

Trade-offs and Issues in Traffic Synchronization

Claus Gwiggner, Masato Fujita,
Yutaka Fukuda, Sakae Nagaoka
Electronic Navigation Research Institute
Tokyo, Japan
Email: claus@enri.go.jp

Tasos Nikoleris
University of California
Berkeley, U.S.
Email: nikoleris@berkeley.edu

Abstract—In traffic synchronization, aircraft will receive traffic windows along their trajectories, such that the resulting traffic flows are guaranteed to be smooth and efficient. While the concept is currently being investigated worldwide, its feasibility is still unclear. In this paper we formulate traffic synchronization as a queueing problem and summarize intuitive results based on analytical and simulation studies. These include insight into the delay propagation in arrival flows, trade-offs between ground and en-route delays, and limitations of speed control due to airspace constraints. All in all, the study clarifies the elementary delay generating mechanisms and opens the door to more transparent decision making in tactical air traffic management.

Index Terms—controlled time of arrival, delay propagation, speed control

I. INTRODUCTION

The International Civil Aviation Organization (ICAO) proposes several components of a modern ATM system that are based on expectations of human capabilities and the ATM infrastructure [1]. One of these components is called ‘Traffic Synchronization’. It is described as the ‘tactical establishment and maintenance of a safe, orderly and efficient flow of air traffic’ [1]. Although this is a very general definition that extends the current function of demand/capacity balancing, its core idea is that future traffic flows shall be sequenced, merged and metered at critical airspace resources in order to avoid congestion. One speaks of traffic windows, or controlled time of arrivals (CTA) [2], [3]. In its simplest form, aircraft will receive traffic windows along their trajectories, such that the resulting traffic flows are guaranteed to be smooth and efficient. These windows will be computed prior to departure and updated during the flight. The concept is general enough to support future system implementation levels, such as time-based operations, trajectory-based operations and performance-based operations [4]. The expected benefit of traffic synchronization is a better usage of the available capacity. This will lead on average to more punctuality, fuel and workload efficiency.

While the concept is currently being investigated worldwide (e.g. Queue Management (Sesar) [2], Time Based Flow Management (NextGen) [3], Calculated Fix Departure Time (Japan) [5]), its feasibility is still unclear. For example, the CTA/ATC System Integration Studies (CASSIS) project conducted flight trial experiments and identified a large number of issues with future scheduling and decision making processes [4], [6]. Likewise, the Contract-based Air Transportation

System (CATS) project runs simulations and window size optimizations, but cannot yet answer what is a reasonable number of traffic windows [7]. Finally, pioneering results were obtained with NASA’s Traffic Management Advisor for single centers (TMA), but the extension to multiple centers is not yet achieved [8], [9].

With higher levels of automation in air traffic control being a high priority for the new generation of ATM systems, we believe that a careful analysis of the concept of traffic synchronization and its limitations may prove a timely and beneficial research effort.

The remainder of the paper is organized as follows. In the next section we describe traffic synchronization in more detail. We then formulate it as a queueing problem and summarize our results from analytical and simulation studies. We conclude with an identification of the most important open problems and give recommendations for future research.

II. TRAFFIC SYNCHRONIZATION

In its core, the traffic synchronization problem can be stated as follows: prior to departure and during the flight, aircraft are assigned and updated traffic windows at each critical resource. A critical resource is a merging point, runway threshold, or similar constrained airspace. A traffic window is a scheduled time of arrival plus a window size. The window size ranges between 0 (during high congestion) and infinity (no demand/capacity imbalance).

Assigning traffic windows potentially creates delays, which can be absorbed either on the ground or during the flight. Moreover, in the case that aircraft miss a window, additional delays may propagate through the airspace. Recent studies suggest that trajectory prediction errors have to be expected in the order of ± 30 sec, so the possibility of missing a window cannot be neglected [4].

The problem of traffic synchronization is not new and researchers have approached it with different goals every time. In a study by Meyn and Erzberger [10], the critical resource was the terminal airspace and the goal was in finding the optimal amount of delay to be absorbed in the TRACON area so as to maximize runway utilization. A simulation tool (STASS) was developed and it was found that part of the delay should be scheduled to be absorbed inside the terminal airspace area.

All in all, traffic synchronization can be seen as an extension of today's departure slot allocation in which aircraft receive several traffic windows whose sizes may adapt to the actual traffic patterns. The major questions are then

- What are feasible sizes and update cycles for traffic windows?
- What is a reasonable balance between ground and en-route delay absorption?
- What is the impact of missed traffic windows on flow performance?

The first question pertains to the scheduling part of traffic synchronization, which investigates the size, number and time-headway of traffic windows. Initial research in this area includes work by [11], [7]. The remaining two questions examine the efficiency of a certain schedule of arrivals that is subject to random events. In the next section we present some recent results, both analytical and from simulation, that investigate these questions.

III. PRE-SCHEDULED QUEUES

The traffic synchronization problem can be stated as a queueing problem, where customers (aircraft) ask for service (traffic window) at one or more servers (critical resource). The main differences to classical queueing models are that the arrival flows are pre-scheduled, but possibly delayed (positive or negative) and that traffic windows have to be computed instead of simple service times. The former implies that arrival flows are serially correlated: the more aircraft miss their window in one time interval, the more will arrive in a subsequent interval [12]. Its full analysis includes the interaction of two queueing processes, the pre-departure plus the delayed one. For such reasons, the problem is since long known to be 'notoriously difficult' [13]. Until now, exact analytical results are not mature enough to be used in applications.

Next, we investigate the trade-offs between absorbing queueing delays at high and low altitudes. First, we tackle this problem by minimizing total fuel consumption and compute the optimal fraction of delay to be absorbed in low altitudes. We then adopt a generalized approach and seek for the buffer between arrivals that results in a schedule that combines efficiency and stability.

A. Trade-offs

In independent studies, we proposed engineering approaches to aspects of the traffic synchronization problem. We focused the analysis on very congested traffic regimes.

a) Trade-off between en-route and descent delay absorption: When queueing delays are absorbed in high altitudes, fuel burn is minimized for individual flights [8]. But due to trajectory prediction errors, there is a risk of under-usage of the runway capacity. Lost landing slots may propagate back to the remaining aircraft, which increases the total delay, and as a consequence the total fuel burnt. This means that queueing delays have to be distributed between the high altitudes (fuel efficient) and low altitudes (fuel inefficient), even when the objective is to minimize fuel consumption. Although recent

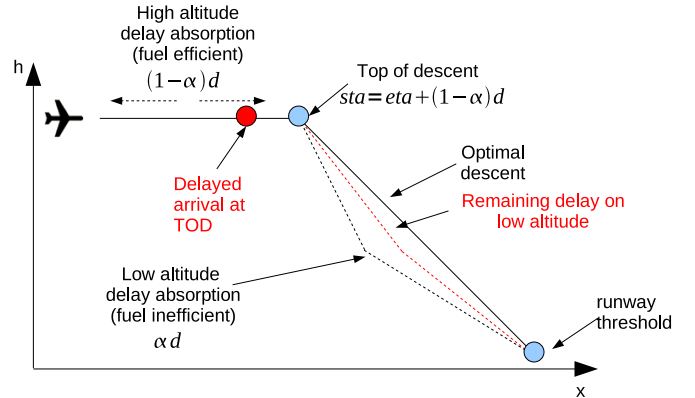


Figure 1. Delay absorption under uncertainty.

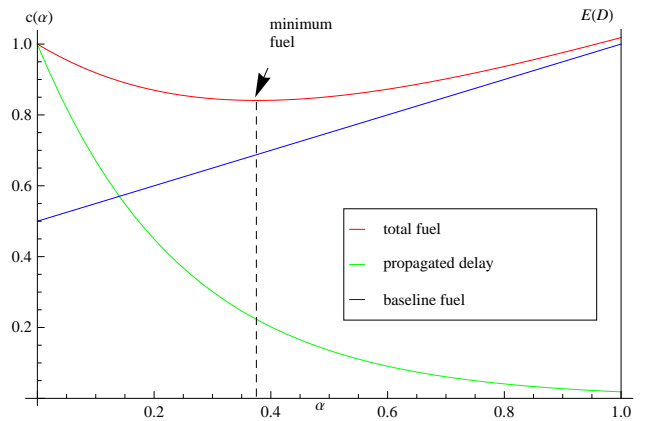


Figure 2. Trade-off in delay distribution.

simulations conclude that elimination of low altitude delays has a modest impact on fuel efficiency [14], the underlying delay propagation mechanism is still badly understood.

As basic model, we consider a single arrival trajectory, as depicted in Figure 1. Given an estimated time of arrival (ηa) at the top of descent, the queueing delay d_i of aircraft i is distributed between high and low altitude

$$sta_i = \eta a_i + (1 - \alpha)d_i, \quad (1)$$

where sta stands for *scheduled time of arrival* and $\alpha \in [0, 1]$ is the *delay balance*. The remaining delay αd_i is included in the sta at the runway threshold. Due to trajectory prediction errors $\epsilon_i \in \mathbb{R}$, the *actual time of arrival* (red point) will be

$$ata_i = sta_i + \epsilon_i. \quad (2)$$

This is similar to [15], except that we do not make assumptions about the service rates.

One can guess from the Figure that delays will propagate when the prediction error ϵ_i is larger than αd_i . We analyzed this process analytically (please see [16] for more details). Our main result can be seen in Figure 2. The horizontal axis is α , the fraction of delay that is absorbed on low altitudes.

The vertical axis has two units: propagated delays and fuel consumption (both are normalized in our illustration). The green line is the average propagated delay that occurs due to trajectory prediction errors, $\mathbf{E}(D(\alpha))$. We obtained an approximation for it, that mainly depends on the probability density functions of the trajectory prediction error ϵ and the queueing delay d [16]. During high traffic densities, the curve always decreases sharply with increasing fraction of absorbed delay in low altitude. This is true for current traffic patterns and for future pre-scheduled ones. As far as the fuel consumption is concerned, we followed the idea of Erzberger [8] and distribute the queueing delay between high and low altitude. This is a simple way to study the average fuel consumption, but more detailed information can be found in [14]. The blue curve is the average fuel consumption in the case that no trajectory prediction errors occur. In this case, the most fuel-efficient strategy is to absorb all metering delays in high altitude ($\alpha = 0$). The red curve is the average fuel consumption under the effect of delay propagation. The trade-off between the low altitude (fuel inefficient) and high altitude (fuel efficient) delay absorption can be seen as its minimum value

$$\min_{\alpha} c(\alpha) = [\alpha c_l + (1 - \alpha)c_h]d(\alpha) \quad (3)$$

$$d(\alpha) = d_o + \mathbf{E}(D(\alpha)), \quad (4)$$

where d_o is the average queueing delay, $\mathbf{E}(D)$ is the expected propagated delay, and c_h, c_l are fuel consumption indices in high (low) altitude in kg per minute. The calculation of the minimum was done by elementary methods.

Our results were validated against several traffic scenarios, including truncated Gaussian, uniform and triangular distributed prediction errors (see [16] for more details). They are in agreement with the simulation studies of Erzberger [8], but our approach is analytical.

b) Trade-off between deterministic and stochastic delay:

The previous analysis computes α , the fraction of delay to be absorbed in lower altitude for fuel consumption minimization. In this section we consider a more general interpretation of the utility to absorb delay at higher altitude and seek for the optimal buffer between successive scheduled arrivals. We consider the case where average demand for service exceeds capacity over a considerable period of time. Aircraft are assigned scheduled times of arrival (sta) at a fix, which they meet with some error (positive or negative). If the minimum required headway between aircraft i and $i - 1$ is h_i , our goal is to find the optimal scheduling (or metering) headway $m_i = h_i + b_i$ between that pair of arrivals. Here, b_i is a *buffer*. To maximize throughput and minimize expected delay, Nikoleris and Hansen [17] show that b_i should be set equal to 0. It may, however, be better to set $b_i > 0$, for at least two reasons. First, as discussed in the previous paragraph, additional flight time can be absorbed in a more fuel-efficient manner. Second, excess time separation b reduces the probability that any late arrivals will propagate backwards. As a result, operations become more predictable and air traffic controllers workload is reduced. A trade-off can be

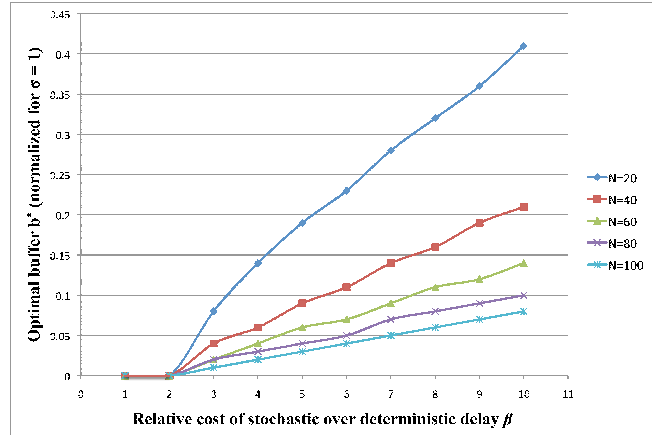


Figure 3. Optimal buffer as a function of β and N (source: [17]).

thus identified between losses in throughput from additional scheduled separation b (deterministic delay) and delays due to imperfect adherence to sta's (stochastic delay).

The deterministic delay for aircraft i is simply ib , while the stochastic delay is defined as aircraft's queueing delay at the fix. Nikoleris and Hansen [17] investigate this trade-off under the following conditions:

- 1) The inserted buffer b between successive aircraft is constant.
- 2) Stochastic errors in meeting sta's are i.i.d normal random variables with zero mean and standard deviation σ .

For a surge of N aircraft arrivals, they express the expected loss $\mathbf{E}(L)$ from the two types of delay as

$$\mathbf{E}(L) = b \sum_{i=1}^N (i-1) + \beta \sum_{i=1}^N \mathbf{E}(Z_i | b)\sigma, \quad (5)$$

where β is the relative cost of stochastic delay over deterministic delay, while Z_i denotes the queueing delay when $\sigma = 1$. One can then find the size of buffer, b^* , that minimizes $\mathbf{E}(L)$. Figure 3 displays values of b^* when $\beta \in \{1, 2, \dots, 10\}$ and $N \in \{20, 40, \dots, 100\}$.

For a given number of aircraft N , the curve of optimal buffer b^* increases with β . This is because, as the unit cost of stochastic delay increases, a larger buffer is required to minimize losses. On the other hand, for a given β , optimal buffer b^* decreases with the number of aircraft N , indicating that the loss from stochastic delays increases at a lower rate than the loss from deterministic delays, as the surge of aircraft becomes larger. That is expected since deterministic delays increase with N^2 (see Equation 5), while stochastic delays increase almost linearly with N as it is shown in [17].

B. Simulation Study

With the simplified queueing models above, we obtained a general understanding of the delay generating mechanisms in traffic synchronization. The underlying sequencing strategies were first-scheduled-first-served, that is, the initial sequence of aircraft was maintained despite delayed arrivals. But for

Table I
AVERAGE WESTERN ARRIVAL FLOW TO TOKYO INT'L AIRPORT (T09).

Origin	Flights	FL_{in} (ft)	v_{in} (kt)	FL_{out} (ft)	v_{out} (kt)	rate (min^{-1})
Central	137 (49 %)	291 (54)	484 (39)	155 (14)	379 (28)	0.15
South	129 (46 %)	357 (44)	507 (39)	157 (16)	382 (29)	0.15
Int'l	13 (5 %)	372 (37)	522 (35)	156 (14)	378 (25)	0.02




Figure 4. Major Japanese traffic flows.

a more realistic view, we developed a fast-time simulator to experiment with new sequencing strategies under trajectory uncertainty. The simulator is trajectory-based and has as major decision variables the traffic sequences at the various merging points. Its main output are visualizations of delay propagation patterns. The simulator was implemented in Java using design patterns. It is freely available and a detailed description can be found in [18]. We report here only one example analysis to illustrate its basic idea.

c) *Priority sequencing*: In today's operations, Japanese arrival flows are merged and metered at the gates between en-route and terminal area. Delays are absorbed on low altitude, which is workload and fuel inefficient. In the future, one wishes to create the conditions for continuous descents: thus metering should be achieved prior to top of descent. A typical problem in this context is the limitation of speed control: the major flows to Tokyo are domestic flights, and some of them have a cruise phase of less than 100 NM. Figure 4 shows these flows and Table I summarizes descriptive statistics of the major arrival flows to Tokyo International Airport. The numbers in parentheses are the standard deviations in the corresponding units.

Aircraft from these flows are sequenced and metered in an en-route sector, called T09. They arrive to this sector roughly in equal number from central and south Japan, and only 5% of the flights are international (column 2). The average ground speed at the sector entry grows with the flown distance, increasing from 484 kt, over 507 kt to 522 kt (column 4). At the sector exit (the gate), the average speed is equally about 380 kt with 27 kt standard deviation, on





Figure 5. Merging strategies.

flight level 156 (columns 5,6). For a typical day, the inflow rates are about 0.15 (ac/min) for central and south Japan and 0.02 (ac/min) for the international flights (column 7), leading to a total arrival rate of about 0.32 aircraft per minute. In current operations, the capacity at the gate is given by a $s_m = 10$ NM spacing requirement. Given the average ground speed of the flow of $\bar{v}_{out} = 380$ kt, this translates into $\mu = \bar{v}_{out}/s_m = 0.63$ (ac/min). In another study, we found

Table II
MERGING STRATEGIES.

Discipline	$E(d_1)$	$E(d_1)$	$P(s > 100NM)$
FCFS	0.78	0.78	0.09
Prio	0.53	1.1	0.02
FCFS	1.79	1.79	0.26
Prio	0.79	2.3	0.05

that the empirical distribution of the corresponding queueing delay can be described by models, in which arrivals occur at random and service times are deterministic [19].

But with a first-come-first-served sequencing strategy, there is a risk that some flights don't have enough time to absorb their delays by speed control. One solution is to change the sequencing strategy from first-come-first-served to a user-defined priority.

Figure 5 visualizes our result. The three Japanese most congested arrival flows will be merged before the top of descent, roughly 150 NM prior to Tokyo International airport (the entry of the T09 sector). The red flow is from central Japan. The yellow from south Japan and the green from International origins, such as South Korea and China. The dots represent the positions at which aircraft will reduce their cruise speed by 10% in order to absorb their queueing delays. This can be computed as follows: When t is the time to fly a distance s at speed v , and t_k is the time at reduced speed kv ($0 < k < 1$) then the absorbed delay is $d = t_k - t = \frac{s(1-k)}{kv}$. Thus, the required distance to absorb d time units of delay at reduced speed kv is $s = \frac{kv d}{1-k}$. For example, aircraft from central Japan have an average cruise speed of 484 kt. Absorbing 1 minute of delay at a reduced cruise speed of 90% takes about 80 NM. The sizes of the dots are proportional to the corresponding number of aircraft. In other words, they represent the spatial distribution of the queueing delays under a 10% speed control rule.

In the upper panel, the distribution for a first-come-first-served scenario is shown. Here, the delays are distributed equally between the three flows. Of particular interest is the tail of the distribution of the red flow: beyond the sector boundary of T21, a few aircraft are concerned by speed control. In Table II we quantify this mass by 0.09 (third column). The average delay of all flights is 0.78 minutes (columns 1,2).

In the lower panel, we used the following priority sequencing rule: if there is a queue, aircraft from central Japan always receive priority over aircraft from the two other flows. We expect from this simple strategy a decrease of queueing delays for the red flow. Looking at the Figure demonstrates the intended effect, because more aircraft than under the FCFS rule have low delays. On the other hand, the distribution has still a long tail. Looking at Table II again, we can read that this time, the proportion of flights having to start their speed control before they enter Tokyo control center is 2%. The average delays for flights from central Japan reduced to 0.53 and those for the remaining flights increased to 1.1 minutes.

We then computed the distributions under a future demand

scenario in which demand for domestic flights increases by 30% and those for International flights doubles. The growth of International flights is predicted for Japan [20]. Average delays and tail probabilities in these cases are in the lower part of Table II. In short, the inequalities between the flight delays become large, as one would expect from such a simple sequencing strategy, and the tails grow with them.

Our current aim is to refine the simulator and identify more flexible strategies that use the available airspace more efficiently. In this context, other statistics, such as fuel consumption and controller workload are also evaluated. Moreover, other absorbing strategies, such as a continuous speed adaptation, instead of a single decrease of cruise speed will be investigated. But we hope that the reader can see that our simulation tool is useful to explore the impact of new traffic synchronization strategies on a system level.

IV. DISCUSSION

It is generally agreed that new navigation technology (flight management systems with required time of arrival function) is the enabler for smoother arrival management and thus for traffic synchronization [6]. Indeed, for one critical resource, such as a runway, the concept of traffic synchronization follows common sense. But already in this simple case one needs to ask for the global goal: is the aim to re-act to the uncertainties that traffic flow management could not predict? This bears the risk to move one bottleneck of the system to another. Or is it desired to close the loop with ATM and provide a system-wide improvement? For example, Japanese flow managers currently distribute 10 minutes of traffic flow management delay in the air, while the rest is absorbed on the ground. In the long-term, their aim is to reduce the amount of en-route delay. Our analysis suggests that the trajectory prediction errors impose a limitation of such a goal. But then, will the technology improvements still be substantial?

The natural next question is if the concept is feasible system-wide. A lot of past research on 'Multi-sector planner' was a failure [9]. Current algorithms for en-route trajectory optimization are at their computational limit. Speed control is available only limited. It seems that identifying the limitations of traffic synchronization is more important than to 'solve' it.

Finally, given a system-wide traffic synchronization algorithm, that has been validated against simulation data. How much confidence can the users put in such an algorithm? Is there a methodology to proof that every traffic pattern will be manageable? Is it satisfying to say that 80% of the traffic patterns will be synchronizable?

At the moment, such questions are far from being answered. But we believe that a good mix of analytical and simulation analysis is the right way to go in order to answer them.

V. CONCLUSIONS AND FUTURE WORK

In the simplest form of traffic synchronization, aircraft will receive traffic windows along their trajectories, such that the resulting traffic flows are guaranteed to be smooth and efficient. These windows will be computed prior to and updated during the flight. Delays will be balanced between ground and air and due to trajectory uncertainties, additional delays will propagate through the airspace. The expected benefit is a better usage of the available capacity. This will lead in average to more punctuality, fuel and workload efficiency. There is worldwide research activity in this concept, but its feasibility is still unclear.

In this paper we formulated traffic synchronization as a queueing problem and summarized intuitive results based on analytical and simulation studies. We found evidence that due to trajectory prediction errors, low altitude radar vectoring (or similar, such as Point Merge) is likely to be necessary in the future. We also found a trade-off between deterministic delays (buffers) and stochastic delays (trajectory uncertainties) that provide useful in increasing predictability and reducing controller workload in future traffic synchronization. Additionally, we developed a new research simulator and explored strategies to distribute arrival delays in a size-constrained airspace. This simulator is particularly useful to analyze time-dependent (transient) phenomena of delay propagation under the impact of trajectory uncertainty.

At the current research stage we want to say that our result are at a fundamental level. We have a certain understanding of the delay generation mechanisms and a new tool to explore sequencing strategies under trajectory uncertainty. Based on this, we believe that a good mix of analytical and simulation analysis is the way to go to provide human-centered decision support in traffic synchronization.

ACKNOWLEDGMENTS

The authors are grateful to the Japan Civil Aviation Bureau of the Ministry of Land, Infrastructure, Transport and Tourism for their cooperation in the provision of radar data.

REFERENCES

- [1] ICAO. *Global Air Traffic Management Operational Concept*. International Civil Aviation Organization, 2005.
- [2] Sesar Consortium. *The ATM Target Concept. D3*. The Sesar Consortium, 2007.
- [3] Joint Planning and Development Office. *Concept of Operations for the Next Generation Air Transportation System*. Version 3.1, 2010.
- [4] CASSIS Project. *CTA/ATC System Integration Studies 2 (Cassis)*. Flight Trials Report. *Eurocontrol*, 2010.
- [5] Study Group for the Future Air Traffic Systems. *Long-term Vision for the Future Air Traffic Systems*. Collaborative Actions for Renovation of Air Traffic Systems (CARATS), Japan, 2010.
- [6] J. K. Klooster, A. Del Amo, and P. Manzi. Controlled time-of-arrival flight trials. results and analysis. In *Eighth USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*, 2009.
- [7] S. Guibert, L. Guichard, C. Rihacek, L. Castelli, and A. Ranieri. Contract-based air transportation system to fulfil the arriving-on-time objective. In *Proceedings of the 26th International congress of the aeronautical sciences. Anchorage. Alaska.*, 2008.
- [8] H. Erzberger. Design principles and algorithms for automated air traffic management. In *AGARD Lecture Series No. 200 on Knowledge-based Functions in Aerospace Systems*, 1995.
- [9] Multi Center Traffic Management Advisor Publication List. <http://www.aviationsystemsdivision.arc.nasa.gov/publications/terminal/mctma/index.shtml>.
- [10] L. A. Meyn and H. Erzberger. Airport arrival capacity benefits due to improved scheduling accuracy. In *AIAA Aviation, Technology, Integration and Operations (ATIO) Conference*. Arlington, VA, 2005.
- [11] L. Corolli, L. Castelli, and G. Lulli. The air traffic flow management problem with time windows. In *4th International Conference on Research in Air Transportation (ICRAT)*, Budapest, Hungary, 2010.
- [12] G. Guadagni, S. Ndreca, and B. Scoppola. Queueing systems with pre-scheduled random arrivals. *Mathematical Methods of Operations Research*, Online first, Sept. 2010.
- [13] D. G. Kendall. Some recent work and further problems in the theory of queues. *Theory Probab. Appl.*, 9/1:1–13, 1964.
- [14] J. E. Robinson and M. Kamgarpour. Benefits of continous descent operations in high-density terminal airspace under scheduling constraints. In *10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Fort Worth, TX*, 2010.
- [15] M.O. Ball and T. Vossen. Analysis of demand uncertainty effects in ground. delay programs. In *Fourth International Air Traffic Management R&D Seminar ATM-2001, Santa Fe (New-Mexico, USA)*, 2005.
- [16] C. Gwiggner. Efficient Delay Distribution in Air Transportation Networks. *To be submitted*, 2011.
- [17] T. Nikoleris and M. Hansen. Queueing models for trajectory-based aircraft operations. In *Manuscript submitted for publication*, 2010.
- [18] M. Fujita. Traffic synchronization simulation tool: Tutorial. Technical report, ENRI, 2011.
- [19] C. Gwiggner, A. Kimura, and S. Nagaoka. Data and queueing analysis of a Japanese arrival flow. *Proceedings of the Asia-Pacific International Symposium on Aerospace Technology. Gifu, Japan.*, 2009.
- [20] S. Nagaoka. ENRI's R&D Long-term Vision. In *Proceedings of the ENRI International Workshop on ATM/CNS (EIWAC). Tokyo. Japan*, 2009.

AUTHOR BIOGRAPHIES

Masato Fujita is a researcher at the Electronic Navigation Research Institute (ENRI), Japan. He received a PhD in Mathematics from the University of Kyoto in 2005. He is a member of the mathematical modeling group at ICAO. His research interests are safety analysis and traffic flow modeling.

Yutaka Fukuda is a researcher at ENRI. He received a Master of Engineering in Japan and was a visiting scientist at NASA Ames in 1998. His research interests are trajectory modeling and conflict resolution.

Claus Gwiggner is a guest researcher at ENRI. He received his PhD in Operations Research from Ecole polytechnique, France in 2007. His research interest are analysis and optimization of transportation systems.

Sakae Nagaoka is a researcher at ENRI. He received a Doctor of Engineering in 1986 from the University of Tokyo. In 1979, he visited the French Air Navigation Research Center (CENA). From 2006-2009 he was visiting Professor at Tokyo University and Tokyo University of Marine Science. He received the Best Paper Award of Japan Institute of Navigation (1993,2007). He is Royal Institute of Navigation Fellow (1986) and IEICE Fellow (2004). His research interests lie in Electronic navigation, safety, and air traffic control systems.

Tasos Nikoleris is a PhD student at the University of California, Berkeley. He received his M.S. degree in civil and environmental engineering in 2006. He worked as a transportation engineer at Dowling Associates, Inc. His research interests include queueing theory and air traffic flow management.