

A Baseline for Terminal Airspace Design Assessment

Tobias Andersson Granberg and Valentin Polishchuk
Communications and Transport Systems, Linköping University
firstname.lastname@liu.se

Billy Josefsson
LFV
firstname.lastname@lfv.se

Abstract—We report on initial steps in our joint work on redesign of Stockholm Terminal Maneuvering Area. We explore possibilities for optimizing assignment of flights to entry/exit points to/from the area under several models of traffic organization on approach/departure routes. Comparison of the distances flown in the optimal flights-to-points matchings under the different routing paradigms allows us to estimate the price of structuring and controlling aircraft within the airspace. In general, our results may serve as a baseline for evaluation of current and future terminal airspace designs in terms of flight efficiency and costs incurred due to the need for control of the traffic flow.

I. INTRODUCTION

Since a great deal of air traffic congestion happens during the initial and final phases of flights, super-dense operations (SDO) in the vicinity of large airports (or air portals – transition airspaces surrounding more than one aerodrome) are a recurring topic in ATM research. Indeed, the complexity of traffic pattern near airports creates higher capacity needs, for the same number of aircraft in the air, than in an enroute setting (e.g., as few as 80 movements/hr already lead to *high* capacity needs in a Terminal Maneuvering Area (TMA), while as many as 160 movements/hr create only *medium* capacity needs enroute [1, p. 20]). In addition, separation standards to avoid wake vortex effects lead to increased sequencing intervals, potentially enforcing time-stretch maneuvers and holding to be executed inside or near the boundary of the TMA (these issues are mitigated by recent efforts in recategorization projects [2], [3]). Last but not least, terminal airspace design needs to take into account ground constraints, such as terrain profiles and noise-sensitive neighborhoods; curved approaches and other techniques are developed in order to smooth air traffic flow in such scenarios.

Stockholm TMA Optimization Project

Stockholm Terminal Maneuvering Area (S-TMA) serves two major airports: Arlanda, with 3 runways and ca. 160000 movements/year (the biggest airport in Sweden and a hub for the national carrier SAS) and Bromma, with 1 runway and ca. 30000 movements/year. In addition, there are 3 small airfields (ESCM, ESOW, ESSU) inside the TMA, that partly contribute to the complex traffic situation, but which will be disregarded in our study. Figure 1 depicts the TMA and its entry/exit points.




Fig. 1. S-TMA boundary is green and entry/exit points are black (note that some points are outside the TMA). The runways at Arlanda and Bromma are thick blue segments. Also shown are tracks of aircraft departed from Bromma (red) and arrived to Bromma (blue) during one day, 2014-05-04 (Arlanda daily flights are too many to show without cluttering the figure).

Luftfartsverket (LFV, the major Swedish Air Navigation Service Provider) manages the TMA. Similarly to many other places around the world, today's S-TMA design is the legacy of a historical development, when expert opinion and rules-of-thumb were used to establish the existing procedures. This was done without a global outlook at the traffic pattern and without using high computing power that has become available since. In 2012 LFV ordered an initial study whose results confirmed the need to investigate possibilities of improving the TMA design with the help of advanced optimization tools. Such investigation is currently underway, performed by LFV in collaboration with Linköping University (LiU) in the course of the ODESTA (Optimal Design of Terminal Airspace) project, funded for the years 2015–2019 by Sweden's innovation agency VINNOVA and in-kind contribution from LFV. A vital component of the project is its reference group, comprised from the industry professionals: Patrik Bergviken (Prod TWR/TMC TMC Landvetter, LFV), Robert Graham (EUROCONTROL), Johan Holmer (Trafikverket, the Swedish Traffic Agency), Anders Ledin (Swedavia), Anne-Marie Ragnarsson (Transportstyrelsen, the Swedish Transportation Authority).

The reference group met in May 2015 at a project workshop organized by LiU, and set up several stages for the project development; in this paper we present the initial results of implementing advices obtained from the reference group

during the meeting. In particular, the group suggested to look not only at the grand challenge of optimally solving the airspace design problem in its full generality, but also to explore possibility of finding "quick-and-dirty" improvements that could be implemented without abolishing the current practices in the overall airspace management. The suggestion is very much inline with the general approach in attacking large-scale optimization tasks: instead of trying to solve the big problem from scratch with some single-shot mega-potent solver, single out several smaller subtasks and optimize each of such components separately (keeping in mind the further opportunities to optimize also interfaces between the components); one classical example of successful industry-wide use of this approach within aviation is splitting the fleet management problem (deciding which plane will fly each link on the schedule) into fleet assignment (deciding which aircraft *type* will fly each link) and aircraft routing (deciding rotations for each aircraft).

From a top view, TMA design falls into the broad class of demand-to-resource matching problems. Specifically, the demand for TMA is formed by the flight plans of the aircraft that intend to land in or depart from the TMA. The resources are the runway(s) at the airport(s), the available fly zones, the surveillance, navigation and control infrastructure, etc. While in principle it is possible to begin solving the airspace design problem from anywhere in the system, one natural approach is to start from a close look at the "outer" and the "inner" boundaries, and gradually expand the optimization frontier, culminating in a "meet-in-the-middle" solution that has optimal designs on both sides. In the TMA case the outer boundary (the "input", the demand) is defined by the flights through the airspace – and this is the focus of the paper. Construction of (parts of) the solution by pushing off the "innermost" structures (the runways), and the design of the "middle-ware" (STARs/SIDs and control sectors within the TMA) are topics of forthcoming work in the project.

In this paper we analyzed the demand for Stockholm TMA based on historical flight data from EUROCONTROL's DDR2 repository. We computed optimal ways to match flights to entry/exit points under two assumptions on the flight paths within the airspace: (1) when the planes follow the currently established STARs and SIDs, and (2) when the aircraft fly directly to the runways. Comparing the assignments between themselves and to the cost of the current operational scheme, we delineated the room for potential improvement in the airspace structure and provided a baseline for estimation of the airspace management cost.

Related work

General guidelines for airspace design are outlined in [4]. Caccavali et al. [5] developed probabilistic models for traffic to a terminal airspace, suggesting alternatives to assuming that the aircraft arrive according to a Poisson process; assigning the arriving traffic to the airspace entry points was not a topic in the paper. Balancing the controllers workload has been the subject in the (re)sectorization research [6]–[12]; in these




Fig. 2. Left: $\mathcal{F} = \{A, B, C, D, E\}$, $\mathcal{E} = \{a, b, c\}$. A 3-matching \mathcal{M} has $\mathcal{M}(A) = \mathcal{M}(B) = a$, $\mathcal{M}(C) = \mathcal{M}(D) = \mathcal{M}(E) = b$. Right: \mathcal{M} becomes an s - t flow of value 5 in the graph where each vertex in \mathcal{E} is triplicated (the edges in the complete bipartite graph between \mathcal{F} and triplicated \mathcal{E} are not shown for clarity).

papers, the ATCOs workload was considered for the entire sector, and not per entry/exit points. Construction of STARs and SIDs within a TMA and TMA capacity estimation was considered in [13]–[20].

II. PRELIMINARIES

This section recapitulates notions from algorithmic theory, relevant for our study. It also describes the information collection and statistical data preprocessing that we performed.

A. Capacitated Matchings in Complete Bipartite Graphs

The *matching* (or *assignment*) is a classical problem in combinatorial optimization. In this paper we will consider matchings in *weighted complete bipartite graphs*. A weighted complete bipartite graph $(\mathcal{F} \cup \mathcal{E}, w)$ is defined by two sets, \mathcal{F} and \mathcal{E} , of *vertices* and the *weight function* $w : \mathcal{F} \times \mathcal{E} \mapsto \mathbb{R}$ that assigns a number, $w(p, q)$ to each pair $p \in \mathcal{F}, q \in \mathcal{E}$; the number signifies the *distance* between p and q , or the *cost* of the "edge" pq (since we consider only *complete* bipartite graphs, which have an edge between every pair of vertices $p \in \mathcal{F}$ and $q \in \mathcal{E}$, we do not explicitly define the set of edges of the graph).

For a number N , a *perfect N -matching* in $(\mathcal{F} \cup \mathcal{E}, w)$ is a set \mathcal{M} of pairs (p, q) , called *edges* of the matching, such that $p \in \mathcal{F}, q \in \mathcal{E}$, for every $p \in \mathcal{F}$ there is exactly one pair in \mathcal{M} that contains p and for every $q \in \mathcal{E}$ there are at most N pairs in \mathcal{M} that contain q . We consider only perfect N -matchings and will drop the modifier "perfect"; also, when N is understood, we call an N -matching simply a *matching*. For $p \in \mathcal{F}$ we denote by $\mathcal{M}(p)$ the (unique, exactly one) vertex of \mathcal{E} to which p is matched, i.e., the vertex such that $(p, \mathcal{M}(p)) \in \mathcal{M}$. Figure 2, left shows an example of a graph and a 3-matching in it. The notion of N -matching that we use is a special case of *b -matching* (in possibly non-complete non-bipartite graph), where b is the vector of vertex capacities (with different vertices possibly having different capacities); a standard matching is a 1-matching.

The *weight* of a matching \mathcal{M} , denoted $w(\mathcal{M})$, is the sum of the weights of its edges: $w(\mathcal{M}) = \sum_{(p,q) \in \mathcal{M}} w(p, q)$. Let $\mathcal{M}^*(\mathcal{F} \cup \mathcal{E}, w)$ denote the *minimum-weight N -matching* (i.e., every vertex of \mathcal{F} is incident to exactly 1 edge of \mathcal{M}^* , every vertex in \mathcal{E} is incident to at most N edges of \mathcal{M}^* , and the total weight, $w(\mathcal{M}^*)$ is the smallest over all subsets of edges

with such properties: $w(\mathcal{M}^*) = \min_{\mathcal{M}} \{\sum_{p \in \mathcal{F}} w(p, \mathcal{M}(p))\}$. Matchings in bipartite graphs are intimately connected to network flows: replicate each vertex in \mathcal{E} N times, add vertices s and t to the graph, connect s to all vertices in \mathcal{F} , and connect t to (all replicas of) all vertices in \mathcal{E} (Fig. 2, right); now there is 1-to-1 correspondence between N -matchings in the original graph and s - t flows of value $|\mathcal{F}|$ in the new graph (here $|\mathcal{F}|$ denotes the size of \mathcal{F} , i.e., the number elements in the set) – the matching edges are exactly those that are used by the flow. In particular, the minimum-weight matching \mathcal{M}^* can be found by computing the *minimum-cost flow* in the new graph, for which efficient algorithms are well known [21].

B. Data Preprocessing

To get input for our problem of optimally matching flights to the entry/exit points, we combined the airspace routes charts obtained from the LFV webpages with flight plans from the EUROCONTROL DDR2 data repository; this section gives the details.

S-TMA Routes

LFV AIS MET and Flight Planning AROWeb [22] contains description of STARs and SIDs for all Swedish airports, including Arlanda and Bromma. We manually extracted the following information from STARs and SIDs e-charts:

- Entry and exit points to and from S-TMA, and
- for each STAR and SID and each entry (resp. exit) point, the distance flown along the STAR (resp. SID) from (resp. to) the entry (resp. exit) point.

The flight tracks in Figure 1 give an idea how some of Bromma’s STARs and SIDs look. Equipped with the above data, we are ready to estimate the distance flown by aircraft along the TMA’s STARs and SIDs, as soon as we know the entry/exit points for the aircraft.

Flight Data

DDR2 repository interface was queried to obtain a historical sample of all flights that originated from or terminated at Arlanda and Bromma in 2014; this resulted in 1Gb worth of data (in the `_m1.so6` format – SAAM 4D flight trajectories last filled flight plans) for the ca. 190000 movements. Thanks to the high quality of the data only a minimal cleanup was needed – we removed the flights that entered or exited the TMA not through the points specified in the LFV echarts (42 flights altogether). We also removed the few flights that had both origin and destination inside the S-TMA (e.g., in 2014 there were 34 circular Bromma–Bromma flights, 119 flights out of Bromma landed in Arlanda, 2 flights landed at ESCM, 21 – at ESOW and 4 – at ESSU; a small number overall). For every flight, we determined the entry (resp. exit) point to (resp. from) the TMA by comparing endpoints of each flight segment with the list of entry and exit points obtained from LFV’s AROWeb; we then followed the flight segments back towards the origin (resp. forward towards the destination), to determine the previous (resp. next) named point of the flight (SAAM special points were excluded from consideration). We

call this point, that precedes (resp. succeeds) the entry (resp. exit) point, the *feeder point* of the flight, or simply the *feeder*. Note that it is more common to use the term “feeder” when speaking about *arriving* flights, since in this case they “feed” traffic into the entry points (e.g., as in “feeder sector”); still we use the term “feeder” also for *departing* flights, to mean the point to which the flight goes after exiting the TMA.

In addition to the main information about the flights (their feeders and entry/exit points), we also retained some other data with each flight: time of flying over the points, fleet type, callsign, flight ID, etc. The preprocessed data was put into a database enabling efficient queries of the type: “Select/Count flights over specified feeders and entry/exit points, within a specified time window, etc.”

C. Statistics of Points Usage

Before turning to the core of our research (matching up entry and exit points with the feeders), we slightly brushed up on the data, removing points with inessential load. Specifically, it turned out that frequency of usage of the points exhibits properties of the Pareto distribution: a small number of points bears the lion’s share of the traffic load.¹ That is, some feeders and some entry/exit points are used very rarely, while others are used quite heavily (Fig. 3). Of course, the fact that a feeder point was used only, say, 700 times during a year does not mean that the point should be treated as inessential to the global traffic picture: even though *on average* the feeder was used less than 2 times per day, it might have been the case that during some periods it was used quite heavily (e.g., 20 times per hour during some 35 one-hour intervals). Therefore we repeated the statistical analysis of the points usage in various time intervals, and identified the set of 155 feeders (out of the total of 195 ever used in 2014) that never carried, altogether, more than 10% of the traffic in any of the intervals. These underutilized points were excluded from further analysis as “outliers” with negligible load.²

We were thus left with 40 feeders to work with (Fig. 4). Removing the low-load feeders excluded 3670 out of the total of 189590 flights, i.e., less than 2%.

III. FEEDERS-TO-ENTRY/EXIT POINTS ASSIGNMENT

This section describes our main results – comparison of distances flown in the TMA under various assumptions on the structure of the flow. We assumed that the flights are not allowed to change their feeders; indeed, flight planning outside of the transition airspace is beyond control of the TMA designer (at least at the tactical level; on the strategic level, the choice of the feeder may depend on the TMA configuration,

¹Note that we do not explore the stochastic nature of the traffic (as was done, e.g., in [5]), and by “distribution” simply mean the traffic statistics, not probabilistic properties of the load as a random variable; in particular, we do not employ p-test or other significance testing techniques to justify exclusion of underused points from the consideration.

²In practice, such removal of points would have to be done with care, e.g., in a situation that flights over them represent some kind of extreme cases (we thank the anonymous reviewer and Patrik Bergviken from LFV for this remark).




Fig. 3. Left: Histogram of the feeder points load (x-axis is the number of aircraft served, y-axis is the number of feeders); majority of feeders served only few flights, while all the essential load was carried by a small number of feeders. Right: On the x-axis are all 195 feeders, from most- to least-used (only some of the feeder names are shown). On the y-axis is the number of flights that passed over the feeders (red) and the cumulative number of flights that passed over the feeders with the smaller load (blue). The red curve drops down fast, and therefore the blue one grows steeply to "saturation" when adding more feeders does not increase the cumulative load by much.




Fig. 4. The 40 feeders and their load (since some feeders are next to each other, their labels overlap); the S-TMA is green.

but we do not model that). Similarly, we do not change the runway for each flight, assuming that the runway assignment has been already made based on time of day, weather, etc. That is, the flights and their assignment to the feeders and runways are given as part of the input to our computations, and we explore the various potential ways of flying between the feeders and the runways.

We use the following notation: A generic flight is denoted by f . The feeder and the entry points through which the aircraft f flew are denoted by $F(f)$ and $E(f)$ respectively; the runway of f is denoted by $RWY(f)$. The great circle distance (GCD) between points p and q is denoted by $GCD(p, q)$. Let $SS(E, RWY)$ denote the distance between an entry/exit point E and a runway RWY along the existing routes in the TMA; when speaking about an arriving flight $SS(E(f), RWY(f))$ is the distance from $E(f)$ to $RWY(f)$ along the STAR, when speaking about a departing flight $SS(E(f), RWY(f))$ is the distance along the SID. We will assume that aircraft fly along great circle arcs between the feeder and the entry/exit point.

We started from computing, for each flight f , the GCD

between the flight's feeder and the runway, i.e., the distance

$$GCD(f) = GCD(F(f), RWY(f))$$

and the total distance

$$GCDF = \sum_f GCD(f)$$

for all flights in 2014. $GCDF$ is a lower bound on the distance that must be flown within S-TMA with the given traffic demand – even without any air traffic control requirements, the planes could not have possibly spent less mileage in the airspace. (In fact, outside 40nm circles centered at the origin and the destination, deviation from GCD is a well established measure of flight efficiency [23].) Of course, such routing paradigm—flying along the great circle arc between the feeder and the runway—is very far from reality, as it completely ignores the potential conflicts due to each aircraft operating in a "FreeFlight" mode.

Therefore, we next looked at the following paradigm: every flight is free to choose the entry/exit point *independently* from all the other flights, and fly GCD from/to the entry/exit point to/from the runway. In this model the flight will choose the entry/exit point E that minimizes the distance between $F(f)$ and E plus the GCD between E and $RWY(f)$; that is, the distance flown by f will be

$$GCD-Greedy(f) = \min_E \{GCD(F(f), E) + GCD(E, RWY(f))\}$$

and the total distance for all flights is

$$GCD-Greedy = \sum_f GCD-Greedy(f)$$

(we denote the distance by $GCD-Greedy$ because each aircraft chooses the entry/exit point *greedily*, without coordination with the other aircraft). Naturally, such routing paradigm does not take into account the possible overload of entry/exit points and, similarly to the above, ignores potential conflicts on the approach/departure due to each aircraft operating in a "FreeFlight" (GCD) mode inside the airspace.

In order to provide some structure to the traffic, we may forbid aircraft to fly GCD between the entry/exit and the runway, and require instead that they use the currently available STARs/SIDs (while still giving the freedom to choose the entry/exit point greedily). Under this paradigm, the flight f will choose the entry/exit point E minimizing the distance $GCD(F(f), E)$ plus the distance $SS(E, RWY(f))$ between E and the runway. The distance flown by the aircraft will thus be

$$Current-Greedy(f) = \min_E \{GCD(F(f), E) + SS(E, RWY(f))\}$$

and the total distance for all flights is

$$Current-Greedy = \sum_f Current-Greedy(f)$$




Fig. 5. Histograms of entry/exit points usage in *GCD-Greedy* (left) and *Current-Greedy* (right): on each histogram, the x-axis shows maximum load (number of flights) of an entry point within 1 hour, and the y-axis shows during how many 1-hour intervals such a load was observed.

(we denote the distance by *Current-Greedy* because each aircraft chooses the entry/exit point *greedily*, without coordination with the other aircraft, but uses current routes through the airspace).

Even though the above routing paradigm avoids the conflicts associated with FreeFlight (GCD routing) within the TMA, it still ignores the fact that each entry/exit point may serve only a limited number of flights during a time interval: Figure 5, right shows that in the *Current-Greedy* scheme, points often have high load assigned to them (interestingly, *GCD-Greedy* features smaller maximum load and, in fact, more points with small load; see Fig. 5, left). To account for this, we split the time into the set I intervals of length $T = 1$ hour and within each interval $i \in I$ computed an N -matching of feeders to exit/entry points, where $N = 7$ is the maximum number of airplanes that were allowed to pass over a point in 1 hour (we chose the value $N = 7$ by looking at the historical data – this was the largest number of flights that went through a single point within an hour in 2014; Section III-A explores matchings using other values of N). We set the cost (weight) of assigning the feeder $F(f)$ of a flight f to an entry/exit point E to be $w(F(f), E) = GCD(F(f), E) + GCD(E, RWY(f))$. This is the same as in *GCD-Greedy*; however now we assign feeders to entry points while respecting the capacity constraint that no entry/exit point is assigned more than N feeders. We computed the optimal matching

$$\mathcal{M}_i^* = \min_{\mathcal{M}} \sum_f GCD(F(f), \mathcal{M}(f)) + GCD(\mathcal{M}(f), RWY(f))$$

within each 1-hour interval $i \in I$ (see Section II-A for definitions and notation related to the weighted capacitated matchings) and the total distance under this model for all flights in 2014

$$GCD-Match = \sum_i \sum_f (GCD(F(f), \mathcal{M}_i^*(F(f))) + GCD(\mathcal{M}_i^*(F(f)), RWY(f)))$$

(we denote the distance by *GCD-Match* because the entry/exit points for aircraft are chosen collectively in the coordinated way according to the minimum-weight matching, but the matching is computed under the assumption of GCD FreeFlight in the TMA).

	GCD	STARs/SIDs
Greedy	<i>GCD-Greedy</i>	<i>Current-Greedy</i>
Coordinated	<i>GCD-Match</i>	<i>Current-Match</i>

TABLE I

THE MODELS: THE COLUMNS SHOW WHETHER GCD FLIGHT IS ALLOWED OR USING CURRENT ROUTES IS NECESSITATED; THE ROWS SHOW WHETHER EACH AIRCRAFT CHOOSES THE ENTRY/EXIT BY ITSELF OR THE GLOBAL MATCHING ALGORITHM (TAKING INTO ACCOUNT THE ENTRY/EXIT POINTS CAPACITIES) IS USED TO ASSIGN THE POINTS. IN ADDITION TO THE FOUR DISTANCES SHOWN IN THE TABLE, WE ALSO COMPUTED DISTANCES *GCDF* (GCD BETWEEN FEEDERS AND RUNWAYS, AVOIDING ANY ENTRY/EXIT POINTS) AND *Current-Current* (THE DISTANCE FLOWN ALONG CURRENT ROUTES AND OVER CURRENT POINTS).

Next, similarly to the Greedy assignments, we computed the optimal matching in the model where the flights follow current STARs and SIDs between the entry/exit points and the runways, i.e., when the weight of assigning the feeder $F(f)$ of a flight f to an entry/exit point E is $w(F(f), E) = GCD(F(f), E) + SS(E, RWY(f))$. As above, we computed the optimal matching

$$\mathcal{M}_i^* = \min_{\mathcal{M}} \sum_f GCD(F(f), \mathcal{M}(f)) + SS(\mathcal{M}(f), RWY(f))$$

within each 1-hour interval i and the total distance

$$Current-Match = \sum_i \sum_f (GCD(F(f), \mathcal{M}_i^*(F(f))) + GCD(\mathcal{M}_i^*(F(f)), RWY(f)))$$

for all flights in 2014 (the name *Current-Match* follows our general convention: current routes are used, and the feeder-to-entry/exit points assignment is optimized globally taking into account the capacity constraints).

Finally, we evaluated the current distance flown in the TMA: for each flight f we computed the distance between the feeder $F(f)$ and the entry/exit point $E(f)$ (that was actually assigned to f , based on the historical data) plus the distance between $E(f)$ and $RWY(f)$ along the current STARs/SIDs:

$$Current-Current = \sum_f GCD(F(f), E(f)) + SS(E(f), RWY(f))$$

The routing paradigms that we explored are summarized in Table I. Figure 6 illustrates the distances that we computed for the flights.

Comparison of the Distances

Table II shows the computed distances. The differences in the flown distance can be attributed to the fact that S-TMA is a controlled airspace. Specifically,

- The difference between *GCDF* and *Current-Current* is the overall cost of controlling the traffic and ensuring separation between the aircraft. (This is reminiscent of the game-theoretic concept of *price of anarchy* – how much equilibrium outcomes of a game differ when the

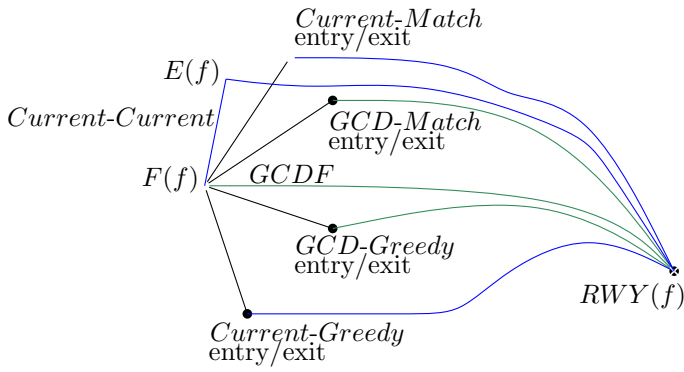


Fig. 6. The 6 compared distances. Green are GCD arcs, blue are current STARs/SIDs.

<i>GCDF</i>	15921330	85.6
<i>GCD-Greedy</i>	15992796	86.0
<i>Current-Greedy</i>	17564138	94.5
<i>GCD-Match</i>	16031315	86.2
<i>Current-Match</i>	17762601	95.5
<i>Current-Current</i>	18799869	101.1

TABLE II

THE DISTANCES (TOTAL FLOWN AND AVERAGE PER FLIGHT), NMI.

players distributively use individually optimal strategies vs. coordinated play [24], [25].)

- The difference between *GCDF* and *GCD-Greedy* is due to the need to enter/exit the TMA through the current entry/exit points. (The difference is small since the entry/exit points are spaced evenly around the TMA, and one does not lose much by going through the points instead of a direct flight between the runway and the feeder; in other words, the graph on feeders and entry/exits points is a good *spanner* having small *geometric stretch factor* [26].)
- The difference between *GCD-Greedy* and *Current-Greedy* signifies the price of keeping the traffic on current STARs and SIDs. The difference between *GCD-Match* and *Current-Match* has similar meaning.
- The difference between *GCD-Greedy* and *GCD-Match* is attributed to the human factors – limited number of flights that may be safely handled during 1 hour. The difference between *Current-Greedy* and *Current-Match* has similar meaning. (Note that we did not take into account the need for separation of aircraft over a point – taking a more profound care of the temporal component is the topic of forthcoming research.)

We remark that even though the differences might not seem large, even small improvements, even implemented on a local scale, may lead to huge savings given the overall large amount of the air traffic.




Fig. 7. The number of 1-hour intervals that have a point whose load is higher than N (y-axis) as a function of N (x-axis). Left: *GCD-Greedy*, right: *Current-Greedy*.




Fig. 8. The dependence of the cost of the matching on N (left: *GCD-Match*, right: *Current-Match*). The dashed horizontal lines show the costs of the greedy solutions.

A. Sensitivity to choice of N

With our choice of $N = 7$, almost half of the 1-hour intervals had an overloaded point (i.e., a point through which more than $N = 7$ flights passed during the hour) in *Current-Greedy*. However, for higher N , the number of overloaded intervals falls down quickly; this holds for *GCD-Greedy* as well (Fig. 7). That is, if a higher load can be safely tolerated, the overload issue in the greedy solutions will have to be fixed in much fewer time intervals.

To quantify the dependence of the solution on N , we computed the matchings for all possible N s – from as few as where needed for a feasible matching to exist up to the maximum load of a point in the greedy matchings. It turned out that the cost drops down with increasing N quite fast (Fig. 8). That is, increasing the control capacity (allowing more aircraft to pass over a single point) indeed helps bringing the distance flown in the airspace down to the values in the greedy solutions.³

IV. CONCLUSION AND DISCUSSION

We singled out a simple, but important subproblem in optimization of terminal airspace management: matching of feeders to entry/exit points. The optimal point assignment constitutes a local modification to the flow management, which can be implemented without changing the current procedures inside the TMA.

³We remark again that delimitation of our model is that we did not take into account aircraft separation: naturally, with 30 flights over a point during 1 hour, the distance between consecutive aircraft would be less than 2 minutes, which may be infeasible for certain categories of aircraft types (even when the reduced sequencing standards [2], [3] are implemented).

Our solution for the redistribution of the demand from the feeder to the entry and exit points can be "propagated" both outward (i.e., over the feeders into adjacent airspaces) and inward (inside the TMA):

- Identifying bottleneck feeders (not simply by counting the load over feeders locally, but also looking at the more broad picture that takes into account the possibility of shuffling the traffic over the available entry/exit points) may suggest looking at alternatives for planning incoming flow before the feeders and outgoing traffic after the feeders (such, more strategic effort is out of the core of the TMA optimization project).
- One possible next step is to optimize locations of the entry/exit points while keeping the existing topologies of STARs and SIDs.

We also believe there exists potential to deepen understanding of traffic even while staying within our framework of "local-scale" improvements, leaving the other parts of the system untouched. One possible enhancement is to weigh distance between the points differently depending on the aircraft type; we included the type information on our data structure, but did not use it during the optimization. Another important extension is to consider 3D ascending/descending flight profiles. Yet another possibility is to evaluate the sensitivity of our output to changes in the parameter T (say, decrease it to $T = 30$ min or increase to $T = 1.5$ hrs) and see the influence on the optimal points matching. Even more importantly, instead of bounding the number of flights passing over a *single* entry/exit, we could have bounded the *total* number of aircraft that fly over the points in a single sector of the airspace; the present study did not take sectorization of the TMA into account, which is a delimitation. Such studies are planned in coordination with the controllers and planners from LFV, as well as with the project reference group.

On a more general note, the work on S-TMA optimization is to continue within the ODESTA project during the coming years. The results reported in this paper will serve as the starting point for future developments.

Last but not least, we note that there are only a few places in our approach where expert (human) intervention is called for: deciding which points are treated as the "outliers" with low load (in the statistical analysis during preprocessing), setting the time horizon T and establishing the limit N on the number of flights over an entry/exit point during the time T ; the rest of the algorithm runs fully automatically. Thus, we envision that the analysis similar to ours can be readily applied to traffic load through TMA boundaries in other airports.

ACKNOWLEDGMENTS

This research is funded by the grant 2014-03476 (ODESTA: Optimal Design of Terminal Airspace) from the Sweden's innovation agency VINNOVA and in-kind participation of LFV. Invaluable input on the subject of the paper has been provided, during a workshop hosted by Linköping University in May 2015, by members of ODESTA project reference group. Many insights were obtained from Patrik Manzi (LFV) during the

authors visit to ATCC Arlanda. We thank EUROCONTROL for granting access to the Demand Data Repository DDR2 which provided the historical data, and the anonymous referees for their useful comments.

REFERENCES

- [1] SESAR, "European ATM Master Plan," 2012.
- [2] EUROCONTROL, "RECAT-EU, optimising Europe's airports capacity," [urlhttp://www.eurocontrol.int/articles/recat-eu](http://www.eurocontrol.int/articles/recat-eu).
- [3] FAA, "Order JO 7110.659B - Wake Turbulence Recategorization," 2015.
- [4] EUROCONTROL, "Manual for airspace design," 2005, section 5, Terminal Airspace Design Guidelines.
- [5] M. V. Caccavale, A. Iovanella, C. Lancia, G. Lulli, and B. Scoppola, "On the statistical description of the inbound air traffic over Heathrow airport," *arXiv preprint arXiv:1302.2027*, 2013.
- [6] A. Yousefi, A. Tafazzoli, B. Khorrami, and T. Myers, "Robust airspace sectorization under weather uncertainties," in *12th AIAA Aviation Technology, Integration, and Operations Conf (ATIO)*, Sep. 2012.
- [7] A. Yousefi and G. L. Donohue, "Optimum airspace sectorization with air traffic controller workload constraints," in *1st International Conference for Research in Air Transportation (ICRAT)*, Zilina, Slovakia, Nov. 2004.
- [8] G. R. Sabhnani, A. Yousefi, and J. S. B. Mitchell, "Flow conforming operational airspace sector design," *Proceedings of the 10th AIAA Aviation Technology, Integration and Operations (ATIO)*, 2010.
- [9] M. Bloem and P. Kopardekar, "Combining airspace sectors for the efficient use of air traffic control resources," *Proc. of AIAA Guidance, Navigation, and Control Conference and Exhibit, Honolulu, HI*, 2008.
- [10] I. Kostitsyna, "Balanced partitioning of polygonal domains," Ph.D. dissertation, Stony Brook University, 2013.
- [11] P. Jägare, P. Flener, and J. Pearson, "Airspace sectorisation using constraint-based local search," in *Proc. 10th USA/Europe Air Traffic Management Research and Development Seminar*, 2013.
- [12] P. Flener and J. Pearson, "Automatic airspace sectorisation: A survey," *arXiv preprint arXiv:1311.0653*, 2013.
- [13] J. Chen, A. Yousefi, S. Krishna, D. Wesely, B. Sliney, and P. Smith, "Integrated arrival and departure weather avoidance routing within extended terminal airspace," in *32nd Digital Avionics Systems Conference (DASC)*, Syracuse, NY, Oct. 2013.
- [14] J. Prete, J. Krozel, J. S. Mitchell, J. Kim, and J. Zou, "Flexible, performance-based route planning for super-dense operations," *AIAA Guidance, Navigation, and Control Conf*, 2008.
- [15] J. Krozel, J. Prete, J. S. Mitchell, J. Kim, and J. Zou, "Capacity estimation for super-dense operations," *AIAA Guidance, Navigation, and Control Conf*, 2008.
- [16] J. Krozel, S. Yang, J. S. Mitchell, and V. Polishchuk, "Strategies to mitigate off-nominal events in super dense operations," *AIAA Guidance, Navigation, and Control Conf*, 2011.
- [17] D. M. Pfeil and H. Balakrishnan, "Identification of robust terminal-area routes in convective weather," *Transportation Science*, vol. 46, no. 1, pp. 56–73, 2012.
- [18] V. Ramanujam and H. Balakrishnan, "Estimation of arrival-departure capacity tradeoffs in multi-airport systems," in *Proceedings of the 48th IEEE Conference on Decision and Control*. 2009, pp. 2534–2540.
- [19] F. Dijkstra, D. Mijatovic, and R. Mead, "Design options for advanced arrival management in the sesar context," *Air Traffic Control Quarterly*, vol. 19, no. 1, 2011.
- [20] H. Lee and H. Balakrishnan, "A study of tradeoffs in scheduling terminal-area operations," *Proceedings of the IEEE*, vol. 96, no. 12, pp. 2081–2095, 2008.
- [21] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*. Prentice Hall, 1993.
- [22] LFV, 2015. [Online]. Available: <https://www.aro.lfv.se/Editorial/View/IAIP?folderId=19>
- [23] EUROCONTORL, "European route network improvement plan part 1: European airspace design methodology guidelines," 2015.
- [24] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria," in *Proceedings of the Conference on Theoretical Aspects of Computer Science, STACS'99*, pp. 404–413.
- [25] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds., *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [26] G. Narasimhan and M. Smid, *Geometric Spanner Networks*. New York, NY, USA: Cambridge University Press, 2007.