AUTOMATED AIRPORT SURFACE TRAFFIC CONTROL

Brian J. Capozzi, Ph.D.*
Metron Aviation, Inc.
Herndon, VA 20170

ABSTRACT

This paper presents a concept for future airport operations founded on the ideas of automated planning of surface operations and active, automated control of surface traffic in order to achieve planned surface movement. We envision such a concept as a stepping stone toward a future NAS air traffic control system which is increasingly automated. The proposed concept of airport operations is designed to dramatically improve the predictability of surface operations while simultaneously supporting increased capacity and delay management in the future National Airspace System (NAS) as traffic demand continues to grow and change over time. Intended to initially provide significant value at major hub airports, this concept can also serve as an enabler for future, point-to-point NAS concepts. The proposed conflict-free, time-based surface trajectories get aircraft to the runway at the time needed to carry out the plan, enabling coordination of runway and taxiway usage between arrivals, departures, and crossings. In describing this concept, we first present a number of factors related to current airport surface traffic control that combine to limit airport throughput. We then introduce a set of core ideas which drive the development of the concept. Based on these core ideas, we present a high-level description of the functionality provided by the concept which we believe will lead to a number of significant benefit mechanisms ranging from capacity to predictability to safety. Finally, we discuss a number of issues related to transitioning from today's airport operations to this future concept, in particular focusing on technology gaps and the proposed shift in roles and responsibilities of controllers due to the introduction of significant automation to the ATC tower.

INTRODUCTION

It is reasonably well accepted across the aviation industry that airports (and the surrounding terminal airspace) represent the primary capacity constraints in the current National Airspace System (NAS). The primary constraints that limit airport throughput are the number of available runways (and procedural constraints related to their use), wake vortex separation requirements, excessive runway occupancy times, downstream departure restrictions, surface congestion, and gate availability.

Several of these constraints (e.g., excessive runway occupancy times) result from the way the airport surface is currently managed. First, the segmentation and distribution of control on the surface introduces a requirement for communication and coordination between different controllers to handle individual flights over different pieces of pavement. This is particularly the case for handoffs between Ramp Control and Ground Control where there is a lack of communication infrastructure to support transfer of preferences and constraints across the "divide" between the movement and non-movement areas of an airport. Often times, procedural constraints are introduced (e.g., one-way taxiway use) in order to reduce the need for explicit coordination between controllers. These constraints can have a negative impact on surface throughput. Other constraints stem from procedures related to safety, such as the requirement to allow only a single aircraft to land or depart a runway at a time. Finally, reliance on visual reference and line-of-sight (e.g., tower view of aircraft position and movement) in the current system introduces reductions in surface capacity during periods of low-visibility, as the situation awareness of both pilots and tower controllers is hindered. Many of these constraints could be relaxed or mitigated through improved sensing accuracy and control authority. However, each of these constraints can play a dominant role with respect to throughput at any given airport or at any given time. Thus, it is not just one constraint that must be relaxed to improve the capacity of airports and, on a larger scale, the NAS.

A number of current research efforts address various aspects of the aforementioned constraints. In terms of technologies for the flight deck, the T-NASA[1], utilized advanced displays such as an Electronic Moving Map (EMM) and Heads-Up Display (HUD) to enhance surface navigation and airport surface position awareness under low-visibility conditions. Similar

---

* AIAA Member, Systems Analyst, Aviation R&D

Copyright © 2003 by the American Institute of Aeronautics and Astronautics, Inc. All Rights Reserved.

flight-deck centric navigation awareness research is presented in Theunissen et al.[2] Another interesting avenue of flight-deck capability enhancement research is that of Cheng[3] who presents the results of simulation studies evaluating the potential for automated taxi control to provide precise time-based surface navigation. Such a capability will enable specification and execution of complex, time-based surface trajectories including required times of arrival at various locations such as active runway crossings. This capability lies at the heart of a proposed future concept of operations for the airport surface described in [4].

Other research efforts focus on providing decision-support for current air traffic controllers. The Departure Planner being developed at MIT[5,6] is designed to provide information to tower controllers to enable the implementation of more efficient departure sequences. MITRE is currently developing its DEPARTS tool[7], which provides similar functionality. Garcia et al.[8] apply both Genetic Algorithm and max-flow/min-cost graph-theoretic techniques to the departure planning problem.

The Surface Management System (SMS)[9], is a decision-support tool (DST) currently under development and operational demonstration testing which provides common situational awareness to all air traffic service providers (ATSP) impacted by or whose actions can influence surface operations. Based on physical modeling of traffic on the airport surface, SMS provides an indication of current and predicted loads on various surface resources, indications of queuing, and a set of "what if" tools to support decision-making. Just as important as the ATSP interfaces, however, is the view that SMS simultaneously provides to NAS users (e.g., airlines) enabling station/ramp operators as well as Airline Operations Control (AOCs) to more efficiently manage their operations. SMS serves as a dynamic information exchange conduit, providing a potential channel for more direct communication of preferences and constraints between NAS users and ATSP.

Although providing decision-support tools aimed at removing constraints within particular NAS domains is a necessary first step to improving NAS throughput, it is also necessary to consider the natural connectivity of the NAS across different domains. This connectivity imposes a set of additional constraints with respect to utilization of scarce resources. Aircraft must be actively and continuously managed not only on the ground, but in the terminal area, en route, and even prior to departure at another airport. The potential benefits related to doing so are discussed in Atkins and Hall[10], which describes the integration of SMS with the Center-TRACON Automation System (CTAS) Traffic Management Advisor (TMA) decision-support tool.

In this paper, we describe a future concept for surface operations being developed under NASA Ames' Virtual Airspace Modeling and Simulation (VAMS) project – a project aimed at developing capacity-increasing solutions to address anticipated demand growth by the year 2020. This paper addresses many similar limitations to those described above. However, the manner in which these limitations are addressed is not merely through the development of DSTs, but rather through a fundamental change in the nature of airport operations. Note that many of these ideas incorporate and extend many of the capabilities of the future NAS as described in various concepts of operations, including the FAA's OEP[11] and NAS Architecture[12] as well as that of RTCA[13]. In particular, we describe a concept for automated airport tower services, providing surface-wide planning and control that reaches outside the airport movement area (to terminal airspace and beyond and to the "ramp") in its scope. Further, we describe the need for explicit integration with NAS TFM via surface capabilities planning in order to provide a seamless progression from strategic planning to tactical implementation. In addition, we propose the need for operational performance tracking to provide an aspect of memory notably missing from the current NAS. Finally, we provide an initial discussion of the roadmap through which we believe such a future system can be achieved, including the need to address a number of socio-political and socio-economic issues related to the fundamental shift in roles and responsibilities associated with our concept.

CONCEPT OVERVIEW

Many of the operational constraints enumerated in the previous section can be traced to a number of root causes, as summarized in **Figure 1**. As indicated, a major contributor to inefficiency is uncertainty related to many aspects of current operations including aircraft current/future position and weather. This uncertainty tends to compound and accentuate the other limiting factors.

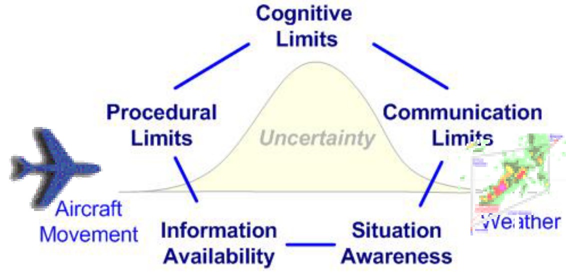

Figure 1: Fundamental Limits on Surface Throughput

The statements above, however, are not intended to imply that current tower and terminal airspace controllers are incapable of providing efficient air traffic services. On the contrary, current ATC provides an astonishingly flexible and exemplary level of service given the tools and procedural infrastructure currently available. The question we ask is: How can we do better? We propose that automation of the planning and execution of surface operations is the way to go for a number of reasons which will be described throughout this paper. Fundamentally, however, we believe that the benefits of automation stem from (a) an increase in the amount and complexity of information that can be used in developing decision-making (planning and implementation) policies, (b) the potential for increasing the extent of look-ahead used in decision-making, and (c) the ability to reason about increasingly complex and multi-dimensional performance objectives.

We describe our concept in terms of a set of Core Ideas aimed at addressing these needs. A particular theme of our concept is the requirement to be non-myopic in the planning and execution of surface movements. The domains of planning and control of our concept are depicted in **Figure 2**. Note, in particular, that the planning aspects extend beyond the physical surface into terminal airspace for both arrivals and departures.

Figure 3 provides a graphical overview highlighting a number of our key concept elements, which are now described in more detail. First, however, we describe a number of key assumptions. We develop our concept under the assumption that on-going research in wake vortex prediction, detection, and visualization (such as the Aircraft Vortex Spacing System[14]) will enable reduction of current minimum required separation standards. By 2020, it is anticipated that wake vortex separation regions will be defined dynamically for each pair of leader/follower aircraft allowing tighter inter-aircraft separation, regardless of the manner in which this is achieved. We also assume infrastructure enhancement providing an increase in the number of runways, the number of high-speed runway exits, and improved sensing, control, and situation awareness displays. These improvements will allow more flexible treatment of runways as independent in all-weather conditions and will reduce runway occupancy time requirements. Such enhancements enable fundamental capacity/throughput limitations to be attributed to actual physical limits rather than the overly conservative procedural constraints currently in place.

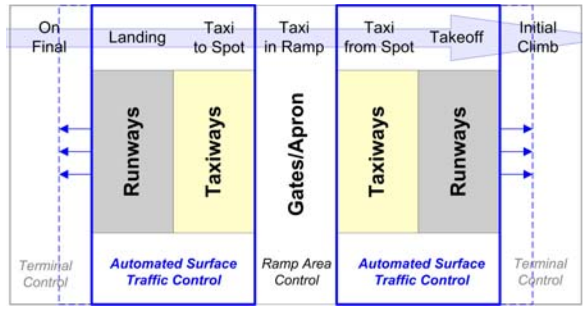

Figure 2: Illustration of the Loci of Control for our Concept

Core Idea 1: Collaborative, Surface-Wide Planning

Since airport capacity is integrally tied to runway throughput, our first core idea is the introduction of automation to plan airport configurations and construct efficient sequences for runway usage. Algorithmic-based configuration planning algorithms will dynamically adapt the airport configuration to anticipated levels of demand, subject to additional constraints arising from weather predictions,

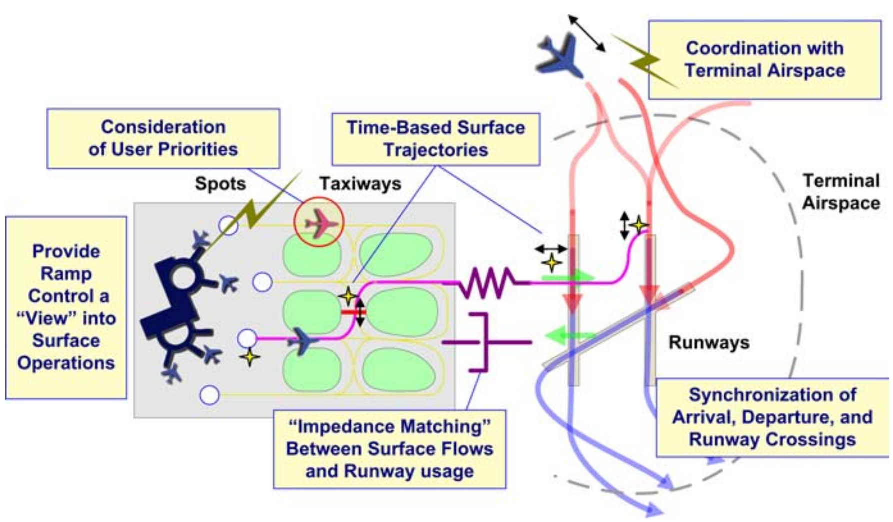

Figure 3: Key Features of our Airport Surface Traffic Control Concept

community noise exposure reduction initiatives, and other considerations. Since changing the airport configuration impacts both airborne flights in the terminal area and flights taxiing or queued on the surface, inputs to these algorithms will come from both the terminal airspace control and NAS users. This will enable coordination between terminal airspace routing and surface traffic routing to discover "preferred" times for planned configuration changes. Similar coordination is required when the active runways are forced to change due to unanticipated events (e.g., blown tire on a runway or sudden, unplanned wind shift).

By looking "beyond the airport" and eliminating the myopic view used for sequencing in the current NAS, our automated runway usage management concept element will enable significantly more efficient sequences than possible today. Through coordination with terminal area control, synchronization of arrival, departure, and runway crossing flows is possible, enabling these streams of traffic to effectively "thread" through one another (extending the ideas discussed in [15]). Sequences will be constructed explicitly considering such things as departure constraints (e.g., MIT, EDCT, APREQ), weight-class, runway slot reachability (based on achievable surface routing), and user priorities. Multiple performance metrics each with its own "weighting" can be defined to control the nature of solution obtained. For example, different parameterizations might favor arrival delay over departure delays or minimize runway-crossing delays. Uncertainty related to pushback times for departures, ETAs at the runway threshold for arrivals, and aircraft taxi performance will be explicitly factored into the solutions to obtain robust runway usage schedules.

It is anticipated that new technologies introduced with respect to control of traffic through terminal airspace will relax the strict mapping between runway and arrival/departure fix often used in today's NAS. More flexible, dynamic generation of terminal area routes, in concert with surface motion planning, will provide a much more flexible departure planning operation. Such technologies will provide a much improved ability to implement efficient sequencing schemes and will dramatically reduce the time required to "turn the airport around" under a much larger set of conditions then is possible in the current NAS.

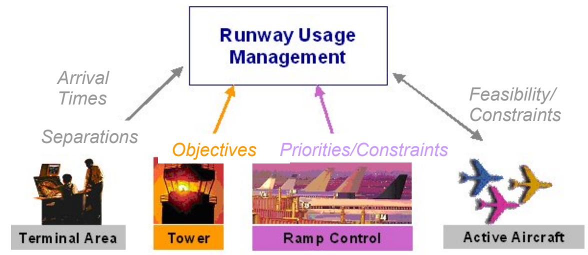

Figure 4: Illustration of the Collaboration Between Different Agents to Determine Runway Usage Schedules

What will enable automation to achieve such improvements is enhanced collaboration and information sharing with NAS users. **Figure 4** highlights the need for two-way flow of information between the various control agents – both on the airport surface and in the terminal area, and with individual aircraft taxiing on the airport surface.

Communication and information sharing will enable the planned runway sequences to be consistent with external flow constraints, environmental concerns, and aircraft-specific constraints. Further, the runway usage plan will be continually adapted based on the evolving likelihood of actually achieving a given plan. The runway usage plan will simultaneously consider arrival and departure flows to match runway throughput to demand while maintaining efficient flow over the airport surface (e.g., avoiding gridlock). The automation will enable higher throughput and make it possible to respond more flexibly and rapidly to unpredictable (or partially predictable) events such as departure fix closings due to weather or a fuel spill on a taxiway.

Core Idea 2: Automated Surface Traffic Control

In order to realize more efficient runway sequences and enable faster response times to changing airport configurations, Core Idea 2 is an improved mechanism for planning and controlling surface movement. In addition to making the runway usage plan of Core Idea 1 possible, this core idea may also provide capacity benefits by enabling reduced separation between events. Without such a mechanism, improved terminal and runway throughput capability would be wasted by taxiway bottlenecks. We thus introduce automation to develop and implement a coordinated motion plan to determine taxi routes consistent with the sequences for runway usage produced by the Collaborative Runway Usage Planner (Core Idea 1).

This core idea includes the functions of negotiating handoffs to/from Ramp Control and the Terminal Area, determining (including timing of movement along routes) and delivering taxi clearances, maintaining separation between aircraft while entering/exiting runways and taxiing, and providing takeoff and landing clearances. Our concept proposes to automate the traffic planning and control functions currently performed by ground and local controllers in the ATC Tower. The primary motivation for introducing automation to this aspect of NAS operations is the cognitive complexity related to producing the time-constrained routes required to implement tightly synchronized runway usage for arrivals, departures, and crossings. The complexity associated with planning and controlling traffic to simultaneously and accurately meet multiple constraints is simply beyond the abilities of human ground controllers. A second motivation is to reduce and/or eliminate the lags in communication and control in current towered operations by providing a tighter connection between the planning and control functions.

The surface routings delivered in the future will be considerably more complex in that they will involve not only 2D paths along taxiways, but temporal constraints as well. These temporal constraints are required to achieve the coordination between flows on the surface required to maximize the efficiency of surface operations. Not only can human controllers not come up with such complex clearances, but the current use of voice communication (and flight deck acknowledgment) between tower and flight deck is incapable of delivering them. Both the time required to deliver and the difficulty in remembering the clearance point to voice communication as infeasible. This core idea thus includes development of alternative clearance delivery mechanisms (see **Figure 5**) such as "out-the-window" surface lighting guidance (green means go, red means stop) and more exotic solutions involving sophisticated heads-up displays and/or Electronic Moving Maps (EMM) and Cockpit Displays of Traffic Information (CDTI) as used in the T-NASA demonstrations[1] and those being developed by Theunissen et al.[2]. Such mechanisms address the need for reducing the lag between the communication and execution of clearances observed today and provide always "on" enhanced vision and situation awareness. They provide pilots with the requisite information to be able to confirm and execute their designated routes in the context of other traffic.

Note that, although we have conceptually separated Runway Usage Management and Surface Traffic Planning/Control as separate processes, they are actually dependent because the runway usage plan must be feasible. A feedback mechanism is included between these processes to resolve situations when the runway usage plan cannot be realized by the Surface Traffic Planner. In such situations, adjustment of the runway usage schedule is necessary. Thus, these processes continually interact with one another in an iterative fashion.

Core Idea 3: Tighter Integration with NAS-wide Planning

Core Idea 1 included coordination between the surface and terminal area for determining how the runways will be used by arrivals and departures. This Core Idea considers how the surface and the TFM system coordinate to plan optimal arrival and departure rates.

Core Idea 3 involves the use of fast-time, airport-specific modeling to establish long-term (2-4 hour) estimates of airport acceptance and departure rates. This modeling will explicitly capture the interaction between arrival and departure flows in sharing gates, runways, and taxiways and will include constraints related to schedule connectivity. Like our other core ideas, this concept element hinges on information sharing and collaboration between NAS Users and automation.

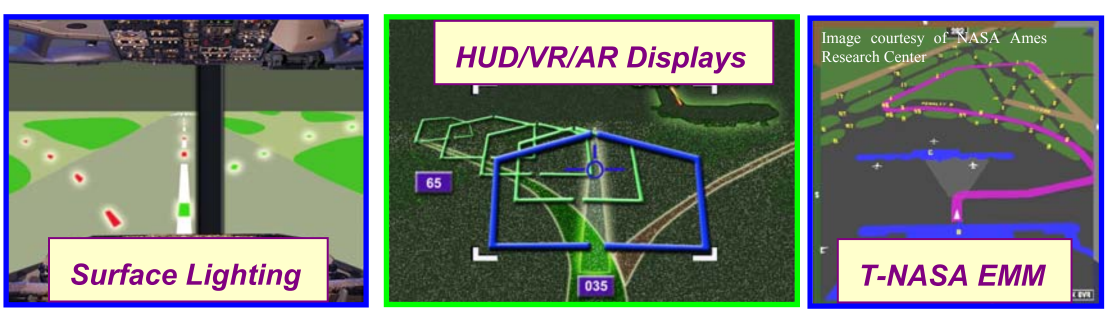

Figure 5: Alternative Clearance Delivery Mechanisms Being Explored

In essence, this core idea can be thought of a separate "application" of our previous core ideas (1 and 2) for a different "customer". Rather than providing real-time planning for actual operations, these tools are instead invoked to provide fast-time investigation of future traffic scenarios. The input data comprising these scenarios will be comprised of probabilistic demand (including effects of cancellations) and probabilistic impacts of predicted weather (e.g., likelihood of fix closure). Feedback from airport surface to NAS TM planning will include the set of feasible arrival/departure rate combinations, the overall arrival and departure delay levels that result from each set point, and the distribution of the delays among different users. Additionally, any particular constraints which may impact utilization of the airport over a given time horizon will be indicated. An example of the latter may be the potential for the airport to reach a gridlock condition or the anticipated lack of gate availability.

In addition to feedback of information to NAS Traffic Management, we propose that each airport maintain a record of the number of operations actually achieved relative to those planned, along with the observed weather conditions and demand. Over time, this capacity-tracking database can be consulted as a mechanism for conditioning the results of the fast-time modeling described above with respect to the likely achievable rates under different conditions. This database can be "mined" to establish key relationships and tendencies that recur – providing feedback useful for adjusting certain parameters which impact surface planning and control performance (**Figure 6**).

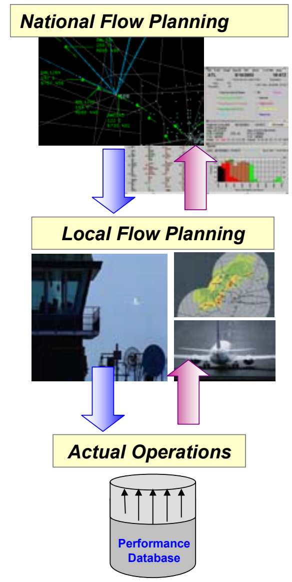

Figure 6: Continual Sharing of Local Information to Aid in NAS-wide Balancing of Capacity and Demand

CONCEPT FUNCTIONALITY

From the perspective of the surface, the rest of the NAS is a single function that interacts with the surface in six ways. The rest of the NAS:

- Delivers steams of arrival traffic to the airport runways
- Provides NAS information (e.g., expected arrival times and trajectories)
- Imposes departure constraints on the airport
- Negotiates with the airport about these departure constraints to ensure a feasible and efficient solution on the surface
- Receives airport status and constraints (e.g., runway configuration and AAR)
- Negotiates with airport about constraints on arrivals (e.g., inter-arrival spacing)

Figure 7 provides an overview of our concept functionality. By defining a "control volume" around the surface domain and specifying the inputs and outputs across that boundary (i.e., how the surface interacts with the rest of the NAS), we express our concept in a way that will facilitate merging with other concepts that address aspects of the NAS Traffic Management box.

Runway Usage Management

Runway Usage Management consists of two parts, Departure Schedule Planning and Airport Configuration Planning.

The output of Departure Schedule Planning is runway assignments and sequence or departure time for each departure, which are provided to the Surface Traffic Planning function. The Departure Schedule Planning function interacts with the rest of the NAS, which we abstract into a single function, NAS Traffic Management (TM), in several ways. First, it receives constraints on the departure plan (e.g., MITs or EDCTs or coordination with en route flows). Second, it negotiates with NAS TM to impose constraints on arrivals (e.g., inter-arrival spacing). Thirdly, it interacts with the NAS TM to ensure that the traffic management solution on the surface that results from the constraints imposed by the NAS TM is both feasible and efficient (e.g., EDCTs are achievable considering where the aircraft is currently located, and departure constraints will not create surface gridlock). Finally, it incorporates any local constraints such as environmental impacts due to emissions or noise (extending noise-avoidance routing ideas discussed in [16]).

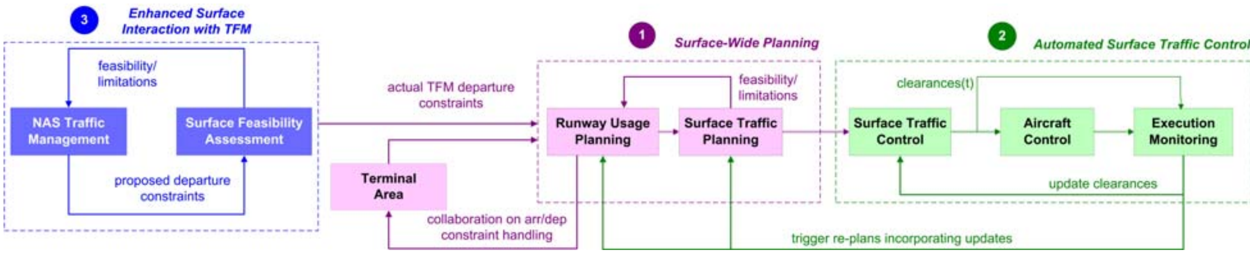

Figure 7: Primary Functions and Information Flow for Automated Traffic Control Concept

Airport Configuration Planning specifies the current airport configuration, planned changes to the configuration, and the planned acceptance and departure rates. These outputs are provided to NAS TM. Primary inputs to this function include the scheduled (for inactive flights) and estimated (for active flights) demand, predicted winds and ceiling, and predicted impacts of weather on terminal area flow rates (e.g., certain configurations may be more compatible with constrained flow rates under some conditions). The output of this function is a list of runways that will be used over the time period.

Note, however, that unlike today's airport usage plans (e.g., Plan X at ORD) in which both runway and nominal operation types (arrivals, departures, or both) is specified, our concept proposes to relax this constraint and leave the operation types on a given runway as flexible as possible. It is anticipated that this function will execute a planning cycle to both (a) establish if a configuration change is warranted and (b) determine the most efficient configuration change time. **Figure 8** provides an illustration of this idea, showing several options related to different potential configurations capable of achieving the required flow rates over the time period of interest. The left-hand figure displays a "work" metric proportional to the difficulty in changing from the current configuration (marked with an "X" on the right-hand Pareto curve) to one with the desired mix of arrivals and departures (marked with an "O").

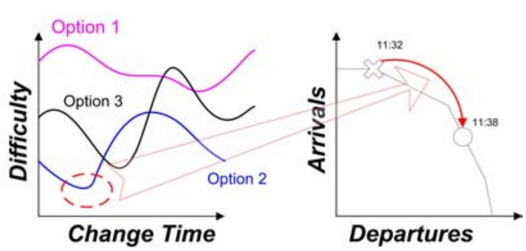

Figure 8: Use of a Work Metric for Determining Changes in Airport Configuration

By comparing work metrics as a function of configuration change time, the Configuration Planner can identify the best change time, as indicated by the dotted circle.

Surface Traffic Planning

Surface Route Planning is done simultaneously for both arrivals and departures and in conjunction with runway assignment and scheduling. **Figure 9** highlights the high-level interactions between these functions.

The output of the Surface Route Planning is a time-based, conflict-free taxi route that meets control time constraints for each specific flight. In aggregate, the set of taxi routes is designed to enable the flow constraints (in terms of flow rate) that are placed on the surface to be met under the assumptions used during the current iteration of the planning process.

Arriving flights whose assigned gate is currently occupied will be either taxied to pre-defined holding locations or, if the gate will be available in short order, the aircraft will be taxied to a "virtual" holding location on a taxiway or put into a surface "holding" pattern which will keep it clear of other surface traffic yet keep the flight moving and ready once the gate becomes available. This capability can be though of as a Holding Area Management functionality.

Similarly, when de-icing operations are required at an airport with defined (possibly remote) de-icing pads, a

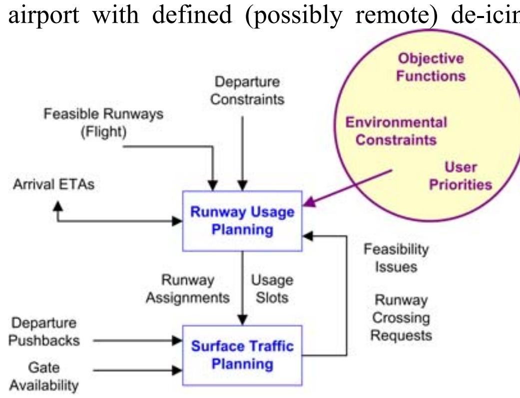

Figure 9: Interaction between Runway Usage and Surface Traffic Planning Functions

De-Icing Management sub-function will be responsible for monitoring and managing the flow of traffic to/from remote de-icing pads and providing information back to the Surface Traffic Planner as to the likely time at which flights will have completed de-icing and be ready for taxi to the runway.

Note that not all taxi routes will by pliable in a given iteration of the planning process. At a certain point, the taxi route will be "frozen" – at least in its spatial degree of freedom. After that time, only time-based adjustments will be made as necessary to remain conflict-free and meet constraints.

Surface Traffic Control

Surface functions related to traffic control correspond to the implementation and actuation of the taxi routes and runway assignments developed by the planning functions of Local Traffic Management.

Departure Clearance/Pushback Clearance

For the purposes of our concept, we assume that departure clearances in the future NAS will be communicated to the flight deck in an automated fashion. Information contained in the departure clearance message will include updated weather information (e.g., Automated Terminal Information System (ATIS) messages), initial runway assignment, and initial terminal area routing. These clearances will result from a negotiation/coordination process between the terminal area and the surface regarding the runway and routing to be used for each departing flight (based on matching overall terminal area and surface throughput to meet NAS constraints). Because we anticipate that future terminal area departure routes/procedures will be more dynamically defined than in the current NAS, however, this initial clearance may actually be updated multiple times while the aircraft is taxiing. For example, the departure clearance received before de-icing might be modified due to events that occur between the time of pushback and the time the aircraft has completed de-icing operations. This again shows the degree of coupling required between the Runway Usage Planner and the Surface Traffic Planner. Receipt of clearance will utilize some form of datalink communication for properly equipped aircraft. For non-commercial users, this might consist of a handheld wireless device similar to a PDA. Note that, under a concept in which NAS TFM constrains the complete 4D trajectory from runway to runway, information regarding terminal area routing and runway assignment is explicitly defined – thus, departure clearances essentially become obsolete.

The function of Pushback Clearance will remain a ramp area responsibility for those airports with dedicated ramp control. For gates that push back onto the airport movement area, pushback will be controlled by the human-staffed position in the tower, working in conjunction with the surface taxi planning automation. This position will be responsible for taxiing the aircraft safely from gate/apron stand to handoff spot on the airport movement area. This function is executed on an event-driven basis, triggered by notification by a given flight that it is "Ready to Push".

Handoff From/To Ramp at Spot

The Surface Traffic Planner automation will use surface surveillance data to determine the presence and anticipated presence of aircraft at handoff spot locations. It will be capable of accepting and initiating multiple handoffs simultaneously. Control of the taxiing of each aircraft will be assumed by the automation at these locations and asserted through datalink communications and flight deck acknowledgment.

Acceptance of handoffs of arriving aircraft from the surface automation to ramp control personnel will be communicated using datalink or equivalent. Ramp personnel will utilize decision-support systems to determine potential ramp area congestion or conflicts and provide an indication to the automated ground controller of the time at which the handoff of specific aircraft will be accepted. This information will be used by the taxi planner to design taxi routes to meet such handoff constraints.

Delivery of Surface Movement Clearances

The output of the taxi planning algorithms will be conflict-free surface trajectories for each aircraft from handoff spot through the runway for departures and from runway threshold through handoff spot for arrivals. These trajectories will be defined as a sequence of intersections as done today, along with a set of controlled times of arrival for critical intersections along the path. In general, not all intersections will be controlled. In some cases, the taxi route may simply be a geometric path with no time constraints while in others, a controlled time of arrival at each intersection might be specified. In general, this trajectory will include clearance of flights across active runways.

Taxi routes and runway clearances will be communicated to the flight deck via a number of mechanisms, depending on aircraft equipage. At one extreme, automated verbal clearances will be generated similar to those used today. Limitations of this delivery technique include difficulty in communication (due to message length), acknowledgment by the flight deck (i.e., read back), and the ability of the pilot to retain in memory the complex clearances required. A potential solution to this difficulty is for the surface automation to instead issue "standardized" taxi clearances (identical to those used today) for non-equipped aircraft. Taxi routes defined in this manner would then serve as additional spatial and temporal constraints for planning the routes of other, equipped aircraft. Alternatively, the taxi route can be communicated via datalink (e.g., extending the capabilities of existing Controller-Pilot Data Link Communications, or CPDLC) wherein a textual representation of the taxi route will be shown.

An additional drawback to each of aforementioned delivery mechanisms is that they hinge on "out the window" navigation by the pilots and orientation of the taxi instructions relative to aircraft position on the airport surface. As such, they can require considerable workload, especially in low-visibility conditions when airport pavement markings and signage become difficult to see. Alternative mechanisms include use of visual taxi clearance representations via either datalink, an electronic moving map, and/or heads-up display (HUD), similar to that used for the T-NASA demonstrations[1]. A heads-up or virtual reality display enables scene-linked symbology and additional situational awareness cues related to other relevant surface traffic to be communicated in addition to the taxi clearance. In addition, these visual enable greater flexibility for the automation to specify the time-sensitive nature of reaching certain intersections and allow for easy visual communication of sequence and spacing information.

One last possibility – should the time-based clearances prove (through human-in-the-loop studies) to be too complex to either communicate and/or execute – is for the surface control automation to send clearances directly to the aircraft's Flight Management System (FMS). This would obviously require an automated taxi capability through which the aircraft control system (throttles, brakes, steering) can effectively track "reference" surface trajectories. Work by Cheng[3] has demonstrated in simulation the potential for such a system to be designed and meet rather stringent timing constraints and navigational accuracy. Such a communication paradigm represents a clear path for integration of UAVs into future surface operations.

Runway Usage Control

Since our Automated Airport Surface Traffic Control concept plans an integrated motion plan for all aircraft from spot through takeoff and from runway threshold to spot, there is no need for "handoffs", per se at the runway threshold or exit, respectively. Rather, the surface trajectories explicitly include runway occupancy times in their definition.

Control of arriving aircraft will be automatically transferred to automated Surface Traffic Control once the aircraft is established on final approach and the runway slot has been finalized through coordination between surface and terminal area control. At this point, the aircraft will be sent a message proving a clearance to land – which must be acknowledged by the aircraft. As the taxi route will have already been determined for each arriving flight prior to its actual landing on the runway, any necessary land and hold short instructions will already have been communicated to the aircraft by the time it touches down. It is anticipated that the proposed integrated approach to surface traffic and runway usage planning should, in fact, minimize the necessity of such operations. As such, landing aircraft will nominally exit from the active runway and immediately commence taxiing on its appointed route toward its assigned parking location.

Conflict Detection and Resolution

By design, the planning components of our concept produce de-conflicted surface trajectories. However, they are only de-conflicted under the assumption that each flight tracks its "reference" trajectory to within a specified degree of conformance. Our concept employs separate "watchdog" logic, driven by live surface position data, to handle situations in which flights are unable to or fail to track their intended taxi route or when unanticipated potential conflicts occur (see **Figure 10**).

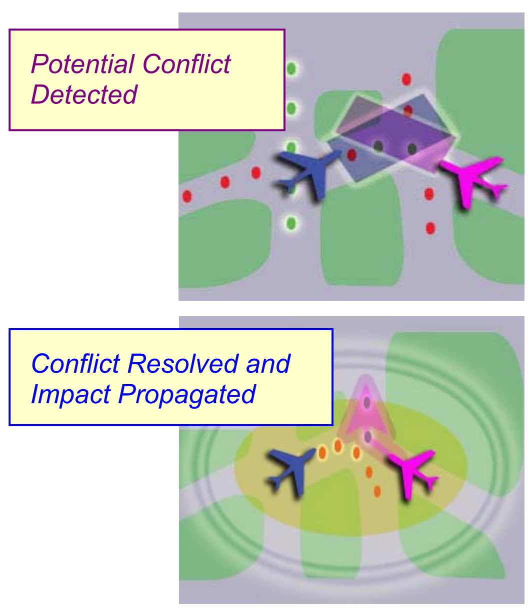

Figure 10: Depiction of Conflict Detection and Resolution Behavior

Thus, as flights are moving along their planned surface trajectories (e.g., time-based taxi routes), separate "watchdog" logic is monitoring their progress and identifying when a potential conflict occurs or when a flight deviates outside of an acceptable margin from its intended course and position (e.g., a missed turn or a pilot that is slow to cross an active runway). This logic will trigger an immediate re-plan of the taxi routes for all impacted flights – with safety as the primary objective function. This re-plan will guarantee a satisficing, safe solution; even if it involves bringing the impacted vehicle to a stop, sacrificing airport performance in a transient sense (see **Figure 10**).

Figure 11 provides an approximation of the relative execution rates and planning horizon (e.g., temporal extent of data) used for each of the high-level concept functions. At this point in time, the actual data rates and output rates for each function are not known. These will be determined through sensitivity studies in subsequent phases of research. The shaded rectangles are meant to indicate the range of look-ahead times over which each function can act, with the intensity of the shading depicting the initial focus areas for the sensitivity studies to follow. The functions are arranged along the vertical axis according to the relative uncertainty of the data used for its execution. The colored rectangles connected to the left of each shaded bar represent initial bounds on the update rates that we will consider.

GETTING THERE FROM HERE

Enabling Technologies

There are several enabling technologies that must exist for the proposed concept to be implemented. Some of these technologies currently exist in forms that will require maturing while others are new technologies that will need to be developed (e.g., technology "gaps"). These technologies and their relationship to our concept are described in this section.

Surface Surveillance

The proposed concept relies on accurate, reliable surveillance of the entire airport surface (i.e., movement area and ramp areas). The surveillance must provide position and identity of all vehicles (aircraft and ground service vehicles) on the movement area. Note that this will require equipping the ground service vehicles with transponders. Separation from vehicles such as baggage carts and catering trucks near the gates will remain the responsibility of the flight crews, ground crews, and

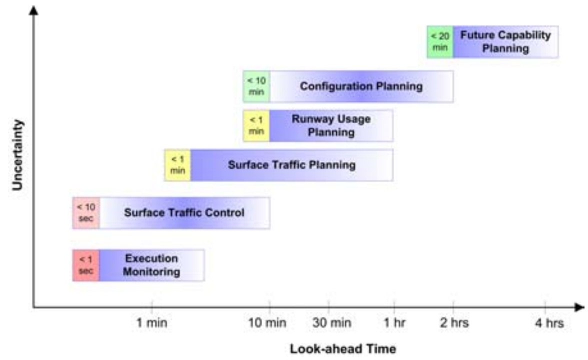

Figure 11: Initial Ranges for Function Planning Horizons and Execution Rates

vehicle operators. The accuracy and reliability required by this concept are the same as for any safety critical system used to maintain aircraft separation. The ASDE-X system currently being developed and deployed will be a first generation of such a system. Subsequent improvements and increased use of ADS-B will provide the needed surveillance in the appropriate timeframe.

Surface Safety Monitor

When the intended taxi paths and clearances are known in advance for every vehicle, the surface safety problem may be decomposed into two parts: checking that the taxi clearances generated by the automation are safe, and taxi route conformance monitoring/alerting. Note that to assure aircraft safety, the taxi path planning automation must be aware of the movement of all vehicles (e.g., snow plows and emergency services) on the movement area. This surface safety automation is an enabling technology for this concept, and especially important during the transition to automated surface ATC.

Surface Lighting System

Rather than requiring aircraft avionics equipage, one option proposed in our concept is the use of airport surface lights to visually deliver taxi and departure clearances to aircraft. While at least parts of this concept will be possible with existing airport lighting, augmented with a new lighting control system, other parts of this concept may require new lights or electronic signs. During later phases of the research, we will develop the requirements on surface lighting to implement various elements of our concept.

Since the lighting system is used to deliver taxi clearances, such as to stop short of crossing a runway or proceed across a runway, it is a safety critical system. Therefore, the lighting system (both the lights and the control system to turn the lights on and off) must be as reliable as other aircraft separation systems.

TRACON Routing

One aspect of our concept is to increase airport capacity by removing current procedural constraints on departure runway assignments. Although controllers make some exceptions, most departures are assigned a runway according to their route of flight. This procedurally provides airborne separation in terminal airspace. Arrivals and departures are similarly procedurally separated. In this concept, automation will assign an aircraft to whichever runway is most efficient (both for that flight as well as the airport as a whole). Before selecting the departure runway, the surface automation will work with terminal airspace (TRACON) control to plan a departure trajectory for the aircraft that will be free of conflicts with arrivals other departures, efficient (i.e., does not deviate substantially from the shortest path to the first fix or desired route to avoid other traffic), and robust to uncertainty in departure times and arrival trajectories (i.e., have sufficient options that are also likely to be conflict free).

However, once the aircraft has departed, the surface automation has no control over whether or not conflicts occur. Therefore, to be feasible, this concept requires that there also be TRACON automation that predicts conflicts and route aircraft to avoid potential conflicts with arrivals and departures. This TRACON automation could be accomplished either by automating ATC in the terminal airspace or by providing controllers with sufficient decision support, such as is generated by the Active Final Approach Spacing Tool (aFAST) and Expedite Departure Path (EDP).

Flight Deck Interface

Although not required by the proposed concept, cockpit displays could improve the flight crew's ability to follow the taxi clearances generated by the automation, further improving both capacity and safety. Aspects of the concept may make use of existing cockpit multi-function displays, which are available on a majority of large aircraft but are not available in most general aviation aircraft. For example, initially the concept will provide 2-D taxi paths for every vehicle. Additional benefits may be achieved by providing a 3-D surface trajectory (including time) for at least some of the aircraft. This function likely requires delivering additional information to the flight crew, beyond what is possible with airport lighting alone. To make use of available cockpit displays, the concept requires information, either via datalink from the aircraft or from the AOC, about what equipment the aircraft has. Aircraft that have sufficient equipment (e.g., multi-function cockpit displays) may be provided timing cues along their taxi path, for example, to increase the efficiency with which the aircraft cross active runways and the accuracy with which the aircraft reach their planned runway slots. Appropriately equipped aircraft, therefore, would be allowed to cross runways in tighter gaps than non-equipped aircraft.

New displays such as HUDs, are currently available in only a very small percentage of aircraft in the NAS, could enable additional surface automation capabilities. However, they are not expected to be available in 2020, since the FAA's technology roadmap through 2015 does not include substantial changes in cockpit displays.

Ramp Area Decision Support Tool

The proposed concept relies on the airlines sharing priorities and constraints for specific aircraft, and airlines using the information from our automation system in their planning process. This will require the airlines to develop some type of ramp tower DST to help ramp provide gate availability (for arrivals) and the when departures will be at the spot to our automation system. This DST will provide the ramp with a "view" into current and predicted future operations (such as anticipated delay for specific flights) including fast-time "what-if" analysis capability (i.e., to check sensitivity to order of pushback on delays for different flights).

Digital communication and display tools would need to be developed to provide our automation planning tools with access to information in a usable form regarding airline priorities and constraints for specific aircraft, and to provide this information to various airline and ATC staff integrated into their own information processing and displays tools.

Transition Plans

Often the term 'transition' and 'revolutionary' don't go together. However, this airport surface traffic control automation concept truly allows an evolutionary approach to a revolutionary concept. This section explains the transition from current operations in today's NAS to a future NAS in which our concept is operational. This will include a discussion of supporting/complementary developments in other NAS domains (e.g., terminal area) required to maximize the utility of our concept.

One of the most challenging aspects of transitioning to automated ATC is the problem of gaining the acceptance of the current controller workforce. One of the primary reasons for the controller union to resist our concept would be its impact on controller jobs. Our transition approach addresses this issue in two ways. First, we assume that the future NAS will include a significant increase in point-to-point travel by small aircraft, as envisioned in the SATS program. This change in the NAS will require improved ATC services at small airports that currently may not have any ATC services. Under our concept, rather than the FAA needing to hire and train a large number of new controllers, the FAA would reassign controllers no longer needed at major airports to these small airports. As long as the FAA continued to pay these controllers on the pay scale of the larger airports, many controllers would likely be willing to move to smaller cities and towns.

In addition, the FAA could offer a "golden handshake" early retirement package as incentive to reduce the controller workforce size. Transitioning to an automated ATC will initially be expensive for the FAA. However, the FAA will realize cost savings in the long run, relative to the NAS without our concept, because fewer controllers will be required at each airport. Of course, if the air traffic control system is privatized as some anticipate, then the resistance to such an automation system might be significantly less if it was proven to be a cost-effective means to improving overall performance of the transportation system.

The transition to automated ATC is also technically challenging. In particular, assurance of safety during such a substantial change to the NAS requires a conservative transition plan with adequate controller training. Our concept would initially be installed at a small airport that currently does not have an ATC tower. After sufficient proving at a small airport, our concept would be installed at a medium size airport for additional hardening prior to installation at the first large airport.

At the first major airport to which our concept is deployed, the FAA would build a second ATC tower that would operate simultaneously with the original tower. Initially, the original tower would remain responsible for controlling aircraft while the second tower, in which our concept had been installed, operates in "shadow mode", staffed with a completely separate set of "controllers" operating the automation. After our concept had been proven in the second tower, responsibility would be shifted to the automation and the original tower closed. Training procedures similar to those in place today will be included during this shadowing at each airport where the system is deployed.

This is not an inexpensive approach. However, automated ATC has the potential to reduce FAA operating costs in the future by reducing the workforce cost. Automating airport surface ATC is a first step toward automating terminal and en route ATC, which would provide additional capacity benefits and further savings in the cost of operating the ATC system.

To make the transition to automated airspace more gradual, various parts of our concept can be phased into operation.

HUMAN PERFORMANCE ISSUES

Our concept proposes a fundamental shift in surface control paradigm from a distributed, human-in-control system to an automation-enhanced, human-guided system. In doing so, our concept potentially reduces the required staffing in a given ATC tower. Under our concept, it is possible for a smaller number of human "collaborators" to work hand-in-hand with the automation and effectively manage the control of surface traffic. As such, new procedures (and staffing with new roles and responsibilities) would need to be developed to make use of these technological capabilities and to provide an equitable operating environment from the perspective of NAS users. Research in subsequent phases of this project will further define the nature of the interaction between the human collaborators and the algorithms planning and controlling surface movements. A particular concern is to make the working of the underlying algorithms as transparent as possible to the human operator. For example, the human operator should be able to quickly determine and modify the "weightings" applied to different performance objectives (e.g., arrival/departure tradeoff, satisfy user priorities) and see, at a glance, any "stumbling blocks" or difficulty in coming up with solutions. Our objective is to keep the human active in the problem-solving loop, but at a level dealing with flows, delays, equity, and overall surface effectiveness rather than the details of moving specific aircraft around the tarmac.

The other major class of human performance issues is related to the Flight Deck and the ability of pilots to actually follow taxi routes with time constraints as well as time-varying taxi-routes. A considerable factor influencing human ability to follow such surface trajectories is the way in which the trajectory is communicated to the Flight Deck. Significant human factors research is required to determine the most appropriate communication medium. There is an inherent tradeoff between providing sufficient situational awareness to insure safety and confidence and simply providing too much information. For example, visual display techniques would require investigation as to the "distance from the user" beyond which changes would be acceptable; changes too "close" to the current aircraft position might be distracting and too difficult to follow. Thus, extensive human-in-the-loop testing will be required to evaluate different communication mechanisms in terms of their effectiveness and usability. In the limit, time-varying surface trajectories may prove too difficult for humans to follow.

User Interfaces

In this section, we provide an initial description of the information to be communicated through user interfaces to be designed in subsequent phases of research. The goal of these interfaces is to maximize human usability of the system and the extent to which the human and automation complement one another's skill sets. Our concept includes two critical interfaces, namely that of the human Surface Collaborator in the ATC tower and that of the Pilot in the Flight Deck. In addition, we aim to design interfaces that both allow for flexibility in the way an operator chooses to approach a given task, and enable the human operator to cope with unanticipated events. The former is accomplished by providing multiple methods of interaction with the system, while the latter is accomplished by providing a continuous and accurate model of the airport surface system and the parameters that affect it.

Surface Collaborator Interface

The objective of the surface collaborator interface is two-fold: to provide a top-down view of the overall surface flow and to provide a "window" into the inner workings of the planning algorithms that are generating motion clearances. Data presented to the Surface Collaborator include:

- Indication of areas with excessive queuing
- Performance indicators related to achieved rates on each runway (vs. planned)
- Display of pertinent surface/terminal Area weather conditions and a depiction of their impact on different areas of performance
- Display of the current performance objectives and constraints driving the planning algorithms
- Indication of the degree to which various performance objectives are being satisfied

Potential inputs expected from the user include:

- Modifications to Constraints
- NAS User Preferences
- Runway Configuration Changes
- Impact of Weather
- Need for De-Icing
- Runway Closures

Pilot Interface

The other critical interface is the pilot interface. Depending on the mechanism for communicating clearances, this interface might consist of one or more of: automated radio clearances, datalinked clearances, "out-the-window" observation of surface lighting, or cockpit-based electronic moving map and heads-up-display. Obviously the complexity of the clearance and the situational awareness afforded by the various mechanisms varies considerably, as indicated notionally in **Figure 12**. At one end of the spectrum, surface lighting guidance can at least partially communicate time-based clearances (e.g., green means go, red means stop) but perhaps conveys the least situational awareness. Essentially, it says "follow the lights and you're safe. Don't, and you're not", without providing the context of own-ship clearances relative to those of other aircraft. Potentially, however, a combination of surface light guidance and automated voice clearances (to provide information regarding other traffic) or datalink (providing countdown to RTA at next intersection or time left in a hold) could be effective. On the other end of the spectrum, Augmented Reality (AR) displays in conjunction with an Electronic Moving Map (EMM) likely offer the highest degree of both situational awareness and the ability to display "followable" time-varying clearances (extending the displays and symbology used by T-NASA[1]). Of course, these delivery mechanisms (particularly the HUD/AR displays) imply considerable cost in terms of equipage. The cost of EMM alone, however, is potentially much lower since it is primarily a software change on many modern glass cockpit aircraft.

Regardless of the physical realization of the interface, the minimal data that must be presented to the user include:

- Taxi route showing intersections
- Holding Locations
- Clearance to Enter Runway
- Clearance to Cross Active Runway
- Clearance to Land

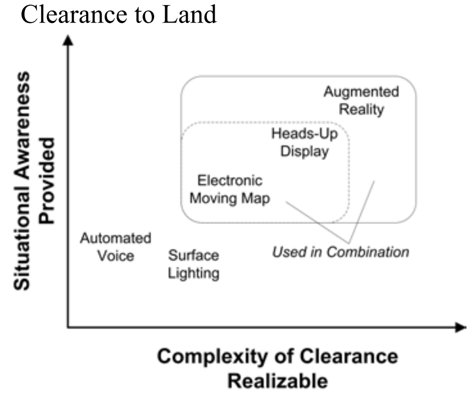

Figure 12: Comparison of Different Clearance Delivery Mechanisms

Additional information to be communicated (as enabled by the form of interface) includes:

- Required Time of Arrival at Intersections
- Countdown to Release from Current Hold
- Command guidance cues (speed, steering, braking, etc.)
- Cues Regarding Status and Motion of Relevant Other Traffic
- Indication of Order in Sequence and Spacing Requirements
- Surface Conditions
- Weather-Related Effects (e.g., wind shear)
- Wake Turbulence Avoidance Regions

As part of our continued research into ATC automation, we intend to develop actual user interface concepts, showing how this information should be accessed, presented and manipulated as part of a candidate situation awareness and decision aid system. Subsequently, we will evaluate the effectiveness of the proposed interface, in terms of capacity and human factors metrics, through both self-evaluation usability studies and VAMS human-in-the-loop simulation.

WHERE DO THE BENEFITS COME FROM?

This section summarizes the mechanisms through which our concept will achieve benefits and defines the metrics of goodness to be used in ongoing efforts to evaluate the extent of benefits actually achieved. In addition, we provide an initial discussion of the costs related to implementation of our concept.

Traffic Flow Management

Improved EDCT/departure constraint compliance will lead to increased predictability in terms of utilization of downstream resources. This will reduce the "waste" of slots at the airport and en route.

By planning continuous changes to the AAR or departure rate to match time-varying demands, automation-enhanced surface operations will reduce the potential for surface gridlock (e.g., TFM system constrains departures and continues to send arrivals).

The availability of improved surface acceptance and departure flow rates (developed explicitly considering uncertainty in traffic demand and weather effects) communicated to NAS TFM through our long-term surface capacity planning will enable improved matching of capacity to demand, reducing the extent of over and under-control of flights (e.g., unnecessary ground holding or the need for unexpected airborne holding, respectively) exhibited in the current NAS.

By employing new technologies, tools and procedures to mitigate weather-related capacity effects, we can reduce the need to adopt less efficient airport configurations and maintain the levels of precision and efficiency afforded by our core ideas.

Controlling Arrivals and Departures

Timing operations on dependent (e.g., crossing) runways requires high workload and coordination between controllers. Consequently, all possible runways are frequently not used. The ability to coordinate arrival times and accurately predict intersection crossing time as well as arrival and departure trajectories allows automation to safely and efficiently use dependent runways – enabling more arrivals to reach the surface in a given period of time.

Improved planning and coordination of runway system use between different types of operations will reduce the number of flights which have to wait for a "gap" prior to crossing an active runway, thus reducing taxi delays. Providing the surface with the opportunity to change arrival runway assignments can lead to reduced arrival taxi times and improved surface operations (i.e., reduced departure delays).

An additional, related mechanism concerns current procedures regarding runway use. Currently, the runway can only be occupied by either a single arriving or single departing aircraft. Automation will not initially do anything about this, but could be a first step toward eventually relaxing other surface procedural constraints such as the single-occupancy requirement. If so, then our concept could support complementary terminal airspace concepts enabling reduced separations and/or formation landing.

Working in concert with the Terminal Area, our concept will relax the "fixed" departure plan mapping from departure fix to runway. This will enable improved, more flexible sequencing that will reduce separations between operations and thus increase runway throughput. An additional impact will be an increase in the ability to satisfy user preferences when constructing departure sequences.

Further, automation will better predict whether there is enough time to safely depart before the next arrival or during a future gap and will inform the pilot to be ready, thus eliminating departure gaps on a mixed use runway which are missed today due to controllers' inability to accurately predict inter-arrival times.

Explicit incorporation of environmental considerations into the performance functions that the planning algorithms use to assess candidate solutions will result in the ability to reduce and/or shape community noise exposure and reduce harmful emissions.

Enhanced vision and situation awareness will provide pilots increased confidence in the separation in space and time between own-ship and other traffic. This will increase the safety of operations by providing pilots visual indications of runway status and the intent of other traffic.

Runway Crossings and Taxi Operations:

Improved planning and coordination of arrival, departure, and taxiing aircraft runway use will enable simultaneous runway crossing clearances – often without the need to wait and stop for a "gap" in traffic. This will result in reduced taxi delays for both arrivals and departures at airport where the geometry requires crossing of active runways to get to and from the gate. Unlike human controllers, our automation concept will issue multiple clearances simultaneously, reducing taxi delay in crossing runways and increasing arrival/departure capacity of the runways. In addition, the automation can inform pilots to "get ready to cross" to expedite crossing.

Enhanced vision and situation awareness will enable increased pilot confidence in knowing their designated taxi routes, time constraints, and relationship to other traffic. This will enable faster taxi speeds in all weather conditions.

SUMMARY AND FUTURE WORK

In this paper we have identified the key focus areas which comprise a proposed future concept for automated airport surface traffic control – a first step towards increased automation throughout the NAS. We have discussed the key functionality as well as many of the transitional issues related to implementation of the concept.

The current concept is the result of the first year of concept development under NASA Ames' Virtual Airspace Modeling and Simulation (VAMS) project. This effort is currently in its second year in which the emphasis is on an initial assessment of the concept from the perspective of establishing "local" benefits (e.g., within a given terminal area or at a particular small set of airports). We anticipate that the proposed concept will provide significant benefits with respect to airport capacity, efficiency, predictability, flexibility (satisfying user preferences), fairness, and environmental impacts. Key aspects of this assessment effort will include simulation and analysis of multi-domain scenarios emphasizing the coupling and interaction of surface traffic with flows aloft. Ultimately, the results of these local concept assessments will drive experiments using the Advanced Concept Evaluation System (ACES) – a fast-time, NAS-wide simulation capability currently under development at NASA Ames[17]. ACES will allow analysis of NAS-wide effects of concept implementation.

In concert with concept assessment activities we are pursuing several critical thrusts including technical feasibility, representation and reasoning with respect to uncertainty, and investigation of the changing roles and responsibilities brought about by the introduction of significant levels of automation to the airport planning and control process.

ACKNOWLEDGEMENTS

This research was supported by NASA Ames Research Center through the Virtual Airspace Modeling and Simulation project, contract NAS2-02071. The authors would like to thank the VAMS project office for valuable feedback in formulating the ideas presented in this paper. In particular, we would like to recognize the efforts of Stephen Atkins in shaping this concept.

REFERENCES

[1] Foyle, D., Andre, A., McCann, R., Wenzel, E., D. Begault, and Battiste, V., "Taxiway Navigation and Situation Awareness (T-NASA) System: Problem, design philosophy and description of an integrated display suite for low-visibility airport surface operations", *SAE Transactions: Journal of Aerospace*, Vol. 105, pp. 1411-1418, 1996.

[2] Theunissen, E., Rademaker, R.M., Jinkins, R.D., and M.U. de Haag, "Design and Evaluation of Taxi Navigation Displays", *Proceedings of the 21st Digital Avionics System Conference*, Irvine, CA, 27-31 October, 2002.

[3] Cheng, V., Sharma, V., and Foyle, D., "A Study of Aircraft Taxi Performance for Enhancing Surface Traffic Control", *IEEE Transactions on Intelligent Transportation Systems*, Vol. 2, No. 2, June 2001.

[4] Cheng, V., "Collaborative Automation Systems for Enhancing Airport Surface Traffic Efficiency and Safety", *Proceedings of the 21st Digital Avionics System Conference*, Irvine, CA, 27-31 October, 2002.

[5] Anagnostakis, I., Clarke, J–P., Böhme, D., and U. Völckers, "Runway Operations Planning and Control: Sequencing and Scheduling", *AIAA Journal of Aircraft*, Nov-Dec 2001, Vol. 38, No.6.

[6] Anagnostakis, I., Clarke, J–P., "Runway Operations Planning: A Two-Stage Heuristic Algorithm", *AIAA's 2nd Annual Aircraft Technology, Integration, and Operations Forum*, Los Angeles, CA, 1-3 October, 2002.

[7] Cooper, W.W., Cormier, R.H., Foster, J.G., Mills, M.J., and Mohleji, S.C., "Use of the Departure Enhanced Planning and Runway/Taxiway Assignment System (DEPARTS) for Optimal Departure Scheduling at Busy Airports", *Proceedings of the 21st Digital Avionics System Conference*, Irvine, CA, 27-31 October, 2002.

[8] Garcia, J., Berlanga, J.M., Besada, J.A., and J. Casar, "Planning Techniques for Airport Ground Operations", *Proceedings of the 21st Digital Avionics System Conference*, Irvine, CA, 27-31 October, 2002.

[9] Atkins, S., and Brinton, C., "Concept Description and Development Plan for the Surface Management System," *Journal of Air Traffic Control*, 2002.

[10] Atkins, S.C., and W.D. Hall, "A Case for Integrating the CTAS Traffic Management Advisor and the Surface Management System", *AIAA Guidance, Navigation, and Control Conf. and Exhibit*, Denver, CO, August 14-17, 2000.

[11] Federal Aviation Administration, *Operational Evolution Plan*, available on-line at: http://www1.faa.gov/programs/oep/.

[12] Federal Aviation Administration, *National Airspace System Architecture Version 4.0*, ASD-110, Washington, DC, Jan., 1999.

[13] RTCA, Inc., *National Airspace System Concept of Operations*, Dec.13, 2000.

[14] Perry, R.B., Hinton, D.A., and R.A. Stuever, "NASA Wake Vortex Research for Aircraft Spacing", *35th Aerospace Sciences Mtg. & Exhibit*, Reno, NV, January 9-10, 1997.

[15] Andrews, J., "The Shortfall in Airport Capacity and Opportunities for Mitigation", *Transportation Analysis: New Concepts and Methods in Air Traffic Management*, pp.75-94, Bianco, Dell'Olmo, and Odoni, Eds., Springer-Verlag Berlin, Heidelberg, 2001.

[16] Capozzi, B., Augustine, S., Thompson, T., and J.E. Robinson III, "An Initial Assessment of Benefits for Noise-Aware Decision Support Tools", *AIAA's 2nd Annual Aircraft Technology, Integration, and Operations Forum*, Los Angeles, CA, 1-3 October, 2002.

[17] Sweet, D.N., Manikonda, V., Aronson, J.S., Roth, K., and M. Blake, "Fast-Time Simulation System for Analysis of Advanced Air Transportation Concepts", *Proc. AIAA Modeling and Simulation Technologies Conference and Exhibit*, Monterey, California, 5-8 August 2002.