CORE
Metadata, citation and similar papers at core.ac.uk
ded by NASA Technical Reports Server

RESOLVING OFF-NOMINAL SITUATIONS IN SCHEDULE-BASED TERMINAL AREA OPERATIONS: RESULTS FROM A HUMAN-IN-THE-LOOP SIMULATION

Joey Mercer, Todd Callantine, Lynne Martin
San Jose State University / NASA Ames Research Center
joey.mercer@nasa.gov, todd.callantine@nasa.gov, lynne.martin@nasa.gov

Keywords: *terminal-area air traffic control, off-nominal recovery, decision support tools*

Abstract

A recent human-in-the-loop simulation in the Airspace Operations Laboratory (AOL) at NASA's Ames Research Center investigated the robustness of Controller-Managed Spacing (CMS) operations. CMS refers to AOL-developed controller tools and procedures for enabling arrivals to conduct efficient Optimized Profile Descents with sustained high throughput. The simulation provided a rich data set for examining how a traffic management supervisor and terminal-area controller participants used the CMS tools and coordinated to respond to off-nominal events. This paper proposes quantitative measures for characterizing the participants' responses. Case studies of go-around events, replicated during the simulation, provide insights into the strategies employed and the role the CMS tools played in supporting them.

1 Introduction

An important research objective of the NASA Airspace Systems Program Super-Density Operations (SDO) research focus area is to safely sustain high runway throughput while minimizing environment impacts through fuel-efficient operations. Scheduling arrivals to fly Optimized Profile Descents (OPDs) along Area Navigation (RNAV) routes in the terminal area is a central element of the SDO concept of operations [1]. Maintaining RNAV OPDs requires managing arrival flows using primarily speed control; tools to aid controllers in this task have been developed in the Airspace Operations Laboratory (AOL) at NASA Ames Research Center [2] as part of a research effort referred to as Controller-Managed Spacing (CMS).

CMS research first focused on controller decision-support tool (DST) development and assessment under nominal operations [3, 4]. An SDO-sponsored AOL simulation conducted in the spring of 2011, named CMS4, broadened the scope of the initial efforts by investigating the robustness of CMS operations to disturbances that may arise due to off-nominal events [5]. In addition, CMS4 afforded the opportunity to examine the potential role of a Traffic Management Supervisor responsible for coordinating with terminal-area controllers to adjust the arrival schedule to support recovery from significant disturbances [6].

This paper examines case studies drawn from replications of scripted go-around events simulated in CMS4. During these events, the supervisor had to adjust arrival schedules to reinsert go-around aircraft into the arrival flows and controllers had to perform control actions to reestablish aircraft on schedule in order to restore nominal operations. In order to better characterize the effects of applying particular recovery strategies, and how well the CMS DSTs supported the efforts in a particular operational context, quantitative measures of the schedule adjustments, the resulting schedule, and the control actions required to achieve it are presented.

First the paper presents some background on schedule-based OPD arrival operations and the CMS DSTs. Second, it introduces the CMS4 operational environment and uses one of the case studies to illustrate how the DSTs might be used to manage a go-around,

and discusses ways in which off-nominal events can impact the DSTs themselves. Third, go-around case-studies are described in detail, together with their quantitative measures. And lastly, the findings and the value of such analyses for formulating effective DSTs and procedures for restoring nominal operations are discussed.

2 Background

When traffic demand is low, managing efficient descents is straightforward because controllers seldom need to intervene. However, maintaining OPDs during sustained periods of heavy demand becomes difficult since current-day control techniques, including altitude level-offs and heading vectors, disrupt OPDs. Speed control is also difficult to apply [7]. DSTs are needed to help controllers primarily use speed control to manage fuel-efficient OPDs for busy arrival flows [8].

DSTs for merging and spacing aircraft in the terminal area have included 'ghosting' displays and/or clearance advisories; Callantine [9], and Kupfer [3] review previous research in this area. The CMS DSTs differ in that they leverage RNAV OPDs for trajectory predictions, and in turn, the arrival schedule, to provide controllers with both temporal and spatial information to support the merging of OPD arrival flows while maintaining high throughput.

2.1 Nominal CMS Operations and DSTs

The schedule-based SDO arrival management concept underpinning CMS assumes all aircraft are Flight Management System- (FMS-) equipped so as to enable Vertical Navigation (VNAV) descents along published RNAV OPDs. En-route controllers are assumed to condition arrival flows so that aircraft enter terminal radar approach control (TRACON) airspace with schedule errors ranging from 60 s early to 30 s late approximately. These errors are small enough that they can be corrected with speed adjustments alone. TRACON Feeder controllers then use schedule information and other DSTs to issue speeds as required for adjusting aircraft toward their Scheduled Time-

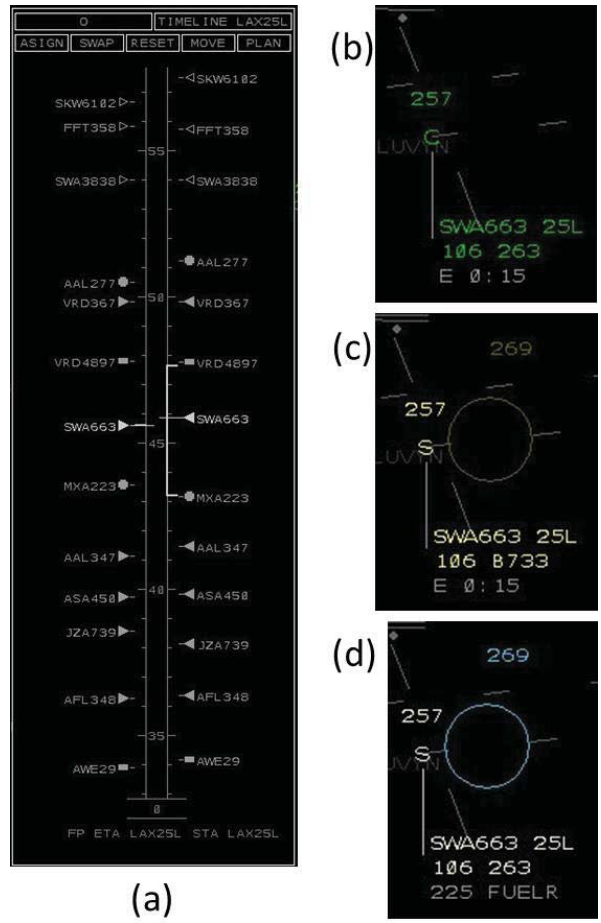

Figure 1. CMS tools.

of-Arrival (STA), while still keeping them on their assigned RNAV OPD. Final controllers issue speeds to remove any residual schedule errors. They also ensure that aircraft are safely merged, established on the final approach, and delivered to the tower such that proper spacing will be achieved at the runway threshold.

The CMS DSTs are designed to provide TRACON controllers with temporal as well as spatial awareness of each aircraft's progress relative to its STA, and also suggest speeds that controllers can issue to correct schedule errors [3] (Figure 1). Schedule timelines show Estimated Times-of-Arrival (ETAs) on the left and STAs on the right for each aircraft that crosses a scheduling point, such as a runway or a merge point (Figure 1a). The scheduling algorithm computes an aircraft's ETA using trajectory predictions based upon its current position and assigned RNAV OPD; it

determines the STAs by applying the required time-spacing between each aircraft in the arrival sequence, as determined by the ETAs. The scheduling algorithm locks the STAs in place when aircraft cross a specified 'freeze horizon,' in order to provide a stable control target. The CMS DSTs also include an early/late indication displayed in the third line of the data block to provide controllers with schedule-conformance information for a given aircraft without diverting their attention from it (Figure 1b).

CMS slot markers convert temporal schedule information into a spatial target controllers can work toward; they display where a given aircraft would be if it were flying its assigned RNAV OPD speed/altitude profile through the forecast wind field and arrived on time at the scheduling point (Figure 1c). Dwelling on an aircraft's data block brightens it and highlights its slot marker and timeline entries (Figure 1a/1d). The ground automation computes each aircraft's slot marker to follow the RNAV OPD assigned to that aircraft in the ground system.

The CMS DSTs also include speed advisories; those used in CMS4 were formulated as a speed that, if flown until decelerating to meet a published speed restriction at a downstream waypoint, then rejoining the nominal speed profile, results in a predicted on-time arrival for the aircraft. When available, the advised speed and fix name are presented in the third line of the data block, replacing the early/late indication (Figure 1d).

Successful CMS evaluations [3, 4] paved the way for testing the robustness of the concept and DSTs during off-nominal events.

3 CMS4 Simulation of Off-Nominal Events

CMS4 is a real-time human-in-the-loop simulation of TRACON operations conducted in the AOL during spring 2011, described in detail in [5]. FMS-equipped, west-flow arrival traffic transited Southern California TRACON (SCT) airspace on RNAV OPDs to runways 24R and 25L at Los Angeles International airport (LAX) (Figure 2). Participants staffed three Feeder sectors (201, 204, and 205) and two Final sectors (202 and 203); in addition, one served as the traffic management supervisor, and a

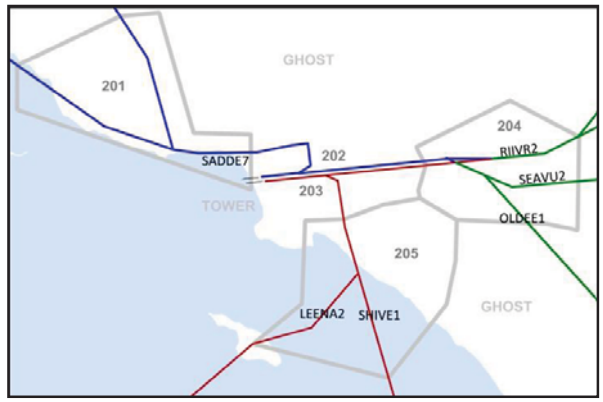

Figure 2. Test sectors and RNAV OPDs to LAX runways 25L and 24R.

confederate staffed the Tower position. Each simulation trial included two off-nominal events, one involving an aircraft assigned to runway 25L and one involving a 24R aircraft (see [5] for a complete description). The off-nominal events were expected to disrupt the arrival flows enough to require schedule adjustments and, as a result, delays that would be too large for controllers to absorb with speed control alone. Path options in the form of named RNAV arrival routes were therefore defined in accordance with controller feedback to help them absorb larger delays (Figure 3). Among the path options were 'long' and 'short' go-around routes designed to provide flexibility in reinserting go-arounds into the arrival flows (green routes that begin at the IGUPE and FUMBL waypoints in Figure 3).

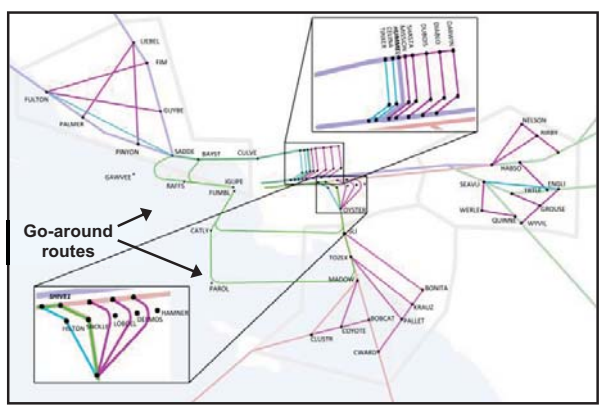

Figure 3. RNAV path options and go-around routes.

Two types of go-around events were scripted to occur in CMS4: pilot-initiated go-arounds and tower-initiated go-arounds. Pilot-

initiated go-arounds were triggered when the designated aircraft was below 4000 ft MSL approximately 15 nmi from the runway. The controlling pseudo-pilot would contact the final controller, declare a landing gear malfunction, and request a go-around to fix the issue. The final controller would cancel the approach clearance, and then issue an altitude to maintain (e.g., "USA395, roger, cancel approach clearance, climb and maintain 7000, remain inbound to runway 24R, expect further instructions shortly."). The final controller would then announce "Go-around!" to alert the supervisor that a go-around was in progress, and notify the tower confederate that the aircraft would be going around. Tower-initiated go-arounds were triggered when the aircraft ahead of the designated aircraft landed and was 'slow to clear the runway.' The tower confederate would contact the final controller (who assumed the role of a departure controller in these situations) via ground-ground voice communication, identify the go-around, and receive instructions about which go-around procedure to assign the aircraft and which feeder controller the aircraft should contact. Similarly, the final controller would announce "Go-around!" to alert the supervisor and other controllers. For both types of go-arounds, the process of recovering from the off-nominal then began.

The tower-initiated go-around event for runway 25L assumed the go-around aircraft would be worked primarily by the feeder controller at sector 205, before returning to sector 203 again, while the intended effect of the pilot-initiated go-around for runway 24R was to have the feeder controller at sector 201 work the go-around aircraft before returning it to sector 202 again. Depending on how the schedule was adjusted in response to the go-around, other aircraft coming through the feeder sectors could be impacted as well.

3.1 Supervisor Rescheduling Support

The controllers and supervisor were expected to coordinate as necessary to resolve situations arising due to off-nominal events while attempting, to the extent possible, to continue OPD operations and sustain high throughput.

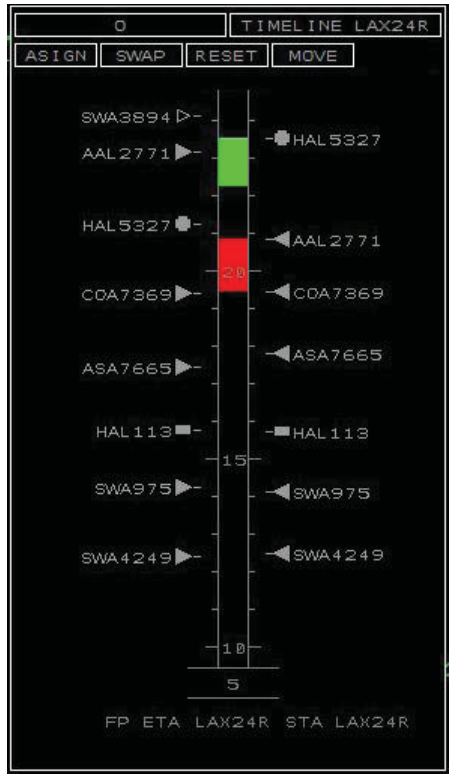

Figure 4. Supervisor timeline rescheduling buttons and bars indicating schedule gaps (green) and insufficient spacing (red).

Controllers were asked to manipulate the traffic using only speed clearances and the pre-defined path options if possible, although vectoring was still a valid option. The supervisor was asked to manage the schedule and attempt to maintain high runway utilization, working the schedule only as far out as the first aircraft that was outside the schedule freeze-horizon (approximate 80 nmi from the runways). The controllers and the supervisor were free to coordinate to formulate off-nominal recovery plans and use the tools to achieve them as they saw fit.

The supervisor staffed a MACS workstation [2] with a traffic display and CMS runway-schedule timelines. The supervisor could adjust the viewable range of the traffic display to visualize aircraft relevant to rescheduling problems, and manipulate the schedule timelines in several ways:

- Re-assign an aircraft's STA.
- Swap STAs for two aircraft.

- Move a specified 'block' of STAs by a specified time.
- Reschedule ('reset') a specified block of aircraft.
- Assign an aircraft to a different runway schedule.

The supervisor performed all of these operations by entering commands in the shortcut window on his workstation. Commands could be composed using a combination of text entries and mouse-selections on the timelines. To aid the supervisor in performing schedule assessments, the timelines were also modified with green and red bars that indicated gaps (or 'slack') or insufficient spacing in a schedule, respectively (Figure 4).

3.2 CMS DSTs under Off-Nominal Conditions

The process of re-inserting the go-around aircraft in the arrival flow required a coordinated set of actions from the controllers and the supervisor. An example of how the controller DSTs could impact the supervisor's DSTs, and vice versa is presented in the following case study T4, a tower-initiated go-around (Figure 5). Tower-initiated go-arounds occurred when the aircraft was close to runway, with the aircraft's ETA and STA near the bottom of the timeline (not shown). After the aircraft passed the runway, its original runway STA would drop off the timeline; without a valid STA, the ground automation is unable to compute speed advisories or a slot marker for the aircraft (Figure 5a). Once the controller assigns a go-around procedure to the aircraft, the scheduling algorithm is able generate a new ETA for the aircraft and display it on the runway timeline; Figure 5b shows the new ETA reflecting the assigned procedure for SWA5488 (go-around) appearing between FFT1134 and AAL5321, with a new STA yet to be assigned. The supervisor could then adjust the timeline to make room for the go-around aircraft. Here the supervisor identifies two areas of nearby slack in the schedule and, using his scheduling tools, creates a new slot for the go-around aircraft. In this case, the supervisor first advanced FFT1134's STA, taking advantage of the slack in front of it (Figure 5c). This effectively combined the slack that was in front of FFT1134 with the additional slack that was behind it, creating an obvious gap in the schedule in which to assign the new STA for SWA5488 (Figure 5d). This produced the final plan (Figure 5e). With a valid STA for the

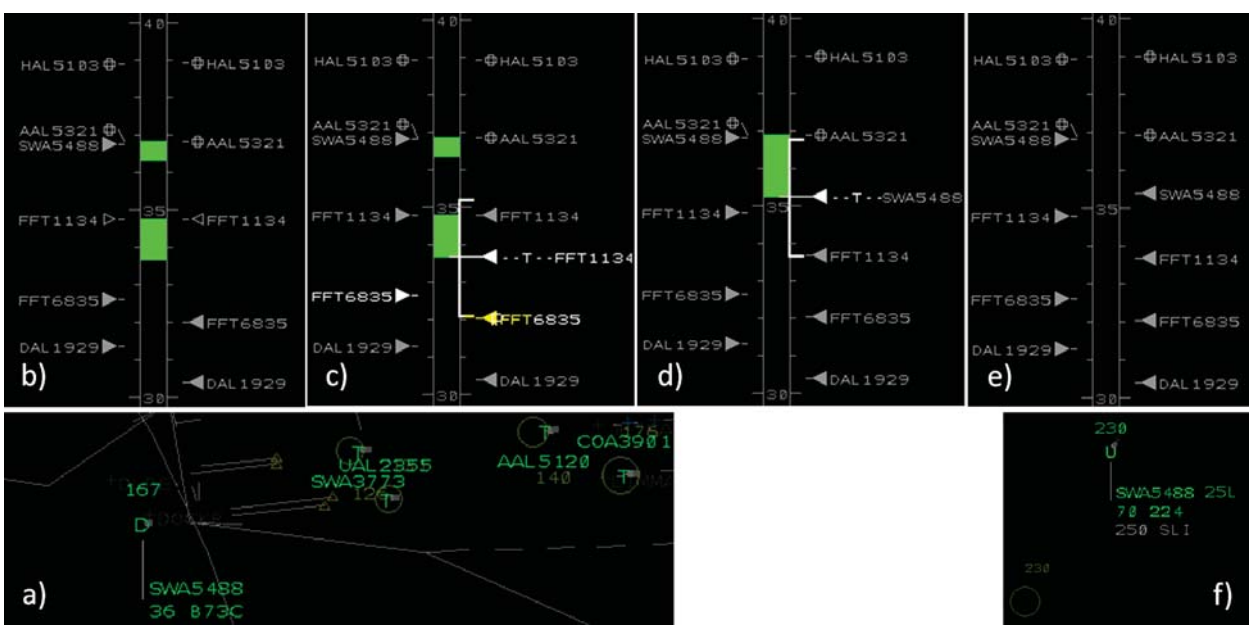

Figure 5. Sequence of events from case study T4 involved with reinserting a go-around aircraft into the arrival schedule for runway 25L (clockwise from bottom left).

aircraft, the automation again computes speed advisories and a slot marker for the go-around aircraft, which the controller can then use to reinsert the aircraft into the arrival flow (Figure 5f).

This example highlights how off-nominal events, and the manner in which they are addressed, impact the DSTs. First, unless an aircraft has a route assigned, and is not vectored too far off it, the ground automation cannot perform the trajectory predictions required to produce a reasonable ETA for the aircraft. Without an ETA, the scheduling algorithm cannot generate an STA for the aircraft, and it becomes difficult to gauge where an STA for the aircraft could be inserted into the schedule. Second, because the trajectory predictions that underlie the CMS slot markers and speed advisories use STAs as a reference, these DSTs require STAs for their associated aircraft, and are most useful for control when the STA is one the aircraft can reasonably achieve along its assigned RNAV route. Additionally, STA changes that occur when the schedule is adjusted cause immediate changes to the slot-marker positions of the affected aircraft on the controller displays. Thus, timely route assignments and judicious scheduling actions bolster the usability and usefulness of the DSTs for restoring nominal operations. The following analyses support these effects.

4 Results

The CMS4 simulation provided a rich set of data on the robustness of the CMS concept and DSTs. This paper extends analyses presented in [5] and [6] by quantifying scheduling and control actions for a particular simulation trial that was replicated four times (with randomized aircraft call signs) during CMS4. The trial included a tower-initiated go-around (denoted 'T') for runway 25L and a pilot-initiated go-around (denoted 'P') for runway 24R. Case studies T1 and P1 are drawn from the first replication, T2 and P2 from the second, etc. Contextual elements behind quantitative measures of the rescheduling and control actions are presented for each.

4.1 Off-Nominal Recovery Metrics

Table 1 presents measures related to recovering from the tower-initiated go-around ('T') events and the pilot-initiated go-around ('P') events. The first three rows of Table 1 provide measures of how the supervisor adjusted the schedules to accommodate the affected aircraft. 'Schedule adjustment time' refers to the difference in time between the supervisor's first and last schedule manipulation when responding to a given off-nominal event. The number of individual schedule-adjustment actions and the number of aircraft STAs affected by the adjustments are shown in rows two and three, followed by the cumulative amount of delay that was added to the schedule. Negative values in this row indicate that, overall, resulting STAs were earlier than STAs before the adjustments; thus, the schedule was advanced.

The next six rows of metrics in Table 1 pertain to control actions taken by controllers to bring aircraft into conformance with the adjusted schedule, starting with the number of unique aircraft that received clearances during the off-nominal recovery period. Subsequent rows tally the types of clearances issued. The numbers of direct-to clearances, (heading) vectors, and 'open' speeds reflect the extent to which controllers used current-day control techniques. Speed advisories and pre-defined paths reflect the degree to which controllers used the CMS speed advisories and assigned aircraft to RNAV routes they could specify to the ground automation. The last row gives the amount of time that elapsed during the recovery process, from the time of the first action related to the off-nominal event to the last. These metrics are used to describe similarities and differences between different replications of the same off-nominal events as they played out during CMS4, in terms of strategies the supervisor applied, salient contextual factors, and the role of the CMS DSTs.

4.2 Off-Nominal Recovery Strategies, Contextual Factors, and DST Implications

During the recovery, the supervisor often used the updated ETA for an aircraft that the ground

		T1	T2	T3	T4	P1	P2	P3	P4
scheduling actions	schedule adjustment time (seconds)	138	217	1202	141	411	54	199	75
	schedule adjustments made	3	11	19	5	3	2	6	2
	STAs affected	1	13	14	5	4	7	9	7
	overall delay added to the schedule (seconds)	30	-166	-273	-47	-154	114	369	387
control actions	aircraft receiving clearances during recovery	2	5	15	2	1	4	8	6
	pre-defined paths issued	0	0	4	0	0	3	5	5
	direct-to's issued	1	3	5	2	1	0	1	1
	vectors issued	1	9	8	2	0	0	0	0
	speed advisories issued	0	2	5	0	1	2	8	5
	open speeds issued	7	12	27	0	0	1	4	4
	recovery time (seconds)	834	828	1617	437	287	603	1073	1021

Table 1. Quantitative metrics for the tower-initiated go-arounds to runway 25L (T1 – T4) and the pilot-initiated go-arounds to runway 24R (P1 – P4).

automation computed when the aircraft was assigned a particular go-around procedure to identify which areas of slack, if any, could be rearranged into a new slot (or partial slot) close to the updated ETA. This strategy was shown to be successful in several case studies; T1, T4, P1, and P2 all had relatively few schedule adjustments, necessitating clearances for small numbers of aircraft during the recovery period. In T1 and T4, the supervisor leveraged slack in the schedule to create a new slot by advancing one aircraft's STA; in P1 and P2, the same adjustment yielded only a partial slot, so the supervisor additionally reset the STAs of the next few aircraft, delaying them as necessary to make room for the go-around aircraft. The effectiveness of this approach in these four cases is reflected in the small absolute values of overall delay added to the schedule shown in Table 1. The supervisor's simple description of his strategy after completing P2 suggests he recognized this: "I assigned a slot, rippled the list, and that was that. I would not have done anything differently."

The supervisor attempted to apply a similar strategy in case studies P3 and P4, which had comparable numbers of schedule adjustments and affected STAs, but it was not as effective in these cases because it added more overall delay to the schedule. The schedule in P3 initially had the least amount of slack, thus adding more delay was unavoidable, while in P4, the supervisor chose to uniformly shift the STAs behind the go-around aircraft's slot back by 15 seconds, impacting the affected STAs more. Table 1 illustrates how the larger impact to the schedules in P3 and P4 required the controllers to issue more clearances (including pre-defined path clearances) to bring aircraft into conformance with the new schedules, which in turn lengthened the required recovery time.

The relationship between schedule impact and recovery time can be characterized by additional explanations, however, as illustrated by the quantitative data in Table 1 for case studies T1, T2, and P1—emphasizing the importance of context in interpreting quantitative metrics. The recovery time in T1 is slightly longer than for T2, but is not due to large schedule impact; rather, it stemmed from a lack of DST support for the sector 205 controller. During T1, the tower confederate was unable to amend the go-around aircraft's route in the ground automation, so the timeline display on sector 205's scope did not have an ETA or STA for the go-around aircraft, and there was no slot marker available to aid the controller. This lack of information about where the go-around aircraft would merge back into the flow of traffic, combined with the geometry and distance between the two merging aircraft,

made the controller's task very difficult. Only as the go-around aircraft neared the location at which the go-around procedure rejoined the nominal RNAV OPD did the controller realize that the go-around aircraft was actually behind where it needed to be, and the spacing to the trailing aircraft was not sufficient to allow the merge. The controller quickly issued a heading vector to ensure separation, followed by a direct-to clearance to shorten the go-around aircraft's route. The controller then had to issue multiple speeds to maintain the required separation.

The recovery strategy applied in case study T2 impacted the schedule similarly to P1, but the recovery time was much longer in T2. The larger numbers of schedule adjustments and affected STAs in T2 is not enough to explain the longer recovery time, which was actually traced to the particular strategy the supervisor attempted to apply. In P1, the go-around aircraft was assigned the 'long' procedure, and the supervisor leveraged the existing slack in the schedule to make that work in a straightforward way, as evidenced by the clearance data for P1 in Table 1. However, in T2 the supervisor followed a more complicated strategy. The supervisor first identified a potential gap in the schedule that was three minutes ahead of the updated ETA that ground automation computed upon assignment of the go-around procedure. Next, the supervisor asked the final controller to maintain control of the go-around aircraft and vector it into the gap he was orchestrating. The supervisor then used existing slack in the schedule to create the earlier slot for the go-around aircraft. Consequently, the sector 203 controller had to issue multiple instructions to the go-around aircraft over the course of several minutes to implement the supervisor's plan. Thus, in T2, the supervisor's apparent desire to resolve the off-nominal quickly added instability to the situation; the resulting plan required more clearances and was ultimately more difficult for the controllers to execute. Upon reflection, the value of those three minutes seemed marginal to the supervisor, who commented, "Thinking back on it, I should have agreed with the [long] procedure instead of experimenting with the radar vectors."

The data in Table 1 also support the value of a simple recovery plan, in case study T3. Whereas the supervisor's decision to try a different recovery strategy led to the complications in T2, complications in T3 arose due to a pseudo-pilot error. T3 was difficult from the outset, because the schedule appeared to have less slack, and the supervisor spent more than eight minutes working with the schedule, trying to best adjust STAs so that he could take advantage of what little slack was available. He accomplished this by first advancing the STAs of three aircraft, then assigning a new STA to the go-around aircraft in the resulting gap. This seemed to work well, but due to a pseudo-pilot error a few minutes later, the go-around aircraft failed to execute a turn and quickly became unable to meet the new time the supervisor had created; the supervisor consequently had to formulate a new plan and re-adjust the schedule. With the go-around aircraft completely off its route, the ground automation was computing unusable ETAs for it, which made it difficult to identify an appropriate time for its new STA. After discussing the issue with the sector 205 controller, the supervisor tried swapping the STAs of the go-around aircraft and the aircraft following it, which did not resolve the issue. The supervisor next identified some slack later in the schedule, and using it as a partial slot, assigned a new STA to the go-around aircraft and reset all the subsequent STAs, again delaying them as necessary to make room for the go-around aircraft. Thus, in case T3, the context of the unexpected pseudo-pilot error brought about the longest observed schedule-adjustment time and the large number of STAs that were adjusted, which in turn directly contributed to the large number of clearances controllers issued.

Contextual factors surrounding the particular off-nominal event being addressed also affect the quantitative metrics and their interpretation. Case study P4 provides an example in which a short schedule-adjustment time suggests a relatively simple, and therefore effective, recovery plan, but this view is contradicted by the number of STAs affected and the amount of delay added to the schedule. The schedule-adjustment time in P4 is much

shorter than that measured for P1, but because more delay was added to the schedule in P4, the STAs of more aircraft were affected and the recovery process required controllers to issue more clearances. However, comparing T2 with P1 requires a different interpretation. The amount of delay added to the schedule is similar for both, and P1 has the longer schedule adjustment period, but given the number of aircraft affected and the number of clearances issued, P1 appears to have been minimally disrupted. While the complicated strategy employed by the supervisor indeed affected the control actions in T2, another possible explanation for this difference is that the pilot-initiated go-arounds were declared earlier than the tower-initiated go-arounds (farther away from the runway), allowing the supervisor more time to formulate a recovery plan. After a tower-initiated go-around was declared, the aircraft was nearing the end of its assigned route (the runway threshold), so there was a certain amount of pressure associated with quickly climbing the aircraft and assigning the desired go-around procedure. Whereas in a pilot-initiated go-around event the aircraft was still several miles away from the runway, allowing the supervisor a few minutes to put a plan in place. Thus, in comparing P1 and T2, the number of affected STAs appears to be a stronger indicator of recovery-strategy effectiveness than schedule-adjustment time; in comparing P1 and P4, the absolute value of the amount of delay added to the schedule appears to be the strongest indicator. However, T1, T2, and P1 all demonstrate how contextual factors can influence the amount of delay added to the schedule.

5 Discussion

The case study analyses illustrate the value of the schedule-adjustment and control metrics in Table 1 for characterizing the effects and effectiveness of particular off-nominal recovery strategies for the schedule-based arrival operations in CMS4, together with the importance of also considering underlying contextual factors. Under different circumstances, different metrics reflect salient aspects of the recovery process. In general, the strategies that the supervisor applied were reasonably consistent, seeking to minimize the overall disturbance to the arrival flow.

The case study analyses highlighted key interactions between the scheduling and control functions central to the CMS concept that could, at times, limit the usefulness of the CMS DSTs and the extensions provided in CMS4. Interdependencies of the DSTs were apparent in T1, when the failure to assign the aircraft's RNAV go-around route in the ground automation left the sector 205 controller without DSTs. Without knowledge of the RNAV route, the automation could not compute a slot marker for the controller to use as a spatial target for reinserting the aircraft into the arrival stream. The problems experienced by the controller in T1 (together with other observations reported in [5]) highlight the utility of slot markers for managing off-nominal recovery, and emphasize the importance of considering such interactions when using the DSTs in off-nominal situations.

T1 also provided an example of how the lack of an updated ETA for the go-around aircraft complicated the supervisor's task. The supervisor could not readily determine where the go-around aircraft could be inserted, which made rescheduling it more difficult and likely increased the schedule-adjustment time. As observed in T3, clear coordination between the supervisor and controllers is also necessary when vectoring aircraft, as vectoring also reduces the accuracy and usability of ETAs the ground automation is able to generate, again complicating the supervisor's efforts to formulate a new schedule.

The case studies also suggest that improvements to the rescheduling functions available to supervisor are warranted. For instance, the supervisor inadvertently reset the schedule during T2, undoing a new sequence that had been determined and causing confusion. This episode made the supervisor hesitant to use the reset command, and he instead used the swap command to re-establish the desired sequence one swap at a time. During P4, the supervisor again had reservations about using the schedule reset function, and instead used the move function, which added unnecessary delay to the resulting schedule. An

'undo' function could give the supervisor another method for correcting any undesired or mistaken schedule changes. In addition, as highlighted in T2, establishing a plan and associated schedule quickly brings stability to the situation; however, due to the control actions that may be required, it may not be the fastest method for resolving an off-nominal event. Decisions regarding such trade-offs could be supported by scheduling DSTs that enable provisional, 'what-if' schedule planning, allowing the supervisor to coordinate with controllers about the implementation details of a proposed solution before putting the plan in place. A related aspect of the current scheduling functions is that any rescheduling actions immediately propagate throughout the ground system; when a new STA is established, the automation re-computes the location of the slot marker associated with that aircraft on the controllers' displays. Linking the provisional scheduling DSTs to provisional controller DSTs would further support coordination.

6 Conclusion

The CMS4 simulation exposed the air traffic control team to off-nominal events that caused large disturbances to the arrival flow of traffic. Their responses to the off-nominal events were analyzed as case studies in order to evaluate the overall effectiveness of different recovery strategies.

This paper presents metrics for characterizing off-nominal recovery strategies in the CMS operational environment. Together with contextual information, the metrics provide insights for understanding how the CMS tools and related procedures support off-nominal recovery, and are useful for understanding the conditions under which particular off-nominal recovery strategies are most effective. Future research should identify additional metrics for improving the characterization of operator strategies in the air traffic control environment.

References

[1] Isaacson, D., Robinson III, J., Swenson, H., and Denery, D. (2010). A Concept for Robust, High Density Terminal Air Traffic Operations. *Proceedings of the AIAA Aviation Technology, Integration, and Operations Conference, Fort Worth, TX., USA.*

[2] Prevot, T., Lee, P., Callantine, T., Mercer, J., Homola, J., Smith, N., and Palmer, E. (2010). Human-in-the-Loop Simulation of NextGen Concepts in the Airspace Operations Laboratory. *Proceedings of the AIAA Modeling and Simulation Technologies Conference, Toronto, Canada.*

[3] Kupfer, M., Callantine, T., Martin, L., Mercer, J., and Palmer, E. (2011). Controller Support Tools for Schedule-Based Terminal-Area Operations. *Proceedings of the Ninth USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany.*

[4] Callantine, T., Palmer, E., and Kupfer, M. (2010). *Human-in-the-Loop Simulation of Trajectory-Based Terminal-Area Operations.* Paper presented at the 27th International Congress of the Aeronautical Sciences, Nice, France.

[5] Callantine, T., Cabrall, C., Kupfer, M., Martin, L., and Mercer, J. (2011). Investigating the Impact of Off-Nominal Events on High-Density 'Green' Arrivals. *Proceedings from the 23rd Digital Avionics System Conference, Seattle, WA., USA.*

[6] Martin, L., Mercer, J., Callantine, T., Kupfer, M., and Cabrall C. (2012). Air Traffic Controllers' Control Strategies in the Terminal Area under Off-Nominal Conditions. *Proceedings of the 4th International Conference on Applied Human Factors and Ergonomics, San Francisco, CA., USA.*

[7] Davison-Reynolds, H., Reynolds T., and Hansman, R. J. (2005). Human Factors Implications of Continuous Descent Approach Procedures for Noise Abatement in Air Traffic Control. *Proceedings of the Sixth USA/Europe Air Traffic Management Research and Development Seminar, Baltimore, MD., USA*.

[8] Becher, T., Barker, D., and Smith, A. (2004). Methods for Maintaining Benefits for Merging Aircraft on Terminal RNAV Routes. *Proceedings from the 23rd Digital Avionics System Conference, Washington, D.C., USA.*

[9] Callantine, T., and Palmer, E. (2009). Controller Advisory Tools for Efficient Arrivals in Dense Traffic Environments. *Proceedings of the AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, SC., USA.*