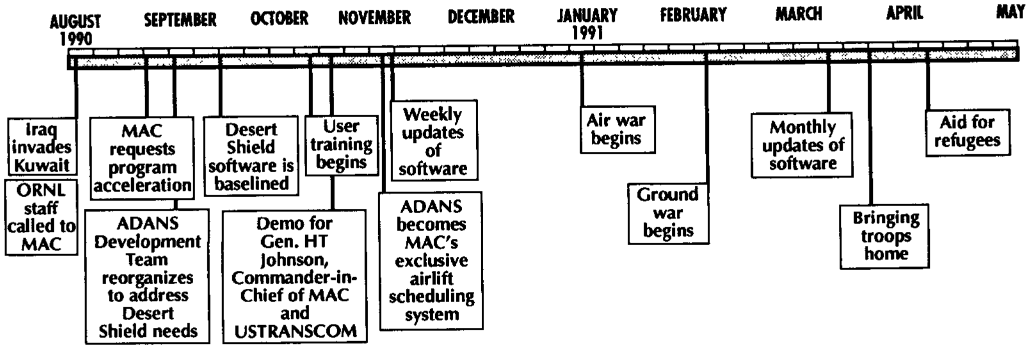

Figure 2: Operation Desert Shield/Storm accelerated and intensified the development of the ADANS software.

ning missions by hand and typing them directly into the MAC command and control system. They had no automated decision support or analysis tools that let them view the overall schedule and address what-if questions.

On August 20, 1990, MAC's crisis action team director requested that the development of ADANS be accelerated to provide an automated airlift scheduling capability as quickly as possible. Knowing that MAC's system would be tested to the maximum, he believed that an improved scheduling capability would help ensure the nation's success. Several ORNL staff members went to Illinois to work with MAC's crisis action team at Scott AFB so that they could better understand the concept of operations for the deployment. The ADANS development effort was reorganized into several "tiger teams" that could confront specific critical issues. These teams addressed airlift network and resource management, communications, reports, data management, schedule analysis, and end-user support. ORNL's operations research staff then accelerated the development of a dynamic programming-based automated airlift scheduling algorithm that would become the centerpiece of ADANS' automated scheduling tools.

The ADANS development staff worked seven days a week with flow planners at MAC, providing continuous participation by ORNL from early August 1990 until mid-January 1991. By working directly with the flow planners, ORNL developers were able to see problems first-hand, to design solutions quickly, and to communicate needs to the software development staff.

On October 19, 1990, ADANS program management staff met with MAC's commander-in-chief, General H. T. Johnson, to discuss the status of ADANS development. The work had progressed much more quickly than had been expected, but ADANS was still a few weeks away from being thoroughly tested and ready for use. General Johnson emphasized that ADANS was needed immediately and that, unless the risk was prohibitive, he wanted ADANS operational by the middle of the following week. The program staff met the challenge; on October 23, 1990, ADANS became MAC's exclusive Operation Desert Shield airlift-scheduling system.

When Iraq invaded Kuwait, ADANS wartime execution tools were still on the drawing board.

The result of the partnership between developers and flow planners was an integrated system that MAC's flow planners could use to enter and evaluate cargo- and passenger-movement requests; to allocate airlift resources, including aircraft availability, aircraft characteristics, crews, airfield resources, and airlift network configurations; to schedule missions; to update mission schedules; to analyze schedules; and to distribute schedules through MAC's command and control system (Figure 3).

Processing Movement Requirements

When a major crisis erupts, MAC organizes a crisis action team composed of a number of specialized cells that work together to track resources, to process requirements, and to develop mission sched-

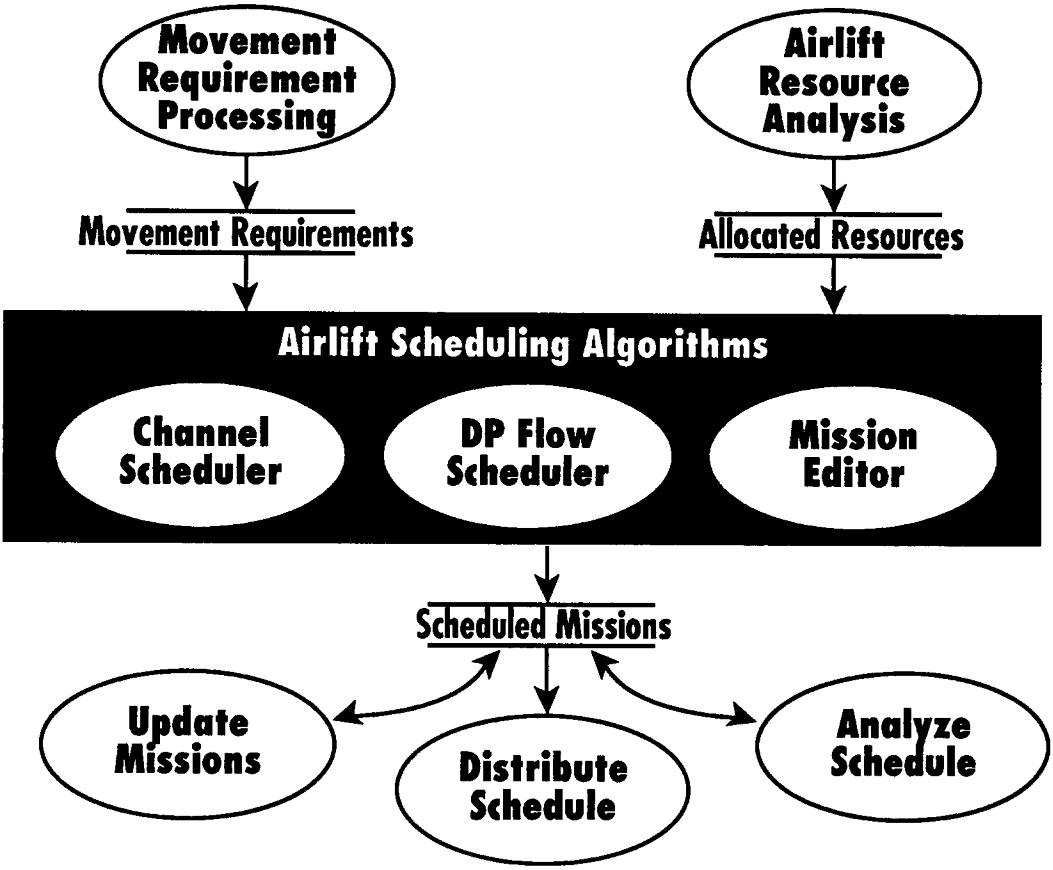

Figure 3: ADANS provides tools for matching movement requirements with airlift resources to create a schedule for an airlift operation.

ules. The first step in developing airlift schedules is to process information associated with movement requirements. This information initially comes to MAC through USTRANSCOM. ADANS software checks the movement-requirement information for errors or missing information. If the information is found to be within acceptable limits, the requirements cell catalogs the movement request, contacts each organization requesting airlift, and solicits specific information, such as the point of contact, details about the cargo, and an estimate of the number and type of aircraft needed. The requirements cell then determines if the cargo and passengers can be moved on commercial or military aircraft. Commercial missions are coordinated with private sector airlines by the civil reserve air fleet (CRAF) cell; as the commercial schedules are developed, they are incorporated into ADANS. If military aircraft are to be used, the movement requirement information is passed to the flow-planning cell for scheduling.

Allocating Airlift Resources

Once the movement requirements have been processed, the flow-planning cell can schedule the airlift. Flow planners specify the number of aircraft of a particular type and configuration that are available from each operator on each day. The system allows flow planners to model the number of aircraft available by day for as many aircraft types as needed. For each aircraft type and configuration, flow planners specify the cargo- and passenger-carrying capacity, the minimum load that can justify a mission, the type of cargo that may be carried, the amount of time required to load or unload the aircraft, the minimum ground time required, and the maintenance requirements.

Each aircraft type requires crews with specific qualifications. A crew can be either basic or augmented. Basic crews normally are limited to 16 hours of duty each day; augmented crews typically have crew duty days of up to 24 hours. However, during wartime, crew duty days can be modified when needed. Within ADANS, flow planners specify the number of basic and augmented crews available at each airfield and the crew duty time limitations for each crew type. The automated scheduler takes into account both the availability of rested crews and the limitations on crew duty time in an attempt to minimize late cargo deliveries.

For each mission, MAC must create a detailed itinerary.

Airfields must have extensive resources to onload, maintain, refuel, and offload aircraft. These resources include fuel, fuel trucks, material handling equipment, parking space, and maintenance support. Instead of trying to model the constraints imposed by each of these limited resources, ADANS models the airfield resources using three major factors: the maximum number of active aircraft that can be on the ground at any one time, the maximum number of tons that can be processed through an airfield each day, and the maximum number of passengers that can be processed through an airfield each day.

To control the timing of arrivals and departures, flow planners can specify operating hours and a minimum length of time required between arrivals (separation time) at each of the major airfields. This modeling of operating hours is essential because some airfields are closed at night, while others have mission support services available only at specific times.

In addition to specifying aircraft, crew, and airfield constraints, flow planners develop a concept of operations that defines how they want to allocate the use of the resources. For instance, they may determine that all missions in the eastern United States should originate from a particular airfield and that they should stop for refueling and crew changes at specified airfields in Europe. A network of flight legs and subnetworks for specific pairs of airfields provides flow planners with total control in routing aircraft through the system. Flow planners define the resource constraints and the network configuration using a set of text-based forms and geographic editing tools.

The airlift network is modeled as a set of airfields, routes between airfields, and permissions for each aircraft at each airfield. Permissions are constraints on the use of airfield resources imposed by a flow planner. Permission or denial of permission may model a physical limitation, such as the lack of appropriate equipment for loading a particular plane; a political concern, such as not landing military aircraft at an airfield; or a logistics need, such as returning aircraft to appropriate airfields for maintenance. Permissions can be specified by aircraft type, configuration, operator, or any combination. Airfield permissions modeled within ADANS include onload, offload, enroute, refuel, and recovery. Onload and offload permissions allow an aircraft to load or unload cargo. Enroute permissions allow aircraft to stop at an airfield to refuel or change its crew even if it is not scheduled to onload or offload at that airfield. Recovery permissions require an aircraft to go to another airfield immediately after an offload instead of waiting at the offload airfield for its next mission. Using a recovery airfield limits congestion, facilitates efficient use of the offload airfield, and provides maintenance facilities.

The ADANS development effort was reorganized into several "tiger teams."

To manage the airlift flow at congested enroute airfields and offloads, flow planners often use a set of prespecified daily departure or arrival times. These prespecified times, or "slot times," control the flow of aircraft into the more heavily used offload airfields by sequencing the flows from the enroute airfields. When flow control can be more relaxed, separation times and parking availability are used to limit the frequency of landings or departures.

Creating Mission Schedules

Missions can be classified into the following four categories:

—*Channel missions* follow regular routes at prespecified intervals to deliver mail, food, and replacement supplies every day. ADANS provides tools to create routes and individual missions. Using ADANS, flow planners can schedule hundreds of channel missions in minutes.

—*Quick response missions* move critical items on extremely short notice. After confirming the need for immediate airlift, flow planners use the mission-editing module to build itineraries and to allocate cargo and passengers to specific aircraft.

—*Civilian aircraft missions* are flown by civilian carriers. The schedules are proposed to MAC by the commercial carriers in response to MAC's stated requirements for passenger or cargo movement. After planners coordinate the mission schedule with the commercial carrier, they use the mission-editing module to incorporate the mission into ADANS to ensure overall efficiency throughout the airlift system.

—*Time-phased airlift flow missions* are planned several days or weeks before their scheduled departures. These are movements of multiple military units from one or more onload airfields to one or more offload airfields. The delivery time windows for these moves may vary. The dynamic-programming-based airlift-scheduling algorithm is designed to schedule these missions either in a single run of the algorithm or through a series of runs.

Each of the four types of missions relies on a different tool that is tailored for its scheduling needs. The channel missions, quick response missions, and civilian aircraft missions are all created with the mission-editing module. The mission-editing module allows flow planners to create, to delete, to copy, and to modify missions, and it provides a means to check that missions are feasible in terms of crew duty restrictions, weight limitations, ground-time requirements, and concept of operations. The mission editor is used to modify missions that are created manually by flow planners or to modify missions produced by the dynamic programming scheduler. It allows flow planners to remain in complete control of the schedule, using the system to provide checks and audits as necessary. The channel mission creation software enables the flow planner to create a series of similar missions offset by a prescribed time phasing (for example, every 24 hours). This tool reduces several days of work each month to a couple of hours of effort.

The dynamic-programming-based airlift-scheduling algorithm iterates through a three-step process to schedule time-phased airlift flows. First, the scheduling algorithm matches aircraft to requirements. Next, it schedules the mission. Finally, it updates the resources and requirements. Using the scheduling algorithm, a planner can schedule missions for thousands of requirements quickly and effectively. It gives the flow planner the option of generating a completely new set of missions or integrating newly generated missions with all currently scheduled missions. The appendix explains its capabilities and the techniques employed.

The airlift-scheduling algorithm is also an excellent tool for analyzing large sets of requirements and providing insights into different proposed concepts of operations. For example, when the flow planners were asked in early November 1990 to determine how long it would take to move 100,000 troops and 100,000 tons of cargo to the Persian Gulf region, they used the airlift-scheduling algorithm to analyze and compare different delivery strategies.

Analyzing Schedules

ADANS also provides integrated analysis tools that allow flow planners to view the schedule from the perspective of the three major components of a mission: requirements, aircraft, and airfields. Each component can be viewed in a summary form. Flow planners can see how much of a movement requirement is not yet scheduled, the number of flying hours that have been scheduled for an aircraft, and the planned peak work loads for an airfield. The components of the summary are decomposed into lists of missions, and the details of any individual mission can be evaluated and modified quickly and easily.

Textual displays are supplemented by graphical displays that show movement requirement deliveries, resource commitments, and aircraft activities. One such display is the rainbow chart, which shows aircraft activity. Adapted from a chart that flow planners drew and regularly redrew manually, the rainbow chart shows the schedule for each aircraft through the use of colored bars. The bars indicate stops and activities at airfields, and the spaces between the bars indicate flying times. The rainbow chart is a powerful interactive decision support tool that permits flow planners to request more information on a stop and to change mission data.

Communicating with MAC's Command and Control System

After a mission has been scheduled, integrated with the rest of the plan, and coordinated with the airlift user, a copy of the airlift schedule is sent to MAC's command and control system. The command and control system provides a worldwide network that follows and manages each aircraft throughout its mission.

System Architecture

The ADANS architecture is based on a relational data-base management system, which operates on a network of powerful, UNIX-based work stations stretching across the United States. ADANS work stations are currently located at Scott AFB, Illinois (HQ MAC); McGuire AFB, New Jersey (21st AF); Travis AFB, California (22nd AF); and Hickam AFB, Hawaii (834th ALD). The configuration includes a data-base management system, a form generation tool, graphical display tools, a report generation system, communication software, a windowing system, and more than 500,000 lines of ADANS-unique code. All ADANS modules exchange data through the data base and run asynchronously; thus, flow planners can use the windowing system to keep track of and to modify multiple pieces of information.

Benefits of ADANS to the Military Airlift Command

It is difficult to quantify the benefits of ADANS for MAC. There was no opportunity to implement the system in a controlled environment or to measure effectiveness during the operation, and there are no comparable activities with which to compare Operations Desert Shield and Desert Storm. MAC had never responded to an airlift challenge of this magnitude before. As MAC commanders have pointed out, the operation moved the equivalent of the combined population of Oklahoma City and Arlington, Virginia, with all of their household goods, vehicles, food, and supplies, from the United States to the Persian Gulf and returned them home.

During the first months of the Persian Gulf airlift, flow planners spent much of their time reacting to events that complicated the scheduling process. The more efficient ADANS scheduling tools enabled the flow planners to make a transition from their reactive posture to a proactive stance by providing them with more time to evaluate the quality of their airlift schedules. This resulted in more efficient airlift operations due to better coordination between the military units being moved and the 21st and 22nd AFs flying the missions.

Being able to create, to modify, and to track channel missions in minutes instead of hours and to manage them in the same system as all other missions increases the flow planners' efficiency and decreases the number of errors. Flow planners can react more efficiently to quick-response missions and, using the airlift-scheduling algorithm, can generate hundreds of missions in minutes either to evaluate courses of action or to develop time-phased airlift flow missions. The schedule-analysis tools and the graphical displays highlight schedule conflicts and overcommitments of resources and allow planners to assess and compare multiple concepts of operation more rapidly by comparing schedules created under different assumptions.

The airlift operation is estimated to have cost between $3.5 and $4 billion. The development cost of the ADANS supporting software was approximately $2 million, a factor of 2,000 to 1. Even a half-percent improvement in efficiency would be a 10-fold return on investment.

Conclusion

Saving money was not the goal of ADANS during Operations Desert Shield and Desert Storm. The primary goal was to deliver the necessary cargo and personnel to the Persian Gulf quickly. ADANS has made scheduling missions more effective at MAC by replacing manual scheduling methods and providing planners with better information so that they can do their job more efficiently.

Major General Vernon J. Kondra, MAC's deputy chief of staff for operations and transportation, visited ORNL on April 3, 1991, to thank the staff members for their work. He spoke of the transition from the early days of the operation when there were no scheduling tools to the advent of ADANS and the beginning of a controlled and managed flow of aircraft and material. He thanked the development staff and said, "I guarantee you that we could not have done that [the deployment to the Persian Gulf] without your help and the contributions you made to ADANS—we absolutely could not have done that."

MAC continues to use the system daily to respond to crises, to provide disaster and humanitarian relief, and to plan military exercises; however, there is still work to be done to implement ADANS fully within MAC. ADANS functionality will be expanded at the 21st and 22nd AFs, and the peacetime tools will be completed and integrated with wartime planning systems to provide smooth transitions as MAC responds to world events. Current research is focusing on specialized tools to assess courses of action and to provide quick estimates of resource requirements and capabilities based on the limited information available in the early phases of a conflict. As MAC and ORNL develop these capabilities, the goal is always to improve MAC's ability to respond to crises quickly and effectively anywhere in the world.

APPENDIX: The Dynamic Programming-Based Airlift Scheduling Algorithm

The goal of the scheduling process is to maximize the amount of cargo and number of passengers delivered on time. However, the complexity introduced by the constraints and the size of the problem preclude the optimization of the schedule as a whole. A typical airlift problem contains several thousand movement requirements to be moved over a network of 50 to 100 airfields using a fleet of several hundred aircraft. Previous work in the related area of cargo-ship scheduling, such as that of Brown, Graves, and Ronen [1987] and Fisher and Rosenwein [1989], produced algorithms for the optimal solution of much smaller problems. These techniques involved generating all (or most) possible schedules for each ship and then solving a set-packing problem to determine the optimal set of schedules. The large number of possible schedules for each of the aircraft in the airlift problem makes such a technique intractable. For this reason, the approach taken by the ADANS scheduling algorithm relies on a heuristic.

The dynamic programming airlift scheduler employs a sequential procedure that schedules one mission at a time, taking into account the resource constraints imposed by previously scheduled missions. Separating the overall airlift problem into a series of smaller problems, each scheduling an individual aircraft, can be justified both theoretically and practically. Current research in routing and scheduling aircraft and crews limited by ground-support resources and time windows has not yielded a computationally feasible optimization technique for real-world problems [Bodin et al. 1983]. MAC flow planners traditionally have constructed overall schedules by piecing together schedules of individual aircraft. On the other hand, optimization does play an important role in the airlift scheduler since minimizing resource consumption and delay for individual missions is essential. Though easier than the overall airlift problem, scheduling a single mission retains enough combinatorial complexity to justify an optimization algorithm. The complexity of a single mission is compounded when a mission includes a second onload or offload (or both) and multiple changes of basic and augmented crews. Each mission must also honor the vast number of routing and timing constraints imposed by the previously scheduled missions, the airfield limitations, and the concept of operations.

Steps in the Scheduling Algorithm

The algorithm iterates over three modules (Figure 4). The first module matches requirements to aircraft, employing a rule-based procedure that suggests pairings of aircraft with requirements and decides the order in which the pairings will be considered. For example, a typical rule considers a requirement available to be moved on January 10 before considering requirements that cannot be moved until January 12. Suitable rules are used to determine compatibility between the aircraft type and the cargo and between the aircraft type and the onload and offload airfields.

The rule-based system provides an aircraft-requirement pair and a network of applicable flight legs to the second module, the dynamic programming module. The dynamic programming module then attempts to schedule a mission, taking into account the available airfield resources.

The solution produced by the dynamic programming module is passed to the third module, which updates the status of resources and the requirements. This process continues until either the requirements or the critical resources are exhausted.

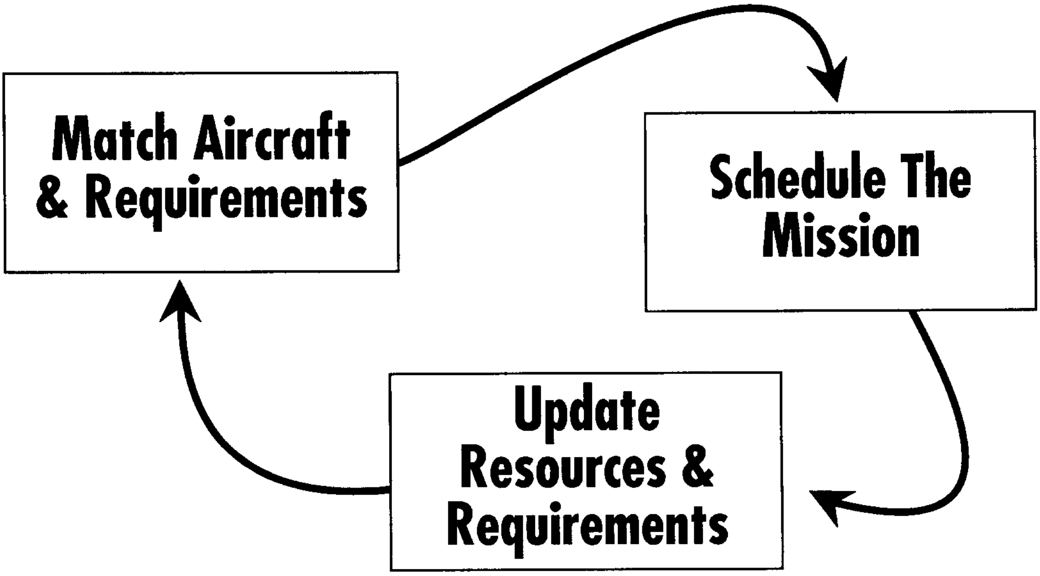

Figure 4: The airlift scheduling algorithm iterates over three modules to create a schedule for an airlift flow.

Dynamic Programming Formulation and Solution Method

The objective of the airlift problem is to maximize the amount of cargo and number of passengers that are delivered on time. However, since the dynamic programming module schedules one mission at a time, it must consider issues related to the efficient utilization of resources along with the maximization of cargo and passengers delivered by a mission. While scheduling a mission, the dynamic programming module seeks to achieve a balance among the following objectives:

(a) To maximize the cargo and passengers delivered by the mission;
(b) To minimize late deliveries;
(c) To minimize the flying time of the mission;
(d) To minimize a weighted sum of ground time where the weights on ground times vary depending on the activity at an airfield (for example, ground time at an offload airfield is penalized more heavily than ground time at a recovery airfield); and
(e) To minimize the number of crew changes.

Objective (a) is actually a constraint in the algorithm. If the amount of the requirement available to be scheduled is enough to fill the aircraft, at least 80 percent of the aircraft's capacity must be used. Often, due to weight limitations, an aircraft can fly longer flight legs only if the weight of the additional fuel is offset by a reduction in the amount of cargo carried. Thus, by permitting the aircraft to carry less than 100 percent of its capacity, the scheduler may find a better schedule that includes longer flight legs but avoids congested enroute stops. The schedule of the mission is obtained by minimizing the weighted sum of objectives (b) through (e) and is subject to the lower bound on the load carried by the mission and the rest of the constraints on the availability of crew and ground-support resources at airfields. The dynamic programming module optimally schedules individual missions based on the above weighted objective function.

The "stages" in the dynamic programming formulation correspond to the airfields in the network of flight legs relevant to the mission (Figure 5). For a given airfield, the states that are included in the corresponding stage are characterized by the departure time from the airfield and the remaining available duty for the current crew. Thus, states can be specified by the triple (s, t, c), where s is the airfield, t is the departure time, and c is the amount of crew duty day remaining. A feasible transition from a state (s, t, c) to a state (s', t', c') requires (at the least) that

(1) airfield s' follow airfield s in the net-

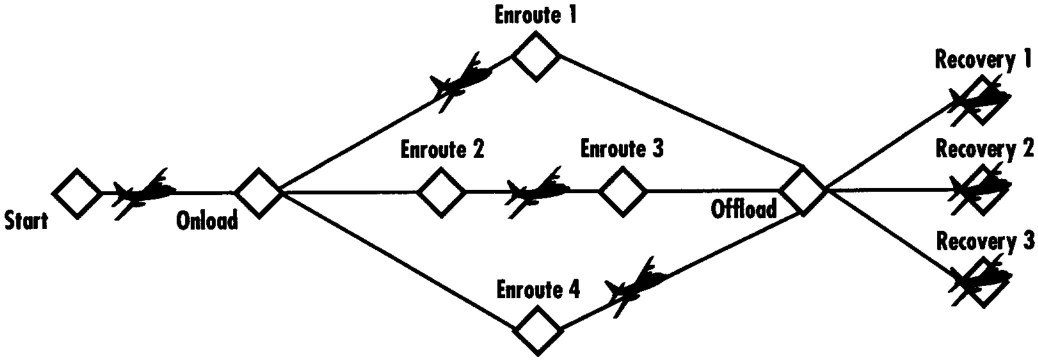

Figure 5: The airlift scheduling algorithm generates the relevant network of possible routings for a mission.

work;
(2) $t' \geq t + f + g$, where f is the flight time from station s to station s', and g is the minimum required ground time at station s';
(3) if a crew change takes place, then c' = crew duty day; if no crew change takes place, then the decrease in the amount of remaining duty day for the crew is exactly the difference in the departure times at the two stations $c' = c - (t' - t)$.

In addition, the needed resources at a station must be concurrently available while the aircraft is on the ground. This is illustrated in Figure 6 by a transition from Frankfurt, Germany, to Dhahran, Saudi Arabia. The highlighted feasible departure interval at each station is the intersection of that station's resource availability intervals.

The optimal objective function value z^* at a state (s', t', c') is given by the recursive relationship $z^*(s', t', c') = \min_{s,t,c}[z^*(s, t, c) + z((s, t, c), (s', t', c'))]$, where the minimum is taken over all states (s, t, c) that allow a feasible transition to state (s', t', c'), and $z((s, t, c), (s', t', c'))$ is the increase in the objective function value corresponding to the transition. Given the states corresponding to the optimal policy, the solution implies the use of station resources and crews.

Streamlining the Dynamic Programming Module

A straightforward dynamic programming implementation would require the computation of the optimal objective function values for all possible combinations of time and crew status at each airfield. The implementation that is used in the ADANS scheduling algorithm uses a streamlined version that requires much less computation and memory by utilizing the following observation: for a given airfield and crew status combination, the minimum value (over all possible routings) of the objective function is piecewise linear over any feasible departure interval.

For example, given the availability of crew and ground support resources at airfields preceding the offload, each possible routing of an aircraft (a sequence of stations along with their crew activities) will result in one or more time intervals for de-

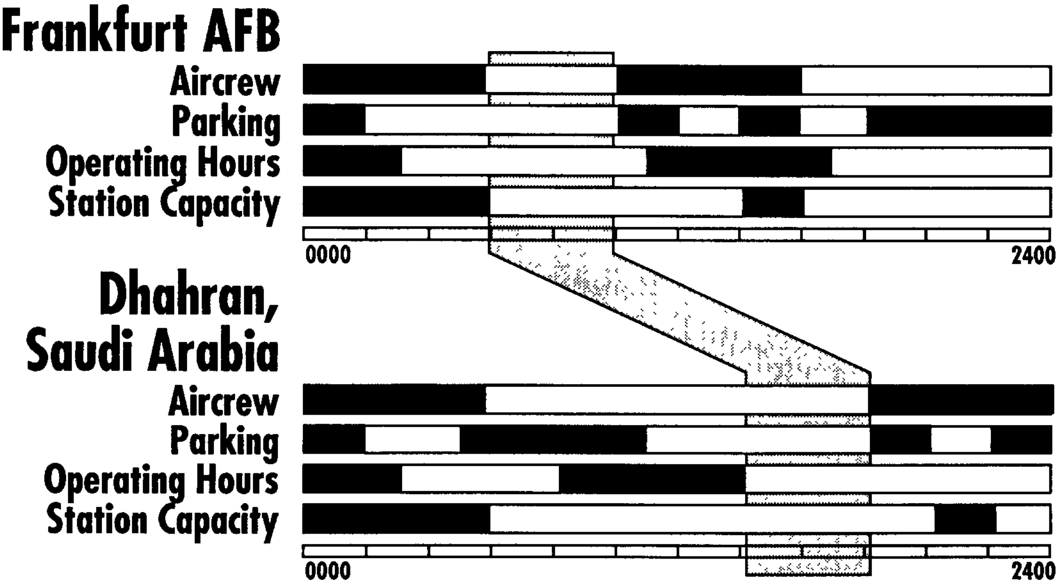

Figure 6: Resource availability limits the feasible departure intervals at each airfield.

parting the offload airfield. For any particular routing, the contributions of flying times and crew changes to the objective function are fixed, while the contributions of airfield ground times and delivery delays are linear with respect to the departure time. Hence, the objective function value varies linearly over each of the departure intervals.

When two different routings result in an overlapping departure subinterval, the routing that achieves the minimal objective function value is retained as the optimal policy (partial routing) for that subinterval. This is illustrated in Figure 7, where the feasible departure interval for routing 1 is $[t_1, t_3]$ and for routing 2 is $[t_2, t_4]$. The departure intervals overlap on the subinterval $[t_2, t_3]$. The objective function values of the two routings for this subinterval intersect at time t^*. Before time t^* routing 2 has a better objective function value; after time t^*, routing 1 has a better objective function value. Thus, routing 1 is retained as the optimal policy for subintervals $[t_1, t_2]$ and $[t^*, t_3]$, while routing 2 is retained as optimal for subintervals $[t_2, t^*]$ and $[t_3, t_4]$. The subintervals preserve the linearity of the objective function. Any change in the slope or the intercept of the objective function leads to a new subinterval. The dynamic programming module operates in the form of Denardo's [1982] forward "reaching" method, projecting subintervals of departure times at a given stage forward in the network and constantly updating the best objective function value for a subinterval. This streamlined implementation achieves a significant reduction in computing time and storage requirements by using intervals of departure time as opposed to the discrete values of departure time. The current implementation of the dynamic programming module typically generates

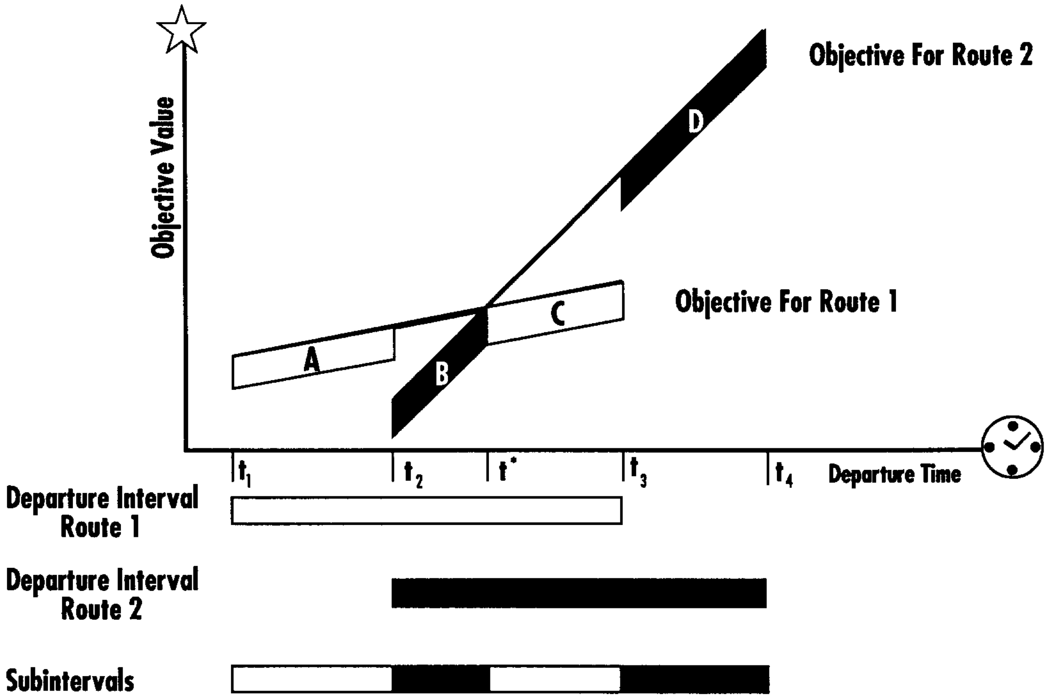

Figure 7: When two different feasible routings result in overlapping departure intervals, the airlift scheduling algorithm creates subintervals to preserve the linearity of the minimum objective function value.

missions at an average rate of two per second. Before we streamlined the algorithm, missions were generated at an average rate of one mission every 10 seconds. Solanki and Busch [1991] provide a more detailed description of the algorithm.

References

Bodin, L.; Golden, B.; Assad, A., and Ball, M. 1983, "Routing and scheduling of vehicles and crews—The state of the art," *Computers and Operations Research,* Vol. 10, No. 2, pp. 63–211.

Brown, G. B.; Graves, G. W.; and Ronen, D. 1987, "Scheduling ocean transportation of crude oil," *Management Science,* Vol. 33, No. 3, pp. 335–346.

Denardo, E. 1982, *Dynamic Programming: Models and Applications,* Prentice Hall, Englewood Cliffs, New Jersey.

Fisher, Marshall L. and Rosenwein, Moshe B. 1989, "An interactive optimization system for bulk-cargo ship scheduling," *Naval Research Logistics,* Vol. 36, No. 1, pp. 27–42.

Hane, Christopher A.; Jarvis, John J.; and Ratliff, H. Donald forthcoming, "A hybrid algorithm for cargo loading problems," *Naval Research Logistics.*

Hooker, John and Natraj, A. N. 1991, "Routing and scheduling by modified chain decomposition and tabu search," working paper 1991-1, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, Pennsylvania.

Rappaport, Harold; Levy, Laurence; Golden, Bruce; and Feshbach, David S. 1991, "Estimating loads of aircraft in planning for the military airlift command, *Interfaces,* Vol. 21, No. 4, pp. 63–78.

Ray, Julian 1990, "The flow resource facility location problem," PhD diss., Department of Geography, University of Tennessee, Knoxville, Tennessee.

Roherkasse, Robert and Hughes, George C. 1990, "Crisis analysis: Operation Desert Shield," *OR/MS Today,* Vol. 17, No. 6, pp. 22–27.

Solanki, Rajendra S. and Busch, Ingrid K. 1991, "A dynamic programming-based scheduling algorithm for the airlift problem," working paper.

Solanki, Rajendra S. and Southworth, Frank S. 1991, "An execution planning algorithm for military airlift," *Interfaces,* Vol. 21, No. 4, pp. 121–131.