Corridor Integrated Weather System

James E. Evans and Elizabeth R. Ducot

■ **Flight delays are now a major problem in the U.S. National Airspace System. A significant fraction of these delays are caused by reductions in en route capacity due to severe convective weather. The Corridor Integrated Weather System (CIWS) is a fully automated weather analysis and forecasting system designed to support the development and execution of convective weather impact mitigation plans for congested en route airspace. The CIWS combines data from dozens of weather radars with satellite data, surface observations, and numerical weather models to dramatically improve the accuracy and timeliness of the storm severity information and to provide state-of-the-art, accurate, automated, high-resolution, animated three-dimensional forecasts of storms (including explicit detection of storm growth and decay). Real-time observations of the Federal Aviation Administration (FAA) decision making process during convective weather at Air Route Traffic Control Centers in the Midwest and Northeast have shown that the CIWS enables the FAA users to achieve more efficient tactical use of the airspace, reduce traffic manager workload, and significantly reduce delays. A real-time data-fusion architecture to assist in national deployment of CIWS is under development, and the CIWS products are being used in integrated air traffic management decision support systems.**

THE CORRIDOR INTEGRATED WEATHER SYSTEM (CIWS) initiative was motivated by increasing delays in the summer months, despite the deployment of a number of convective weather decision support systems since 1997. Figure 1 compares the delays measured by the Federal Aviation Administration (FAA) Air Traffic Operations Network (OPSNET) as a function of time for an eight-year period. In the years of higher traffic (2000, 2001 prior to September 11, and 2003 through 2005), delays are noticeably higher in the summer months, when convective weather is most prevalent. The number of delays in May, June, and July of 2004 exceeded the number of delays in the same months for any of the six previous years, while July of 2005 had the highest number of delays in any month to date. As discussed in the sidebar, "National Airspace System Network Congestion Arising from Convective Weather," much of the increase in summer delays is the result of increased congestion in high-altitude en route airspace, making the National Airspace System (NAS) much more sensitive to the loss of capacity from convective weather.

The increase in delays due to thunderstorms in the summer of 1999 led the FAA to consider new approaches for managing the impacts of convective weather on the en route system. The Collaborative Convective Forecast Product (CCFP), a multi-hour forecast of convective weather, was developed and used as the input for FAA/airline strategic planning discussions during the convective season of 2000 [1]. Collaborative traffic routing plans were established via teleconferences between the FAA traffic flow managers and airline system operations center (SOC) personnel, based upon two-, four-, and six-hour collab-

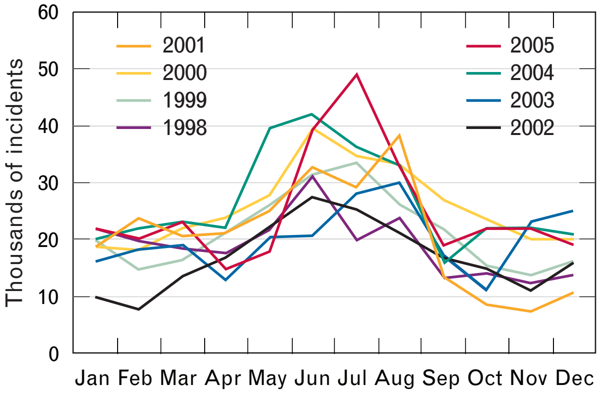

FIGURE 1. The Federal Aviation Administration's (FAA) Air Traffic Operations Network (OPSNET) delays by month for an eight-year period. OPSNET delays are reported delays of fifteen minutes or more. These delays are attributable to individual FAA facilities that assign causality to the events. Typically, approximately 70% of the OPSNET delays are attributed to weather (e.g., wind, rain, snow/ice, low cloud ceilings, low visibility, tornados, hurricanes or thunderstorms). Note that the delays are more numerous during the summer months, when thunderstorms are most frequent.

orative convective weather forecasts [2]. Studies of the collaborative planning process [3] have found that the strategic plans typically have to be modified significantly after they are generated, because of difficulties in accurately forecasting the impacts on capacity four to six hours in advance. Forecasting the capacity has two major challenges: (1) it is difficult to forecast the fractional coverage of convective weather in a region hours in advance [4], and (2) there is considerable uncertainty in the airspace convective capacity even if the fractional coverage of convective weather could be accurately forecast [5].

NATIONAL AIRSPACE SYSTEM NETWORK CONGESTION ARISING FROM CONVECTIVE WEATHER

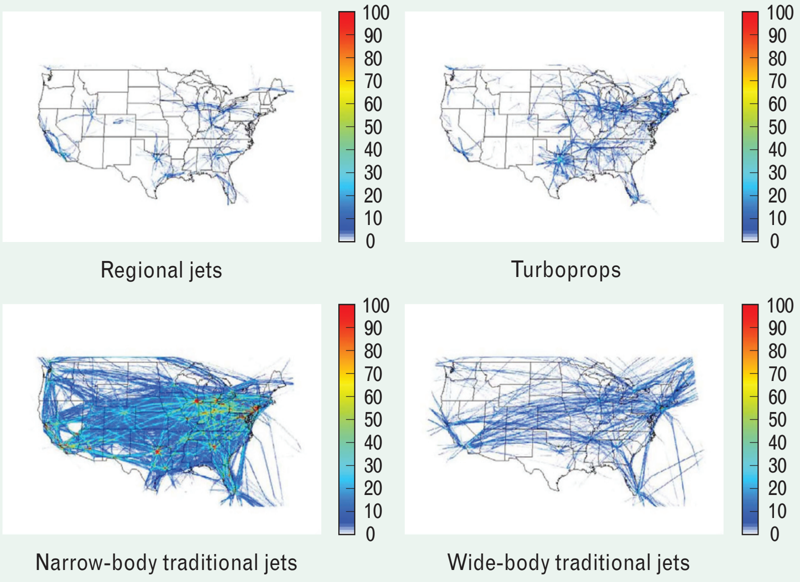

Regional jets Turboprops

Narrow-body traditional jets Wide-body traditional jets

FIGURE A. Number of flight plans in a twenty-four-hour period that pass through various regions as a function of aircraft types in January 1998. At this time, there was minimal impact observed from regional aircraft on the higher-flying jets.

In Federal Aviation Administration (FAA) publications such as the Aviation System Capacity Plans [1] and major FAA initiatives such as the Operational Evolution Plan [2], weather delays in the National Airspace System (NAS) have been attributed largely to insufficient major airport capacity when there are instrument meteorological conditions (IMC) arising from weather phenomena such as low-ceiling and/or visibility conditions. The principal approach to addressing the IMC capacity problem has been to build additional runways at major airports. However, it has become apparent from both delay statistics (such as are shown in Figure 1 in the main text) and major FAA initiatives to reduce delays [1, 2] that convective weather is an increasingly significant cause of delays.

When Lincoln Laboratory first started addressing convective weather delays in the early 1990s with the Integrated Terminal Weather System (ITWS) [3], it was believed that the principal cause of convective delays was terminal impacts. However, it became clear from the operational experience with Lincoln Laboratory's ITWS demonstration system between 1997 and 2000 that en route constraints due to convective weather were becoming more important than the terminal constraints.

Three major changes in the usage of the NAS have significantly increased the demand in en route airspace: (1) the use of secondary airports such as Manchester, New Hampshire; Providence, Rhode Island; Chicago Midway, Illinois; and Long Beach, California, by low cost carriers such as Southwest Airlines and Jet Blue; (2) the transition from turbo-prop planes that fly at relatively low altitudes in en route airspace to regional jets that fly at the same altitudes as narrow- and wide-body jets; and (3) the rapid growth in the use of business jets. In Figures A and B, we show the density of various types of air carrier aircraft, as determined by A. Mozdzanowska and R.J. Hansman of the MIT International Center for Air Transportation [4].

The spatial distribution of flights is another key factor. In Figure C, we show the density of all aircraft tracked by the FAA's Enhanced Traffic Management System (ETMS) for a twenty-four-hour period in September 2002, along with the coverage of the Corridor Integrated Weather System (CIWS) in 2005. We observe that the majority of the aircraft are flying along well-described routes.

In the area west of the New York and Philadelphia airports, the traffic is particularly constrained to closely spaced parallel routes, so as to yield a manageable method of transitioning traffic from en route airspace to a number of major airports in close proximity to one another. One of the very important elements of the highly congested airspace inside the CIWS coverage contour, shown in Figure C, is that there is often little or no excess capacity available when severe weather occurs. For example, rerouting aircraft around areas of actual or predicted weather can be very difficult when we must be concerned about controller overload in the weather-free sectors.

To illustrate, when convective weather occurs in Ohio, traffic from the western portion of the United States to the New York, Boston, and Philadelphia airports may need to be rerouted into Canada (north of Toronto) and/or via Georgia. When aircraft are rerouted through Georgia, it then becomes difficult to find a space in the overhead stream of aircraft for Atlanta departures. The result is then significant delays to Atlan-

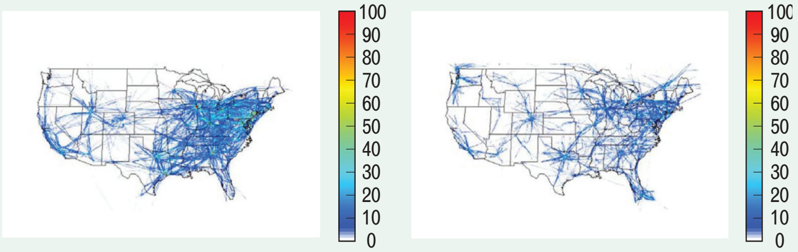

Regional jets

Turboprops

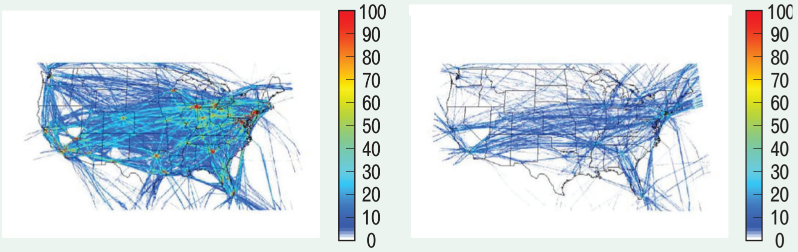

Narrow-body traditional jets

Wide-body traditional jets

FIGURE B. Number of flight plans in a twenty-four-hour period that pass through various regions as a function of aircraft types in January 2003. The increase in density of regional jets was highest in the northeast quadrant of the United States, i.e., the current Corridor Integrated Weather System (CIWS) domain.

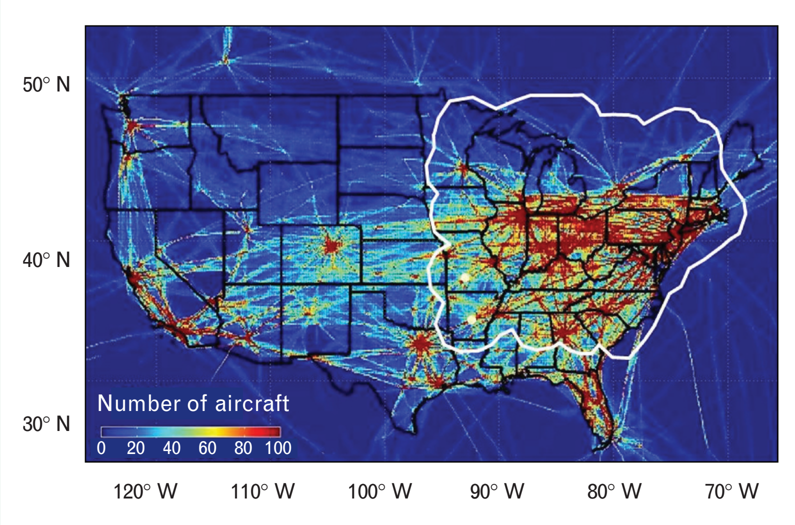

FIGURE C. Density of traffic in the United States over a twenty-four-hour period starting at 1000 GMT on 12 September 2002, with an overlay of the spatial coverage of the 2005 CIWS concept-exploration system. Aircraft are concentrated on particular routes as opposed to a random distribution in space.

ta flights even though there may never have been convective weather near Atlanta on those days. The use of Canadian airspace is somewhat limited, since Canadian Air Traffic Control (ATC) is not prepared to handle significant increases in traffic north of Toronto.

When rerouting is not feasible, the alternative is to delay departures, which results in large departure queues. Queue delays have a very nonlinear dependence on demand, capacity, and time duration of events. In the simplest case, the fair-weather capacity (C_v) of airspace under the control of an ATC facility (e.g., an airport or an en route sector) is reduced by convective activity to a lower convective weather capacity (C_w) for time duration t. Typically, C_v is greater than the demand D, but $D > C_w$. For this case of constant capacities and constant demand, the accumulated delay AD for all the aircraft involved in the queue can be shown [5] to be*

$$AD = \frac{0.5t^2\left(D - C_w\right)\left(C_v - C_w\right)}{\left(C_v - D\right)}.$$

Analysis of operations in the northeast portion of the United States [6, 7] has shown that queues (e.g., holding patterns aloft and aircraft held on the ground) are very common during convective weather events.

References
1. Federal Aviation Administration (2001): Operational Evolution Plan version 3.2, Washington, http://www.faa.gov/programs/oep.
2. Federal Aviation Administration (2005), Operational Evolution Plan version 7.1, Washington, http://www.faa.gov/programs/oep.
3. J.E. Evans and E.R. Ducot, "The Integrated Terminal Weather System (ITWS)," *Linc. Lab. J.* 7 (2), 1994, pp. 449–474.
4. A. Mozdzanowska and R.J. Hansman, "Evaluation of Regional Jet Operating Patterns in the Continental United States," MIT International Center for Air Transportation Report ICAT-2004-1, May 2004.
5. J.W. Andrews, "Impact of Weather Event Uncertainty upon an Optimum Ground Holding Strategy," *Air Traffic Control Q.* **1** (1),1993, pp. 59–70.
6. J. DeArmon, C. Wanke, E.K. Beaton and S. Miller, "Complex, Congested Airspace and Impacts on Airline Operations," chap. 34 in *Handbook of Airline Operations,* 1st ed. (McGraw Hill, New York, 2000).
7. M. Robinson, J. Evans, B. Crowe, D. Klingle-Wilson, and S. Allan, "Corridor Integrated Weather System Operational Benefits 2002-3: Initial Estimates of Convective Weather Delay Reduction," Lincoln Laboratory Project Report ATC-313 (9 Apr. 2004).
8. S.S. Allan and J.E. Evans, "Operational Benefits of the Integrated Terminal Weather System (ITWS) at Atlanta," Lincoln Laboratory Project Report ATC-320 (15 July 2005).

---

* The generalization of the above equation to consider the case where C_w and D change with time was validated with Atlanta thunderstorm event data [8].

The CIWS demonstration began in 2001. The CIWS objective is to provide improved tactical air traffic management (ATM) decision support, via short-term accurate convective weather forecasts, as a complement to strategic traffic flow management (TFM) accomplished through the CCFP strategic planning teleconferences.

In this article, we describe the current status of CIWS, including initial operational results of Air Traffic Control (ATC) use of the CIWS weather products. We begin with some CIWS background, describing the motivation for the program, the role of CIWS products in the overall convective weather planning process, and the functional domains in which CIWS products can provide operationally significant benefits. We then review the current CIWS capabilities, spatial coverage, sensors used, products, operational users, and integration with ATM systems. Next we discuss the real-time demonstration system, which currently provides the products, and describe the evolution of this prototype toward a technology-transfer package for CIWS operational implementation. Then the results of detailed CIWS operational benefits studies carried out in 2003 and 2005 are summarized. Finally, we discuss the FAA plans for CIWS and near-term enhancements to the system.

Background

Figures 2 and 3 show examples of convective weather events in 2005 that led to long delays in the NAS. We see that delays arose from both disorganized convective weather (characterized by weather cells scattered apparently randomly over a large area) and from organized convection, such as squall lines. Note that in both figures, the storms are far from impenetrable. The various Air Route Traffic Control Centers (ARTCC) in the CIWS domain, shown in Figure 4, typically experience convective weather impacts on 40% to 75% of the days in the period of May through August.

The objective of the CIWS is to assist the decision makers in rapidly determining and executing an appropriate weather impact mitigation plan that will take advantage of the available capacity. Two key factors had to be considered in designing the system: (1) development of weather products that could assist in the identification of usable en route airspace, and (2) identification of the key users and an appropriate training paradigm for CIWS products.

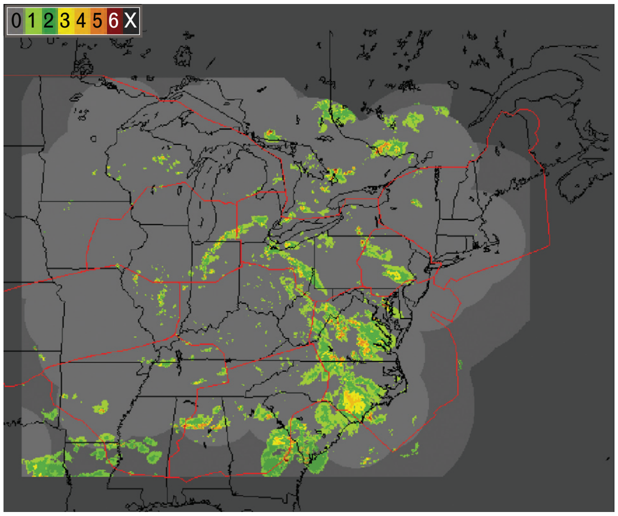

FIGURE 2. Convective weather at about 8 p.m. EDT on 13 July 2005. This day set a new all-time record for the number of national OPSNET delays. The convective weather never did impact the Chicago terminal areas on this day, but led to substantial delays at these facilities nonetheless. The New York City airports were directly impacted by storms for about one hour of the fifteen-hour period of significant delays. In the figure, light green denotes regions of light rain; dark green is light to moderate rain; yellow, orange, and red regions denote increasingly heavy rain, corresponding to a rain rate greater than 0.4 inches per hour. Red regions may include three-quarter-inch hail as well as very heavy rain. Pilots will typically avoid flying through storms indicated by the yellow, orange, and red regions.

Lincoln Laboratory studies of en route pilot decision making in convective weather have shown that both the radar reflectivity and a storm's vertical structure are important in determining whether a pilot will fly through or over a convective storm [6, 7]. These studies determined that pilots will typically avoid storms with an equivalent reflectivity greater than 40 dBZ. In en route airspace, however, pilots will generally fly over storms if the aircraft altitude is at least 5000 feet above the storm radar echo tops. Thus highly accurate representations of the three-dimensional (3D) storm information are needed as input to the traffic management decisions.

The speed of execution of the operational decision loop, shown in Figure 5, must be commensurate with the time scale over which the weather changes and with the ability to accurately forecast the weather im-

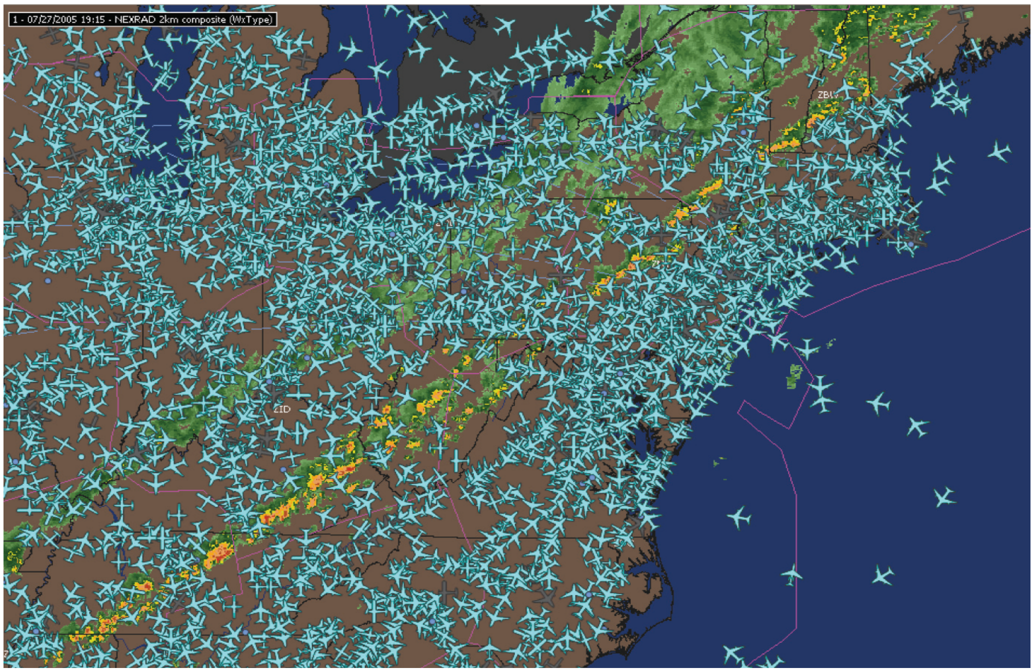

FIGURE 3. Convective weather and aircraft locations about 3 p.m. EDT on 27 July 2005. This day had a very high number of OPSNET delays. Note that the aircraft are penetrating a long squall line at a number of locations. (This figure was generated with Flight Explorer software.)

pact. If this cannot be achieved, then the plans that are executed will no longer be an appropriate solution for the weather situation.

As noted in the sidebar, the NAS in the northeastern U.S. is a highly congested network in which disturbances arising from a loss of capacity due to adverse convective weather rapidly propagate throughout the network. A consequence of the highly coupled nature of the NAS is the need for extensive collaboration and coordination in addressing convective weather impacts. In Figure 6, we show the results of an analysis of interfacility coordination issues by the MIT Center for Air Transportation [8].

Figure 6 suggests that the CIWS products need to be available at towers, terminal control facilities (TRACON), en route ARTCCs, the ATC System Command Center (ATCSCC), and airline dispatchers. In the ARTCCs, we have found that both the traffic flow managers and the area supervisors are important. The traffic management coordinators (TMC) play a key role in addressing NAS network problems, but they must coordinate with many other potential aviation weather forecast users. Note that airline dispatchers are

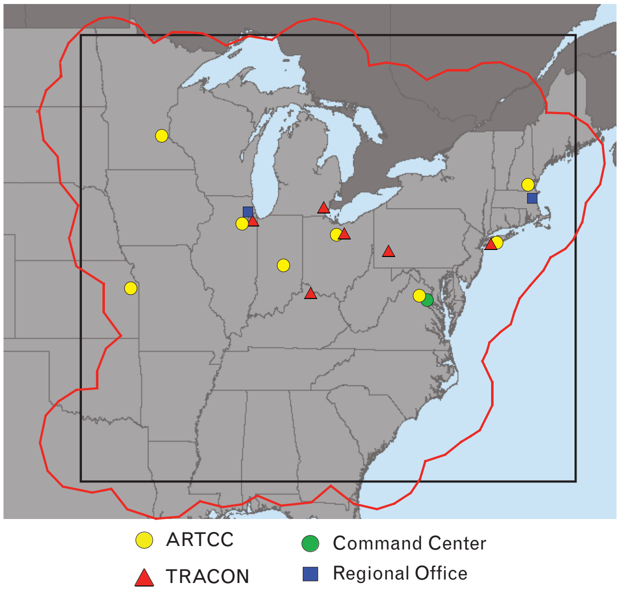

FIGURE 4. Corridor Integrated Weather System (CIWS) domain and FAA users for 2006. The black outline indicates the size of the current CIWS grid, the red outline depicts the outer boundary of the radar coverage contributing to the generation of CIWS products within the grid, and the symbols refer to the locations of FAA facilities (within the CIWS domain) that currently have dedicated CIWS displays.

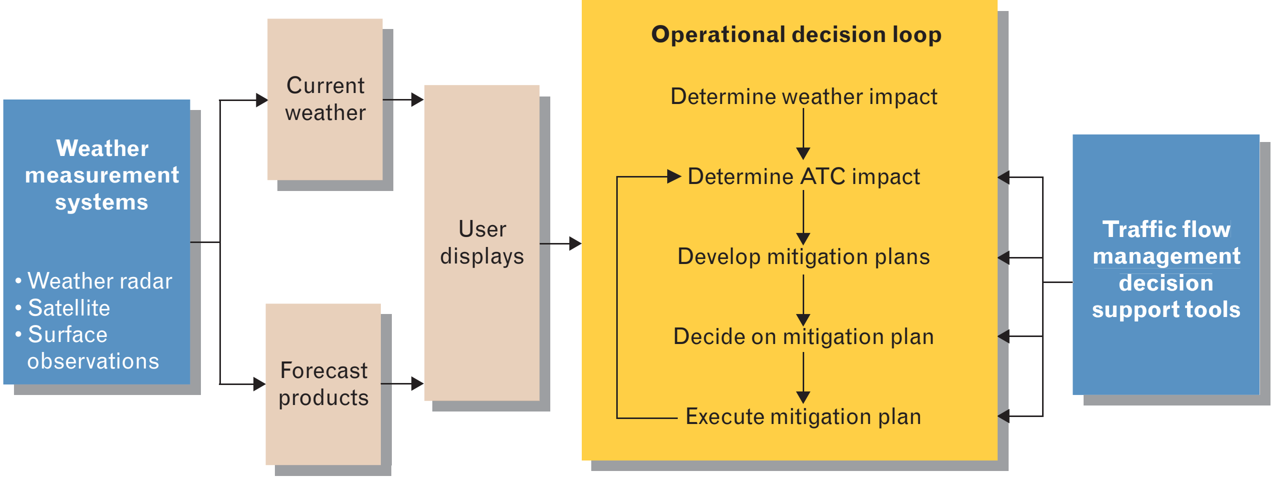

FIGURE 5. Overall convective weather impact mitigation process. Timeliness and accuracy of weather predictions and air traffic control (ATC) impacts are critical for an effective execution of the decision loop.

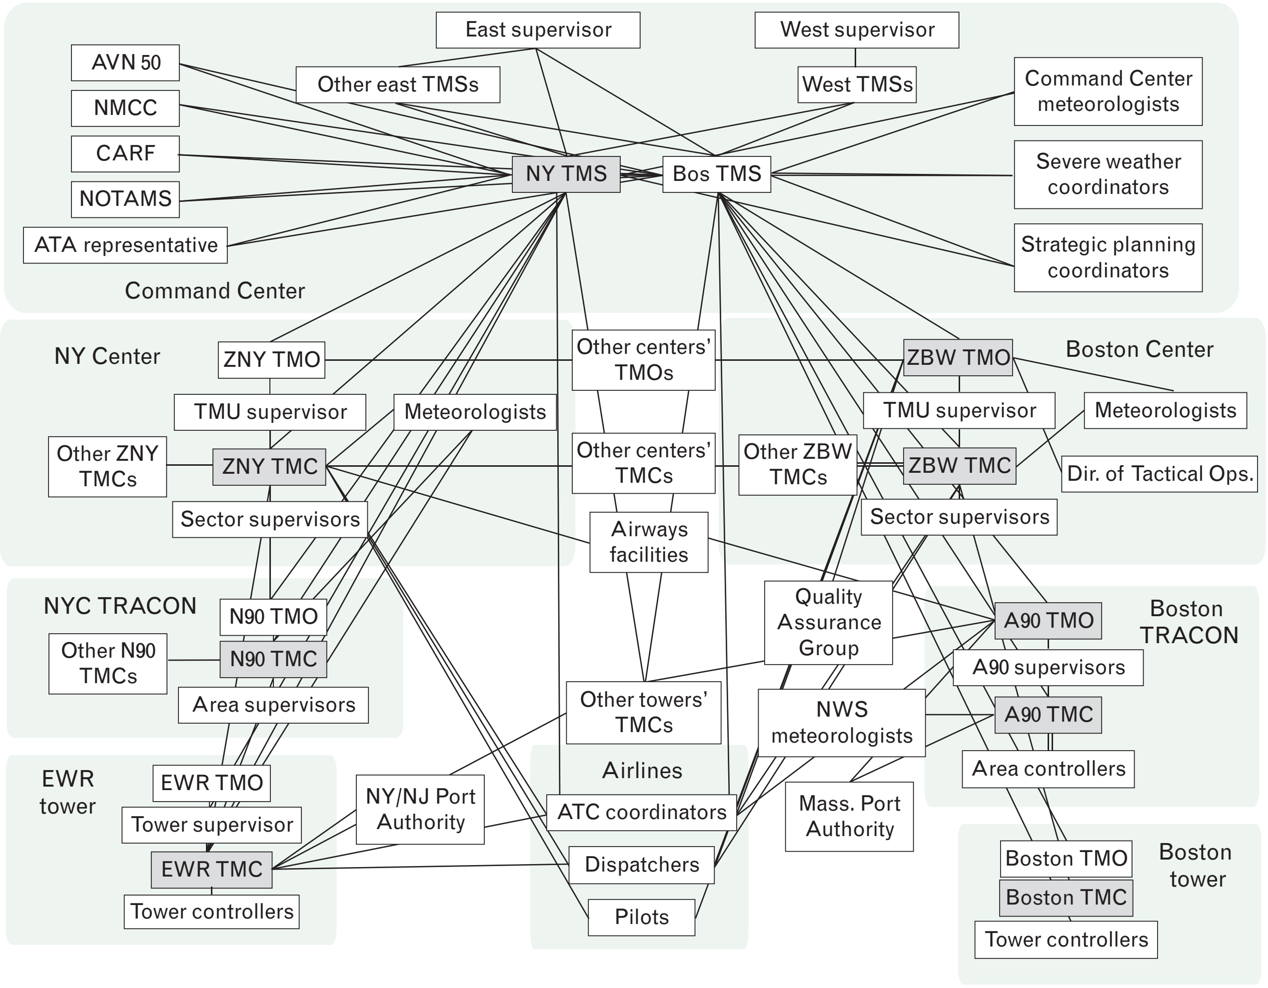

FIGURE 6. Interactions between various FAA facilities and airlines in addressing congestion problems related to Newark International Airport (EWR), identified by H.J. Davison and R.J. Hansman [8]. The dark boxes indicate the personnel who principally manage the coordination with other FAA facilities and airline ATC coordinators. This figure shows that coordination is a major factor in accomplishing the convective weather impact mitigation process shown in Figure 5.

an important component of the coordination process, because rerouting and other adjustments to air carrier flight plans may be necessary to address the combination of weather and congestion problems.

Active research is under way [9] to integrate the CIWS three-dimensional convective weather forecasts with the NAS airspace structure information (including aircraft trajectory information) to provide an integrated decision support system that will further enhance the FAA traffic manager's ability to develop and execute mitigation plans. A companion paper [5] in this issue of the *Journal* describes the operational use of the CIWS products by the Route Availability Planning Tool (RAPT) to increase the rate of departures from the New York City airports during en route convective weather.

Generating the CIWS Products

The CIWS takes advantage of the high density of existing FAA, National Weather Service (NWS), and Environment Canada weather sensors, shown in Figure 7, and the FAA-funded research conducted on thunderstorm evolution. Figure 8 illustrates the principal CIWS products. Real-time observations of product usage have consistently shown that an accurate depiction of the current three-dimensional weather situation is at least as important as the three-dimensional forecasts in the development and execution of appropriate aviation weather mitigation plans.

Lincoln Laboratory's past experience in developing fully automated wind-shear systems [10] and integrated terminal weather decision support systems [11]

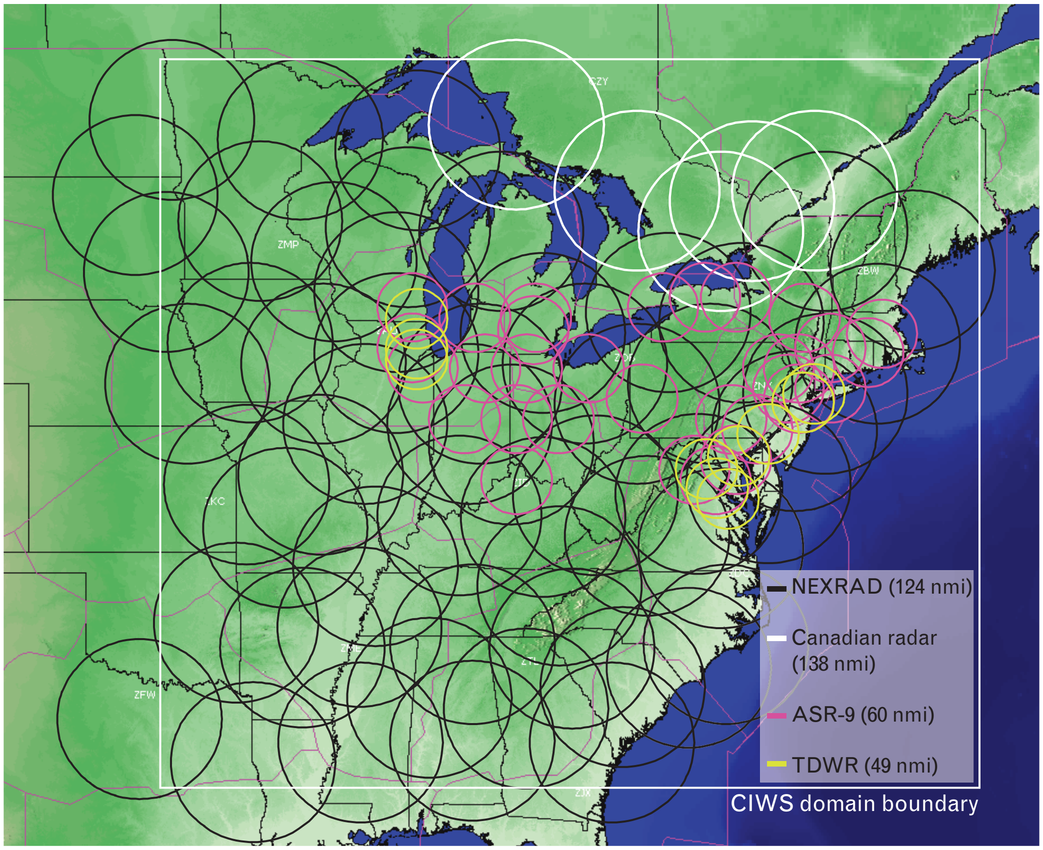

FIGURE 7. Terminal and en route weather sensors utilized to create the CIWS products in 2006. The thirty-second update rate of the Airport Surveillance Radar Model 9 (ASR-9) radars (shown in pink, range of 60 nmi) is utilized to detect rapidly growing cells, while the Next Generation Weather Radar (NEXRAD, shown in black, range of 124 nmi), Terminal Doppler Weather Radar (TDWR, shown in yellow, range of 49 nmi), and Canadian radars (shown in white, range of 138 nmi) provide information on three-dimensional storm structure and boundary layer winds. Data from lightning sensors and the Geostationary Operational Environmental Satellite (GOES) are also integrated with the radar data.

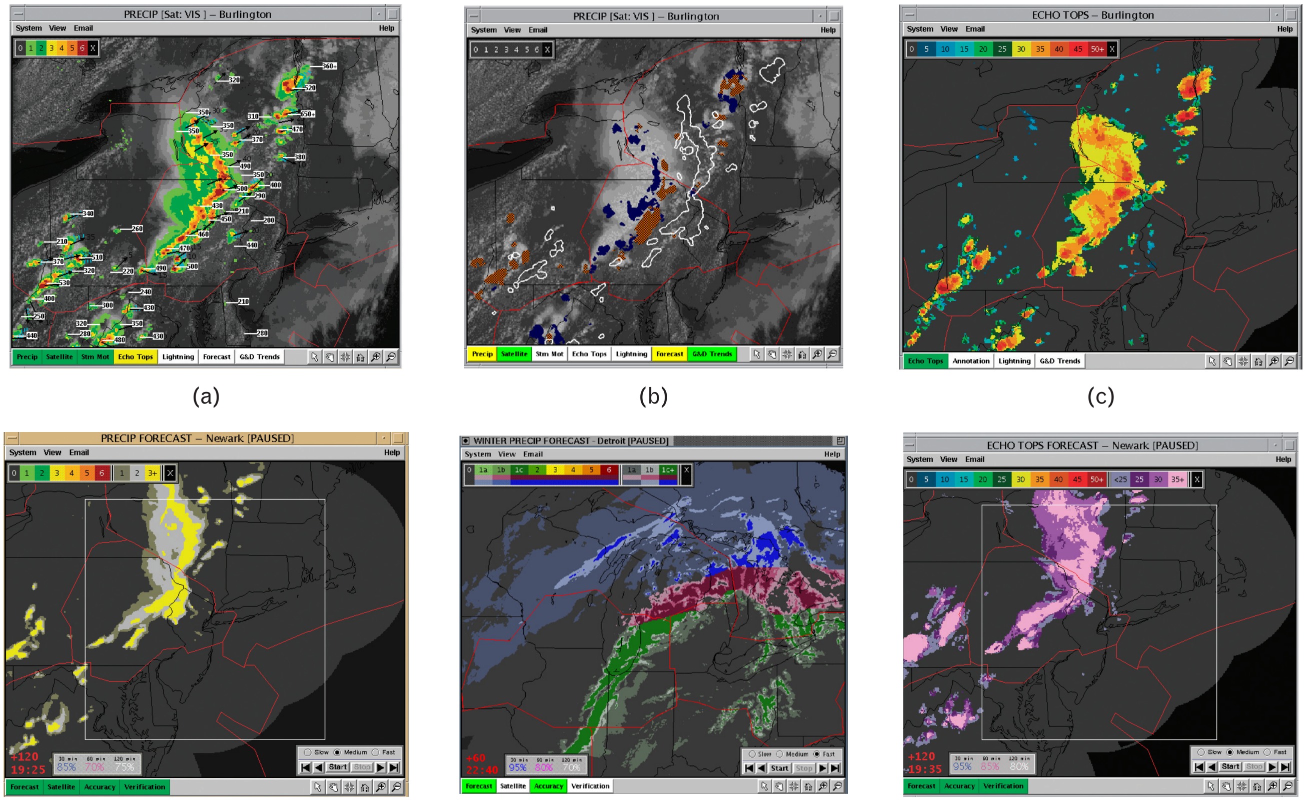

(a) (b) (c)

(d) (e) (f)

FIGURE 8. Examples of the 2006 CIWS products: (a) precipitation mosaic with echo tops tags, storm motion, and satellite data; (b) growth and decay trends and two-hour precipitation forecast contours (in white) overlaid on satellite data; (c) echo tops mosaic; (d) convective two-hour precipitation forecast; (e) two-hour winter-mode forecast with precipitation phase; and (f) two-hour echo tops forecast.

has shown that successful data fusion typically requires considerable attention to addressing data quality problems of the individual sensors, as well as the use of innovative data-fusion algorithms. The CIWS uses the raw product data from radars such as Next Generation Weather Radar (NEXRAD) as opposed to the current products generated by those radars. The acquisition of raw NEXRAD data allows the use of improved product algorithms to generate the NEXRAD products used by CIWS.

For example, reflectivity measured on a single elevation scan of the radar and/or the maximum reflectivity seen at any altitude (used on en route controller displays) are susceptible to anomalous propagation contamination, ground clutter, and spurious bright-band high-reflectivity returns from melting snow or ice particles. A Lincoln Laboratory study [12] showed that vertically integrated liquid (VIL) is a better indicator of storm severity and new growth and is less susceptible than other precipitation representations to anomalous propagation and other anomalies. Lincoln Laboratory has developed improved anomalous propagation recognition algorithms for individual NEXRAD sensors, as well as an improved VIL mosaic algorithm, which is used to generate the CIWS precipitation field.

A second significant data quality issue is the need to mitigate the underestimated storm echo tops from the current NEXRAD echo tops algorithm. An accurate measure of radar echo tops is particularly important in the en route domain with the rapid transition to regional jets for commuter operations. However, the current operational NEXRAD products underestimate tops by as much as 12,000 feet. An improved echo tops algorithm has been developed at Lincoln Laboratory that more accurately estimates the true storm echo tops [13]. This algorithm is used within CIWS and has also been transferred to the NEXRAD program to improve that system's operational products.

The CIWS Product Components

Users who are principally concerned with the current locations of high-reflectivity storms can focus on Figure 8(a), which contains the current precipitation (VIL) mosaic product with storm motion vectors and storm-top height tags (in hundreds of feet) all overlaid on the visible satellite image. Figure 8(b) is concerned with storm trends: the regions of storm growth are shown in red/brown, decay trends are in blue, and the two-hour forecast of high-reflectivity storm locations is shown as a white contour overlay. These growth and decay trends show very readily the areas of expected convective development and decay. Users who are assessing where it is safe to fly over storms that would otherwise block major jet routes need accurate radar-based echo tops information. In a manner similar to the VIL mosaic, Figure 8(c) exhibits a mosaic of echo tops to show the current heights (in kilofeet) of the convective weather.

The CIWS three-dimensional forecast is composed of two main products: the precipitation (VIL) forecast and the echo tops forecast. Details on the algorithms used to generate the forecasts are provided in a companion paper in this issue of the *Journal* [14]. Our discussion here will focus on the presentation of the products to the operational decision makers.

Figures 8(d) and 8(f) show the forecasts of regions of high-reflectivity returns and echo tops in time animation. In both cases, the animation shows sixty minutes of past weather, then advances the forecast in fifteen minute increments to the maximum forecast lead-time horizon of 120 minutes. A winter mode, shown in Figure 8(e), can be enabled to display the weaker VIL values in much greater detail and convey more information about cold-season aviation impacts such as heavy snow bands. The winter mode also includes a forecast of the phase (snow, ice, rain) of the precipitation.

In Figures 8(d) through 8(f), forecast accuracy scores are provided on both the precipitation and echo tops forecast animation displays so that the user can determine how much credence to give to the forecast information at various forecast lead times. Additionally, the user can view the past thirty-, sixty-, and ninety-minute forecasts for the current time overlaid on current weather to better understand the spatial distribution of the forecast errors.

The CIWS Demonstration System

The real-time product-generation and display functional prototype is critical to the CIWS evolutionary development. By design, the functional prototype provides a rapid update capability that allows us to be responsive to operational user feedback. Figure 9 shows how the demonstration system, after it was fielded in 2001, quickly evolved, both in coverage and set of products provided.

In building the system, we were able to draw heavily on the successful prototype architecture developed for the Integrated Terminal Weather System (ITWS) [11], which allowed us to produce a running version of CIWS in less than six months, from the time of initial concept to real-time operation. The CIWS prototype offers the type of flexibility we associate with a research and development activity, but which is not normally available as part of an operational system. This flexible infrastructure, in which the software can transition rapidly from a desktop testing environment to the real-time system without re-engineering, and computing resources can be configured into the system as needed, is critical to our ability to support the rapid-development timelines shown in Figure 9.

The CIWS real-time demonstration system, which runs round the clock, has served the program very well thus far, with a time availability of greater than 99%. However, at this juncture, the system faces three end-of-life issues, which are being addressed in the design of the next-generation CIWS prototype. The first is the obsolescence of the existing infrastructure and the related operational overhead associated with the aging prototype. The second deals with both product-generation and display scalability concerns. Unless major modifications are made, the current system is at its computational limits with the existing CIWS grid. The third is the need to support rapidly emerging standards, in the software development itself, in the definition of the interfaces, and in the techniques for handling the movement of the large data fields.

In the sections that follow, we begin with an overview of the CIWS prototypes, both existing and planned. We then discuss the end-of-life issues identi-

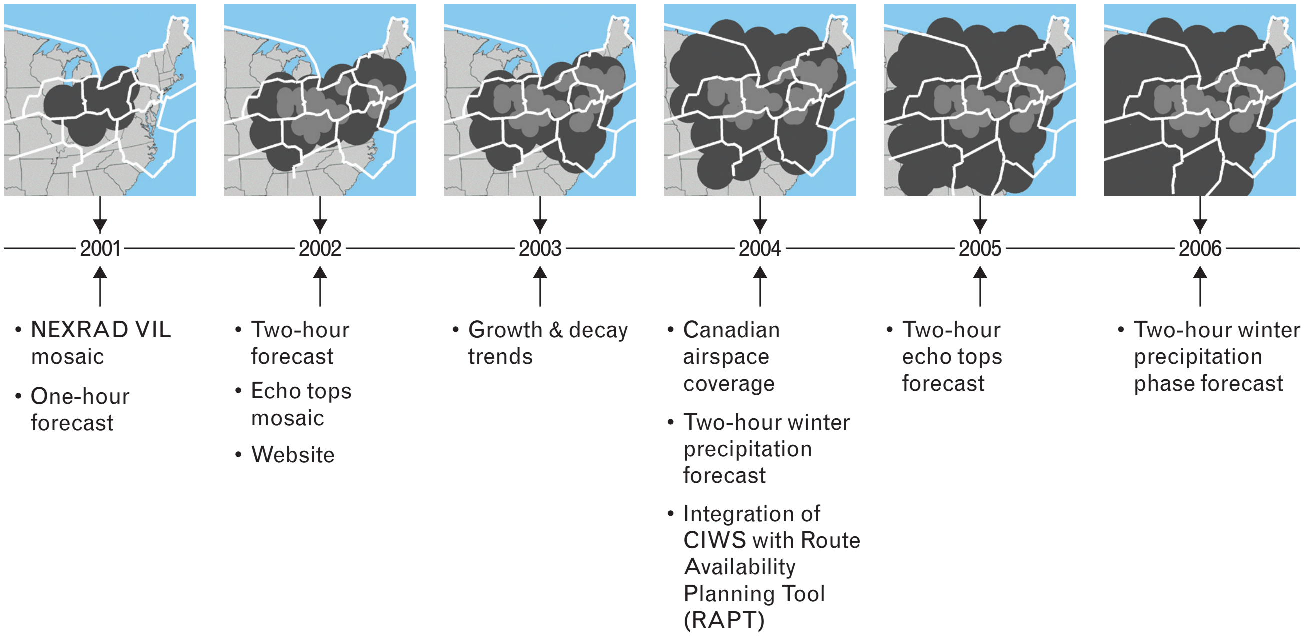

FIGURE 9. Milestones in the evolution of the CIWS demonstration system. The evolution of the system included both expansion in the region covered and in the products provided to the users. The lighter brown regions are those covered by ASR-9 radars, and the darker brown areas indicate those covered by NEXRAD and Canadian radars.

fied above from the perspective of how they manifest themselves in the existing system and how the new prototype will address the concerns.

Overview of the Prototypes—Existing and Planned

The core CIWS product-generation system, as it exists today, consists of a heterogeneous set of networked Solaris and Linux servers. The software is organized as a configuration of loosely coupled, asynchronous data processing modules, distributed across the set of available compute nodes. In the current CIWS prototype, over 1300 process instances are distributed across more than sixty-five compute nodes. Output products from the product generator are sent to more than sixty CIWS situation displays (SD) over a dedicated private frame relay network. The current CIWS displays are placed in a number of FAA operational facilities (eight ARTCCs, six TRACONs, and the Command Center) and nine airline SOCs. An even larger community of users (including key users from the FAA's counterpart in Canada) accesses the CIWS products in real time over the Internet through a standard web browser.

Figure 10 depicts the structure of the new CIWS prototype, which is intended to serve over the next few years as a testbed for demonstrating concepts pertaining to the evolution of integrated aviation weather decision tools, as well as a vehicle for providing part of the CIWS technology-transfer package to the FAA for operational implementation of the CIWS capability. One of the important features of the new CIWS system is the connection between the design of the new prototype and the long-term plans within the FAA for managing the dissemination of sensor data and products. The FAA's future net-centric infrastructure, System Wide Information Management (SWIM), is shown on both sides of the CIWS product generator in Figure 10. Once SWIM is in place, it will deliver observations and sensor and model data to CIWS through a single network point of presence and will manage the dissemination of the products from CIWS by exploiting emerging standards. In the meantime, judicious use of the dissemination concepts that SWIM will eventually support provides architectural benefits to the FAA's NAS system designers, which go beyond the value of the CIWS products to the users. The current prototype performs similar functions, but uses custom interfaces and makes point-to-point connections to the individual sources and sinks over the CIWS frame relay.

The design and development of the new prototype

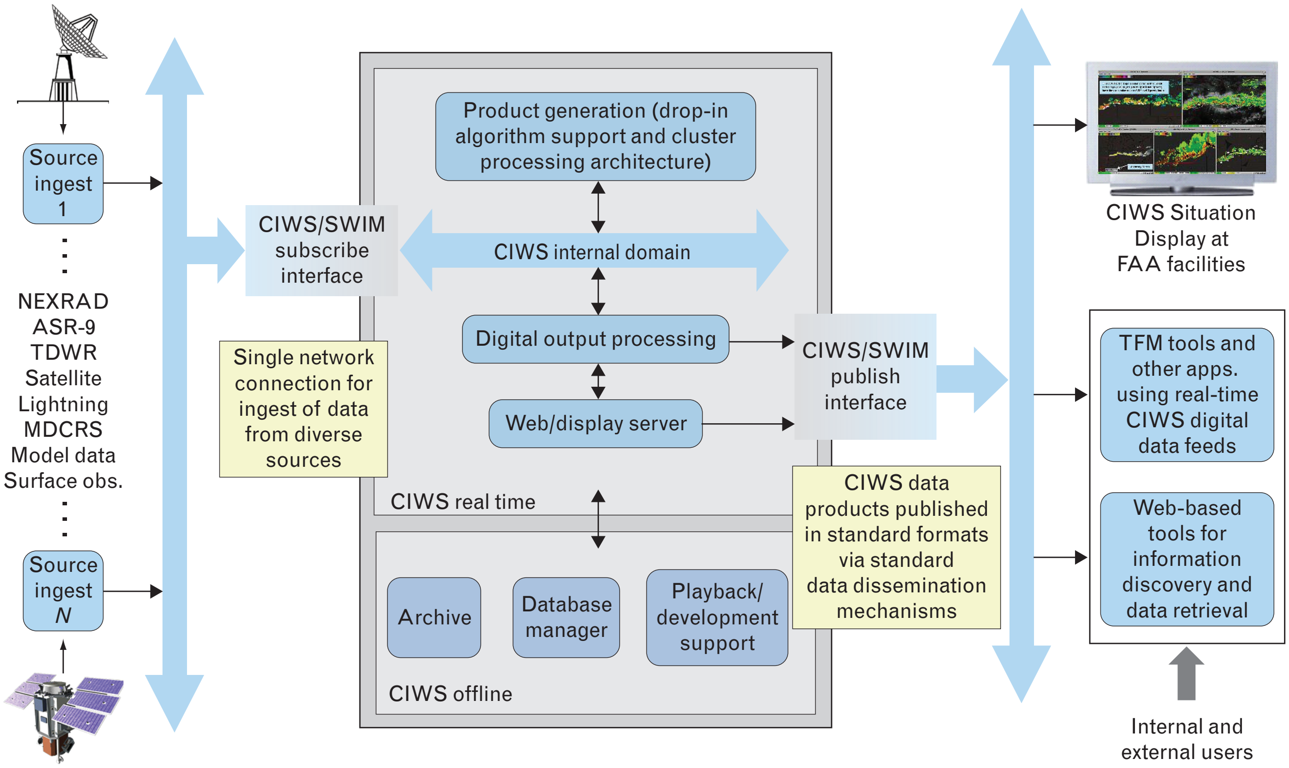

FIGURE 10. Structure of the new CIWS prototype. The addition of SWIM, the net-centric infrastructure, will simplify the delivery of sensor and model data to CIWS, and will manage the dissemination of products to the users.

began in earnest at the beginning of 2005. The transition from the existing prototype to the new system for live operations is expected to take place in time for the 2008 convective season.

Obsolescence

The heterogeneous network currently in the field has been built up over time. Many of the oldest of its commercial-off-the-shelf (COTS) hardware elements, having been reclaimed from other projects, predate the first CIWS prototype. Over the last year, these elements have begun to fail. Unfortunately, when this happens, we are often unable to replace the hardware with systems that provide a comparable software environment. A number of the software libraries used by the algorithms are no longer well supported. In addition, a number of components have been made obsolete or redundant by the advances in the ANSI C++ standard library implementations. In many cases, the oldest software components cannot be rebuilt on the new systems available from the vendors. The logistical issues posed by the aging infrastructure are felt in both the management and operation of the prototype, leading to significant extra administration and monitoring in order to meet the high availability requirement associated with the demonstration.

The new prototype will be largely homogeneous in both hardware and software. A 100-node Dell cluster forms the core of the product-generation system, and the software is being rewritten on the basis of unified software standards in a single higher-level language (C++ for the product generation and Java for the displays). The individual cluster nodes are Dell PowerEdge 1855 'Blades,' each configured with dual 3.2 GHz Xeon processors, 4 to 6 GB of memory, and dual gigabit Ethernet. The nodes are installed in ten rack-mountable chassis, with ten nodes per chassis. Nodes within each chassis are hot swappable, and each chassis is configured with dual power supplies, which are also hot swappable. Dual uninterrupted power supply units are used to provide redundant power to each chassis. Figure 11 shows the cluster nodes connected by a group of ten Cisco high-speed gigabit Ethernet switches, organized into two switch stacks of five switches each. The cluster is highly scalable in both compute and input/output bandwidth, which allows for a potential increase in both the complexity of the algorithms supported and the richness of the product suite offered.

The use of the cluster not only addresses the immediate system obsolescence issues, but also mitigates some of the system administration and real-time monitoring load associated with the existing prototype. The hardware is inherently reliable, due to the form factors used, the dual hot swappable power supplies, and unlimited power supplies. In addition, open-source cluster management tools (e.g., ROCKS and Ganglia) provide much-improved capabilities for system management and simplify the system control and monitoring tasks.

Although Ganglia can be leveraged to provide a certain level of fault tolerance, none of the existing tools provide all the functionality needed to achieve the fault-tolerant scheduling capabilities desired for the system. Open systems support for fault tolerance in a cluster environment is still in the developmental stage, with no freely or commercially available robust implementations to date. One general approach that we are pursuing is centered on the concept of cluster virtualization, in which a cluster appears as a set of virtual nodes rather than physical nodes [15]. This approach allows a set of processes running on a single virtual node to be automatically rescheduled to another physical node should a non-recoverable fault occur.

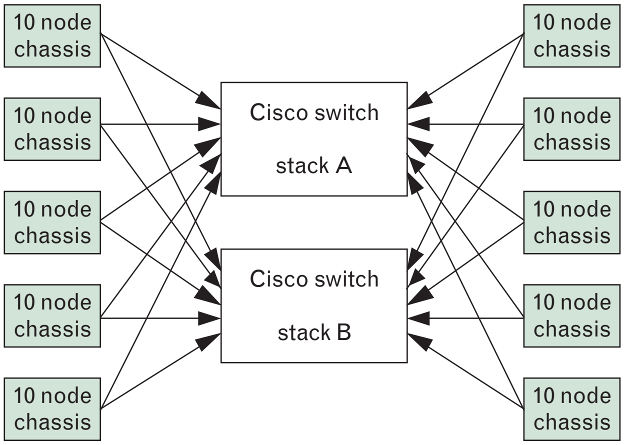

FIGURE 11. CIWS cluster interconnect configuration. This configuration of node-switch interconnections is highly scalable.

Given the requirement for the CIWS system to run around the clock with minimal manual intervention, the CIWS design will incorporate a cluster-virtualization layer and virtual-node process scheduler. Figure 12 shows the basic design of the system control and monitoring subsystem. Processes are first assigned to a virtual-node group, using a straightforward file-based mechanism. The master process manager then assigns each virtual node to an online, healthy, physical node and issues the remote commands necessary to start the processes on the physical node. In the event of a non-recoverable failure, error logging daemons supported by the system will detect the failure and report it to the master system monitor node, and any affected processes will be remapped to an available spare cluster node.

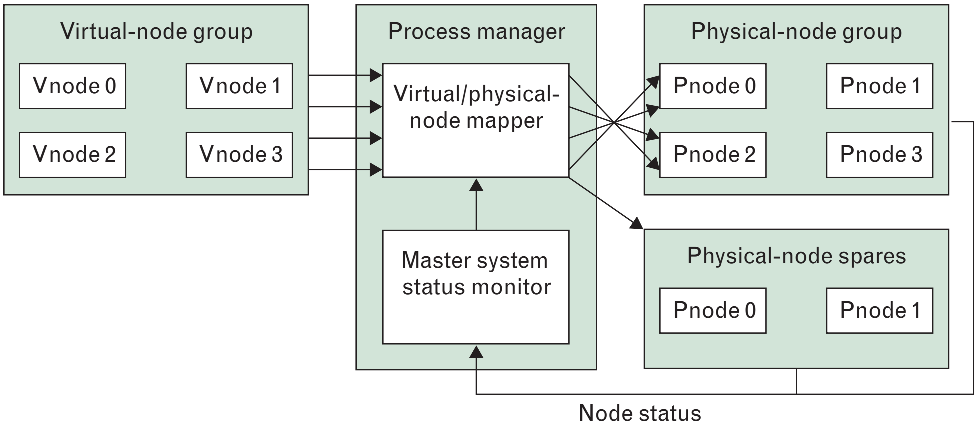

FIGURE 12. CIWS virtual-node fault-tolerant scheduler. The process manager connects virtual nodes to available physical nodes.

Scalability—Product Generation

As indicated above, the existing hardware and software infrastructure of the CIWS demonstration system had reached its computational limits. With the implementation of the new prototype, the future scalability of the hardware system is assured. However, software scalability issues remained. The long-term vision is for CIWS to cover the continental United States (CONUS). As a result, there is a significant challenge in scaling up from the current CIWS grid, which covers about one-third of the eventual total, to the end state. The existing CIWS prototype is designed around the assumption that algorithmic results, whether intermediate or final, can be generated on a single processor for the entire grid within the required product update cycle.

For many of the CIWS products, a number of sequentially generated intermediate results might be needed to build a product that updates every five minutes. For the grid to grow beyond its current size, changes to the computing paradigm are required, that is, a decomposition of the grid into subgrids, parallelization of the computation, and efficient recombination of the components into final results. For this new paradigm to be efficient, a common representation of the grid over the entire domain is needed. Moreover, this common representation should be one that supports both forward and inverse mapping, is computationally simple, and yields low processing errors. Selecting such a grid representation presented a number of technical challenges.

The CIWS image processing algorithms make extensive use of matched-filter techniques to generate both intermediate and final output products. For fixed-size matched-filter kernels to be used, the spatial distortion of the underlying grid must be minimized to ensure that the filter response is consistent across the entire processing domain. In the ITWS system, this domain was limited to a few hundred kilometers in size, and a simple modified flat-earth projection was sufficient. When we use the ITWS mapping to represent a grid the size of the current CIWS system, we approach, or in some cases slightly exceed, the acceptable error bounds for the image processing computations. The shape distortion that would occur if the ITWS legacy mapping were used for a CONUS-sized grid is shown in Figure 13(a), which depicts the percentage error in the shape of the cell at a given location of the grid centered at the middle of the United States. For a CIWS CONUS domain, the shape errors become unacceptably large and a different projection must be used. Because there is no mapping of the lines of latitude and longitude onto the *x-y* (row-column) space that can preserve both north-south direction and shape (area), a compromise between area distortions and directional errors must be reached. The more critical criterion for CIWS is the preservation of shape.

A projection was sought that would eliminate as much of the shape distortion as possible, while at the same time producing a representation that could be provided to the users. The resulting map should be one that is recognizable, pleasing to the eye, and useful in dealing with other geographic tools. The azimuthal equal-area mapping presented by J.H. Lambert in 1772 [16] meets all of these requirements, providing a near-optimal balance of very low shape distortion coupled with a directional error that is well within acceptable bounds for a two-hour forecast. Figure 13(b) shows the grid-cell distortion-error characteristics for the Lambert projection. Comparing the custom ITWS projection in Figure 13(a) and the new standard CIWS projection in Figure 13(b) illustrates the positive attributes of the new grid in terms of scalability and integrity across the domain. The added appeal of this representation, in addition to its efficiency of use, is that it is the same mapping used by the aviation community in general and in particular by the Enhanced Traffic Management System (ETMS), a traffic flow management system that is heavily used operationally in conjunction with CIWS. Moreover, because the mapping is a recognized standard, any of a number of standard data representations can be used for easy export.

Scalability—Dissemination and Display

The larger CIWS domain provides challenges for the dissemination of the products and the management of the products on either a dedicated display or Web site. Whether the products are to be viewed on a dedicated display or accessed via the Web, the system needs to be responsive to the users' requests. With the larger do-

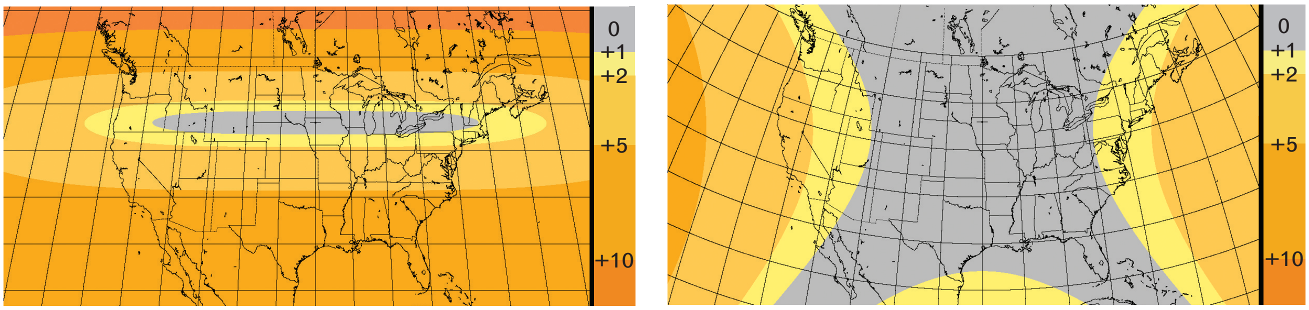

(a) (b)

FIGURE 13. Representations of a continental United States–sized grid: (a) Legacy grid representation; and (b) Lambert azimuthal equal-area grid centered at the middle of the United States. This grid significantly reduces the shape error of the legacy grid, while retaining recognizability in the image. The colors in these figures represent the percentage error in the shape of a cell in a given location.

main and the underlying 1 km resolution of the products, the size of the objects that the displays must ingest, store, process, and render is significant. Handling these objects in real time stresses the performance of the display as the users select from among the products they want to see—navigating across the domain and focusing in on specific areas of interest. Due to screen resolution constraints, the display is not able to render the entire national grid at the full 1 km resolution. The display design utilizes a tiling technique to help manage the large suite of products available at 1 km, allowing the display to subscribe to and receive only those portions of the grid visible to the users at any given time. In addition, the displays will use tiled images at various resolutions to speed the access and display of the data and will manage the choice of display products at 4 km, 2 km, and 1 km resolution, depending on the need. The resolution of the tiles requested will be determined by the window size and the zoom level, so that the display will deliver higher-resolution images as the user zooms in. In addition, compression of the data flows will take place as part of the transmission to the displays.

Need to Support Emerging Standards

In recent years, standards for the development of systems and dissemination of information have gained increasing support from research and commercial organizations alike. The power of these standards has been demonstrated many times, allowing the exchange of software and the leveraging of shared utilities and libraries in the development of object-oriented processing and display systems in a cost-effective manner. The existing prototype relies heavily on custom software developed over a number of years. By contrast, in tackling the design for the updated CIWS prototype, we are committed to the use of standards for development and documentation, as well as relying wherever possible on COTS libraries for processing and display. This commitment has been most apparent in the selection of the grid representation, in the use of publicly available software support for grid manipulation, and in the development of the displays. The rewards from embracing dissemination standards are already being felt in the ability to provide the weather products using an accepted format, such as HDF5 [17], which is designed for efficiently handling large grids. An experiment with this dissemination mechanism led to the successful exchange of products between CIWS and the National Aeronautics and Space Administration (NASA), without requiring an interface description or exchange of decoding software. Use of the shared libraries that support this and other emerging standards reduces the cost of development in the near term and system maintenance in the future. Reliance on the technologies for display development coming from the commercial sector is speeding the development of the user tools for this project as well. Support for emerging standards such as Web Services and other related Web-based technologies also allows CIWS to keep pace with the evolving SWIM infrastructure discussed earlier.

Operational Benefits of the CIWS

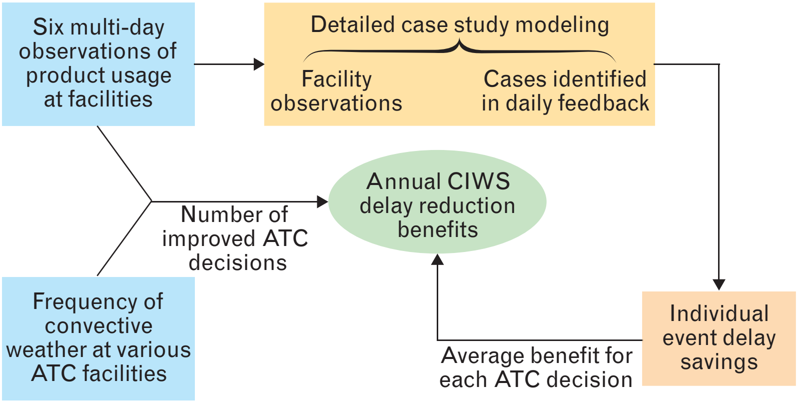

FIGURE 14. CIWS convective weather delay reduction benefits analysis [18].

The principal focus of the CIWS operational benefits assessment has been on assessing the ability of the CIWS users to achieve better air traffic management (ATM) in convective weather, as manifested by a greater number of capacity enhancing/delay reduction decisions. There are formidable challenges in assessing the benefits of a decision support system such as CIWS by comparing delays before and after the system has been deployed [18]. The major problems in use of delay statistics include normalizing the delays to account for (a) differences in the severity and frequency of convective weather events, (b) changes in the demand (e.g., replacement of turboprops by regional jets), (c) changes in FAA strategic planning (e.g., four to six hours in advance) using the CCFP, and (d) determining to what extent CIWS was responsible for changes in the delays as opposed to other weather and ATM systems. The normalization problem is particularly difficult when the NAS is operating in a highly nonlinear mode, as is discussed in the sidebar.

Given these difficulties in ascertaining the CIWS benefits by comparison of delay statistics, the CIWS operational benefits assessment has been accomplished by direct observations of FAA personnel using the CIWS to make various ATM decisions during convective weather events. The results of these observations were then used to determine the parameters of ATC models, which in turn were used to estimate the delay reduction benefits.

The basic assumption is that a weather product or ATM tool is useful only to the extent that it changes user decisions. Thus, we can analyze the various decisions that the users indicated were improved as a result of having access to the ATM decision support system under study.

The CIWS benefits analysis methodology, shown in Figure 14, calls for deployment of Lincoln Laboratory and FAA observers to a number of ATC facilities during periods when significant convective weather was expected. These extensive observation periods were treated as a sampling of the population of significant convective weather events at a facility. On the basis of both the real-time observations of users interacting with CIWS displays and the users' real-time verbal descriptions of ATC decisions they made while using the CIWS products, detailed statistics were generated on the number of times per significant convective weather day a given beneficial ATC decision was made by using CIWS products. Table 1 lists the frequency of various beneficial uses of the CIWS products per convective weather day for various ARTCCs [19]. The ARTCCs included in the analysis were as follows: Minneapolis (ZMP), Chicago (ZAU), Cleveland (ZOB), Washington (ZDC), New York (ZNY) and Boston (ZBW). The results shown in Table 1 are useful both as a metric for assessing the effectiveness of the CIWS and in computing the CIWS delay reduction benefits.

To convert the benefits frequency to delay reduction, we need estimates of the average number of convective impacts in an ARTCC and an estimate of the average beneficial decision within the ARTCC. The average annual frequency of a given benefit was estimated by using the long-term climatology for convective weather impacts within an ARTCC* together with the results of Table 1.

---

* Estimating the convective-weather impacts within an ARTCC requires a regional climatology that is quite different from the point climatology that is normally obtained. The regional climatology was determined by M. Robinson of Lincoln Laboratory by using the CIWS observations of ARTCC convective weather impacts together with the point-location thunderstorm reports for those same days to obtain an area climatology scaling factor (similar to earlier Lincoln Laboratory work [20] that determined terminal area climatology scaling factors).

Table 1. Frequency of Beneficial Uses of CIWS Products per Storm Day at Various Air Route Traffic Control Centers [19]

Benefit Category	*ZMP* *Minneapolis*	*ZAU* *Chicago*	*ZOB* *Cleveland*	*ZDC* *Washington*	*ZNY* *New York*	*ZBW* *Boston*
Keeping routes open longer and/or reopening closed routes earlier	2.0	1.2	5.2	5.2	1.5	2.8
Closing routes proactively	0.2	0.2	0.9	0.5	1.5	0.8
Proactive, efficient reroutes	3.8	2.2	4.9	4.8	2.0	2.8
Improved Ground Stop program management (shorter/fewer stops, ground stops avoided, more efficient use of ground stops)	0.5	0.5	1.1	1.2	0.5	1.2
Reduced Miles-in-Trail (MIT) restrictions (proactive management of routes in use)	0.3	0.3	1.3	2.2	0.5	1.0
Traffic directed through gaps	0..5	0.5	0.6	0.2	0.5	0.5
Improved management of weather impacts on terminal arrival and departure transition areas (ATA and DTA)	1.5	0.7	4.3	0.8	0.0	0.3
Improved Ground Delay program management	0.0	0.0	0.6	0.2	0.0	0.0
Greater departures during Severe Weather Avoidance Plans (SWAP)	0.5	0.5	0.9	0.3	0.0	0.3
Directing pathfinders	0.7	0.3	0.6	1.2	0.3	0.3
Interfacility, intrafacility coordination assistance	9.0	4.7	16.3	14.5	2.5	7.3
Improved safety	0.2	0.0	0.4	0.7	0.0	0.0
Reduced workload (including proactive impact mitigation planning)	5.8	2.5	11.7	11.5	3.0	5.0
Federal Aviation Administration (FAA) facility staffing assistance	1.2	0.2	0.7	0.2	0.0	0.7
Situational awareness	9.5	9.5	13.0	13.3	4.5	16.8
Number of days that Lincoln Laboratory/FAA observers present	6	6	9	6	4	6

The most labor intensive element of the analysis depicted in Figure 14 was the task of estimating the average delay reduction associated with a given benefits category for each ARTCC. This estimation was accomplished by detailed analysis of actual cases in which a given benefit was identified, using archives of CIWS weather products, flight tracks, and Aviation System Performance Metrics (ASPM) data. The cases analyzed were chosen randomly from all of the specific situations in which a given benefit was observed (that is, by randomly sampling the population of all benefits cases, we sought to generate a realistic estimate of the average benefit) [21]. Two significant results of this analysis are notable.

(1) The delay reduction for a given benefit (e.g., keeping routes open) varied significantly among different ARTCCs and within an ARTCC. For example, in the Cleveland ARTCC individual case benefits associated with *keeping routes open* varied from 189 hours to 4 hours. This variability was handled by using median

or mean case benefits (depending on the number of cases analyzed).

(2) Queues were common in the operational situations analyzed: eleven of the fifteen randomly chosen specific instances of *keeping routes open longer* involved reduction of queue delays. Also, the highest benefits generally were associated with queue delay reduction.

The case benefits results were then combined with benefits frequency data and the climatology of ARTCC impacts by convective weather to generate detailed estimates of the annual hours of delay reduction for two key ATM decisions: *keeping routes open longer/reopening closed routes sooner* and *proactive reroutes of aircraft* by FAA facilities. Based on the 2005 frequency of benefits decisions in Table 1, the 2003 case analysis [21], and the 2005 price of jet fuel, our first-order estimate of the delay reduction per year is over 90,000 hours with an airline direct operating cost savings in excess of $90M and a jet fuel savings in excess of 4.8M gallons.

Subsequent to the completion of the study [21], it has been necessary to create estimates of the CIWS convective weather delay reduction benefits for the period from 2004 to 2023, in which both demand and en route capacity are projected to increase. The approach presented in Figure 14 has facilitated such studies, since there were a number of concrete cases that could be reexamined to determine the impact of higher demand and/or greater effective capacity on delay reduction. Additionally, the case study results (e.g., an assessment of the number of aircraft utilizing a route when it was reopened sooner) could be independently analyzed by other organizations, since the case study results were based on analysis of archived data.

Analyzing the benefits for individual cases as part of the overall process shown in Figure 14 is very time consuming. The users identified how CIWS was used in real time to make a beneficial decision. However, the case analysis methodology then requires that we determine what would have happened to the affected aircraft if these same users had not had CIWS and therefore not taken the actions they did. Assessing these alternatives is quite challenging. For example, one of the *keeping routes open* decision cases (Reference 21, case B-7) found that CIWS allowed eleven aircraft, whose plans might otherwise have been altered, to fly along their desired route. For each of these, we must determine the possible alternative routing, considering the convective weather impacts on capacity of the NAS, and the demand on alternative routes at that time. Currently, the process of determining a reasonably optimal allocation of the aircraft between alternative routes and/or holding them on the ground, and determining the magnitude of the resulting delays, is largely manual. As a result, the number of specific cases analyzed was fewer than desired. Optimization tools

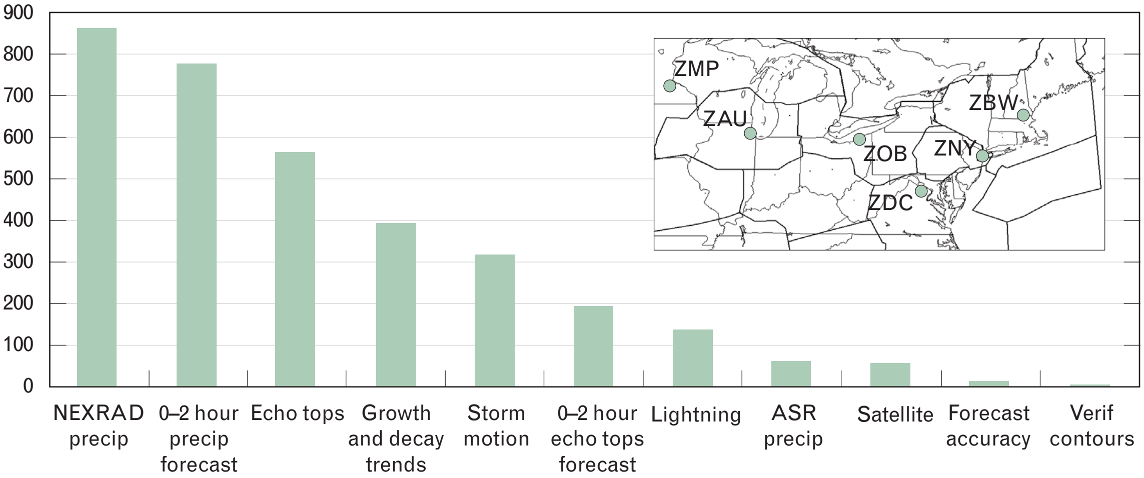

FIGURE 15. Observed use of the various CIWS products in the real-time observations at various Air Route Traffic Control Centers (ARTCC) in 2005 [19]. Note that the users relied heavily on both current weather state and forecast information. The echo tops forecast was new in 2005. On the basis of past experience with the usage of a new CIWS product after introduction, we anticipate that usage of the echo tops forecast product will increase in the future.

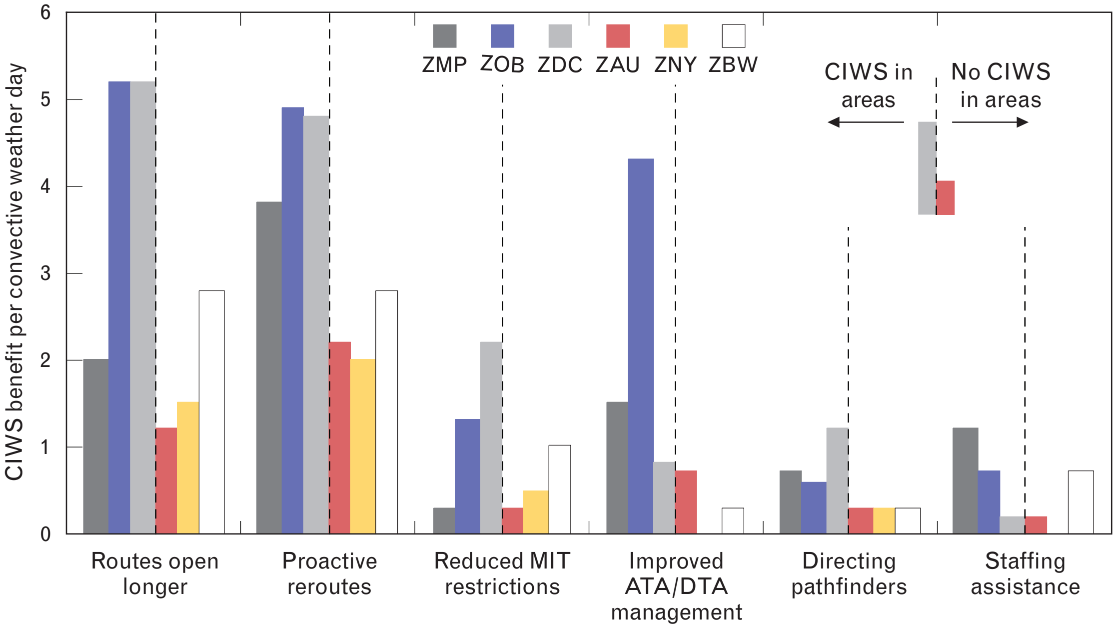

FIGURE 16. Frequency with which an improved Air Traffic Management (ATM) decision was made per thunderstorm day for the most significant CIWS en route delay reduction benefits categories and the FAA staffing assistance category [19]. ZMP was a new user of the CIWS, while ZBW had several very intensive users of CIWS in the Traffic Management Unit (TMU). Note that the experienced CIWS uses facilities (ZOB and ZDC) had a much higher rate of capacity-enhancing uses of CIWS than the other ARTCCs. The much higher rate of management of weather impacts on the ATA and DTA at ZOB reflects the availability of CIWS at the major terminals within ZOB.

of the type under development by Lincoln Laboratory [5] show promise as potentially effective methods of increasing the productivity for the analysis of convective weather cases.

Another significant benefit of the methodology used in the CIWS benefits study is the ability to better understand which products are used operationally, and the importance of product availability to key decision makers in determining the overall benefits. In Figure 15, we show the usage of the various CIWS products at the ARTCCs where intensive real-time observations were made in 2005. Figure 16 shows that the availability of the CIWS products to the ARTCC area supervisors* and to TRACON traffic managers is an important factor in the operational benefits achieved by using CIWS. We conclude from Figure 16, and the results in Table 1, that the frequency with which various capacity enhancing decisions are made per thunderstorm impact day can be significantly increased by providing key decision makers with the CIWS products.

In Figure 5, we noted that reducing the time to develop and implement an appropriate convective weather migration plan is an important factor in the overall productivity of the CIWS users. In Figure 17, we show the savings in the time to make and implement decisions with CIWS. As was the case with the frequency of capacity-enhancing benefits decisions, the availability of CIWS products to the area supervisors appears to be a significant factor in the time savings achieved.

Summary

Convective weather impacts on high-altitude en route airspace have become a major factor in weather-related delays. This increased sensitivity in high-altitude en route airspace is a consequence of significantly increased traffic loads at high altitudes due to the increased utilization of smaller terminals by low-cost

---

* The ARTCC airspace is divided into a number of 3D sectors, each of which has several controllers. Several adjacent sectors form an area that has an area supervisor whose responsibility includes insuring that a manageable number of aircraft will be sent into the area. The traffic flow managers for the ARTCC must coordinate with the various area supervisors in their ARTCCs in developing and implementing convective weather impact mitigation plans.

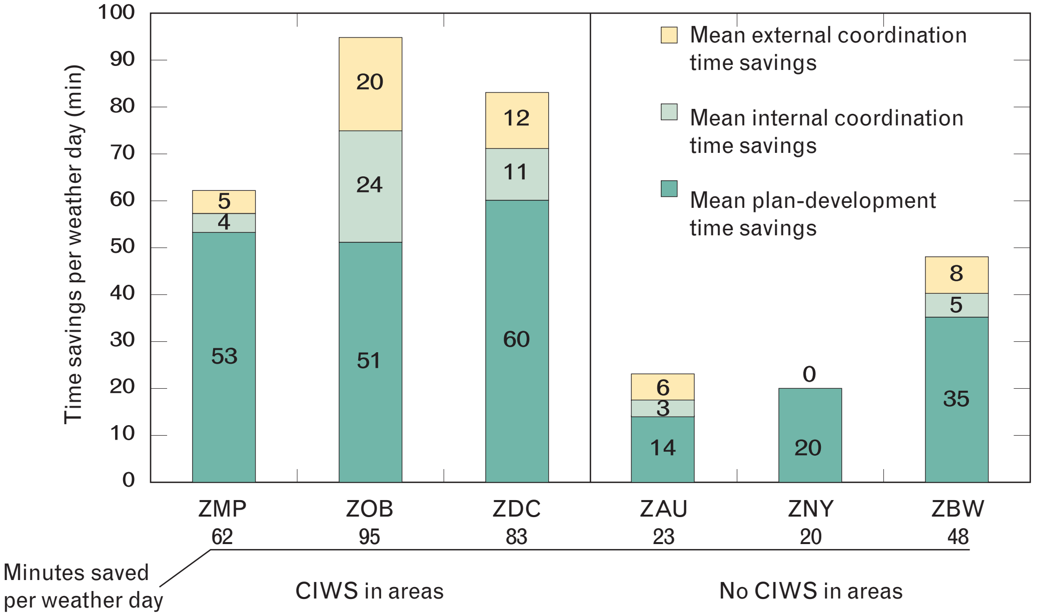

FIGURE 17. Estimated TMU time savings at various ARTCCs per convective weather day as determined from the 2005 Lincoln Laboratory/FAA study [19]. The time savings is the convective weather impact mitigation planning time savings attributed to CIWS.

carriers, the transition from turboprops to regional jets, and the expansion in the use of business jet aircraft. The management of convective impacts in the highly congested airspace in the northeast quadrant of the United States is particularly challenging because rerouting aircraft away from the regions of severe convective weather can cause queues to arise in many other regions of the NAS.

The CIWS takes advantage of the high density of existing FAA and NWS sensors in the congested en route corridors and the FAA-funded research conducted on thunderstorm evolution to dramatically improve the accuracy and timeliness of the storm severity information. The CIWS also provides state-of-the-art accurate, automated, high-resolution animated 3D forecasts of storms (including explicit detection of storm growth and decay). These CIWS *tactical* traffic flow management products complement the longer-term (four to six hour) national CCFP forecasts that are also needed for overall flight planning and traffic flow management.

The CIWS products generated by a demonstration system have been used for real-time decision making in the Northeast and Great Lakes corridors since 2001. Real-time observations of FAA decision making during convective weather at ARTCCs have shown that the CIWS enables the FAA users to achieve more efficient tactical use of the airspace, reduce traffic-manager workload, and significantly reduce delays relative to what would have occurred in the absence of CIWS. The significant improvement in the frequency of making capacity-enhancing decisions when area supervisors in the ARTCCs and traffic flow managers at TRACONs have access to the CIWS products has demonstrated the importance of providing common situational awareness among all of the key decision makers in the NAS.

The FAA is in the process of determining how the CIWS technology can best be incorporated into the NAS. Lincoln Laboratory's research work on a real-time CIWS processing architecture that would facilitate the expansion of the CIWS to the lower forty-eight states is expected to be a key element of the technology-transfer package for the CIWS capability.

Active research is under way to determine how the CIWS products should be enhanced to better support the development of integrated convective weather decision support systems. These support systems must assist in the development and execution of ATM plans that more fully utilize the available en route capacity during severe convective weather. A companion paper [5] in this issue of the *Journal* describes the use of the CIWS products as an input to capacity models and strategy generation tools.

Such integrated applications will require modifications to the current CIWS products so that the integrated system can achieve its objectives. For example, the method used to represent forecast uncertainty for integrated applications may well need to be application specific. Lincoln Laboratory's initial operational utility assessments of the RAPT have shown the current models for pilot preferences in avoidance of convective weather also need to be refined. Such refinements to the models for pilot preferences are likely, in turn, to necessitate the development of alternative storm structure representations. Hence it is important that the overall CIWS system architecture be highly flexible and easily expanded.

Acknowledgments

This work was sponsored by the Federal Aviation Administration under Air Force Contract FA8721-05-C-0002. The majority of the Lincoln Laboratory Weather Sensing group, which consists of more than sixty people, has contributed to the CIWS development described in this article. However, the authors would like to make special mention of the contributions of Michael Robinson to the section on CIWS benefits and Oliver Newell to the section on the demonstration system.

REFERENCES

1. T.H. Fahey III and D. Rodenhuis, "Continual Evolution of CCFP-User Needs for Extended Range Prediction," *Proc. 11th Conf. on Aviation, Range, and Aerospace Meteorology, Hyannis, Mass., 4–8 Oct. 2004.*
2. A.N. Sinha, "FAA/Industry Spring 2000 Initiative," presentation at *All Weather Decision Support Planning Workshop, 31 May 2000,* www.mitrecaasd.org/library/documents/3_Sinha.pdf.
3. M. Huberdeau and J. Gentry, "Use of the Collaborative Convective Forecast Product in the Air Traffic Control Strategic Planning Process," *J. Air Traffic Control,* Apr.–June 2004, pp. 9–14.
4. M.P. Kay, J.L. Mahoney, and J.E. Hart, "An Analysis of CCFP Forecast Performance for the 2005 Convective Season," *12th Conf. on Aviation, Range, and Aerospace Meteorology, Atlanta, 29 Jan–2 Feb. 2006.*
5. J.E. Evans, M.E. Weber, and W.R. Moser, "Integrating Advanced Weather Forecast Technologies into Air Traffic Management Decision Support," in this issue.
6. D.A. Rhoda, E.A. Kocab and M.L. Pawlak, "Aircraft Encounters with Thunderstorms in Enroute vs. Terminal Airspace above Memphis, Tennessee," *10th Conf. on Aviation, Range, and Aerospace Meteorology, Portland, Oreg., 13–16 May 2002.*
7. R. DeLaura and J.E. Evans, "An Exploratory Study of Modeling Enroute Pilot Convective Storm Flight Deviation Behavior," *12th Conf. on Aviation, Range, and Aerospace Meteorology, Atlanta, 29 Jan–2 Feb. 2006.*
8. H. Davison and R.J. Hansman, "Identification of Inter-Facility Communication and Coordination Issues in the U.S. Air Traffic Control System," MIT International Center for Air Transportation paper ICAT 2001-11-21, http://icat-server.mit.edu/Library.
9. M. Weber, J. Evans, M. Wolfson, R. DeLaura, B. Moser, B. Martin, J. Welch, J. Andrews, and D. Bertsimas, "Improving Air Traffic Management During Thunderstorms," *12th Conf. on Aviation, Range, and Aerospace Meteorology, Atlanta, 29 Jan.–2 Feb. 2006.*
10. J.E. Evans and D.M. Bernella, "Supporting the Deployment of the Terminal Dopper Weather Radar (TDWR)," *Linc. Lab. J.* 7 (2), 1994, pp. 379–398.
11. J.E. Evans and E.R. Ducot, "The Integrated Terminal Weather System (ITWS)," *Linc. Lab. J.* 7 (2), 1994, pp. 449–474.
12. M. Robinson, J.E. Evans, and B.A. Crowe, "En Route Weather Depiction Benefits of the NEXRAD Vertically Integrated Liquid Water Product Utilized by the Corridor Integrated Weather System," *10th Conf. on Aviation, Range, and Aerospace Meteorology, Portland, Oreg., 13–16 May 2002.*
13. J.E. Evans, K. Carusone, M.M. Wolfson, B. Crowe, and D.J. Smalley, "Multi-Radar Integration to Improve en Route Aviation Operatons in Severe Convective Weather," *19th Conf. on Interactive Information Processing Systems (IIPS), Long Beach, Calif., 8–13 Feb. 2003.*
14. M.M. Wolfson and D.A. Clark, "Advanced Aviation Weather Forecasts," in this issue.
15. D. Becker, "Beyond MPI," *Linux Mag.,* Nov. 2005, pp. 28–30.
16. J.P. Snyder, "Map Projections: A Working Manual," U.S. Geological Survey Professional Paper 1395, U.S. Government Printing Office, Washington, 1987, pp. 182–190.
17. M.Q. Yang, R.E. MCcGrath, and M. Folk, "HDF5—A High Performance Data Format for Earth Science," *21st Int. Conf. on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology, San Diego, 9–13 Jan. 2005.*
18. J. Evans, M. Robinson, and S. Allan, "Quantifying Convective Delay Reduction Benefits for Weather/ATM Systems," *Air Traffic Control Q.* **14** (1), 2006, pp. 69–93.
19. M. Robinson, J. Evans, and T. Hancock, "Assessment of Air Traffic Control Productivity Enhancements from the Corridor Integrated Weather System (CIWS)," Lincoln Laboratory Project Report ATC-325 (June 2006).
20. P. Bieringer, D.W. Miller, and D. Meyer, "A Refinement of Thunderstorm Climatology for the Terminal Radar Control Airspace," *8th Conf. on Aviation, Range, and Aerospace Meteorology Society, Dallas, 10–15 Jan. 1999.*
21. M. Robinson, J. Evans, B. Crowe, D. Klingle-Wilson, and S. Allan, "Corridor Integrated Weather System Operational Benefits 2002-3: Initial Estimates of Convective Weather Delay Reduction," Lincoln Laboratory Project Report ATC-313 (9 Apr. 2004).

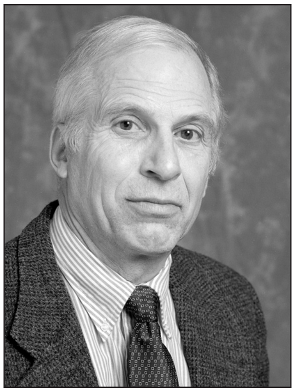

JAMES E. EVANS
is a senior staff member at Lincoln Laboratory who is responsible for initiating and contributing to research programs in improved aviation weather decision making. He joined the Laboratory in 1967 and commenced work in air traffic control in 1970. He has led the Lincoln Laboratory programs to develop the Terminal Doppler Weather Radar (TDWR), the Integrated Terminal Weather System (ITWS), and the Corridor Integrated Weather System (CIWS). His current research includes improving Air Traffic Control and airline decision making to mitigate the impacts of adverse weather, developing integrated weather air traffic management (ATM) systems, and assessing operational benefits for deployed systems. He was presented with outstanding paper awards at the last two USA/Europe ATM R&D Symposia. He was honored with an MIT Lincoln Laboratory Technical Excellence Award in 2002. He received the S.B., S.M., Engineer and Ph.D. degrees in electrical engineering from MIT.

ELIZABETH R. DUCOT
is an associate leader in the Weather Sensing group, where she is currently serving as the program manager and project engineer for both ITWS and CIWS. She received a B.A. degree in physics from Smith College. Her graduate studies in computer science were conducted at the University of California. In 1988, she transferred to Lincoln Laboratory from MIT, where her research focused on large-scale distribution systems. Shortly after joining Lincoln Laboratory, she received the IEEE Control Systems Society Distinguished Member Award.