FAA TACTICAL WEATHER FORECASTING IN THE UNITED STATES NATIONAL AIRSPACE*

W.J. Dupree, M.M. Wolfson, R.J. Johnson Jr., R.A. Boldi, E.B. Mann, K. Theriault Calden, C.A. Wilson, P.E. Bieringer, B.D. Martin, H. Iskenderian

MIT Lincoln Laboratory,
Lexington, Massachusetts

1.0 INTRODUCTION

During the past decade the United States (US) National Airspace System (NAS) has become significantly more congested. When weather perturbs the NAS, already operating at near or maximum capacity, the air traffic system can become chaotic resulting in massive delays that can propagate throughout the country. Studies have shown that roughly 75% of these delays are a result of convective weather. As a result, the US Federal Aviation Administration (FAA) has assembled a convective weather product development team (PDT) as part of the Aviation Weather Research Program (AWRP), made up of a collaboration between several institutions to help in air traffic delay reduction. These institutions include the Massachusetts Institute of Technology Lincoln Laboratory (MIT LL), the National Center for Atmospheric Research (NCAR), the National Oceanic and Atmospheric Administration (NOAA) Forecast Systems Laboratory (FSL), and several universities.

Air Traffic control planning is primarily split between two planning methodologies: strategic (2-6 hour forecast lead time) and tactical (0-2 hour forecast lead time). Strategic planning occurs daily and is often only useful when the weather is highly predictable. When convective weather impacts both enroute and terminal settings, the airspace capacity drops below a sustainable threshold and the planning switches to tactical. In this case, it is critical that reliable and accurate short-term 0-2 hour forecasts of precipitation and storm tops be available for tactical maneuvering in both terminal and enroute airspace.

*This work was sponsored by the Federal Aviation Administration under Air Force Contract No. FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

The development of the Tactical Convective Weather Forecast (CWF) algorithm grew out of the need for better situational awareness during weather events and has been primarily driven by the needs of the air traffic management community and hence has evolved as the user needs have become better understood. We will briefly review the development of the CWF system, however a more in depth review is given in Wolfson et al. (2004-a).

The “corridor” between New York and Chicago covers one of the busiest air traffic regions in the world (Figure 1). In 2001, a study of weather delays in the New York Terminal Radar Approach Control (TRACON) (Allan et al., 2001), showed that two-thirds of delays were caused by convective weather in the terminal areas and one-third of the delays were due to weather in enroute airspace. This study showed that terminal operations were highly coupled to the enroute operations in the corridor. As a result, in 2001 a radar-mosaic version of CWF was released for use by the enroute controllers at the Air Route Traffic Control Center and the US Air Traffic Control System Command Center. This concept exploration demonstration weather system is called the Corridor Integrated Weather System (CIWS) of which the CWF is a key component. CIWS has proven to be an extremely valuable tool for the air traffic community (Robinson et al., 2004).

To extend the forecast out to 2 hours, a few new concepts were added in 2002. These included the addition of the weather classification algorithm and the multiscale tracking technology. In parallel, a growth and decay trending algorithm was being developed. The weather classification algorithm (weather type) provided information by which precipitation trends could be applied to meteorological features of interest. Fuzzy function methods allowed for specific storm models to be constructed that contained growth and decay

trends, weather type, storm phase and time horizon of interest. This technology was added in 2003 to the terminal and corridor versions of CWF (ITWS and CIWS respectively).

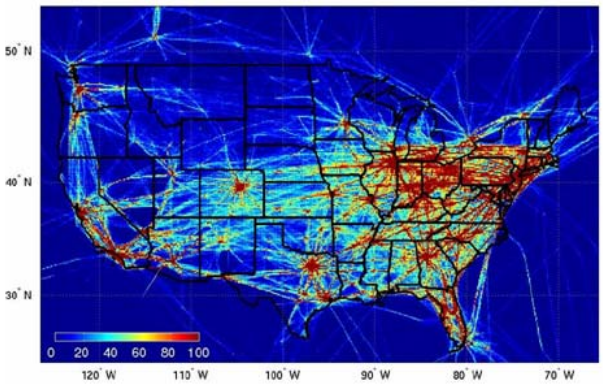

Figure 1: 24-hr traffic counts over the continental United States on a clear weather day (12-13 September 2002) showing the high traffic density in the Northeast Corridor, where several major terminals are located. The strong coupling between terminal and enroute delay in this area motivated the 2-hour forecast horizon for CIWS, to help improve tactical management of the congested enroute traffic.

Even though the CWF was primarily developed for use on convective weather, our aviation users continued to use the precipitation forecast during the winter season, primarily for situational awareness. In order to provide more utility of the forecast in the winter, a new VIL color scale was introduced that had greater resolution in weaker precipitation levels. See Table 1 for the correspondence of levels in Vertical Integrated Liquid Water (VIL) products. The winter mode display split the color level 1 into three categories, adding finer resolution required to resolve weak winter features like snow bands. The winter product was released in 2004.

Evans et al. (2003) conclude that one of the key problems that needs to be addressed to help improve the capacity of the NAS is the introduction of a storm echo tops structure. In addition to the VIL mosaic, CIWS provides a high resolution mosaic of echo tops with 1000 ft (305 m) vertical resolution and 2 km horizontal resolution. This product proved to be one of the most valuable to users in the 2002-2003 seasons (Robinson et al., 2004). As a result, in the spring 2005 an Echo Tops Forecast (ETF) was added to CWF. This forecast is coupled with the CWF technology and uses much of the existing methodology.

Table 1
Correspondence of Levels in VIL Products.

Color Level	dBZ	VIL (kg/m^2)	Scaled VIL
1a*	-10.0	0.03	5
1	18.5	0.15	16
1b*	22.7	0.28	34
1c*	27.4	0.53	59
2	30.2	0.75	74
3	41.1	3.54	133
4	46.2	7.08	160
5	50.1	12.2	181
6	57.2	32.3	219

**These color levels were added by MIT LL to further divide the level 1 category for winter.*

This paper describes the Tactical 0-2 hour Convective Weather Forecast (CWF) algorithm developed by the MIT LL for the FAA. We will address the algorithm and focus on the key scientific developments. Future directions will also be discussed.

2.0 FORECAST DISPLAY

The forecast display is designed to graphically present predicted future weather in a manner that is easily understandable and requires no further interpretation. The CWF forecast is composed of two forecast products: the Precipitation Forecast (VIL), which has two modes (standard and winter), and the Echo Tops Forecast. The Precipitation Forecast is a forecast of the intensity and location of VIL given in operationally significant levels. The Echo Top Forecast is a forecast of the maximum altitude where the radar reflectivity drops below 18 dBZ and is a proxy for the cloud “precipitation” top. The echo top forecast levels depicted are of significant operational use. Each of these products is displayed to the users in separate windows but the windows can also be displayed side by side on the same situation display. We will describe these windows separately.

2.1 VIL Forecast Display

The VIL forecast can be displayed to users in two ways. The first display method is an animated loop that shows past weather, transitions to the forecast after the current time,

and then increments the forecast out to the maximum time horizon. For the CIWS, the length of the past weather loop is 60 minutes, the maximum forecast time horizon is 120 minutes and the forecast increment is 15 minutes. Satellite data are selectable as the background on the past weather portion of the loop.

For the second display method, the forecast can be displayed as contours overlaid on the current weather. The past weather VIL uses the six standard National Weather Service (NWS) color levels.

Since the CWF is a continuous automated system and must operate in winter conditions, a method to convey more information about snow bands containing moderate to heavy snow is provided. There are two choices for forecast display: standard and winter. In standard mode the 3 levels represent color levels 1, 2, and 3 (Figure 2). The solid yellow represents high intensities of level 3 and above, the lighter grey represents moderate intensities of level 2, and the dark grey represents the lowest intensities of level 1.

When the winter forecast mode is selected, the lower VIL forecast values appear in much greater detail (Figure 2). The winter color scale replaces and extends the standard level 1 precipitation interval with three new intervals (1a, 1b and 1c) (see Table 1). The precipitation levels 2 through 6 are unchanged. The forecast portion of the animation loop will also change in color when the mode is changed to winter. With the winter color scale, level 1c and above is forecast at the high intensity, and this level is displayed in solid green to alert the users that the winter forecast is being displayed. Additionally, level 1b is displayed as light grey and level 1a is dark grey.

The forecast also displays a forecast accuracy scoring metric (Figure 3). This is a "User Confidence Score" and provides past performance of the algorithm based on the location of level 3 for standard and level 1c for winter modes. See section 5 on Forecast Performance.

2.2 Echo Tops Forecast Display

The ETF is displayed as an animated loop from 60 minutes in the past to 120 minutes in the future in 15 minute increments (Figure 3). The past echo tops are shown in 5 kft (1.5 km) increments from 0 to 50+ kft (0 to 15.2+ km) above mean sea level, while four colors are used to indicate the forecast heights. [<25, 25, 30 and >= 35 kft (<7.6, 7.6, 9.1, 10.7 km)] These colors are grey, dark purple, purple and light purple respectively. The forecast accuracy numerical scores indicate the past performance of the ETF based on predicting areas of 30 kft (9.1 km) or higher echo tops.

3.0 FORECAST SYSTEM

It has long been known in nowcasting that scaling of storm physics is key to predicting future behavior. Wilson (1966) showed that spatially large scale features (lines) mapped to longer time scale persistent features whereas spatially small scale features (airmass cells) were typically short lived. Contributions by Marwitz (1972), pointed out that multicellular storms often grew and moved with the mean wind whereas the initiating disturbance moved with a distinctly different motion coincident with the large scale forcing. Classification of the cellular nature of storms has been made by Weisman and Klemp (1986), where three distinct groupings were made: airmass cells, multicellular (lines and squall lines), and super cells. In such a classification, airmass cells are typically short lived and it is difficult to predict where new cells will regenerate after collapse. Super cells and line storms, on the other hand, tend to grow new cells along a preferred flank and produce a new grouping of cells. More recent work has shown that scale extraction of radar imagery can lead to better tracking predictions of large scale motion (Bellon and Zawadzki, 1994; Wolfson et al., 1999). Attempts to predict motion on multiple scales have been made by Seed and Keenan (2001), Dupree et al. (2002), and Lakshmanan et al. (2003).

The CWF system is built upon the underlying assumption that weather can be classified based on scale and texture. This classification is used to extract motion and trending behavior.

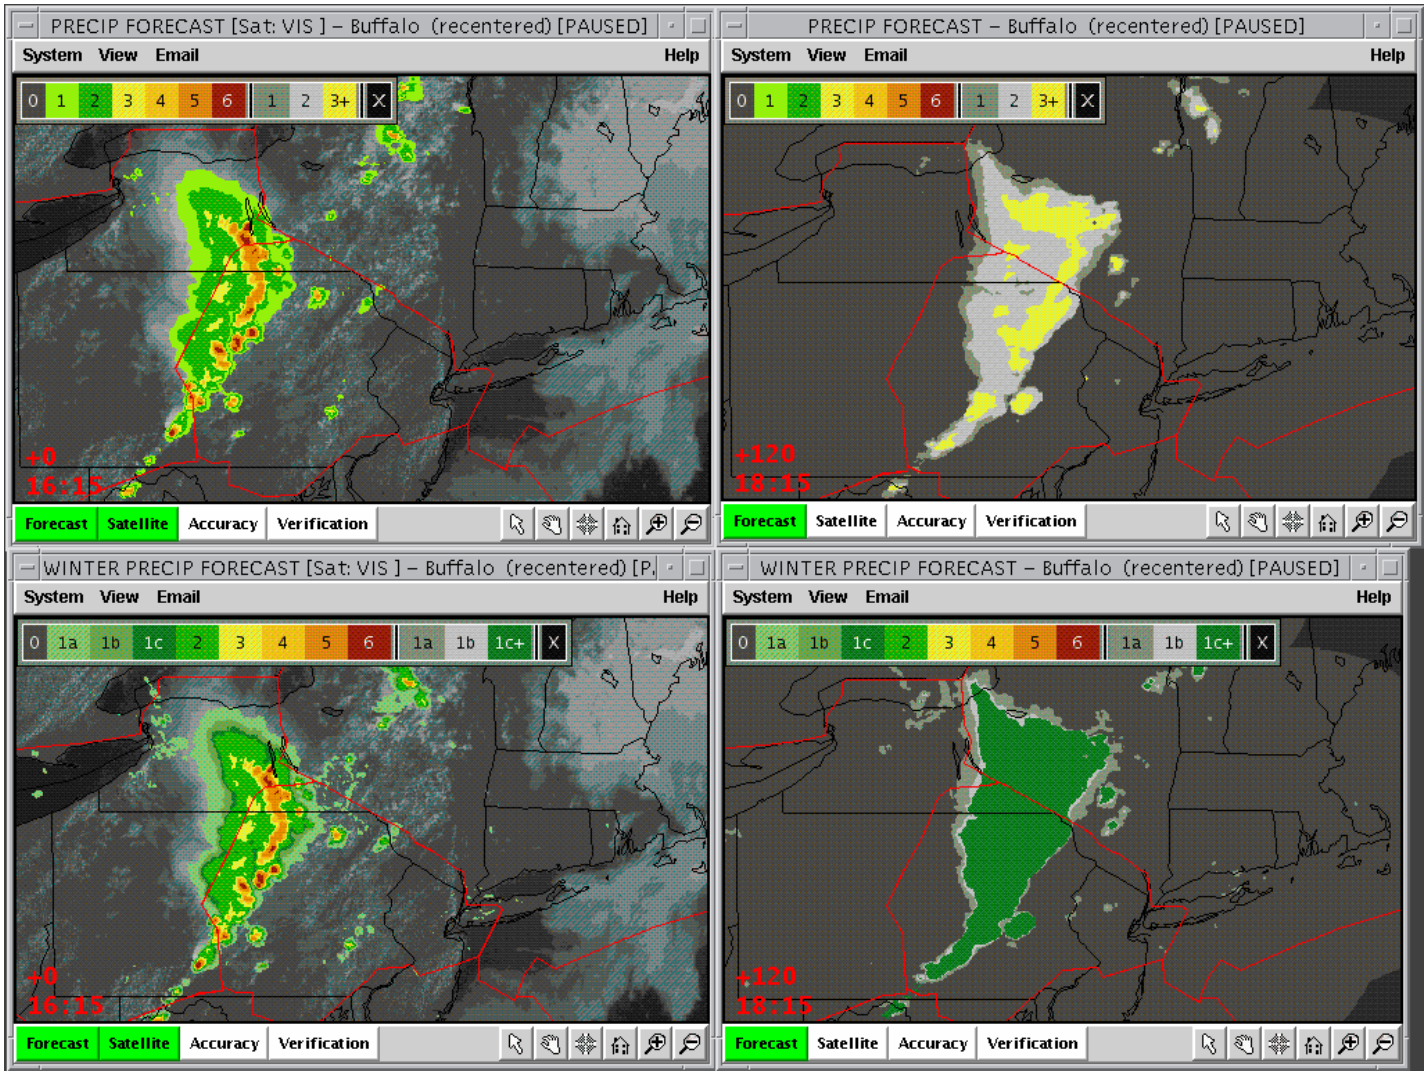

Figure 2: CIWS Forecast Display Standard and Winter Modes: The upper left figure is the current VIL in standard mode with the six NWS color levels overlaid on top of the satellite image. The upper right shows the standard mode VIL forecast loop. The lower left figure is the current VIL in winter mode with the satellite image. The lower right shows the winter mode looping VIL forecast.

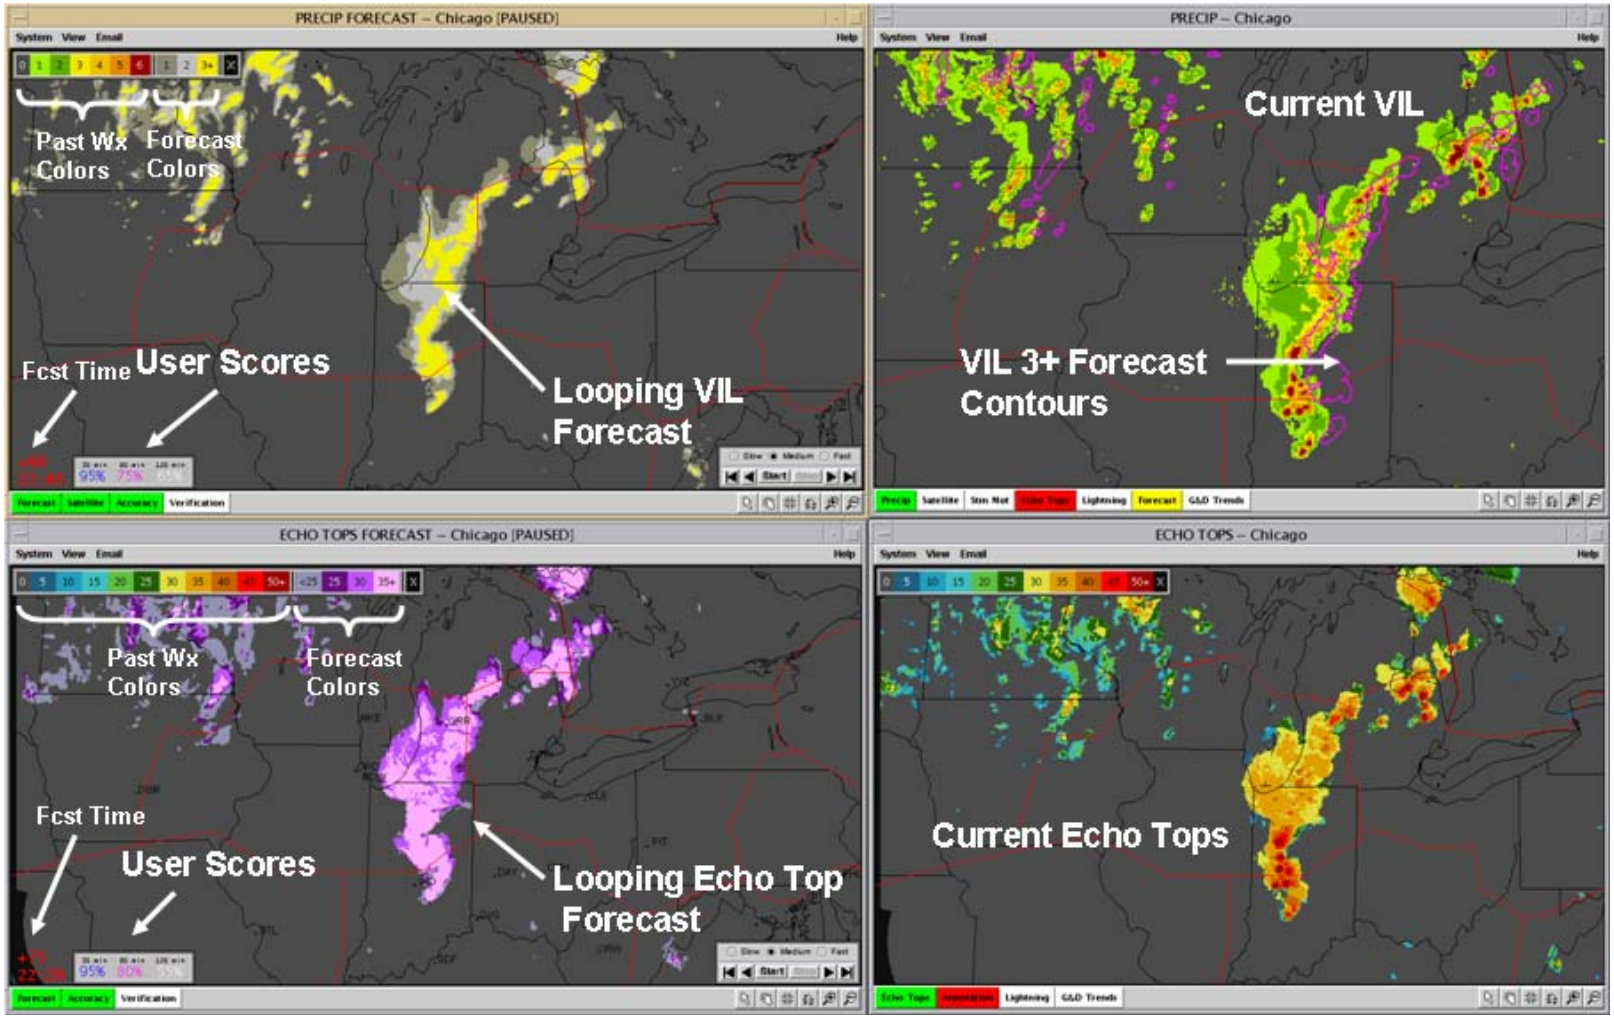

Figure 3: CIWS Forecast Display: The upper left window in the looping VIL forecast. The lower left is the looping echo tops forecast. The upper right is the current VIL with the 60 minute forecast level 3 contours overlay. The lower right side is the current echo tops.

Furthermore, the classification can be used to construct heuristic models of storm behavior where this scaling and trending behavior are relevant. In the next section we will follow the CWF processing chain and discuss each section in terms of the underlying scientific rationale.

4.0 PROCESSING OVERVIEW

In this section we discuss the main areas of processing for the Forecast System. Figure 4 gives an overview of the processing chain. This figure depicts four main processing blocks in the system:

1) Per radar processing (pink): Processing on a single radar image.
2) Mosaicing processing (blue): Mosaicing radar data onto a large grid encompassing the domain of interest.
3) Full grid processing (green): Feature extraction, advecting, and model application.
4) Post-processing (grey): Scoring, contouring, preparing data for display, and display of the grids.

4.1 VIL

We selected VIL as the aviation precipitation map and forecast product for two reasons: one statistical and one meteorological. The first reason is that VIL is a vertical integral or a central measure and is not overly susceptible to outlier observations such as ground clutter or bright band contamination, as shown in Figure 5. The second reason for choosing VIL is that it is indicative of the amount of water aloft, and at higher values it is related to the strength of the updrafts and hence storm intensity. This makes VIL highly relevant to aviation decisions regarding areas to avoid.

4.2 Radar Data Quality Editing

Before image processing occurs it is critical that each data set undergo a series of data quality editing (DQE) steps. Two types of radar data are ingested into the system: data from the United States network of NEXRAD radars and data from the neighboring Canadian network of SIGMET radars. The NEXRAD data are edited as follows. First, a strong point target detector (clutter editing) is applied in the NEXRAD Radar Data Acquisition (RDA), followed by application of the Open Radar Products Generator (ORPG) DQE module. The ORPG applies FAA–specified low velocity and low spectrum width editing, constant power editing (including, test pattern, sun strobe and star burst editing) and velocity dealiasing. At this stage, the data are converted into VIL and echo tops and sent out as NEXRAD Level III data products. Once data are available we apply the MIT LL data quality editing routine. This routine is a lint detection algorithm that detects and removes small, thin and isolated regions (Wolfson et al., 2004-b).

For the SIGMET radars from the Canadian Radar network we edit out returns under a 0.5 km floor and apply a strong point target detection algorithm to the reflectivity data. Next, the radar data are converted to high resolution VIL and echo tops using the same methods as the NEXRAD ORPG.

4.3 Scale Separation and Tracking

The observation that line storms are persistent and follow the envelope motion and not the cellular motion was the driving force for developing the Growth and Decay Storm tracker (Wolfson et al., 1999). A fundamental and critical step is the extraction of scales of interest from the precipitation images and has come to be known as the multiscale tracking method. In this method, the VIL images are correlated with a series of elliptical-mean filters using an array of orientations. Then taking the point-wise maximum of the resulting filtered images, one can construct an envelope or “line” detector. With application of the cross-correlation tracking (XCT) method (Chornoboy et al., 1994) to successive images one can deduce the motion of the propagating line. Cartwright et al. (1999) showed that the most optimum settings, i.e. the filter that produced the best tracking results, was a 13x69 km elliptical filter rotated at 5 degree increments. Furthermore, the optimum scale used to track airmass cells was found to be a 13 km diameter circular mean filter. We chose to split the tracking into not just distinct scales but distinct phenomenon.

Early analysis of our envelope forecasting algorithm showed that occasionally the tracking motions were in error due to the over-rejection of cross-correlation vectors (Theriault et al., 2000 and 2001). Early versions of the XCT allowed one to compare a tracking vector against the scene average or the full NEXRAD domain. If the vector did not fall between parameterized limits (+/- 75 degrees) then the vector was removed. However, this average was not

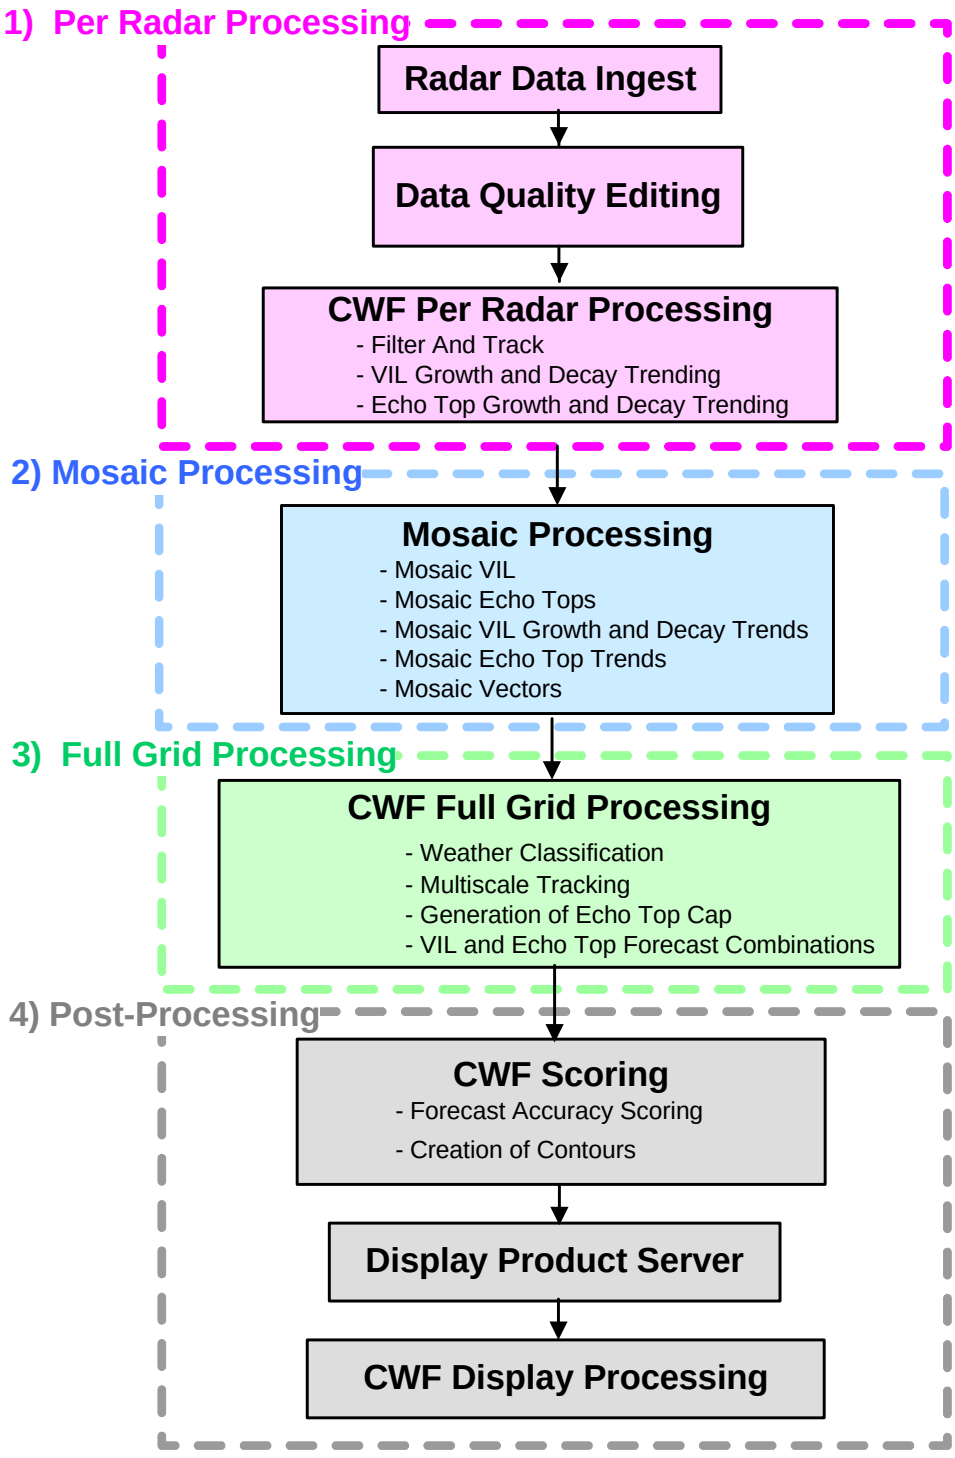

Figure 4: CWF Algorithm Processing

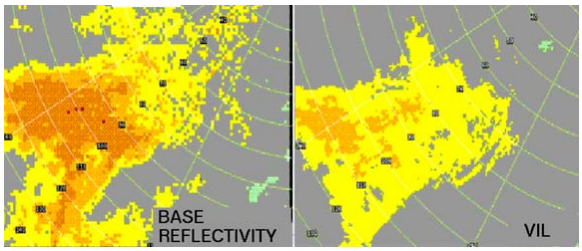

Figure 5: A typical "Bright-Band" situation. This occurs when thin layers of precipitation at the melting level have enhanced reflectivity.

indicative of some local motions and important motion data was being rejected, so XCT was modified to determine the average motion locally in spatially smaller boxes. Currently we use a restriction of +/- 45 degrees compared to the average motion in a 56 km box size for envelope motion and a 28 km box size for cell motion. This change provided much better motion estimates in airmass, meso-lows, and hurricanes situations.

Additionally, the envelope motion vectors are temporally averaged in an Eulerian sense where the most recently updated vector is weighted at 30% and the previous running average is weighted at 70%. For cell tracking, the newest vectors are weighted at 90% and the previous average at 10%, since cells are typically short lived. Correlations of the filtered images are done on the same temporal scale of 12 minutes for envelope and cell tracking.

The echo top data basically provides a measure of the top of the VIL image. For this reason, we chose to use the VIL tracking motion as a proxy for the echo top tracking motion.

4.4 VIL and Echo Top Trending

Trending of precipitation rates in radar data has received much interest in the past. Wilson (1998) and Tsonis (1981) concluded that essential physical processes that dictate the change in rainfall with time are not necessarily observable in the past history of a particular echo development. In an attempt to alleviate this problem the forecast generation algorithm uses the weather classification image in conjunction with the trend interest images to determine the expected behavior of a storm cell over time. The weather classification image dictates how to model the trend information for a particular storm type.

The Growth and Decay Trend (GDT) algorithm consists of a suite of image processing feature detectors that produced interest images containing growth and decay information on a per pixel basis. Both the VIL and echo tops radar data are processed using similar methodology, however, there are differences in the fundamental output of each data type. Common processing will be discussed as will fundamental differences.

The VIL Growth and Decay Trend (VGDT) algorithm ingests VIL and the envelope and cell vectors generated in the scale separating and tracking section. These detectors create short-term trends, long-term trends, isolated cell, and boundary growth images. The resultant interest images can be used to (1) modify the VIL forecast values, (2) add new regions of interest to the VIL forecast and (3) provide subtype classifications in the weather classification algorithm.

The Echo Top Growth and Decay Trend (ETGDT) algorithm is similar to the VGDT algorithm in the initial stages. The echo top trend process requires echo top images and cell vector images as input. The output from the ETGDT is a short-term trend image where the pixel values represent a physical rate in units of meters/min. This image is used to modify the echo tops forecast.

4.4.1 Trending

This section covers detectors that are common to VIL and echo top trending. The first feature detector to be executed is the VIL variance detector. This detector calculates the standard deviation under a 15x15 km circular kernel over the entire VIL image. The standard deviation values returned are used by subsequent feature detectors to identify possible convective areas. High standard deviation values are indicative of convective areas while low values indicate stratiform areas.

The fundamental image processing step for the trend feature detectors is the differencing of prior images with the current image. A prior image is advected to the current time with a set of vectors which capture the desired scale of motion. The cell vectors are used for the short-term VIL and echo top trend images while the envelope vectors are used for the long-term VIL trend image. The cell vectors capture the

motions of individual cells within a storm complex while the envelope vectors capture the motions of the large scale storm structure. Once the prior image is aligned in time with the current image the two images are subtracted. This difference image represents the change in VIL or echo tops in the given time period.

The short-term trend image is generated from the difference between two consecutive VIL or echo top images. Typically this time difference is the radar scan update rate of about 5-6 minutes (Figure 6).

The long-term trend interest images are generated similarly by differencing two VIL images separated by a wider time difference. Alignment of the prior time images differs between the short-term and long-term trend processes. A quasi "Lagrangian" advection technique is used in the long-term process. Since the advection of the prior image spans several radar updates a new set of vectors is used at each radar update timestep. A history of vectors is saved for each radar update. As the prior image is advected to the current time the new vectors are used at each timestep. A difference interest image is created for each radar update.

A single difference image contains high frequency noise. By temporally averaging several of these difference images together much of the high-frequency noise is mitigated. Fortunately, this averaging still allows the persistent trends to remain.

Once the averaged VIL and echo top trending images have been calculated, the processing diverges for the two input images. We will discuss the final processing steps for VIL and echo tops separately.

4.4.2 VIL Trending

Filtering of the averaged VIL difference images is done using Functional Template Correlation (FTC; Delanoy et al., 1992) and produces an interest image that represents the short-term and long-term trends of VIL.

The short and long-term VIL trend images each have one final processing step performed prior to application in the VIL forecast. The trend images have been scaled such that the interest values range from 0-255. The value of 128 is indicative of a neutral trend, or no change in VIL, while values below 128 indicate the increasing magnitude of decay. Values above 128 indicate increasing magnitudes of growth. An important image called the trend modified VIL image is used to generate the VIL forecast. These images are created by making a modification to the VIL image using the growth and decay trend values and parameterized scaling factors. This modification creates an interest image by applying the magnitude of the trend interest image to the current VIL pixel value. That is,

$$V' = V + \delta V \qquad (1)$$

$$\delta V = TD_f \qquad T < 128 \qquad (2)$$

$$\delta V = TG_f \qquad T > 128 \qquad (3)$$

$$\delta V = 0 \qquad T = 128 \qquad (4)$$

T = pixel value of trend image
G_f = growth factor trend image
D_f = decay factor trend image.

Here V' represents the modified growth and decay field. The resultant interest image is the trend modified VIL interest image. These feature detectors output the trend modified short-term and long-term interest images in 0-255 scaled VIL units.

Two additional feature detectors are used in the VGDT algorithm only: the boundary growth feature detector and the isolated cell feature detector. By utilizing the short-term VIL trend, standard deviation interest image, and the current VIL image, the boundary growth feature detector returns an interest image that represents regions of linearly aligned growth (see Figure 7). Specifically, the detector assigns high interest to areas that show thin bands of moderate to strong growth surrounded by no radar returns. These are areas where rapid convective development occurs not only in intensity but in aerial extent. Additionally, the detector outputs the orientation of the major axis of the interest region. The boundary growth feature detector tries to capture the initiation phase of the storm from the growth of the low-moderate strength radar returns. The feature detector is tuned to capture only the initial growth phase of the storm cells.

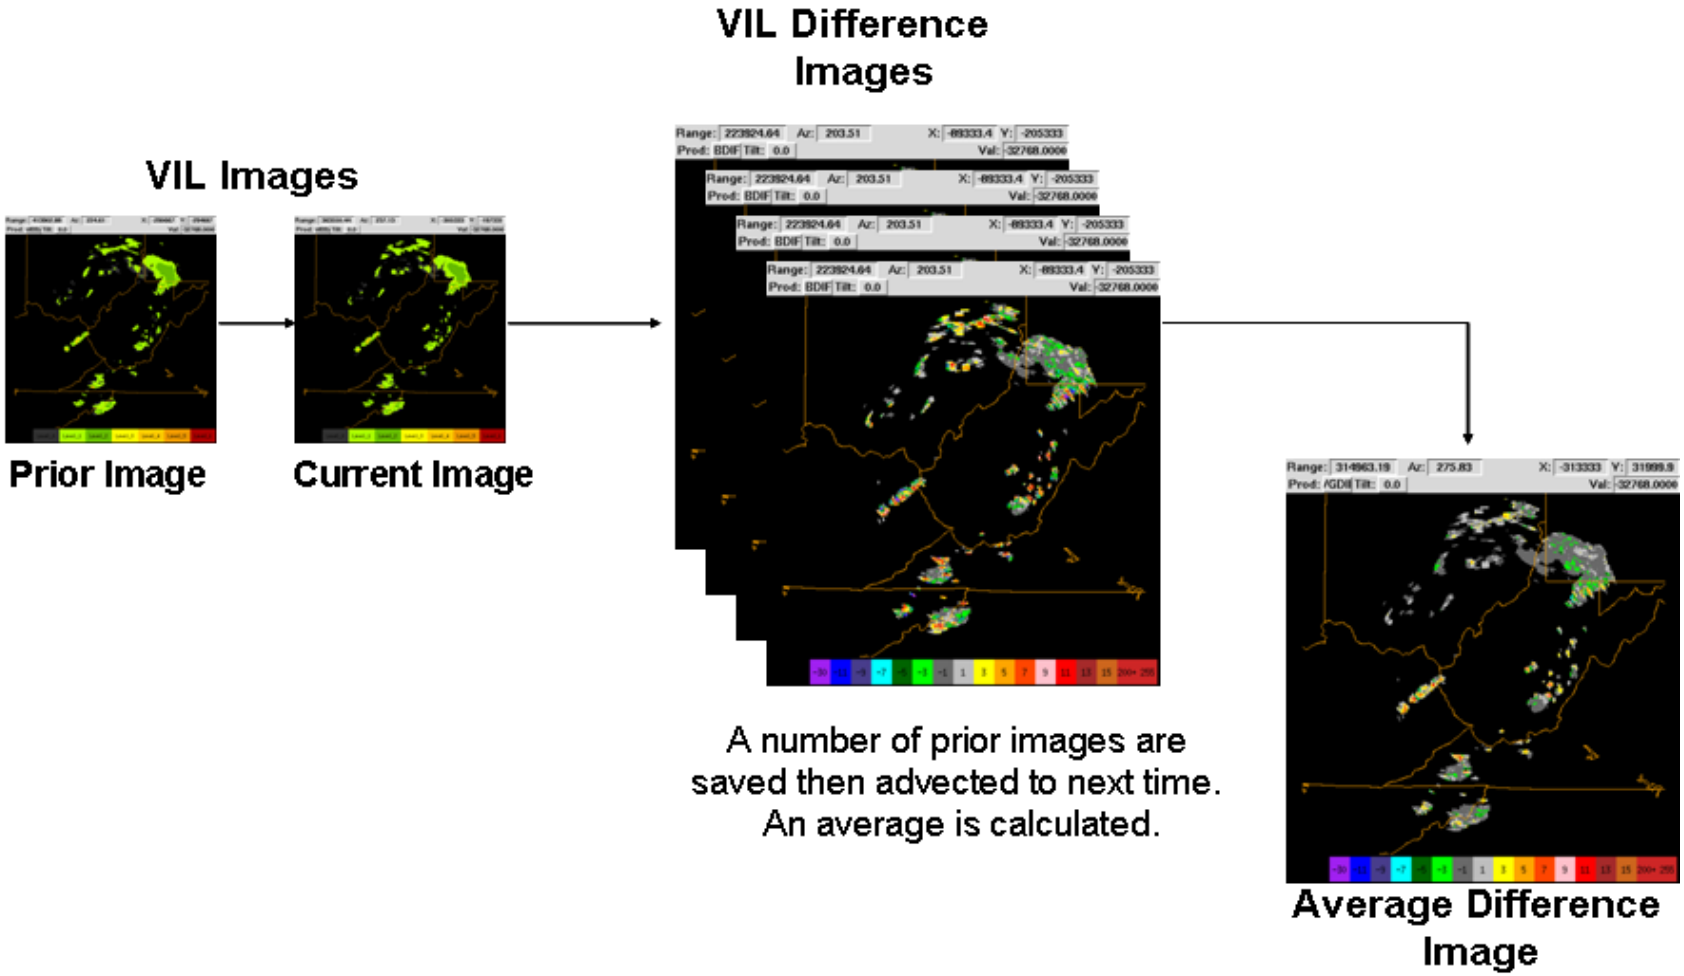

Figure 6: VIL Differencing: This figure shows the process of differencing prior VIL precipitation images and averaging to produce short-term and long-term growth and decay trends.

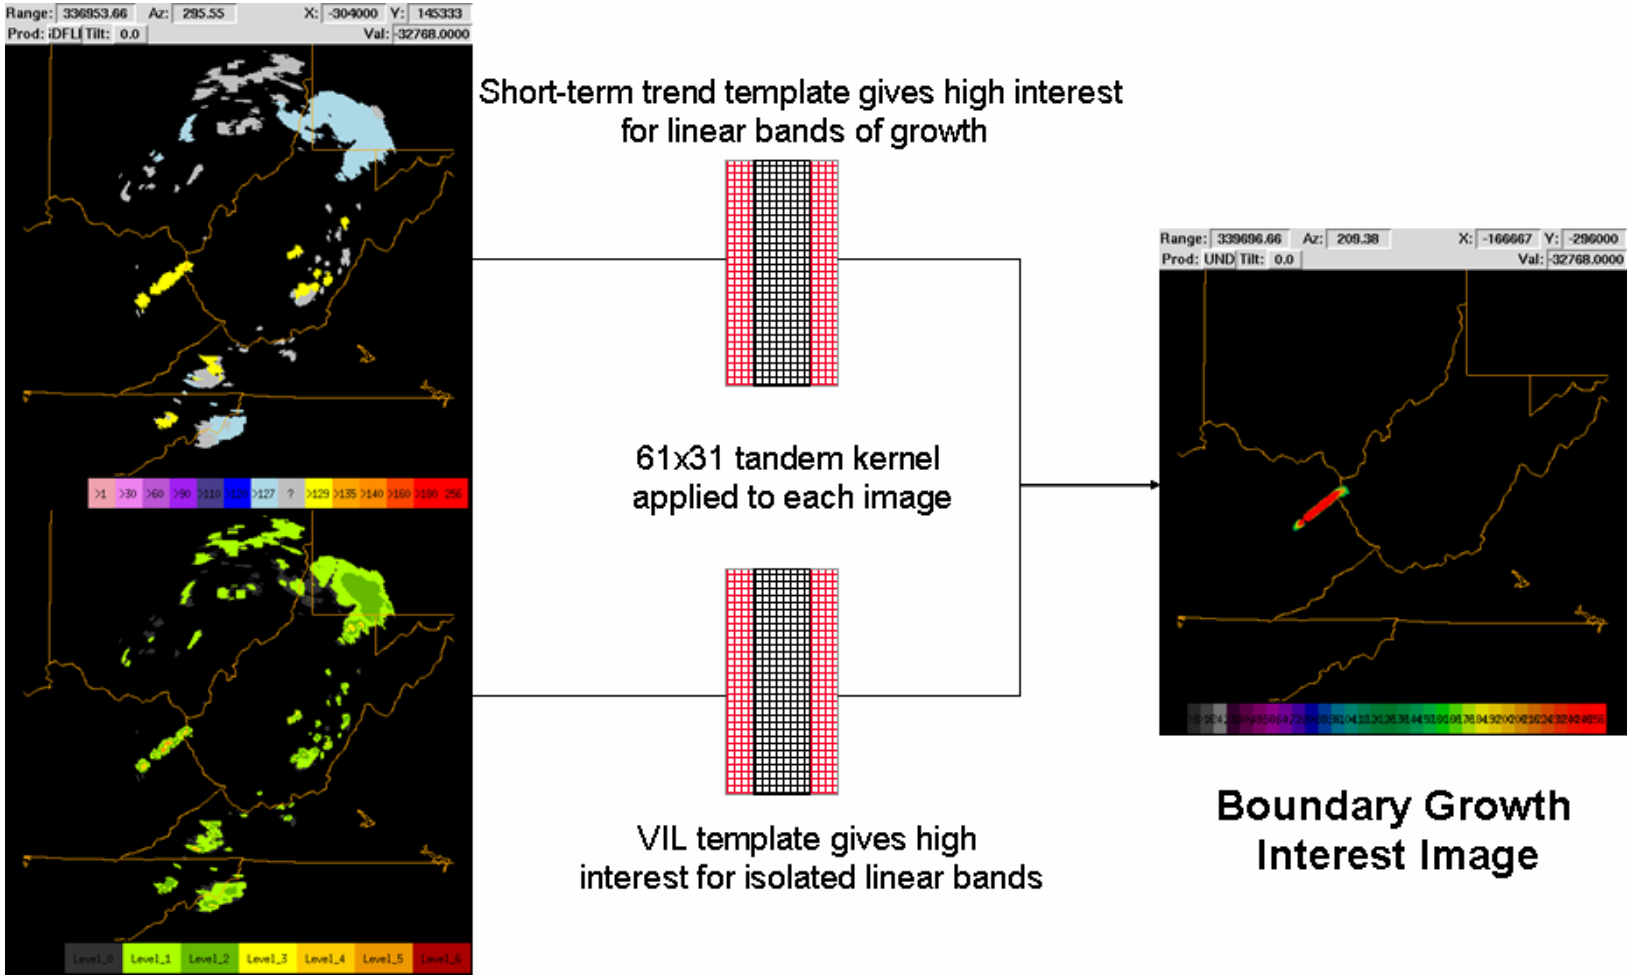

Figure 7: Boundary Growth: The boundary growth detector showing the functional template correlation kernel.

Another important VIL feature detector produces the Isolated Cell interest image. This feature detector uses the short-term trend interest image to assign interest to small storm cells that show a growth signature. These small cells have been empirically shown to be responsible for forecast false alarms when isolated from any surrounding growing cells. The isolated cell interest image is used to suppress the forecasting of interest in these cells. The FTC kernel is structured such that high interest is given only to small (< 20 km diameter) storm cells that are surrounded by no radar returns.

4.4.3 Echo Top Trending

The only image filtering that occurs in the echo top trends is a simple 7x7 km mean filter. This filtering is done to reduce some of the high-frequency noise. The processing chain for the echo top trend algorithm is shown in Figure 8.

4.5 Mosaic Processing

The per-radar processing steps discussed in the previous sections generate a set of maps or images of (1) VIL, (2) echo tops, (3) VIL growth and decay trends, (4) echo top growth and decay trends and (5) motion images (envelope and cell). These images are generated asynchronously at the update of each individual radar system (nominally once every 5-6 minutes). The VIL, VIL trends, and motion data from individual NEXRAD radars extend out to 460 km and for the Canadian SIGMET radars out to 256 km from the radar location. The NEXRAD echo top data are used out to 230 km from the radar and the SIGMET echo top data out to 256 km. Due to the average spacing of the radar systems, there is overlapping coverage over most of the Continental United States (CONUS), especially the eastern half. The overlapping coverage is shown in Figure 9. The figure shows that over the CONUS, there are typically between three and six values at each pixel for the 230 km range.

There are different approaches to mosaicing data from overlapping radar systems. For example, one can consider overlapping data as point estimates and interpolate (Zhang, 2002) or select from a fixed pre-evaluation of the neighboring radars' data quality (Lang, 2003). The essential features of our approach are to (1) time align or advect the individual asynchronous images to a common time, and (2) combine them on a product-by-product basis in a conservative or worst-case representation that still rejects ground clutter or other forms of Anomalous Propagations (AP) by use of the "maximum plausible" rule detailed below.

4.5.1 Time-Alignment

In order to synchronize the mosaic data, we generate a time strobe at every 2.5 minutes so that the mosaiced data are broadcast in regular time increments. Using the motion vectors, the VIL, echo top, and trend images and the vectors themselves are advected to the next time strobe. We use, in essence, a forward, or downwind advection scheme. In the case of VIL, we additionally select the largest value which lands on a grid location. In addition to time alignment, a time check is made. Data will not be advected more than a given amount of time (nominally 15 minutes); if a data set is too old, it is dropped out of the mosaic. Locations in the mosaic where there is no radar coverage are flagged with a no coverage value; hence a radar drop out may change the appearance of the overall coverage pattern.

4.5.2 VIL and VIL Trend Mosaic

The VIL images are combined using the "Maximum Plausible" rule. This rule, in effect, chooses the highest value that can be confirmed by a (slightly smaller) value reported from a neighboring radar. Whether another radar has the ability to confirm a particular value (the test value) at a particular location (the test location), depends on the distance from the test location to the radar that is trying to confirm the value (the confirming radar). If the highest value is not deemed plausible, it is removed and the process is repeated until a value is found which can be confirmed, or there is only one value left. This procedure reduces clutter and AP breakthrough in the VIL mosaic output, yet provides the maximum in normal situations, as shown in Figure 10.

Once a value of VIL has been found, the VIL trend value from the same radar at the same location as the VIL is chosen for use. This keeps the VIL trend data (the time derivatives of VIL) consistent with the VIL data.

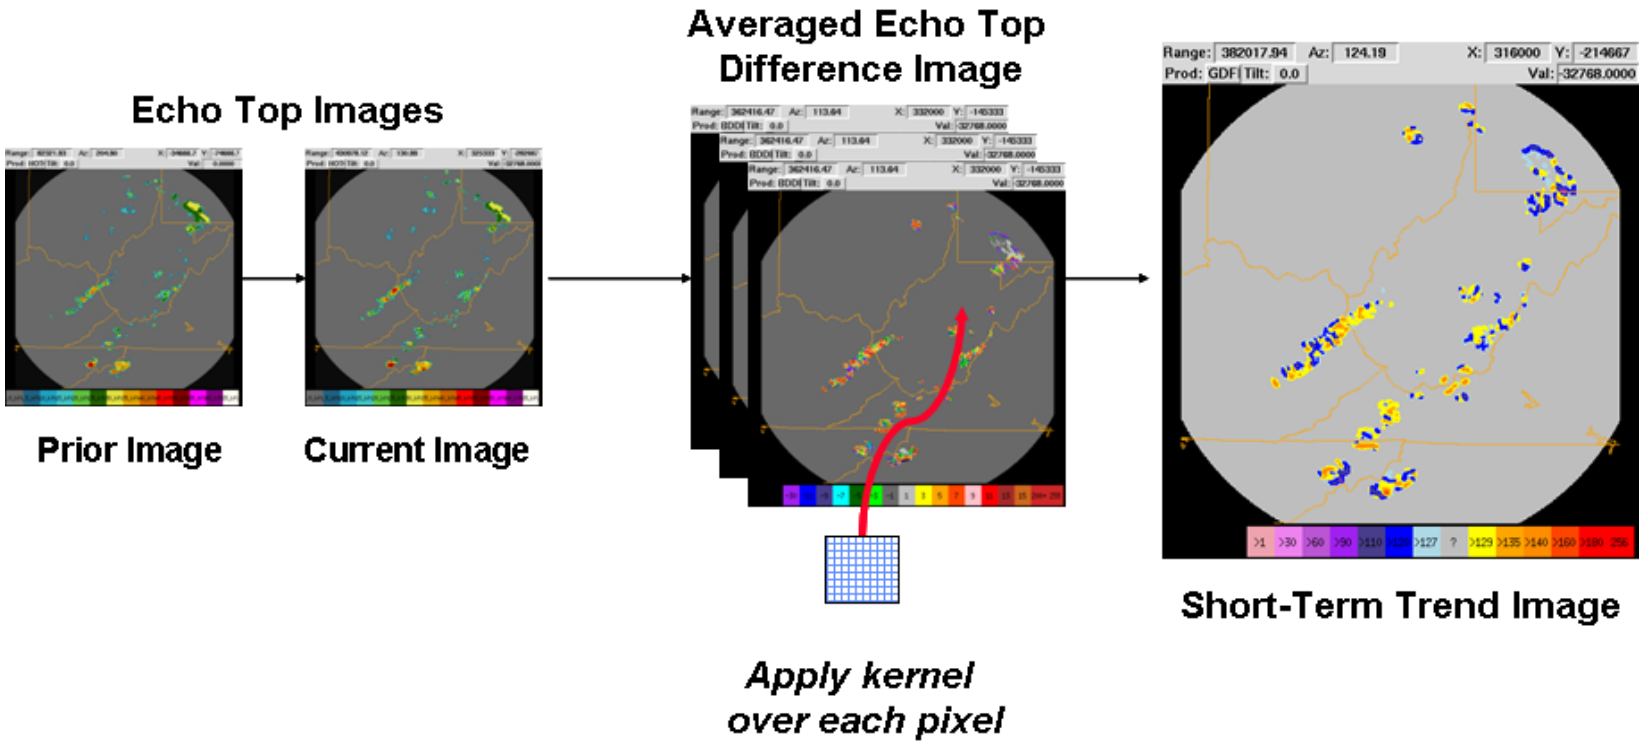

Figure 8: Echo tops short-term growth and decay: This figure shows the algorithm steps used to create the echo tops short-term growth and decay interest. The first step shows the differencing of prior images, the middle graphic is the step where previous images are temporally and spatially averaged, the right image is the final short-term echo top growth and decay image.

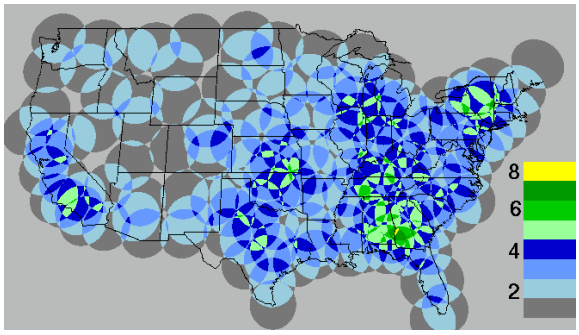

Figure 9: Number of NEXRADs within 230 km radius coverage. This figure does not consider terrain blockage which serves to reduce the range of NEXRADs in mountainous regions.

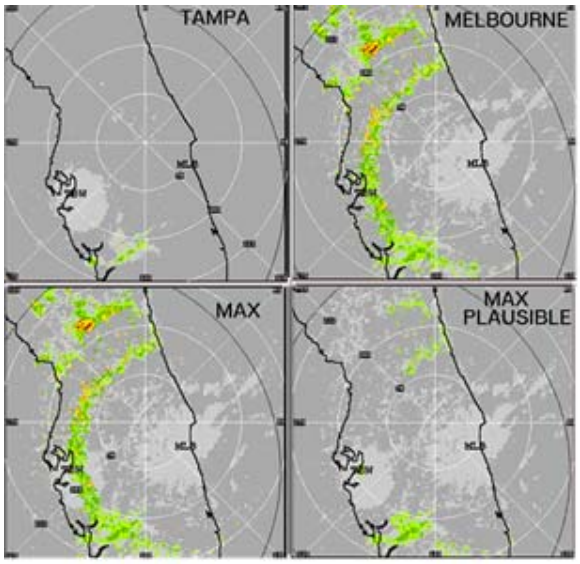

Figure 10: Example of AP contaminating a radar image (MELBOURNE) and failing to be confirmed by a neighboring radar (TAMPA) Both the max and max plausible mosaic results are shown.

4.5.3 Echo Top and Echo Top Trend Mosaic

For the echo top mosaic, we simply select the maximum value provided by all the radars at the time and location of interest. For the echo top trend mosaic we use the same radar from the echo top mosaic to determine which echo top trend data value gets mosaiced.

4.6 Weather Classification

Weather Classification is a key component of the forecast system in that it serves as the underlying classification scheme used to assign specific phenomenological behavior in subsequent models.

Convective and stratiform partitioning classification schemes based on radar precipitation images have been addressed by many researchers (Anagnostou, 2004, Biggerstaff and Listemaa, 2000, Awaka et al., 1997, Steiner et al., 1995). However, these schemes limit the classifications to three basic categories: convective, stratiform, and other.

Dupree et al. (2002) introduced the convective weather classification scheme that extracts lines, cells and stratiform precipitation regions from VIL images. This approach aims to classify the radar returns not only as convective or non-convective but assigns them a distinct phenomenological class. This algorithm was later enhanced to use additional input fields, and to provide growing and decaying sub-type categories (Wolfson et al., 2004-a). Weather types are constructed from the VIL, echo tops, and VIL growth and decay trend images using FTC and image processing region analysis. Figure 11 shows an example of the weather classification image for a single radar mosaic.

A more in-depth explanation of the weather classification algorithm is given in Wolfson et al. (2004-b).

First a spatial standard deviation image ("variability") is used to differentiate between convective and non-convective regions. This image is constructed by running a 15x15 km kernel over the input VIL image and calculating the spatial standard deviation of the radar returns. Following this step, the standard deviation images are filtered and smoothed with a 15x15 km mean filter producing an image in which higher values represent convective elements and lower values are stratiform in nature.

In order to construct a "line" detection, defined as a strong linear precipitation area, convective weather regions are filtered with a 13x69 km rotated elliptical filter to highlight the interest. Region analysis is further applied to each of the non-convective and convective images to generate specific classifications. Embedded cells are defined in this context as any strong cells above VIL Level 2, with echo tops above 26 kft (7.9 km) and located in a VIL

region greater than 70 km in size. The echo top requirement was added to reject those regions with strong and variable precipitation that can form in stratiform weather often due to bright band contamination. Isolated convective regions (< 70 km in size) are sorted into sizes from 4-20 km (small cells) and 21-70 km (large cells).

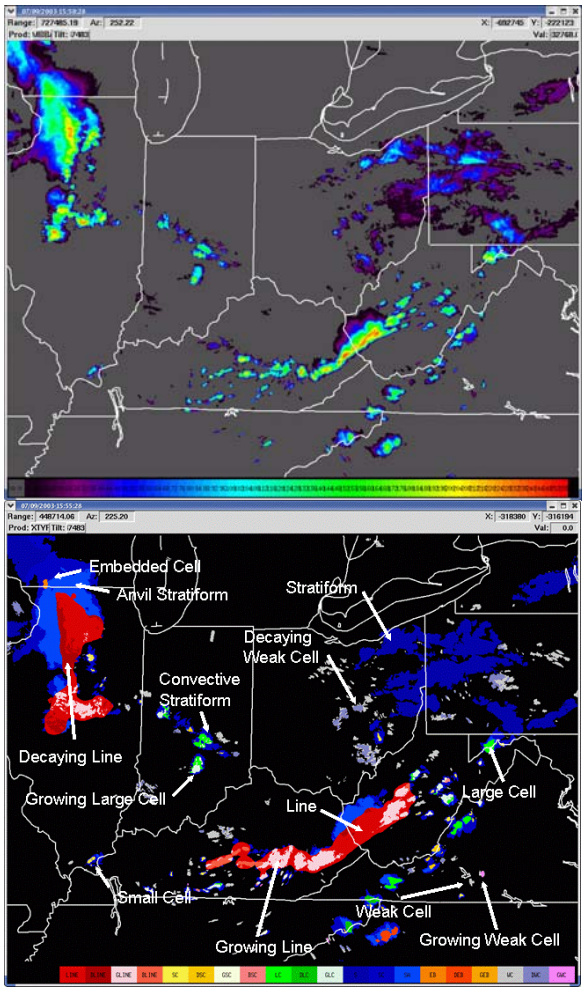

Figure 11: The VIL and weather classification image showing examples of weather features identified in the algorithm.

Non-convective elements are classified into stratiform and weak cells. Weak cells are simply regions with low variability that are less than 70 km in size and have precipitation less than Level 2. At this stage of development it is difficult to distinguish a growing cell from an isolated small patch of stratiform weather. It has been noted that weak cell types contain the first signatures of growing airmass cells. Future work is required to find a classifier that can identify proto-airmass cells.

All remaining stratiform regions are considered next. These are divided into convective and non-convective stratiform even though variability is low throughout. There are two types of convective stratiform pixels that we assign: anvil stratiform, or those with high echo tops values that form around convective cores, and convective stratiform, or any stratiform that is in proximity to convective regions.

To add further sub-classification, types are split into growing, decaying or neutral sub-types based on input from the VIL short-term growth and decay trends. These sub-types have been valuable in identifying the stage of some classes.

The final step assembled all the sub-classified images into a single weather classification image. This is done with a precedence order (see Wolfson et al., 2004-a). In most cases classification pixels are mutually exclusive, but there are a few cases where pixels can be classified as more than one weather type. For example, line regions can contain embedded cells and convective stratiform pixels and therefore we present line as the class for these pixels based on the greater precedence order.

4.7 Multiscale Tracking

Tracking and producing forecasts on multiple scales remains an area of active research (Bellon and Zawadzki, 1994, Wolfson et al., 1999, Seed and Keenan, 2001, Dupree et al., 2002 and Lakshmanan et al., 2003). All of these studies apply some type of classifier to separate features and some method of tracking the feature, usually either cross-correlation tracking or mean-square error tracking. A key result is that all show certain features are more predictable than others. Specifically, large scale features are more predictable than small scale features. We have chosen to simplify this scale extraction step and to only consider two scales or objects of interest, those indicative of line storms (which can also be applied to stratiform weather and “super cells”) and those that depict single cell motion.

Track vectors are calculated on individual radar grids and then time aligned and merged into data streams with vectors for all the radars. Because tracking of different radars can result in different motion vectors for the same patch of

weather when placed on a common grid, it is necessary to look for locations where the vector fields deviate. We calculate the standard deviation of vectors from regions under a 57 km diameter circular kernel and reject all vectors in the top 15% of deviants.

Once deviant track vectors have been rejected, the envelope and cell vectors are sorted according to weather type. We group cell vectors as all vectors that occur under the "small cell" weather type and no larger than 20 km region. All other weather classification types receive envelope track vectors. The merging of the envelope and cell vectors is simply a mapping of the envelope and cell vectors according to the weather classification onto a single grid and can be thought of as a phenomenological weather classification mosaicer.

The mapped vectors are clustered in locations where weather exists and are only located where a correlation track was determined, so the data are irregular and sparse. For the advection procedure, a vector is required at every pixel. To fill this motion field we apply a two part scheme. First, the weighted average of vectors under a 1/r shaped weighting function is found at a user-specified spatial frequency (50 km), and the remaining unfilled pixels are found using bilinear interpolation with the previously calculated 1/r values. The result is the multiscale track vector field which is used to advect both the VIL and echo tops forecasts for each time horizon.

4.8 VIL Forecast Model

The VIL forecast model creates a separate, independent forecast for each forecast time horizon. Each forecast time horizon is created by a two step process. First, we construct an initial "Time 0" forecast (F0) from the current interest images (VIL trends, VIL, isolated cells, weather type, boundary) using a weighted averaging method described in the next section. Because these images portray information in locations representative of conditions at analysis time ("Time 0"), the resulting F0 is therefore also located at a position reflecting conditions at the analysis time. Second, the F0 forecast is advected to the position it is predicted to occupy at the desired forecast time horizon ("Time T") using the multiscale vectors. Third, we apply a Time T suppressor. Fourth, we apply a final median filter to the forecast.

4.8.1 Construction of VIL Forecasts

The VIL Forecast Model combines the current interest images into a single image representing a preliminary forecast for a particular time horizon using a weighted average of the inputs. The weighted average can be expressed as the summation over the i interest images of a pair of functions: the scoring function and the weighting function.

Specifically, there is one scoring function and one weighting function for each time horizon, weather type, and input interest image type. In a system with n_t time horizons, n_w weather types, and n_i input interest image types, there will be two arrays of functions, each array having dimensions [n_t by n_w by n_i]. For each time horizon and input interest type, the particular weather type-specific scoring and weighting functions chosen are based on the weather type at the pixel in question; i.e. weather types are mutually exclusive, and each pixel has exactly one weather type. The scoring functions are represented here as S_{twi} and a pixel in an interest image i has the value v_i. The scores s_i returned for a value v_i at a pixel in an interest image of type i having a weather type w at time horizon t is then represented as:

$$s_i = S_{twi}(v_i). \qquad (5)$$

In a manner exactly similar to the scoring functions, the weights (ω_i) used to average the scores s_i are represented as:

$$\omega_i = \Omega_{twi}(v_i), \qquad (6)$$

where Ω_{twi} represents the weight scoring functions. The value of weighted average (f) can then be expressed as the summation over the i interest images or:

$$f = \frac{\sum_{i=1}^{n_i} \omega_i s_i}{\sum_{i=1}^{n_i} \omega_i}. \qquad (7)$$

This forecast is for the forecast time t at a particular point in the forecast grid.

The scoring and weighting functions were constructed based on an *a-priori* expectation of how the different types of weather classifications evolve over time and how the relative importance of each of the interest images were in that evolution. These were then adjusted based on empirical results.

4.9 Echo Top Forecast

In the NAS, if echo tops exceed flight levels of about 30 kft (9.1 km) the aircraft will be diverted around the high top cloud. If this weather is also spatially extensive it can cause an entire route to be closed. A reliable and accurate rapid update of echo tops is therefore highly desirable to the air traffic flow management community. We chose to use portions of the existing CWF technology to construct a forecast of echo tops.

4.9.1 Echo Top Cap

When forecasting echo top heights it is important to have a good estimate of what height one might expect a storm to achieve. Estimates of the maximum tops are made by calculating the distribution of echo tops under a 51 km circular kernel that is applied to the current and two previous mosaic images (up to 10 minutes in the past) in an Eulerian sense (Figure 12). For each pixel in the domain, the 98^{th} percentile is found and placed on the echo top cap image. We take this height to be the maximum growth height in the echo tops storm model described in the following section. The echo top cap image is currently only used in the echo top forecast.

4.9.2 Echo Top Storm Growth Model

The echo top storm growth model is distinctly a different evolution model than the VIL model. For the echo top model we ingest VIL growth and decay forecasts and use them in conjunction with the other interest images to create a trending and advection echo top forecast. The echo top forecast engine differs from the VIL engine in that interest is combined in a post-advection model. Post-advection combination can be achieved efficiently by pre-calculating a set of advection maps. These advection maps contain the indices of the grid locations for each time horizon. Additionally the echo top cap, weather type, echo top trends, and VIL are ingested. For each time horizon these images are advected. The echo top trends are applied to convective elements, weather types with line, and large, small and embedded cell interest (see Figure 13). These trends are in units of meters/min and we assume a linear growth model for the initial growth phase. Once the echo top has grown to the echo top cap, the top is held at this level for all forecast time horizons. For the remaining types, we advect the existing echo tops. The advected VIL serves as a baseline forecast that can be compared to the growth and decay forecast. Difference images between the baseline forecast and the growth and decay are created and areas of growth and decay can be identified. The echo tops are then dilated to match the growth signatures. In a final step, a 5x5 km median filter is applied to the final forecast.

5.0 FORECAST PERFORMANCE

The CWF algorithm performs a quality assessment of the VIL forecast and the echo tops forecast in realtime. For research and development, binary Critical Success Index (CSI; Theriault et al., 2001) scores are generated in addition to other scoring statistics. Performance is assessed in several ways including the generation of histograms, difference images and images of hits, misses, and false alarms. Binary scores are meaningful for research; however, they do not portray the value that the user gains from forecasts that are only slightly off due to location errors. Users desire a 1 hour forecast that is accurate to within 5 nm (10 km) of the actual weather (Hallowell et al., 1999). As the forecast time horizon increases to 120 minutes, the accuracy restriction is relaxed to 10 nm (20 km). The "user" score that is displayed on the situation display incorporates this methodology. In realtime, user scores are available for the 30, 60 and 120 minute forecasts, denoted by a CSI percent accuracy number in the lower left corner of the forecast window (Figure 3). Forecast verification contours are also displayed on past and current weather images. Further details of the scoring techniques are described in Theriault et al. (2001) and Wolfson et al. (2004-b).

Each of the forecasts have corresponding thresholds that are deemed critical to flow management operations. Interviews with Air Traffic Control (ATC) users have shown that during convective events, color intensity levels 3 and above are the critical levels of interest. In winter events, studies at MIT LL show that color

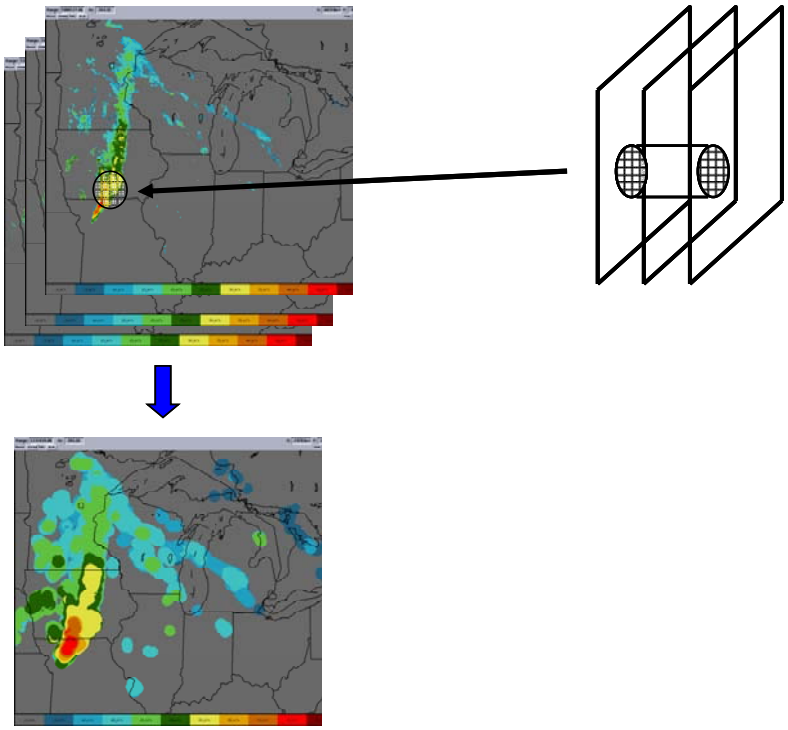

Figure 12: Echo top cap image: The echo top cap is estimated by finding the 98th percentile of an echo top distribution from previous echo top images, 5 minutes apart and under a 51 km circular kernel.

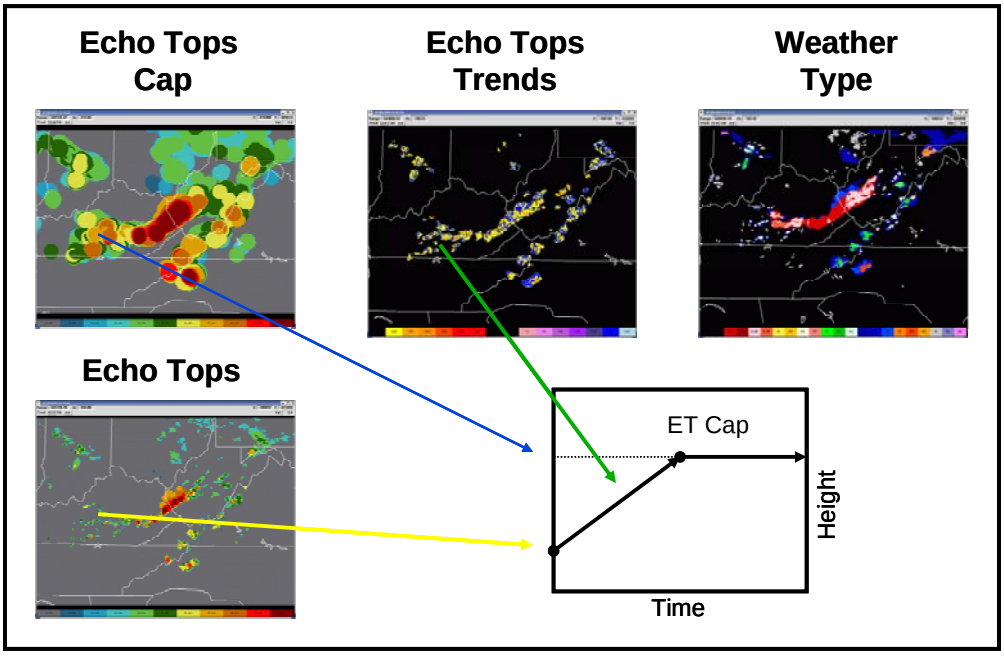

Figure 13: Echo Top Forecast Engine. The echo tops forecast engine applies a simple growth model to the current echo top height field. Short-term echo top growth rates are applied as a function of forecast time horizon in regions of convective weather, until the echo top height reaches the maximum cap as measured by previous echo top distributions in the area. Non-convective weather is simply advected and growth and decay trends are not used.

level 1c is a good predictor of winter weather that is of concern to aviation. Forecast performance is calculated using these critical levels. These confidence scores indicate past forecast performance and may not represent how the current forecast will verify. Table 1 shows the relationship between display color levels and scaled VIL. Table 2 shows the scoring parameters that are used by the algorithm.

Separate CSI performance scores for the 30, 60, and 120 minute forecasts are available for both qualitative analysis (binary scoring) and for display to the users (user scoring).

5.1 Binary Forecast Accuracy Scores

A 1 km resolution pixel by pixel comparison is used to calculate the binary CSI. CSI scores are calculated for a performance evaluation of the data, but can be deceiving without further understanding of the usefulness of a slightly missed forecast in either time or space (Wilson et al., 1997). Examples of both VIL and echo top forecasts are presented. Each section will show CSI curves in addition to a binary verification image that captures the pixel by pixel specifics of the forecast performance.

5.2 VIL Performance

In order to illustrate the VIL forecast performance, we provide a data set of forecast performance over approximately one month. Figure 14 shows the 30, 60, and 120 minute standard VIL forecast performance scores over that time period. The levels scored correspond to the high forecast threshold as noted in Table 2. The results show the dependencies of CSI upon the number of forecast pixels present at the threshold being scored.

The case of 9 July 2003 was studied to assess the performance of the growth and decay trending algorithm. Two modes were created: one with growth and decay trending and one with just advection. Figure 15 below shows the CSI comparisons for both modes for 30, 60 and the 120 minute VIL forecast. It is important to note the marginal improvements realized by examining the data in this way. The blue curve represents VIL forecast performance with the addition of growth and decay trending, while the green curve shows performance without growth and decay trending.

Table 2
Thresholds for Forecast Performance Algorithm

	Color Level	Standard Forecast (scaled VIL)	Echo Top Forecast (km)
HIGH forecast threshold	3	133	9.1 (30 kft)
HIGH truth threshold	3	133	9.1 (30 kft)
Moderate forecast threshold	2	73	6.1 (20 kft)
Moderate truth threshold	2	73	6.1 (20 kft)

Figure 16 illustrates the binary forecast verification image with and without growth and decay trending during two time periods for the 9 July 2003 case. False alarms are shown in red, hits are shown in green, and misses are shown in blue. At time 14:30 there is a large region of forecast false alarms (shown in red) in the top left image, but note the proximity to the misses. The next five minute time period (top right) better reveals the forecast improvement with the addition of growth and decay trending. It captures the full extent of the line that was rapidly developing. Examining cases in detail is crucial in understanding the benefit that the growth and decay trending adds to the algorithm.

5.3 Echo Tops

Interviews with ATC personnel have shown that the critical flight level of 30 kft (9.1 km) is extremely important to enroute operations. For this reason, the ETF scores the performance of the 30 kft forecast level against truth. Knowledge of echo top forecasted heights assists ATC in guiding planes to fly around or above the storms. Air traffic routes may remain open longer in convective situations given this information. Figure 17 shows the binary scores of the 30 kft threshold for the first month of operation.

An echo top performance analysis was also performed on the 9 July 2003 case we studied for the VIL forecast. CSI curves for 30, 60, and

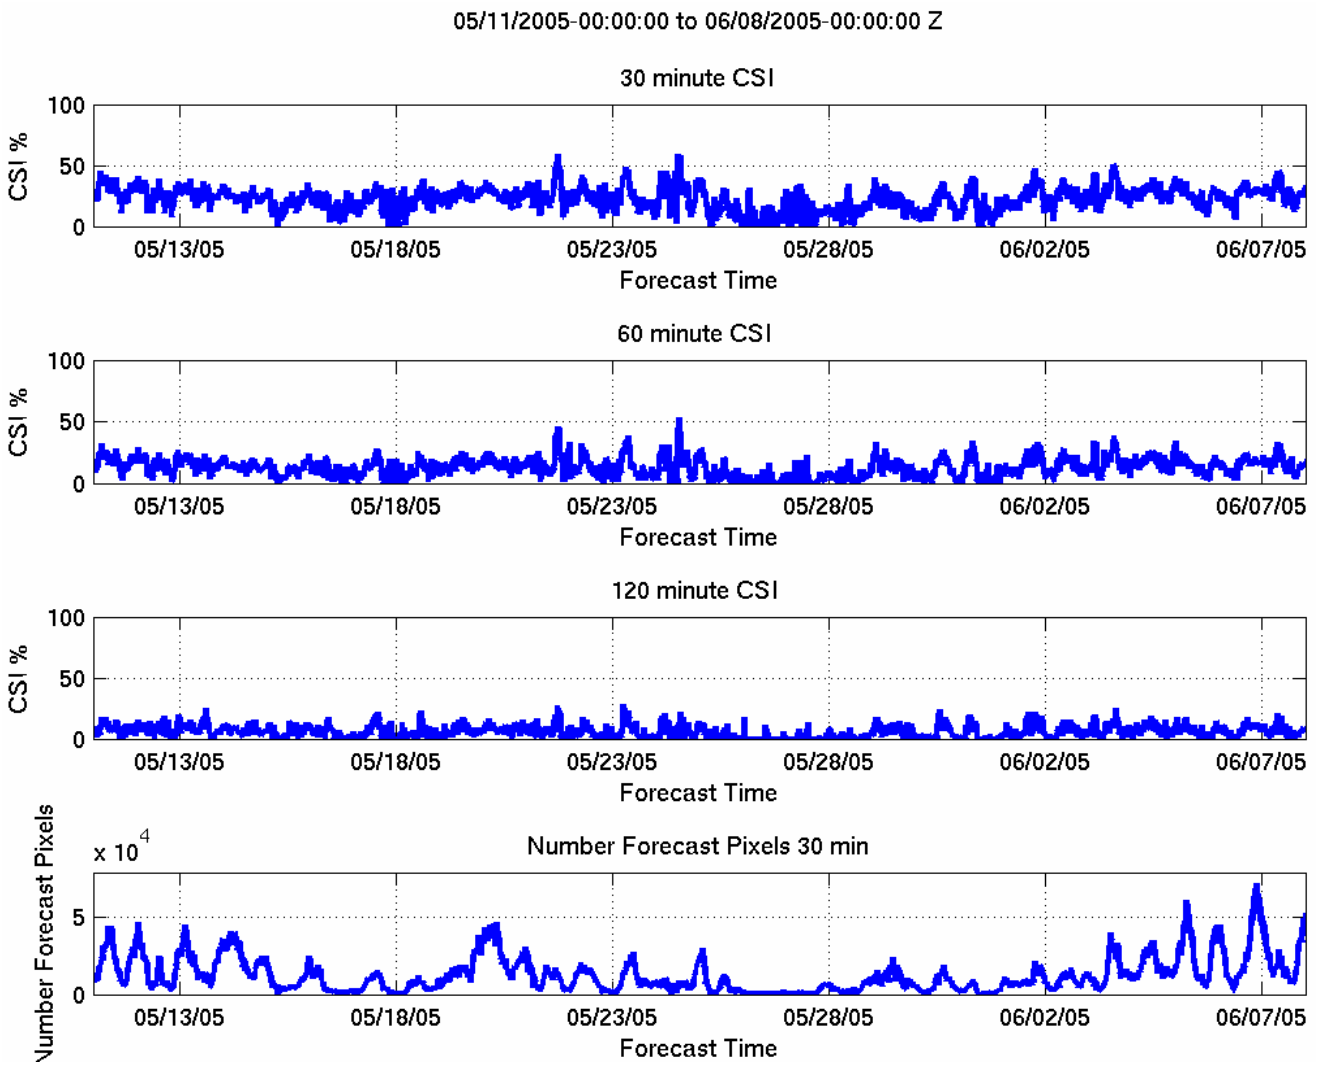

Figure 14: This figure shows the binary 30, 60, and 120 minute CSI scores for the VIL forecast over approximately a one month period. Forecast is for VIL level 3 and above.

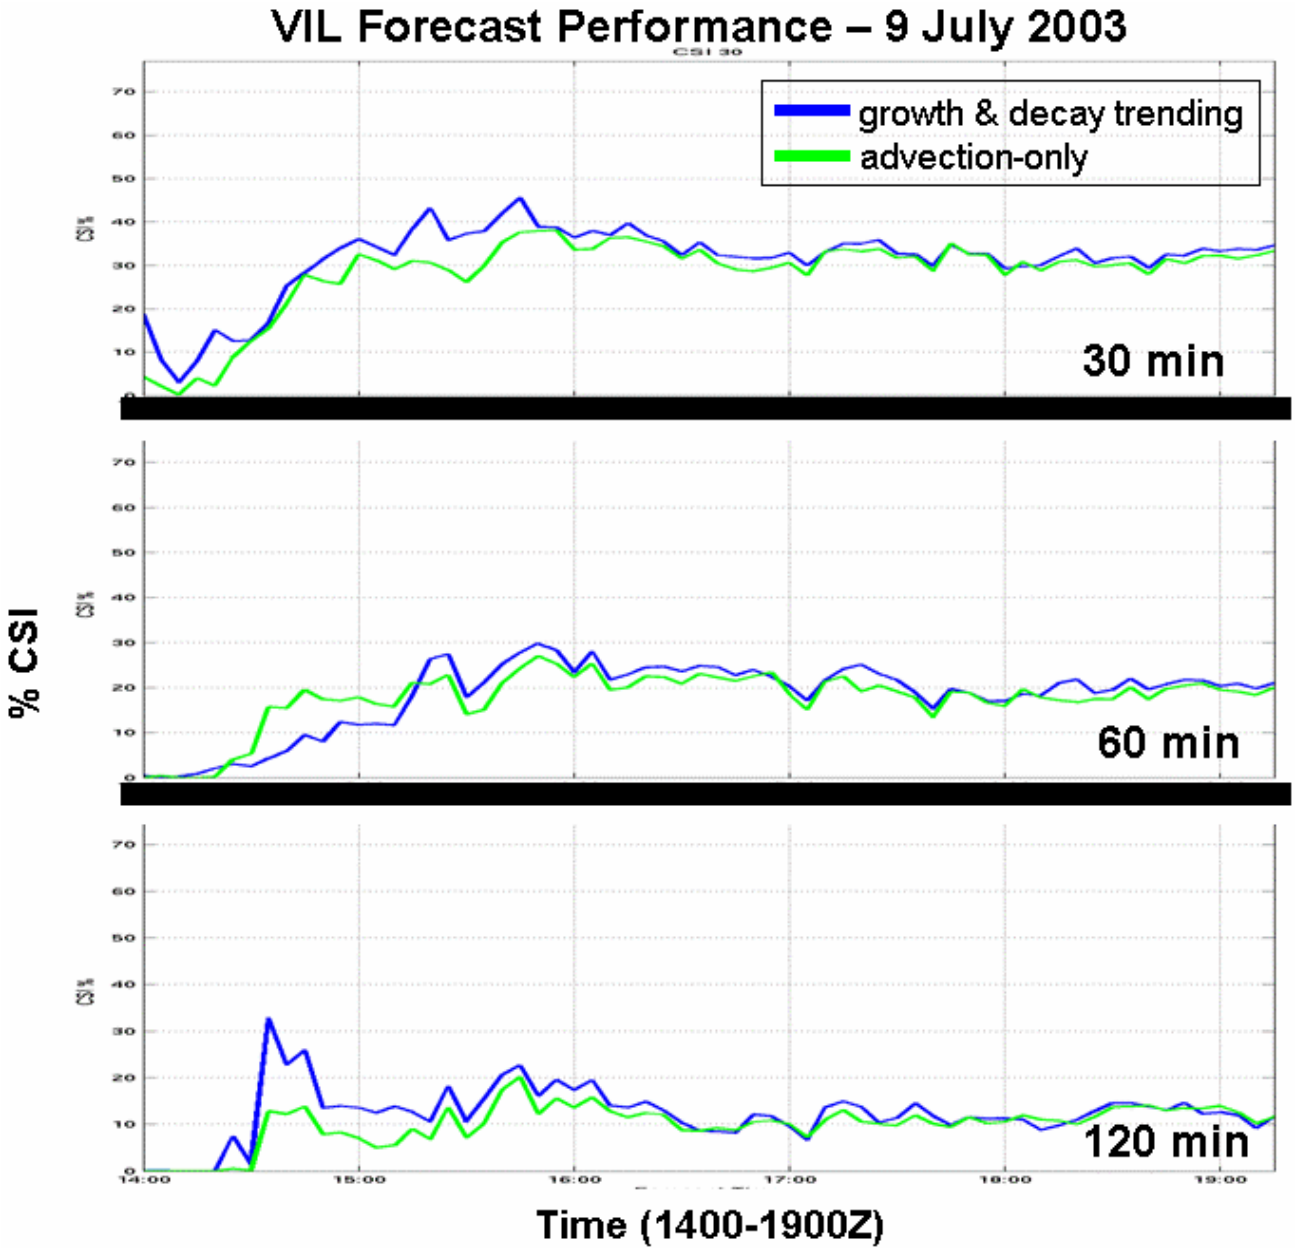

Figure 15: The CSI comparisons between the growth and decay VIL forecasts and the advection only VIL forecasts on 9 July 2003. The blue curves represent the forecast performance with the addition of growth and decay trending, while the green curves show performance without growth and decay trending.

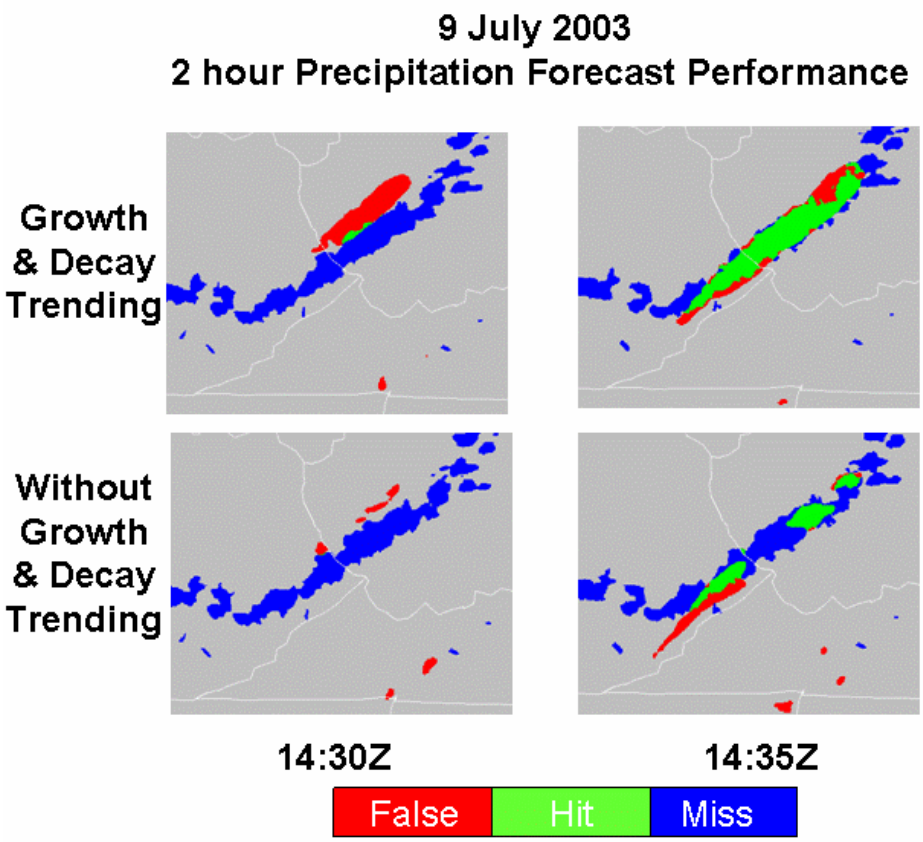

Figure 16: Shown is the superposition of the 120 minute VIL forecast at the level 3 threshold. False alarms are shown in red, hits in green, and misses in blue. The upper left side shows where the forecast has false alarmed in close proximity to the newly grown line. The top right shows the next time period where the growth and decay forecast provides significant benefit.

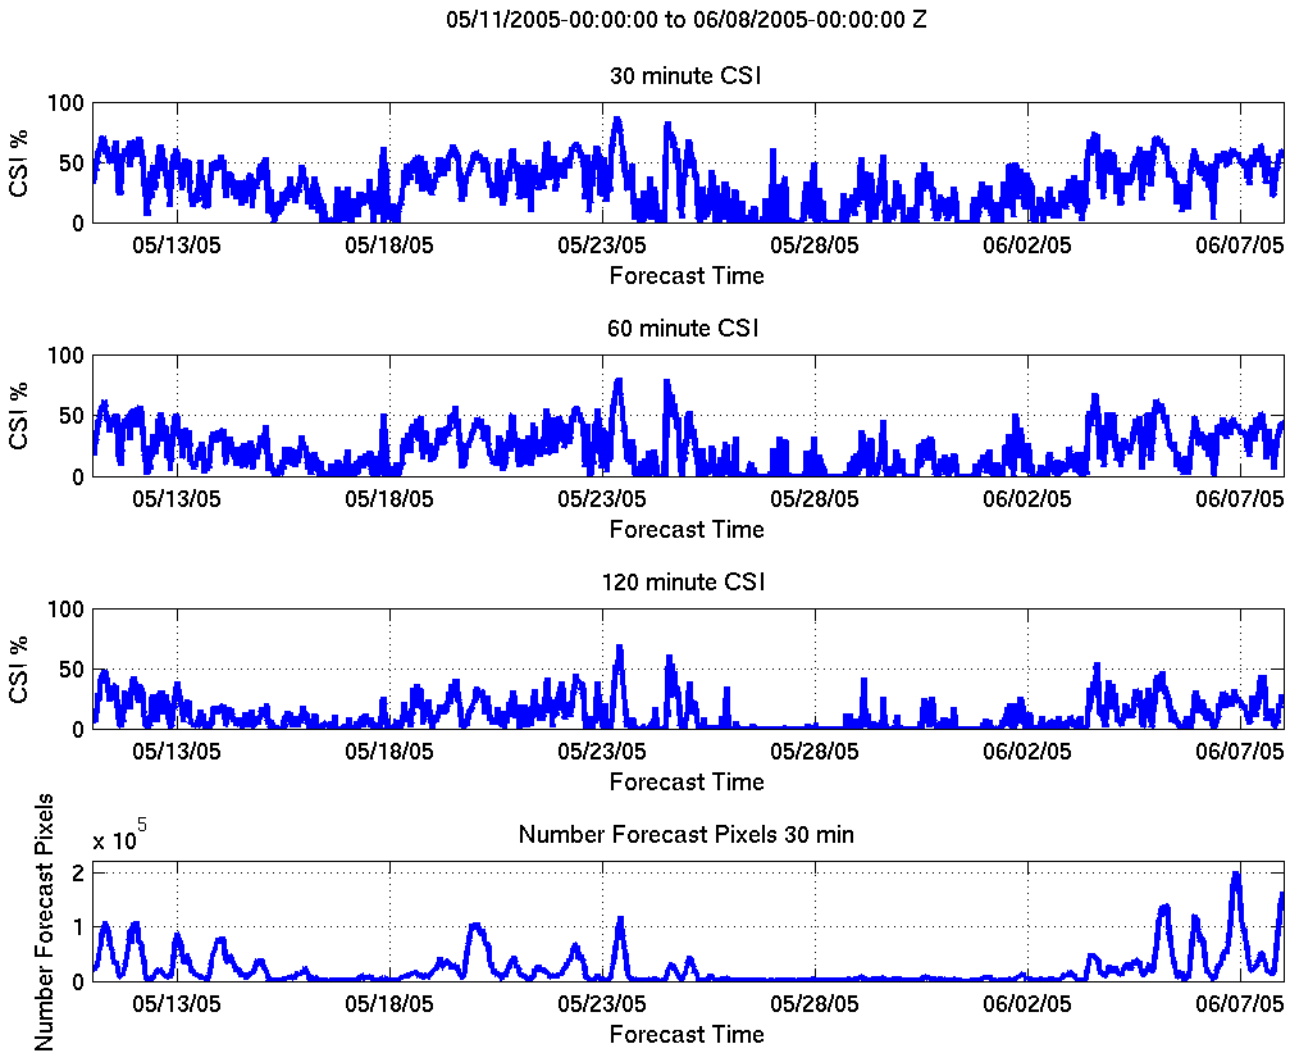

Figure 17: This figure shows the binary 30, 60, and 120 minute CSI scores for the echo top forecast over approximately a one month period.

120 minute echo top performance are shown in Figure 18. The addition of growth and decay trending adds significant benefit to the forecast as shown by the blue curves, while the green curves show performance without growth and decay trending. For this case we see clear performance gains for the echo top trend forecast.

Similar to the VIL forecast performance scores, it is important to not only look at the statistical diagnostics of the scoring, but understand the specific regions of benefit. Figure 19 shows an example of the additional benefit gained using echo top growth and decay trending on the 9 July 2003 event.

Overall, both the VIL and echo top forecasts perform well on a variety of storm classifications. The echo top forecast recently added to CIWS works well and will be analyzed throughout the next several months for possible enhancements. It has been shown that the addition of growth and decay trending adds benefit above and beyond an advection-only forecast.

6.0 FUTURE WORK

6.1 Convective Initiation

Convective weather, and in particular the initiation of new thunderstorms, makes the efficient management of air traffic in the NAS difficult. Prior knowledge of when and where new convection will develop, when it will decay, and the type of storms that will form are key elements in producing effective convective weather forecasts. The CWF system currently produces forecasts of VIL based primarily on the characteristics of existing storms. Since the CWF system is unable to account for convective initiation, the development of new storms is one of the largest sources of error in the CWF system. Examples and a further discussion of this shortcoming can be found in Wolfson et al.(2004-a).

In many situations, convective initiation is preceded by low altitude convergence in the horizontal winds (Wilson and Megenhardt, 1997). These regions of low altitude convergence, often referred to as boundaries, are typically associated with synoptic scale fronts, drylines, and thunderstorm outflows. Gridded wind analyses that utilize Doppler weather radar, surface, and aircraft measurements are one of the best sources of low altitude winds that can be used to identify wind boundaries over large domains. Here a “large domain” represents an area greater than or equal to the continental US east of the Mississippi river.

Two gridded wind analysis systems are being evaluated by MIT LL for use as a data source for a “large domain” system capable of providing realtime detections of wind boundaries. The first is the Space Time Mesoscale Analysis System (STMAS) operated by NOAA FSL (Koch et al., 2004). STMAS produces a surface wind analysis from a large network of surface anemometers. The STMAS product is a 5 km grid point resolution analysis that updates every 15 minutes, with a 45 minute latency. The second system being evaluated is the Corridor Boundary layer wind analysis system (CBOUND). CBOUND is a prototype wind analysis system based on a modified version of the ITWS Terminal Winds (TWINDS) system that utilizes the 20 km RUC model wind analysis, 5 minute Aviation Surface Observation System (ASOS) measurements, and all of the available WSR-88D (NEXRAD) and Terminal Doppler Weather Radar (TDWR) data in the domain. The final products from this analysis are two 9-layer wind analyses that extend from the surface to 800 hPa at 1 and 5 km horizontal resolutions. The prototype system currently covers a domain centered near Chicago, IL and produces an updated wind analysis every 5 minutes. The long-term goal of this development effort is to integrate data from all of the available FAA wind sensors in the CIWS domain into a single high resolution wind analysis.

Automated techniques for the detection of synoptic scale fronts are also currently under development at MIT LL. Lagrangian Scalar Integration (LSI) is the technique being used to identify these fronts. The purpose of automation is to provide a rapid update of the location of synoptic airmass boundaries. When ingested into CWF, this can improve the automated forecast of convective initiation and growth.

Developed by MIT LL, LSI is applied to the gridded surface wind analyses for atmospheric feature detection (patent pending) (Winkler, 2001). A grid of tracers is specified over the wind analysis at a resolution consistent with

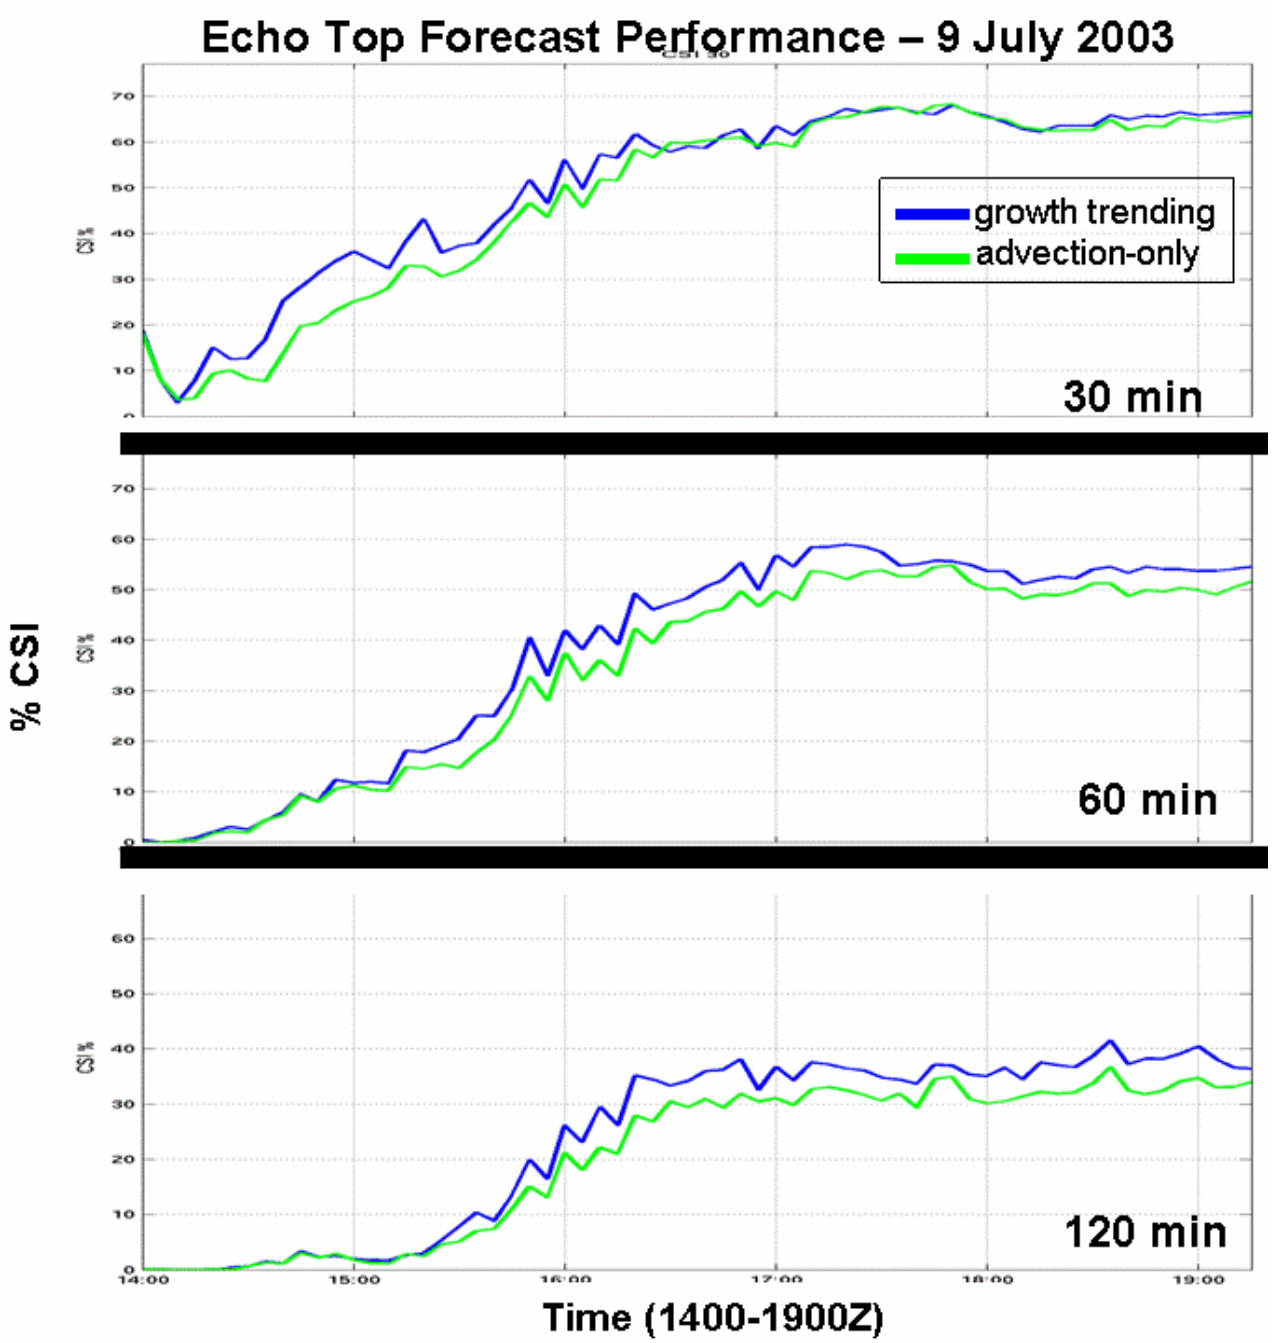

Figure 18: The CSI comparisons between the growth and decay echo top forecasts and advection echo top forecast for the 30, 60 and 120 minute forecasts for 9 July 2003 data. The blue curves represent the trending forecast performance, while the green curves show performance without growth and decay trending.

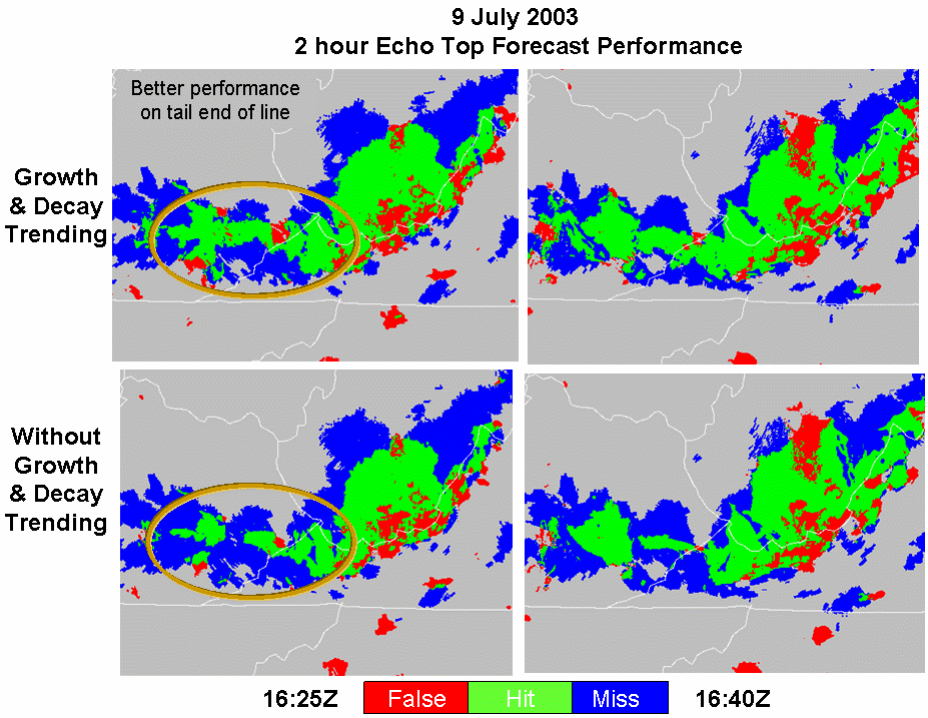

Figure 19: Shown is the superposition of the 120 minute VIL forecast at the 30 kft (9.1 km) threshold. False alarms are shown in red, hits in green, and misses in blue. Highlighted under the ellipse are areas where the echo top trend forecast showed improved performance over the advected forecast.

features of interest (synoptic or mesoscale) and advected following the horizontal winds. A scalar quantity such as temperature, moisture, wind divergence, etc., can be measured at each time interval and integrated. This technique provides a measurement of the dispersive properties of the synoptic scale motions as well as an estimate of the conditions experienced by the tracer particles as they are advected in time.

When parameterized for the detection of synoptic scale boundaries, LSI effectively sharpens the gradients of the scalar quantity while at the same time reducing noise generated by an analysis. This is evident in Figure 20 when comparing the gridded Eulerian divergence field (far left image) to the LSI divergence field (center image) during the passage of a synoptic cold front over the Western Ohio River Valley. Here, artifacts and topographic effects present in an STMAS wind analysis as revealed by the Eulerian divergence calculation are filtered by the LSI divergence calculation and the convergence signature associated with the front is more coherent.

The sharpened scalar fields produced by the LSI computation are then processed with a modified version of the Machine Intelligent Gust Front Algorithm (MIGFA) (Troxel et al., 2002) developed by MIT LL. The modified version of MIGFA utilizes multi-dimensional image processing and fuzzy logic techniques to identify synoptic fronts in the LSI data. In the far right image of Figure 20, the LSI relative dispersion provides an interest field for the modified version of MIGFA to make a front detection. Overlying the interest field (magenta lines) is the location of the surface front as detected by the automated algorithm. While the position of the front becomes discontinuous in the southwest quadrant of the grid, the detectable interest can be enhanced by layering other scalar quantities output by the LSI calculation, which can lead to a continuous automated detection.

Automated boundary detections of synoptic and mesoscale fronts can also be made using the LSI-processed CBOUND data. Figure 21 illustrates the Eulerian, Lagrangian, and LSI relative dispersion fields derived from the 1

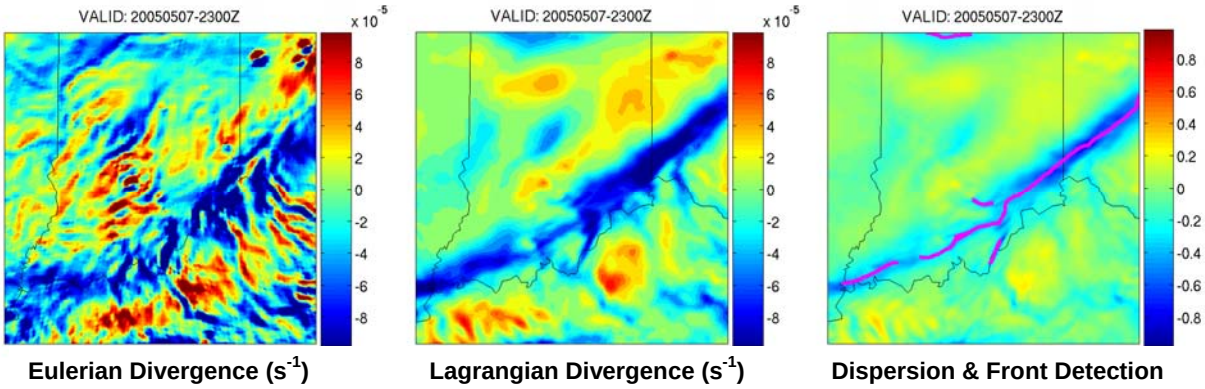

Figure 20: This image shows the automated detection of a synoptic scale cold front utilizing STMAS surface analysis and machine intelligent front detection algorithms developed by MIT LL. The far left image illustrates the noise generated by an Eulerian divergence calculation. The center image shows the smooth fields resulting from a Lagrangian scaled integrated divergence calculation developed by MIT LL. The far right image illustrates Lagrangian Scalar Integrated Relative Dispersion used in automated front detection with the calculated position superimposed.

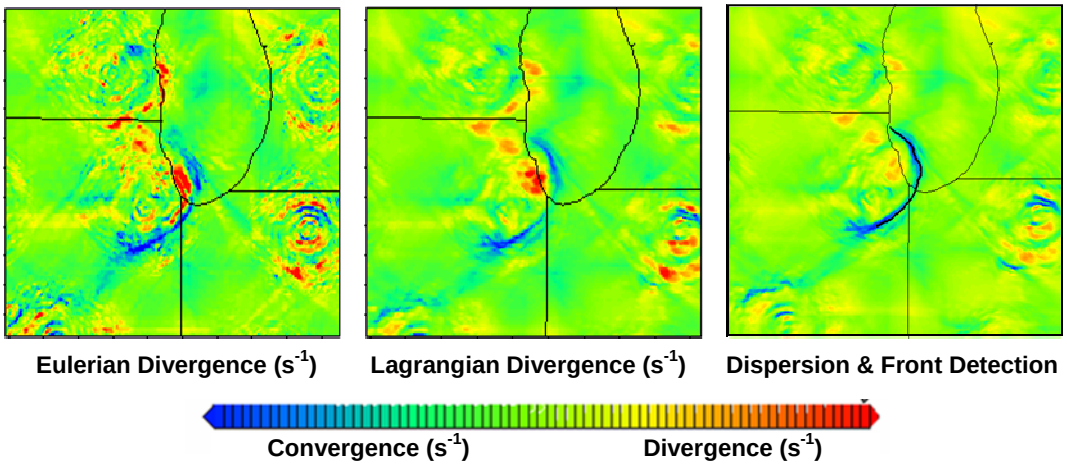

Figure 21: This image shows automated detection of a mesoscale convective outflow boundary utilizing the CBOUND surface analysis and machine intelligent front detection algorithms developed by MIT LL. The far left image illustrates the noise generated by an Eulerian Divergence calculation compared to the smooth fields resulting from a Lagrangian integrated divergence calculation developed by LL in the center image. The far right image illustrates Lagrangian Scalar Integrated Relative Dispersion used in automated front detection with calculated position superimposed as a black line.

km resolution CBOUND wind analysis of a thunderstorm outflow. The automated frontal detection of a thunderstorm outflow using MIGFA is overlaid as a black line on the LSI relative dispersion analysis. When parameterized for the detection of mesoscale boundaries, LSI effectively sharpens the gradients of the mesoscale scalar quantities and reduces the noise generated by an analysis. A comprehensive review of these wind analysis systems will be conducted during the summer convective season of 2005. This review will provide a recommendation of which system can best meet the boundary detection capability required to make successful automated convective initiation forecasts.

7.0 SUMMARY

This paper has reviewed the scientific rationale and development path for the MIT LL automated Convective Weather Forecast (CWF) system. The CWF system provides 0 to 2 hour forecasts of precipitation, and a novel forecast of storm height called the Echo Top Forecast, which is currently undergoing operational evaluation. The CWF system operates in a large part of the US NAS and has proven to be extremely beneficial to air traffic flow management operations. The forecast shows consistently good performance over long periods of operation and during periods of growth and decay. We recognize that integrating convective initiation into the system can provide further benefit in aiding with weather related decisions and we are actively working towards creating initiation prediction tools.

8.0 REFERENCES

Allan, S., S. Gaddy and J. Evans, 2001: "Delay Causality and Reduction at the New York City Airports Using Terminal Weather Information Systems" Massachusetts Institute of Technology, Lincoln Laboratory, Project Report ATC-291 (available at http://www.ll.mit.edu/AviationWeather/reports.html)

Anagnostou, E. N., 2004: A convective/stratiform precipitation classification algorithm for volume scanning weather radar observations, Meteorol. Appl., **11**, 291-300.

Awaka, Jun, 1997: Rain Type Classification Algorithm for TRMM Precipitation Radar: Proceedings of the IEEE 1997 International Geosci. Remote Sen. Symposium. 3-8 August, Singapore, 1636-1638.

Bellon, A. and I. Zawadzki, 1994: Forecasting of hourly accumulations of precipitation by optimal extrapolation of radar maps, J. of Hydrol., **157**, 211-233.

Biggerstaff, M. I. and S. A. Listemaa, 2000: An Improved Scheme for Convective/Stratiform Echo Classification Using Radar Reflectivity, J. Appl. Meteor., **39**, 2129-2150.

Cartwright, T.J., M.M. Wolfson, B.E. Forman, R.G. Hallowell, M.P. Moore, and K.E. Theriault, 1999: The FAA Terminal Convective Weather Forecast Product: Scale Separation Filter Optimization, 29th International Conference on Radar Meteorology, Montreal, Quebec, 852-855.

Chornoboy, E.S., A.M. Matlin, and J.P. Morgan, 1994: Automated Storm Tracking for Terminal Air Traffic Control, MIT Lincoln Laboratory Journal, **7**, 427-448.

Delanoy, R. L., J.G. Verly, and D. Dudgeon, 1992: Functional Templates and Their Applications to 3-D Object Recognition, Proc. Int. Conf. On Acoustics, Speech, and Signal Processing (ICAASP), IEEE, **3**, III-141-144.

Dupree, W.J., R.J. Johnson, M.M. Wolfson, K.E. Theriault, B.E. Forman, R.A. Boldi, and C. A. Wilson, 2002: Forecasting Convective Weather Using MultiScale Detectors and Weather Classification – Enhancements to the MIT Lincoln Laboratory Terminal Convective Weather Forecast, 10th Conference on Aviation, Range, and Aerospace Meteorology, Portland, Oregon, 132- 135.

Evans, J. , M. Robinson, B. Crowe, D. Klingle-Wilson and S. Allan., 2003, Reducing Severe Weather Delays In Congested Airspace With Weather Decision Support for Tactical Air Traffic Management, 5th Eurocontrol/FAA ATM R&D Seminar, Budapest, Hungary, http://atm2003.eurocontrol.fr/

Hallowell, R.G., M.M Wolfson, B.E. Forman, M.P. Moore, B.A. Crowe, T.M. Rotz, D.W. Miller, T.C. Carty and S.F. McGettigan, 1999: The Terminal Convective Weather Forecast Demonstration at the DFW International Airport, 8th Conf. On Aviation Meteorology, Dallas, TX, 200-204.

Koch, S.E., Y. Xie, N. Wang, J.A. McGinley, P.A. Miller, and S. Albers, 2004: A new approach for mesoscale surface analysis: The Space-Time Mesoscale Analysis System. 22nd Conf. on

Severe Local Storms, Hyannis, MA, Amer. Meteor. Soc.

Lakshmanan, V., R. Rabin, and V. Debrunner, 2003: Muliscale Storm Indentification and Forecast, J. Atm. Res., 367-380.

Lang, J., 2003: Radar Mosaic Generation Algorithms 20th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology, Seattle, WA, Amer. Meteor. Soc., paper 12.10.

Marwitz, J.D., 1972: The structure and motion of severe hailstorms. Part II. Multicell storms. J. Appl. Meteor., **11**, 180-188.

Robinson, M., J. Evans, B. Crowe, D. Klingle-Wilson and S. Allan, 2004: CIWS Operational Benefits 2002-3: Initial Estimates of Convective Weather Delay Reduction, MIT Lincoln Laboratory Project Report ATC-313, 326.

Seed, A. W. and T. Keenan, 2001: A Dynamic Spatial Scaling Approach to Advection Forecasting. 30th International Conference on Radar Meteorology, 19-24 July 2001, 492-494.

Steiner, M., R. A. Houze Jr., and S. E. Yuter, 1995: Climatological Characterization of Three-Dimensional Storm Structure from Operational Radar and Rain Gauge Data, J. Appl. Meteor., **34**, 1978-2007.

Theriault, K.E., M.M. Wolfson, B.E. Forman, R.G. Hallowell, M.P. Moore, R.J. Johnson, Jr., 2000: FAA Terminal Convective Weather Forecast Algorithm Assessment. 9th Conf. on Aviation, Range, and Aerospace Meteorology, Orlando, FL.

Theriault, K.E., M.M. Wolfson, W.J. Dupree, B.E. Forman, R.G. Hallowell, R.J. Johnson and M.P. Moore, 2001: TCWF Algorithm Assessment–Memphis 2000, MIT Lincoln Laboratory Project Report, ATC-297, 109.

Tsonis A. A. and G. L. Austin, 1981: An Evaluation of Extrapolation Techniques for the Short-Term Prediction of Rain Amounts, Atmos.-Ocean **19(1)**, 54-65.

Troxel, S.W., R.L. Delanoy, J.P. Morgan, W.L. Pughe, 1996: Machine intelligent gust front algorithm for the Terminal Doppler Weather Radar (TDWR) and Integrated Terminal Weather System (ITWS), AMS Workshop on Wind Shear and Wind Shear Alert Systems, Oklahoma City, OK, 70-79.

Weisman, M.L., and J.B. Klemp, 1986: Characteristics of isolated convective storms, Mesoscale Meteorology and Forecasting, P. S. Ray Ed., Amer. Meteor. Soc., 331-358.

Wilson, J.W., 1966: Movement and predictability of radar echoes. NSSL Tech Memo, **28**, Norman, OK., 30 p.

Wilson, J.W. and D.L. Megenhardt, 1997: Thunderstorm Initiation, Organization, and Lifetime Associated With Florida Boundary Layer Convergence Lines. Mon. Wea. Rev., **125**, 1507–1525.

Wilson, J.W., Crook, A., Mueller, C.K., and Dixon, M, 1997: State-of-the-art of Nowcasting Thunderstorms; 7th Conf. on Aviation, Range and Aerospace Meteorology, Amer. Meteor. Soc., Long Beach, CA.

Wilson, J.W., Crook, N.A., Mueller, C.K., Sun J. and M. Dixon, 1998: Nowcasting Thunderstorms: A Status Report, Bulletin of the Amer. Meteor.l Soc., **79**, 2079-2099.

Winkler, S, 2001: Lagrangian Dynamics in Geophysical Fluid Flows, PhD Thesis, Brown University.

Wolfson, M.M., B.E. Forman, R.G. Hallowell, and M.P. Moore, 1999: The Growth and Decay Storm Tracker, 8th Conference on Aviation, Range, and Aerospace Meteorology, Dallas, TX, 58-62.

Wolfson, M.M., B.E. Forman, K.T. Calden, W.J. Dupree, R.L. Johnson Jr., R. A. Boldi, C. A. Wilson, P.E. Bieringer. E. B. Mann, and J.P. Morgan, 2004-a: Tactical 0-2 Hour Convective Weather Forecasts For FAA, 11th Conference on Aviation, Range and Aerospace Meteorology, Hyannis, MA.

Wolfson, M.M., R. Boldi, K. Calden, W. Dupree, B. Forman, R. Johnson, Jr., E. Mann, and C. Wilson, 2004-b: ITWS Pre-Planned Product Improvement Algorithm Description for Convective Weather Forecast, Project Report ATC-316, MIT Lincoln Laboratory, 476 p.

Zhang, J., J. J. Gourley, W. Xia, and K. Howard, 2002: 3D multi-radar reflectivity mosaic for the CIWS domain. Preprints, The 10th Conference on Aviation, Range, and Aerospace Meteorology. 13-16 May 2002, Portland, Oregon, J107-J109.