MODELING AN AIRLINE OPERATIONS CONTROL

Nicolas Pujet
pujet@mit.edu

Eric Feron*
feron@mit.edu

International Center for Air Transportation
Massachusetts Institute of Technology,
Cambridge MA.

Abstract

A discrete event model of the Airline Operations Center (AOC) of a major airline is proposed. The model represents each agent in the AOC as a multi-class queueing server, and the complete AOC as a multi-agent, multi-class queueing system. Model parameters include the paths and time constants of information and decisions flowing through the AOC. The identification of these parameters is carried out by combining direct observations in the AOC and statistical analysis of computer transactional data. A computer implementation of the resulting model is described.

Introduction

A key challenge for major US airlines is to achieve efficient information management to alleviate the impact of unforeseen schedule disruptions. In addition to planning up to 2,500 flights a day, the operators in the Airline Operations Center (AOC) of a major airline adjust in real time the movements of the hundreds of aircraft and thousands of crewmembers of the airline to minimize costly delays and cancellations, while complying with complex contractual and maintenance routing constraints[1]. The benefits of the evolution of the National Airspace System towards more user operational flexibility (as embodied in the Free Flight concept) will depend in part on the ability of the airlines to make operational decisions quickly[2]. Thus the economic impact of the AOC decisions and its real-time nature motivate an in-depth analysis of its performance and dynamic characteristics. In particular, such an analysis would help evaluate the performance limits of the current AOC and predict the impact of future AOC decision aids and processes (such as Collaborative Decision Making). This paper presents an approach to modeling an AOC as a network in which each operator is a queueing server.

Queueing network models were initially developed for telecommunication systems[3] and have been analyzed and applied extensively in academic research[4] and industry, in particular in the context of multiprocessor computers[5] and Flexible Manufacturing Systems[6], but to our knowledge they have not been applied to modeling a real-time manned control center.

The first section of the paper presents the structure of the model and the modeling assumptions. The second section details how the model parameters can be identified using a combination of on-site observations, archived operational data and computer transactional data. The third section briefly presents a computer implementation of the resulting model.

Section 1. Structure of the model

The AOC of a major airline is composed of 50 to 100 operators working in three shifts around the clock. These operators are responsible for “dispatching” flights (i.e. preparing and filing flight plans and following flights from departure to destination) and for adjusting the airline schedule (flight schedule, departure slot assignments, aircraft assignments and crew assignments) in response to external perturbations such as thunderstorms, airport capacity restrictions, equipment failures, etc. The control authority of the AOC ranges from a few hours before scheduled flight departure to one hour before departure (see Figure 1). Within one hour of departure, most of the control of the flight is handled by the departure airport (or “station”) control center of the airline.

* Assistant Professor of Aeronautics and Astronautics

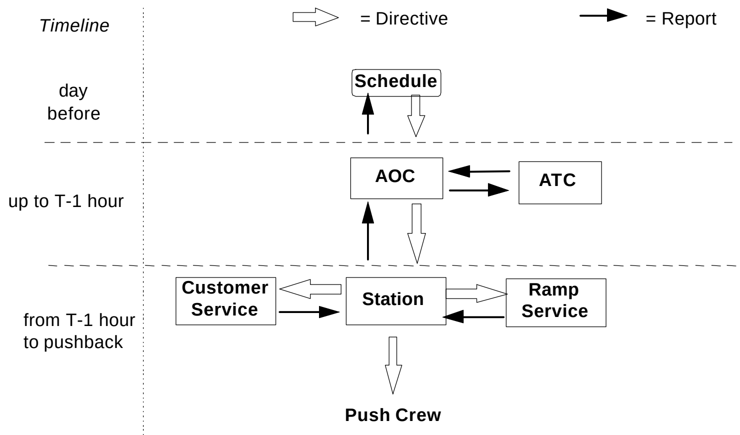

Figure 1. Airline operations timeline (T = scheduled departure time)

Most of the real-time activity in the AOC involves the following operators:

- System Operations Controllers (or SOC) who oversee the operations of the airline and are responsible for the major adjustments to the schedule, such as cancellations.
- Aircraft Routers who are specialized in monitoring and adjusting the routing of aircraft through the network of flights, while complying with aircraft maintenance routing constraints.
- Crew schedulers who monitor and adjust the assignment of crew members (pilots and flight attendants) while complying with contractual constraints.
- Dispatchers who are responsible for preparing flight plans (i.e. deciding on fuel loads, alternate destinations, etc.), releasing and following flights.
- ATC coordinators who are the interface between the AOC and the FAA, particularly during ground delay programs

The model is based on the following hypotheses:

a. Operator model: Each operator is only working on one task at any given time and can therefore be modeled as a single multi-class G/G/1 queueing server[7] (see figure 2).

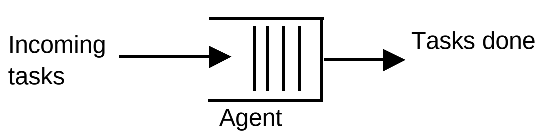

Figure 2. Modeling an operator in the AOC

The “customers” of this queueing server are tasks waiting to be worked on by the operator (e.g. writing a flight plan, finding a replacement for an aircraft which has broken down, etc.). Observations made in the AOC of a major airline have confirmed this hypothesis. The identification of the queueing characteristics of each operator is addressed in section 2. Since each operator in the AOC is modeled as a queueing server, the AOC is now modeled as a network of queueing servers.

b. Processes: The order in which tasks are performed by the operators is quite constant, so that “process maps” could actually be charted to represent the order in which these tasks occur. Figure 3 shows an example of such a process map. Observations in the AOC and extensive interviews conducted with the operators have confirmed that a small number of “processes” is sufficient to give an almost exhaustive description of AOC activities. Each process is triggered by an external event (e.g. mechanical failure of an aircraft, announcement of a Ground Delay program by the FAA, etc.) and includes some queueing any time a task is given by one operator to another. Table 1 (shown at the end of the paper) presents the processes that were identified under irregular operations, while table 2 lists the “background” processes which take place on a more routine basis (regardless of weather and traffic conditions).

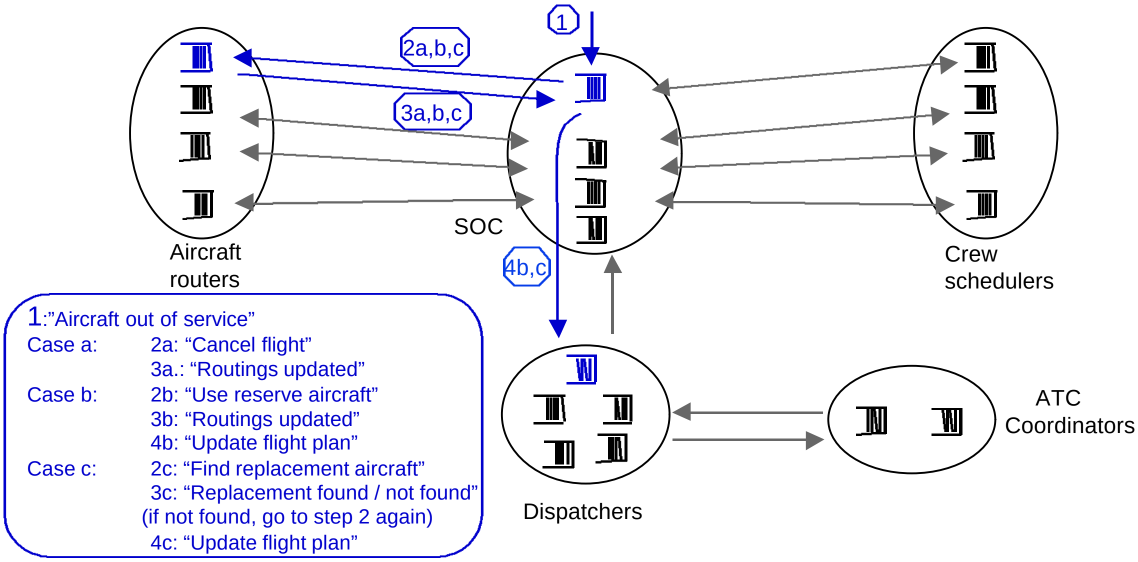

Figure 3. An irregular operations process: recovery from mechanical failure and ensuing flight plan revision

c. <u>External world</u>: the position, status and scheduled assignments of the hundreds of aircraft and thousands of crewmembers, and the hundreds of operational constraints and decision variables that the AOC operators are working on need not be included in the model. Indeed a tractable and useful model of the AOC can be obtained by focusing on the states and transitions of the queueing servers representing the AOC operators. This assumption implies that the model does not account for the "slow" feedback loops from the AOC decisions to its future inputs (in the context of this model, "slow" means taking more than a few hours). Instead, the AOC inputs are modeled by deterministic or stochastic arrival processes, whose parameters are chosen to match statistical measurements in the real system (see figure 4 and section 2).

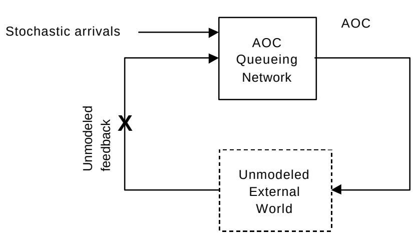

Figure 4. The model replaces slow feedback loops with stochastic arrivals

This assumption is justified by the observation that the effects of these slow feedback loops (e.g. the exact effects of an adjustment to several aircraft and crew routings) are influenced by a large number of uncertain variables (such as future decisions and weather events).

<u>Section 2. Identification of model parameters</u>

The objective of the identification effort is to obtain from the real AOC system some statistics on the parameters of the queueing network model:

- "Customer arrival process", i.e. frequency of occurrence of background and irregular operations processes;
- "Task priorities and service times", i.e. which tasks will be given priority by each operator, and how long they will be worked on;
- "Branching probabilities", i.e. the probability that a process will follow one of several alternative task sequences (see example below).

Example: consider the simplified model of the Mechanical failure process shown on figure 5 (with standard queueing network notations [8])

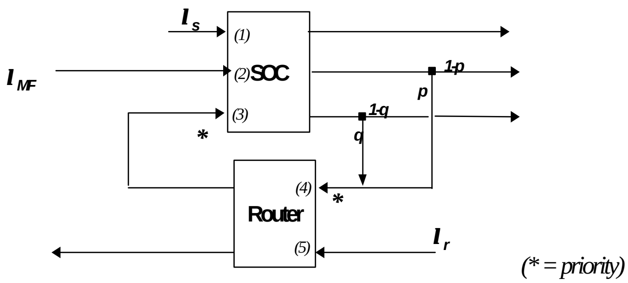

Figure 5. A multi-class queueing network representation of the Mechanical failure process

- the occurrence of mechanical failures and other processes is assumed to follow Poisson processes with the following parameters:
  - ⇨ λ_{MF}: mechanical failures
  - ⇨ λ_s : other SOC tasks
  - ⇨ λ_r : other Router tasks (e.g. routing checks)
- the task service times are the random variables:
  - ⇨ $S_1, S_2, \ldots, S_5$
- the branching probabilities are:
  - ⇨ *p*: probability that the SOC controller calls on the Aircraft Router to help solve the problem
  - ⇨ *q*: probability of a new iteration

An initial on-site identification and calibration effort was carried out in order to obtain reasonably accurate estimates of the parameters of the model. Interviews with operators and on-site manual measurements of service times were conducted in order to build a baseline quantitative model of the current AOC processes and dynamics (see paragraph a. below).

It appeared however that a systematic, objective calibration effort was needed in order to obtain reliable and useful predictions from the model. This led to the collection and analysis of operational and transactional data from the main AOC computer system (see paragraphs b. and c. below).

The transactional data was also combined with Airline Service Quality Performance (ASQP) data to examine the relations between scheduled flight operations, AOC processes and actual flight operations. (see paragraph d.).

a. Identification through on-site manual measurements

The most straightforward approach to obtain identification data is to observe the operators on-site, to conduct focused interviews and to make manual measurements. Table 3 illustrates this approach by showing some of the process parameters of the "Airborne Holds and Diversions" process. These values were averaged over a small number of samples or were estimated by the dispatchers during focused interviews.

Similarly, a first estimate of service times and queueing disciplines was obtained by manual observations and interviews with AOC operators. Some standardized tasks were found to have a fairly constant service time (e.g. preparing a flight plan) while some other tasks were found to have a widely varying service time (e.g. finding a replacement aircraft).

Queueing discipline rules were also inferred from direct observations and interviews. The priority of a task was found to depend on :

- the estimated time until departure
- the type of market and route : international flights having the highest priority, followed by hub-to-hub flights. Simple out-and-back flights have a lower priority.

However, this data collection method presents many shortcomings:

- the results of focused interviews may be biased by the perception that the operators have – or feel they should have – of their functions.
- rare events such as major irregular operations, which could reveal a lot about the performance limits of the AOC, are rarely witnessed, and their size and complexity make it very difficult for a single observer to capture significant data.
- only a few samples of each parameter are collected when many would be needed to obtain useful statistics.
- the number of parameters that is required for a realistic simulation of an AOC on a typical day is quite large. Even without accounting for individual differences between operators of the same function, more than 120 parameters are to be identified to run the processes in tables 1 and 2. A rigorous identification of these parameters would require a large number of samples in different representative conditions (regular day, bad weather day, etc.) and thus appears quite intractable by hand.

b. Identification using irregular operations data

Archived irregular operations data (i.e. data on cancellations, Ground Delay Programs, etc.) can be used to complement on-site observations, since they provide many objective samples.

Thus the statistics of occurrence of some of the irregular operations processes (from table 1) were estimated using archived operational data. For example, the occurrence of Ground Delay Programs (GDP) can be quite accurately modeled as a Poisson process. The parameter of this Poisson process was estimated for each hub of the airline from the list of all GDPs experienced in a given month of interest. The extent of the GDP (duration of the delay program and size of the delays) was also estimated from historical data.

However, the archived data only show the end results of the AOC interventions but give little indication of the dynamic aspects of information flows and decision making which led to these results.

c. Identification using transactional data

Since most of the AOC activity is computer based, the previous data collection methods can be complemented by an automatic, computer-based data collection scheme which resides in the computer system of the AOC.

Operator terminal transactions are continuously recorded in the AOC for auditing purposes. Until now, these logs have apparently never been used to study the large scale dynamics of the AOC. They contain the times at which each operator started (or finished) work on each task (but they do not contain the time at which each task arrived in the operator's queue). These logs can be used to estimate process occurrences and priorities, task service times and branching probabilities for different tasks.

***Example**:* Transactional data was used to study the work patterns and service times of four dispatchers during the same period of time: an afternoon shift in March 1998. Dispatchers #1 and #2 were working on sectors in the western US, while dispatchers #3 and #4 were working on sectors in the eastern US.

Ø *Work patterns*

The distribution of two tasks over time is presented on figure 6.

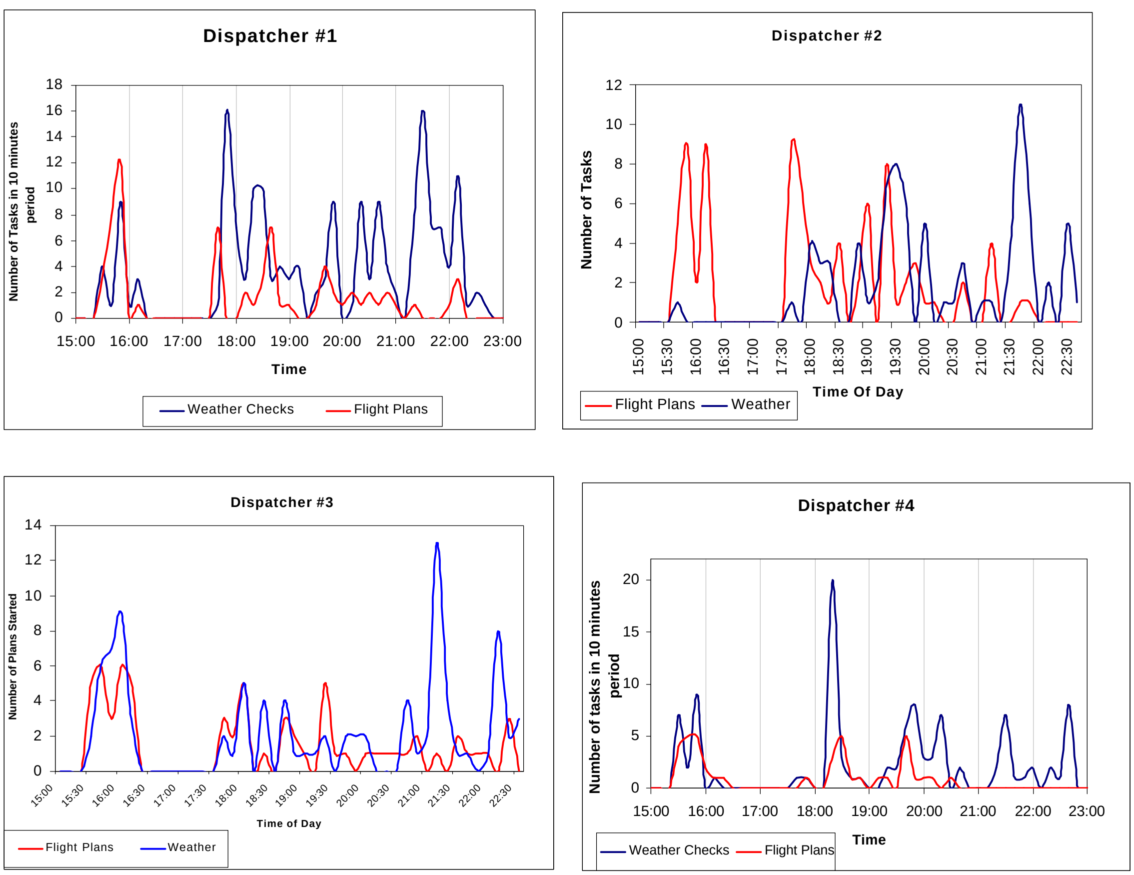

Figure 6. Distribution of tasks over time for four dispatchers
(no data available from 16:25 to 17:30)

These charts display the distribution of the flight plan service start times over the course of the shift along with the distribution of weather checks.

There appears to be a high frequency of weather checks and flight plans at the beginning of the shift (around 15:30), as the dispatchers are trying to get a picture of the situation in their sectors. Other peaks in weather check frequency appear as dispatchers check the new forecasts. The precise time of these peaks varies somewhat between dispatchers. These plots also show that the number of flight plans which are started during the clusters of weather checks (typically just after a new weather forecast is issued) is usually small, which indicates that the two activities are usually not mixed but rather are worked on in batches. This fact had been observed on site, but only qualitatively.

In queueing server terms, since flight plan tasks arrive in the dispatchers' queues quite regularly (following the airline schedule), this means that weather checks have priority over flight plan tasks during these periods.

It appears that dispatchers #1 and #2, whose sectors in the western US were affected by severe weather disruptions, spent more time checking the current state of flight plans in the beginning of their shift than dispatchers #3 and #4, and performed more frequent weather checks in the second half of the shift, alternating with batches of flight plan updates.

In our model, weather checks and flight plan tasks arrive according to a deterministic process; the rate of arrival depends on the time of day and the type of operations (severe weather day or routine day). Note that the charts shown above give the times at which tasks are started ("service commencement epochs"), but not the rates at which these tasks come due ("arrival rate"). One of the next steps of this research will be to obtain these rates from other sources (e.g. respectively, the Meteorology department and the Official Airline Guide – or OAG), so that these charts can then be used to estimate the dispatchers' workloads (and therefore their availability to take over new tasks) and the potential benefits of decision aids and/or improved communications.

Ø *Service times – weather checks*

The computer transactional data can also be used to evaluate task service times, under certain assumptions (namely, that the operator idle time is small). Figure 7 shows an estimate of the service times of the weather checks for the same 4 dispatchers.

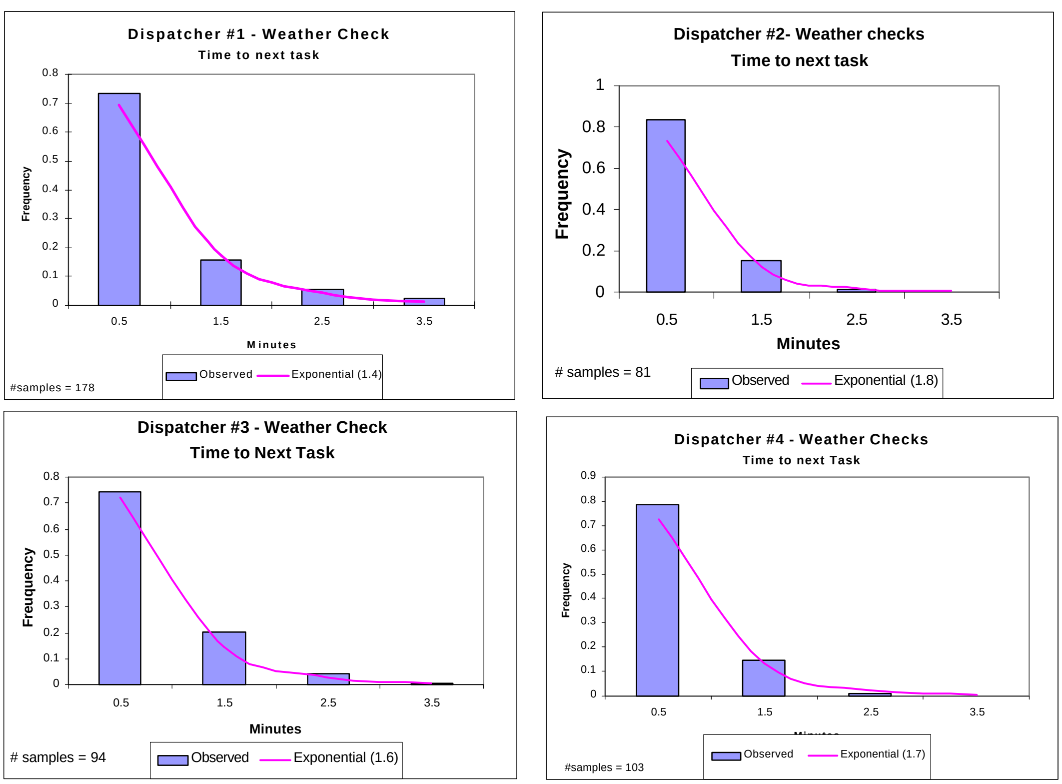

Figure 7. Distribution of weather check service time estimate for four dispatchers

The observed distributions appear to be well approximated by exponential distributions, and the parameters of these exponential distributions seem to vary between 1.4 and 1.8.

Note that dispatcher #1, who appears to have the fastest service rate, also performed twice as many weather checks during his/her shift than the other dispatchers did. This could be a result of the severe weather situation in his/her sector and/or of personal work habits.

The data set indicates that a good model of the service time for weather checks would be an exponential distribution with parameter between 1.4 and 1.8, depending on the weather conditions and the individual operators.

Ø *Service times – flight planning*

Figure 8 shows the distribution of an estimate of the service time for flight planning that was observed for the same 4 dispatchers.

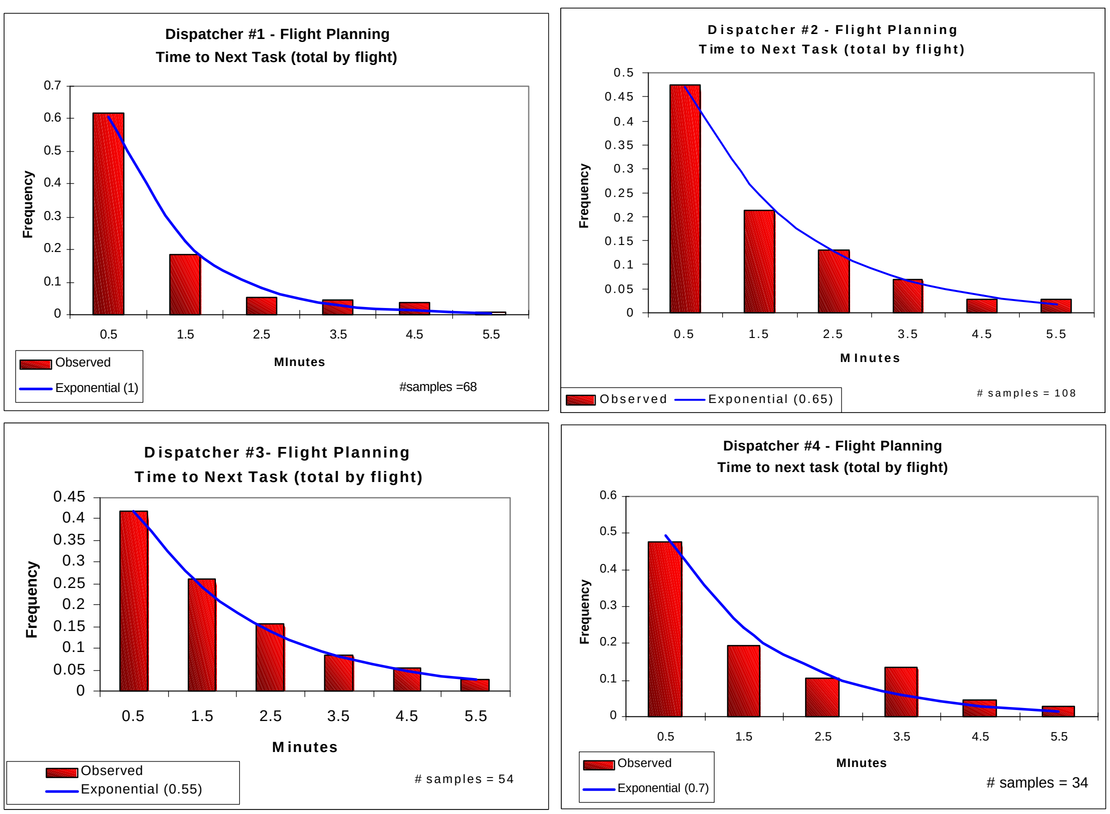

Figure 8. Distribution of flight planning estimated service time for four dispatchers

Once again, the observed distributions appear to be well approximated by exponential distributions. The parameters of these exponential distributions seem to vary between 0.55 and 1.

Once again, dispatcher #1 demonstrates the fastest service rate, but dispatcher number 2 processed more flight plans at a slower rate. Again, this could be a result of the weather conditions and/or of personal work habits.

This data set indicates that a good approximate model of the service time for flight planning would be an exponential distribution with parameter between 0.55 and 1, depending on weather conditions and the individual operators.

Similar statistics are currently being compiled for all the other AOC operators and tasks to obtain an extensive numerical calibration of the model.

Automatic search programs are also currently under development to extract and recognize rare irregular operations processes from the very large amount of transactional data available. The analysis of these processes will give statistics on the occurrence of these events and on the associated “branching probabilities” introduced earlier.

d. Combining transactional and ASQP data

The set of transactional data was also combined with the Airline Service Quality Performance (ASQP) database. This database consists in ACARS reports from the scheduled domestic flights of the ten major U.S. carriers (except Southwest, which collects this data manually) and is publicly available. [9]

The two following examples demonstrate that combining these two data sources provides insight in the relation between the timing of tasks within the AOC and the timing of events outside of the AOC.

Ø *Example 1: timing of flight planning*

For a given flight, define:

- FP_{start} = the time at which the dispatcher started to work on the flight plan
- FP_{finish} = the time at which the dispatcher entered the last flight planning command into the computer system
- SDT = scheduled block departure time of the flight
- ADT = actual block departure time of the flight

Figure 9 shows the distribution functions of SDT - FP_{start} and SDT - FP_{finish} over 4180 domestic flights of a major US airline on March 7th and 8th , 1998.

Figure 9 shows, for example, that:

- for 50% of the flights, the flight plan was started less than 180 minutes before scheduled departure;
- about 60% of the flight plans were started between 1.5 hours and 4 hours of scheduled departure;
- for 50% of the flights, the flight plan was finished less than 90 minutes before scheduled departure;
- for 10% of the flights, the flight plan was finished *after* scheduled departure.

Note that if a flight is subject to changes in crew and/or aircraft assignments, the flight plan has to be revised, which increases the dispatcher's workload. Hence, knowing when a flight plan is typically started would allow to predict the increase in workload that aircraft/crew changes would bring to the dispatcher. For instance, note that for about 35% of the flights the flight plan was started less than 2.5 hours before the scheduled departure time; hence if the aircraft/crew changes are made 2.5 hours before the scheduled departure time, we could estimate at 35% the chance that the dispatcher has not started working on the flight plan, and thus that the changes will not increase his/her workload. Conversely, we could estimate at 65% the risk that the dispatcher has already worked on the flight plan, and that he/she would have to go back to this flight plan and revise it. If the dispatcher is expected to experience a high level of workload in the next 2 hours (e.g. due to airborne holding and diversions, other flight plan revisions, or many airborne flights to follow), a decision could be made to assign the flight plan revision to another dispatcher.

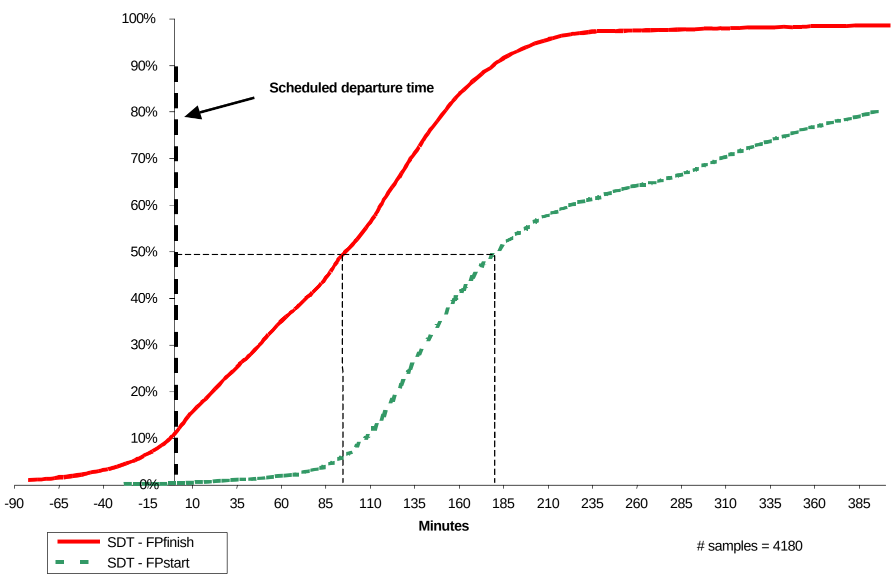

Figure 9. Distribution of flight planning start and finish times relative to scheduled departure

Figure 10 shows the distribution functions of ADT - FP_{start} and ADT - FP_{finish} over the same flights. This figure shows that:

- for 50% of the flights, the flight plan was started less than 195 minutes before actual departure;
- for 50% of the flights, the flight plan was finished less than 105 minutes before actual departure;

An interesting feature of figure 10 is that many flight plans are worked on almost until the aircraft leaves: about 17% of the flight plans are finished less than 25 minutes before actual departure (i.e. while the crew is already on board and the bags and passengers are being loaded into the aircraft).

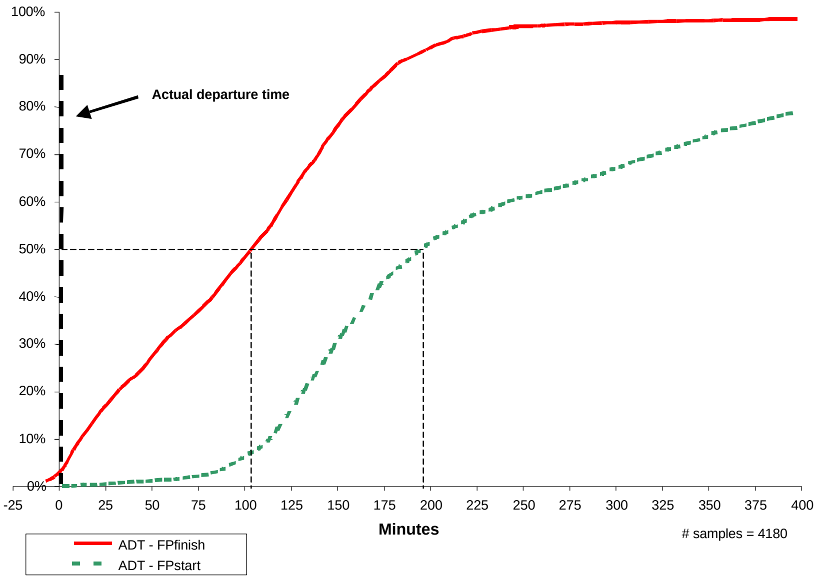

Figure 10. Distribution of flight planning start and finish times relative to actual departure

At this point we can formulate the hypothesis that dispatchers tend to work closer to actual departure on delayed flights. The rationale is that a delay at the departure gate introduces more uncertainty about passenger and cargo loads, since the load planner and customer service agents may take advantage of the delay to add last minute passengers and cargo loads onto the aircraft. If these weight additions exceed a given threshold, they have to be reviewed and approved by the dispatcher, which could explain the last minute flight planning tasks.

The data actually confirms this hypothesis: figure 11 shows the distribution of ADT - FP_{finish} separately for "on-time 5" flights and for flights with more than 5 minutes of delay. This figure shows that dispatchers do work closer to actual departure on delayed flights. For instance, 24% of the delayed flights (compared to only 14% of the "on-time 5" flights) were worked on in the last 25 minutes before actual departure. Some delayed flights apparently suffered from a "vicious circle" in which an initial delay enticed OCC and airport employees to increase the aircraft load, which resulted in additional load planning and eventually a few minutes of additional delay. A possible way to break this circle would be to "freeze" the loads some time before the desired departure time (10 minutes seems realistic). The benefit would be a reduction in last-minute replanning delays. The drawback would be the loss of some passengers and cargo load (some of which could be recovered on later flights). An analysis is underway to study correlations between last-minute load planning and delays.

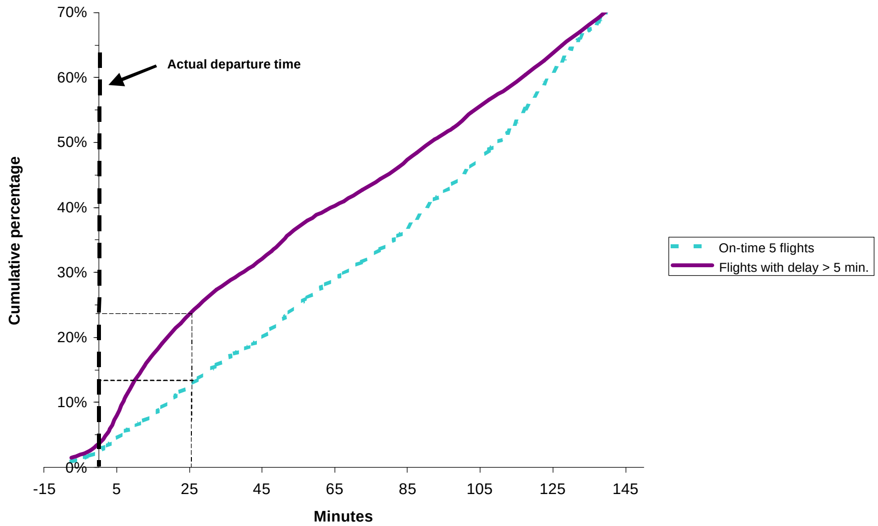

Figure 11. Comparing the ADT – FP_{finish} distributions of "on-time 5" flights and delayed flights

Ø *Example 2: Dispatch reaction to aircraft change*

If the aircraft assigned to a flight is changed (e.g. because the aircraft that was originally assigned is suffering a mechanical failure), the dispatcher has to revise the flight plan. Let $D_{scheduled}$ be the time interval between the new aircraft assignment decision and the scheduled departure of the flight (i.e.how much time is available for the dispatcher to respond)

Let R be the time that the dispatcher takes to finish the flight plan revision after the new aircraft assignment decision (i.e., the dispatcher's actual response time). Figure 12 shows the relationship between R and $D_{scheduled}$. Each dot represents one irregular flight.

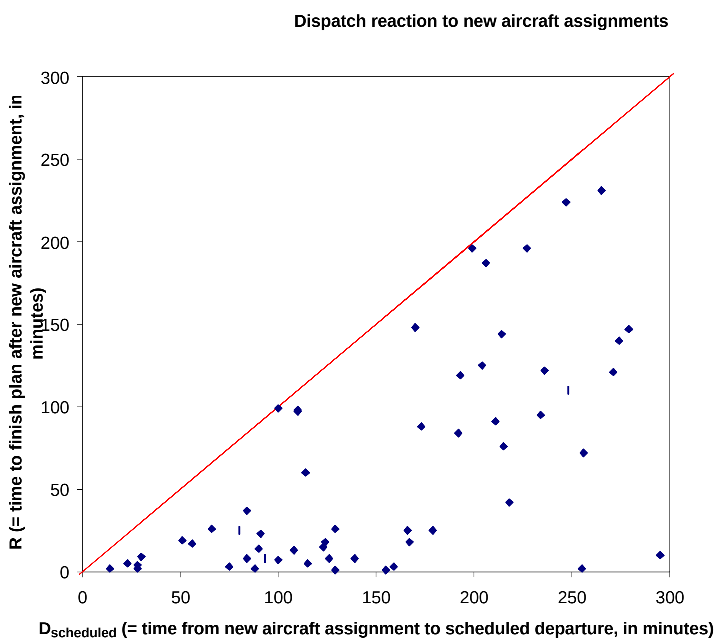

Figure 12. Dispatch reaction to new aircraft assignments

The closer a dot is to the horizontal axis, the faster the dispatcher has finished responding to the new aircraft assignment. The closer a dot is to the diagonal, the later the dispatcher keeps updating the flight plan. A dot on the diagonal means that the last flight plan task was done at the time of scheduled departure (a dot above the diagonal would mean that the dispatcher was still looking at the flight plan after scheduled departure time – these flights are not shown here). This plot suggests that the dispatcher has a different strategy depending on $D_{scheduled}$. If $D_{scheduled}$ is smaller than 100 minutes, the dispatcher tends to revise the flight plan very quickly; if $D_{scheduled}$ is larger than 100 minutes, the dispatcher is in no hurry to revise the flight plan since he/she knows that he/she will have to look at the flight later on anyway.

A more detailed demonstration of these statements is given on figure 13, where $R/D_{scheduled}$ is plotted against $D_{scheduled}$. This ratio represents the fraction of the available response time that the dispatcher actually used.

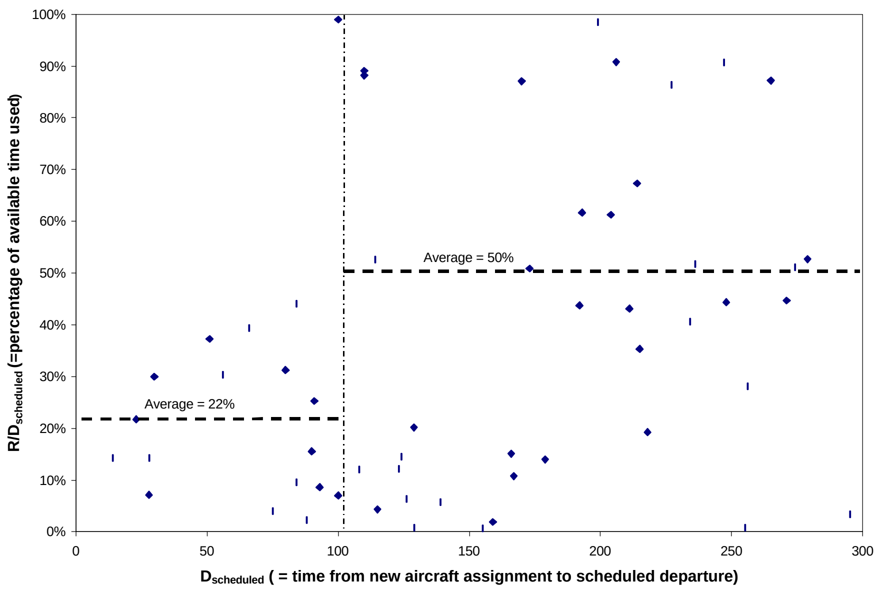

Figure 13. Dispatch reaction to new aircraft assignments: percentage of available response time used

The average actual response time of the dispatcher is on average 22% of the available response time when $D_{scheduled}$ is under 100 minutes, while it is on average 50% of the available time when $D_{scheduled}$ is over 100 minutes.

These figures show that in the AOC model, a new aircraft assignment should result in a high priority flight plan revision task in the queue of the dispatcher if the flight is scheduled to depart in less than 100 minutes, but only result in a low priority task if the flight is more than 100 minutes away. Note that this conclusion rests on the particular data set used in this study and could be refined as additional data becomes available.

Section 3. Computer Implementation of the model

The model described in this paper has been implemented on a computer as a "discrete-event" simulation (i.e. this simulation does not advance time by constant increments, but is controlled by an event calendar which keeps track of when events, such as the launch of a process or the completion of a task, are due to happen). A Graphical User Interface (GUI) was written in Java to allow easy visualization and interpretation of the simulation results.

These simulation tools are currently used to investigate further the dynamics of the AOC and strategies to monitor and control the workflow in the AOC network.

Conclusion

The economic benefits of ensuring timely decision making in the Airline Operations Center motivate an in-depth study of its dynamics. To this end, this paper proposed a modeling approach which represents the operators of the AOC as queueing servers and charts the flow of tasks in this network as fixed process maps. A preliminary calibration effort used focused interviews and direct observations in the AOC to estimate the parameters of the model. A more complete identification was performed using archived operational data and computer transactional data. A computer simulation and a graphical user interface have been built to further investigate the dynamic behavior of the model.

Acknowledgements

We would like to thank Honeywell and NASA Ames Research Center for their support in this research.

.

Table 1. Irregular Operations processes

Process	**Launched by**	**Involves the following operators**
Ground Delay program	Message from ATC	System Operations Controllers Aircraft routers Crew Schedulers Dispatchers ATC Coordinators
Cancellation Plan	Severe weather forecast	System Operations Controllers Aircraft routers Crew Schedulers
Recovery from long delay	Expected delay > 30 min	System Operations Controllers Aircraft routers Crew Schedulers
Recovery from mechanical failure	Aircraft breakdown	System Operations Controllers Aircraft routers
Crew problem	Scheduled crew is unavailable	Crew Schedulers
Flight Plan Revision	Change of aircraft or crew	Dispatchers
Severe weather pilot report	Call from pilot	Dispatchers
Airborne Holds and Diversions	Pilot reporting airborne hold/diversion due to bad weather at destination or airport closure	Dispatchers ATC Coordinators
Assist flight	Flight having lost contact with ATC	ATC Coordinators

Table 2. Background processes

Process	Launched by	Involves
Flight Plan	Schedule	Dispatchers
Checking weather conditions	Weather forecast updates	Dispatchers
Checking aircraft routing	Schedule	Aircraft routers

Table 3. Airborne Holds/Diversions process parameters

Parameter	Value	Comments
Time to read/hear pilot report of first holding pilot	30 seconds 2 minutes	If pilot sends electronic message If pilot calls
Time to build list of future affected flights	60 seconds	Need only one command to print all flights
Minimum inter-arrival time of affected flights	3 minutes	Typical of busy hub airport at rush hours (e.g. San Francisco)
Maximum inter-arrival time of affected flights	10 minutes	
Maximum airborne holding time of each flight	20 minutes	Typical fuel reserve allocated by pilot/dispatcher to airborne holding
Minimum duration of event	10 minutes	Shortest time to get airport arrival process back to normal
Time to call an individual flight and give update	3 minutes	About 1 minute to establish radio patch and 2 minutes to give update

References

[1] "Airlines Operational Control Overview " (Draft 1.4). Airline Dispatchers Federation and Seagull Technology, Inc. May 1995.

[2] "Flight Plan Optimization In Flexible Air Traffic Environments". N. Pujet, E. Feron. 1996 AIAA Guidance, Navigation and Control Conference.

[3] E. Gelenbe, G. Pujolle. "Introduction to Queueing Networks". John Wiley & Sons. 1987.

[4] L. Kleinrock. "Queueing Systems: Theory", vol.1. John Wiley & Sons, New York, 1975.

[5] T.G. Robertazzi. "Computer Networks and Systems: Queueing Theory and Performance Evaluation". Springer-Verlag. 1990.

[6] Stanley B. Gershwin. "Manufacturing Systems Engineering". Prentice Hall. 1994.

[7] D.K.Schmidt. "A Queueing Analysis of the Air Traffic Controller's Work Load". IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-8, No.6, June 1978.

[8] J.R. Morrison, P.R.Kumar. "New Linear Program Performance Bounds for Queueing Networks". Unpublished.

[9] B. Delcaire, E. Feron. "Dealing with Airport Congestion: Development of Tactical Tools for the Departure Flows from a Large Airport". MIT ICAT report 98-3. June 1998.