Developments in echo tracking - enhancing TITAN

Mike Dixon[1] and Alan Seed[2]

[1]National Center for Atmospheric Research, Boulder, Colorado, USA, dixon@ucar.edu
[2]Centre for Australian Weather and Climate Research, Melbourne, Australia, a.seed@bom.gov.au
(Dated: 26 August 2014)

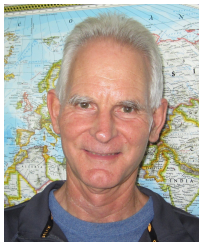

Mike Dixon

1 Introduction

The TITAN algorithm was developed over 20 years ago (Dixon and Wiener 1993) and a number of changes have been incorporated since then. Most recently, three major improvements have been investigated: (a) identifying convective regions prior to storm identification, to avoid tracking large stratiform areas; (b) improving the forecasting for larger scale features by performing 2-D spatial filtering on the reflectivity field to enhance those features; and (c) adding the option of using a field tracker (such as optical flow) to provide motion estimates for storms that have limited history.

For background, in section 2 we will describe the basic capabilities of TITAN. In sections 3 through 5 we will introduce the new features that are the focus of our current research.

2 Basic TITAN processing

2.1 Storm identification – single dBZ threshold

The most basic storm identification in TITAN comprises finding contiguous regions of reflectivity that exceed a defined threshold. In Figure 1 below, the white outlines depict the extent of storms identified using a threshold of 35 dBZ. For clarity we sometimes depict storms using an ellipse, since these are easier for the eye to interpret. The panel on the left uses ellipses for storm shape, while the right-hand panel uses polygons.

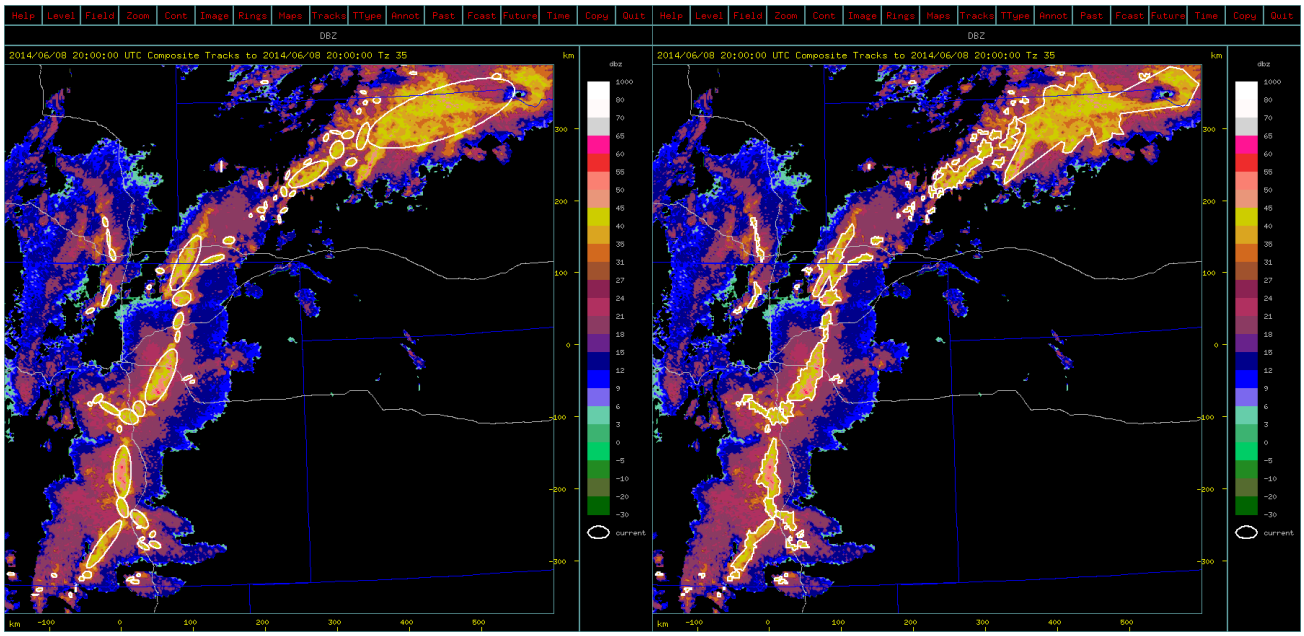

Figure 1: Storm outlines at a 35-dBZ threshold. Left: storms as ellipses; Right: storms as polygons.

These figures depict the storm shapes in 2 dimensions. However, if the input data is 3-dimensional, TITAN will identify the storms as 3 dimensional entities.

2.2 Storm identification – dual thresholds

As you can see from Figure 1, some storms just seem to 'touch' each other at a point, while others are truly merged into a single entity. In order to distinguish between the two situations, the dual-threshold identification technique was developed.

In figure 2, on the left you see the outlines of storms at the lower threshold (35 dBZ), along with dividing lines that show how the storms are broken up into sections using the dual-threshold method. On the right you see, as depicted colored areas, the sub-regions identified at the higher of the two thresholds, in this case 40 dBZ.

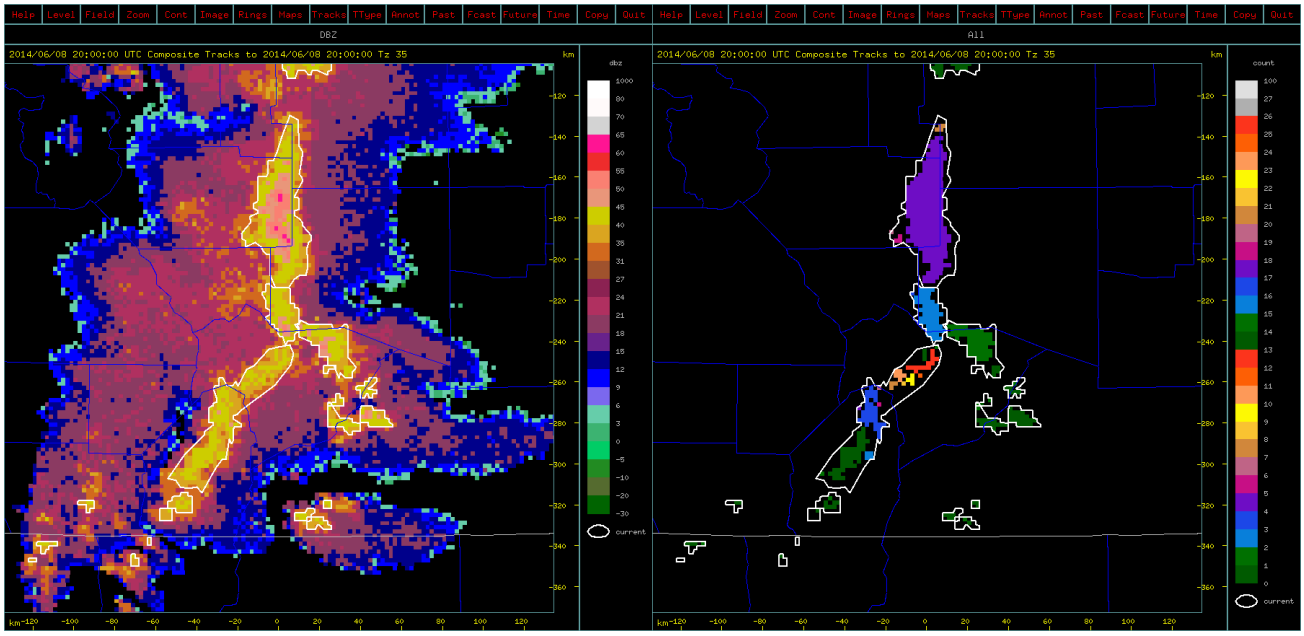

Figure 2: Left: storm outlines at a 35-dBZ threshold. Right: reflectivity clumps at a higher threshold, in this case 40 dBZ. Each colored region shows a separate sub-region at the higher threshold.

The colored sub-regions in figure 2 (right) are tested to determine whether they represent separate storms, or merely different parts of the same storm. We check that (a) the sum of all sub-regions exceeds 50% of the area at the lower threshold, (b) that each individual sub-region exceeds 5% of that area, and (c) that each sub-region exceeds some specified size (in this case 20 km^2). If these conditions are met, we treat these as 'valid' sub-regions, as depicted in Figure 3 left. We then 'grow' the sub-regions out to the original threshold outline, as shown in Figure 3 right.

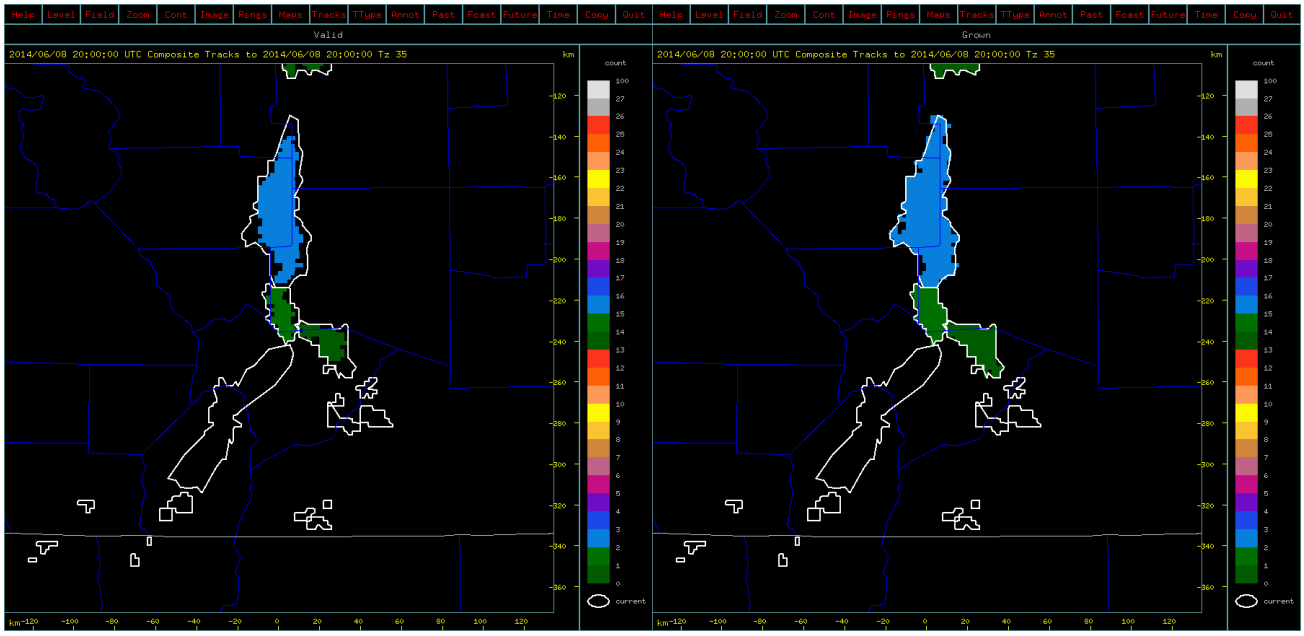

Figure 3: Left: 40 dBZ valid sub-regions. Right: sub-regions 'grown' out to the original 35-dBZ envelope.

2.3 Storm properties

Having identified the storms using the methods described above, we can then compute the properties of each storm. These include the following:

- centroid location (x,y,z)
- area and volume
- top and base
- mean and maximum reflectivity

- reflectivity profile with height, and height of the maximum reflectivity
- VIL (Vertically-integrated liquid), from the maximum reflectivity value at each height
- estimated precipitation flux and storm mass

2.4 Basic tracking

Once the storms have been identified, at sequential times, we can pair them up from one time to the next in order to associate them into tracks.

The primary tracking step searches for overlapping regions from one time to the next. Figure 4 below shows the outlines of the current storms in white, and those from the previous scan in yellow. In this particular case, with slow-moving storms, there is ample overlap between the storms from successive scans, and matching is relatively easy. In other, more challenging situations, storms may move faster and/or the time between tracks may be longer, leading to more ambiguity in the tracking decision.

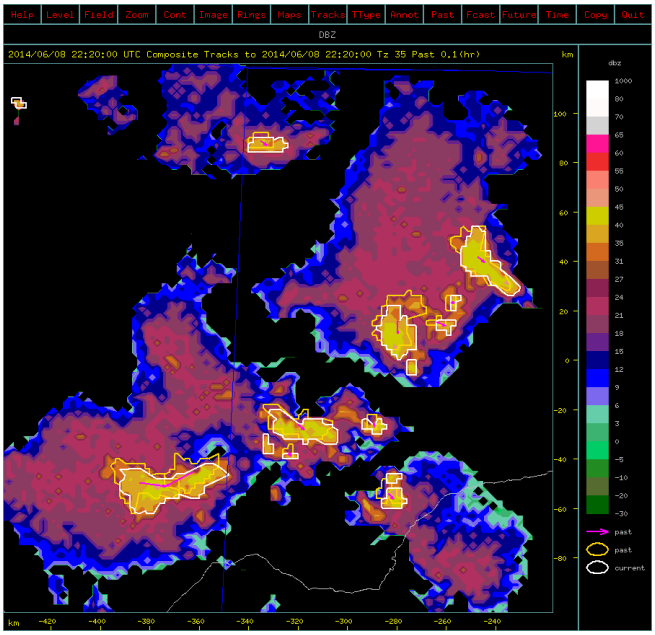

Figure 4: pairing storms for tracking using overlaps

Generally only some of the tracks will be matched using the overlap procedure, so we need a secondary step to track the storms unmatched after the first step. Figure 5 below depicts a situation in which 4 storms at time t1 (unshaded) are to be matched with 5 storms at time t2 (shaded). Some of the possibilities are discarded because they are too far apart and would therefore require motion that exceeds some set limit (say 100 km/hr) - these are shown dotted. The solid lines that remain depict the possible matches.

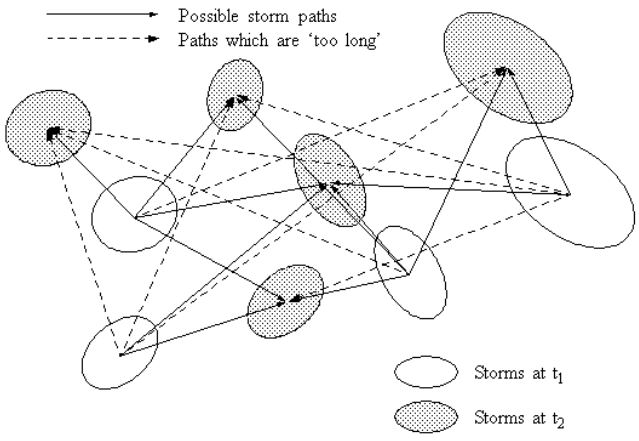

Figure 5: secondary step: matching storm pairs from time t1 to time t2 (from Dixon and Wiener 1993).

We find the best match using an optimization scheme that minimizes a cost function defined as follows:

cost = (distance between storm centroids) + (difference in (storm volume)$^{1/3}$)

This simple sum has the advantage that the terms both have units of distance. The Hungarian optimization method is used to find the optimum match.

2.5 Handling mergers and splits

As viewed by a radar, storms often merge or split. For accurate tracking and forecasting it is essential to handle mergers and splits correctly.

Figure 6 below shows diagrammatically how we identify mergers and splits. A *merger* occurs when the forecast centroids of two or more tracks lie within the observed envelope of a single storm. A *split* occurs when the observed centroids of two or more storms lie within the forecast envelope of a single storm.

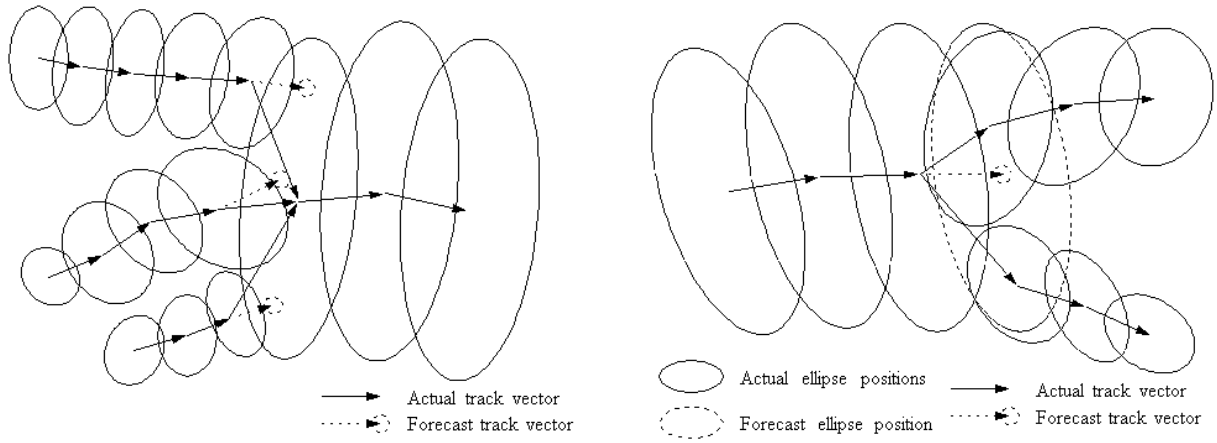

Figure 6: handling mergers (left) and splits (right) (from Dixon and Wiener 1993)

Figure 7 below shows an example of a storm that splits into 4 parts from one scan to the next.

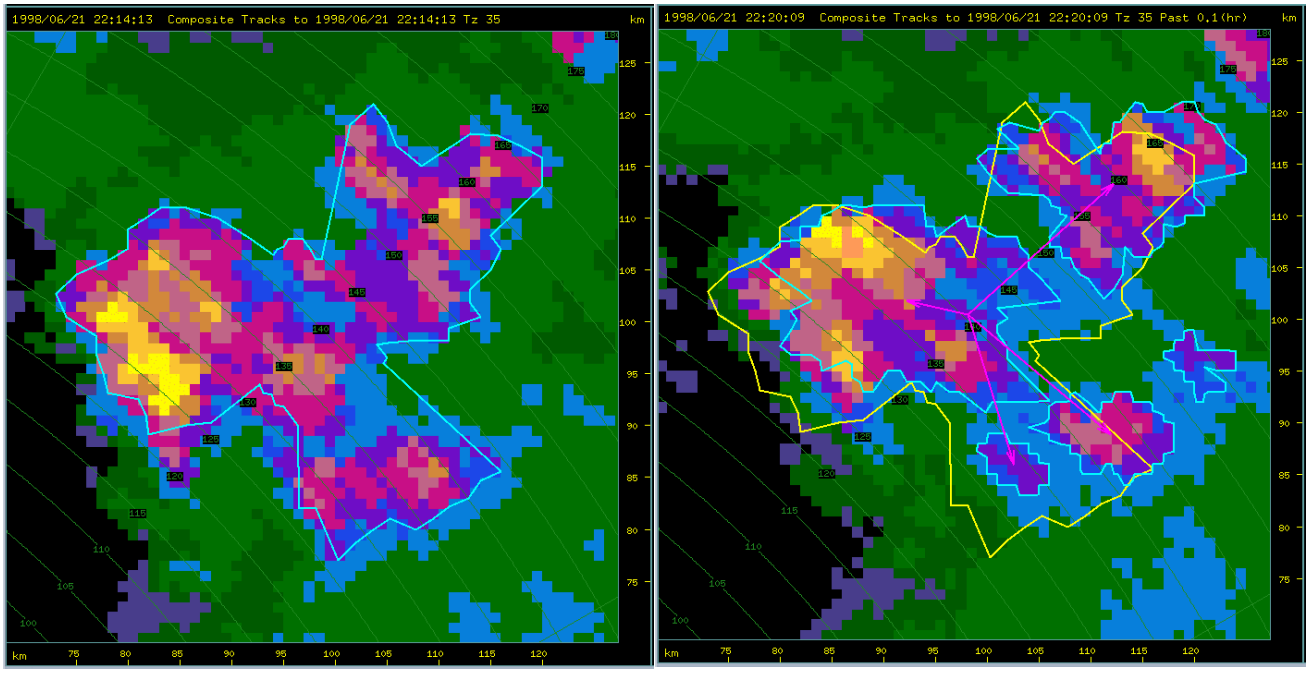

Figure 7: example of a splitting storm. The storm shown on the left breaks up into 4 parts, as shown on the right. The storm outline at the previous time is shown in yellow.

2.6 Trend-based forecasting

In TITAN, all forecasts are based on extrapolating trends in the observed storm propertiess. We perform a weighted linear fit to past observations, with the weights decreasing the further into the past we go.

Figure 8 shows examples of storm location forecasts.

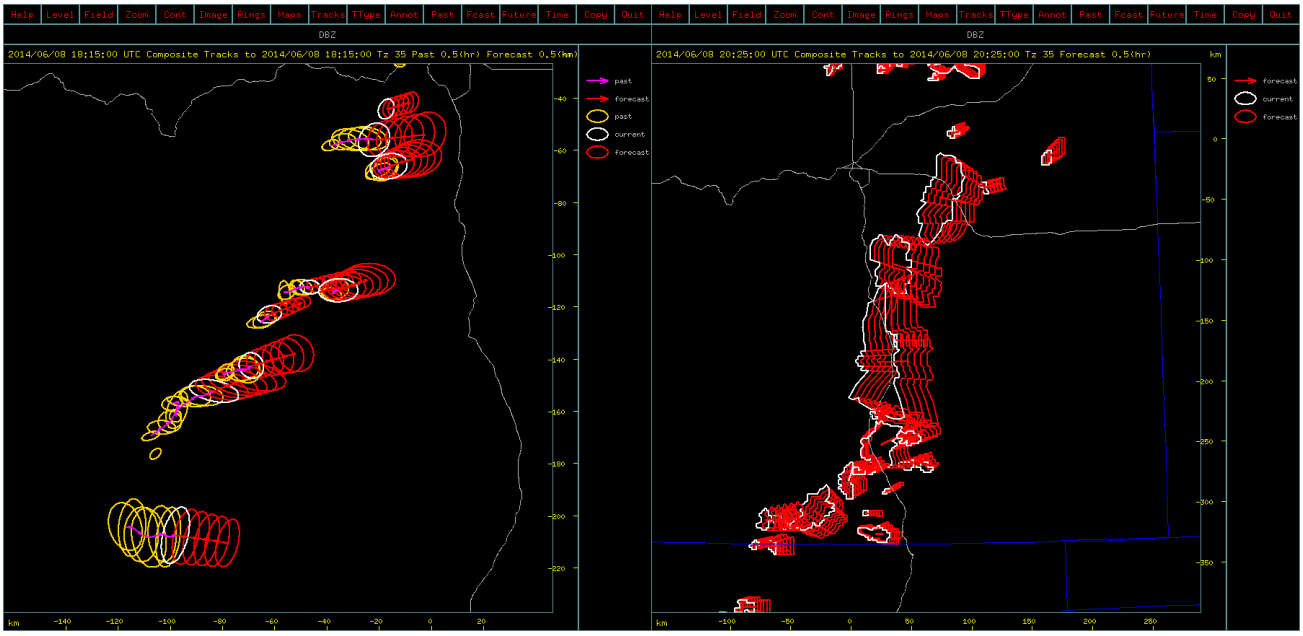

Figure 8: example of extrapolation forecasts for storm location, using ellipses (left) and polygons (right). The past is shown in yellow, the current location in white, and the forecasts locations in red. There are 5 forecasts, 6 minutes apart, out to a total of 30 minutes.

We can also apply trend-based forecasting to storm properties such as area and volume. In figure 9 left, we show the trend forecast for each time step in a complete track history, while in the right panel we show the observed time-height history of the storm maximum reflectivity.

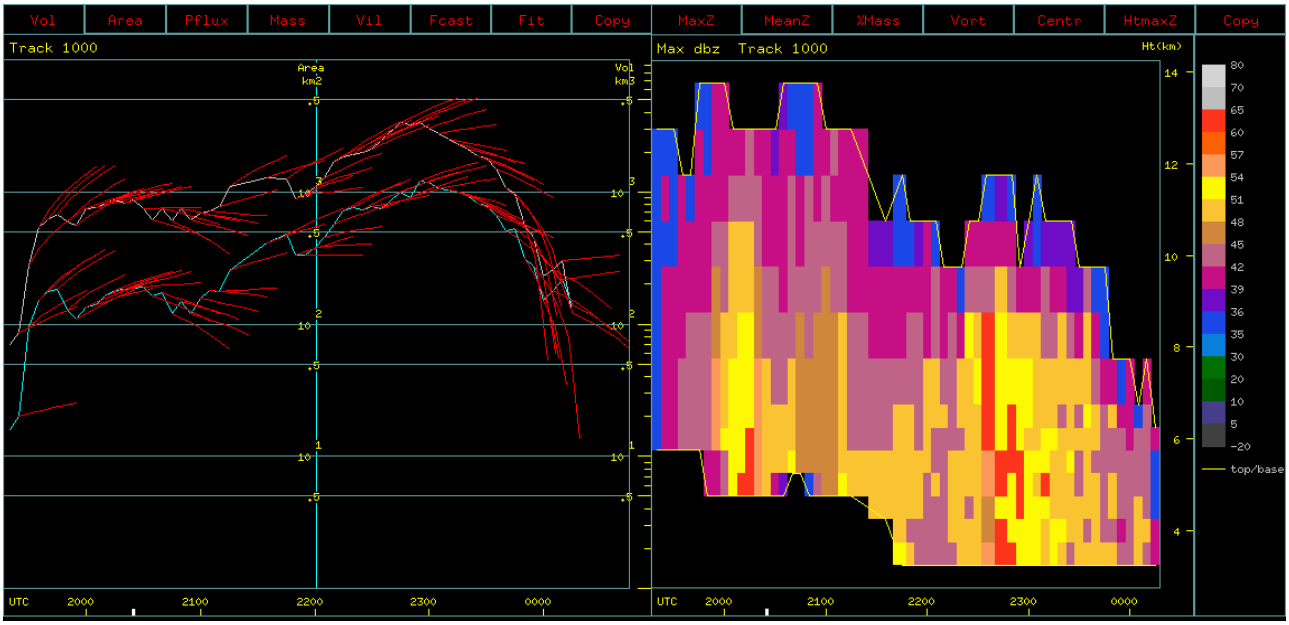

Figure 9:Left: the observations for volume (upper line) and area (lower line), along with the forecast at each scan time shown in red. Right: the observed time-height profile of maximum reflectivity.

3 Isolating convective regions for identification and tracking

Storm events are frequently made up of a combination of convective areas and stratiform areas. TITAN is generally better suited to handling convection, and is not as useful when tracking stratiform regions.

Figure 10 below shows an MCS in eastern Colorado, with a line of convection moving towards the east, trailed by a large region of stratiform rain with some embedded convection. The bright-band in the stratiform regions raises the reflectivity to above 35 dBZ, so TITAN tends to identify the stratiform region as connected to the convective line, creating one large blob. Figure 11 demonstrates this problem.

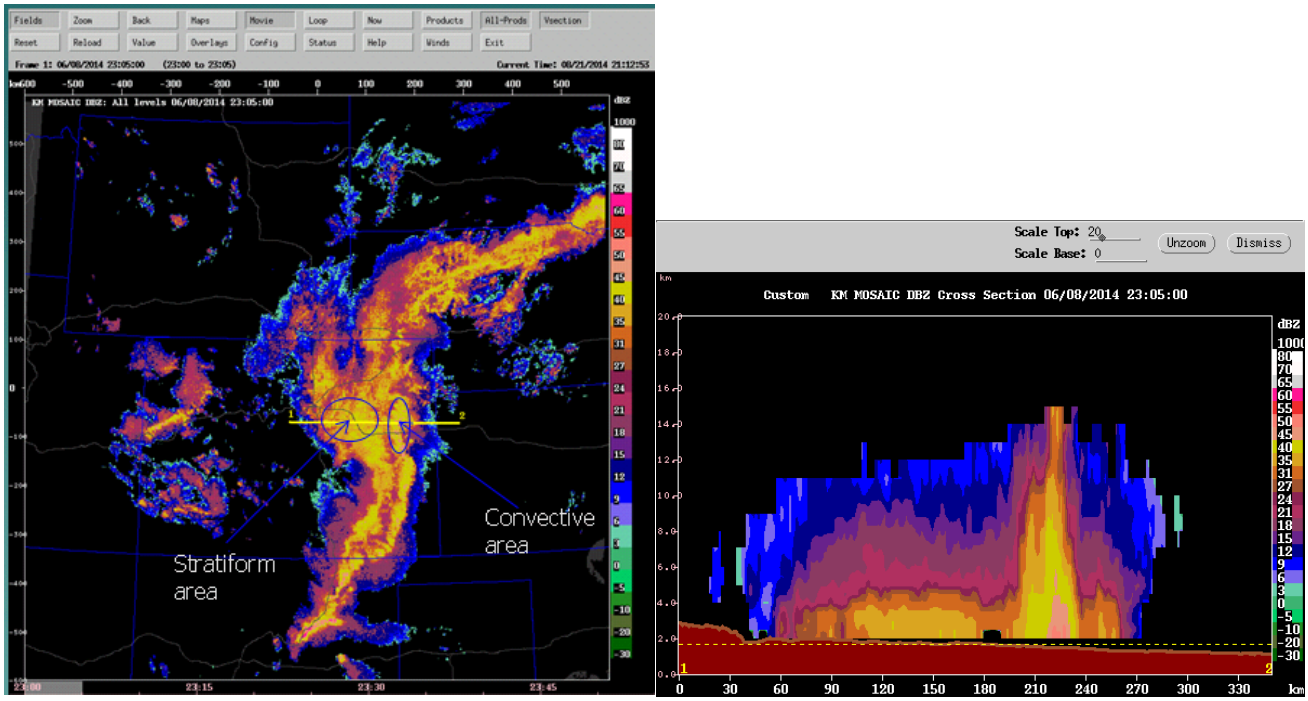

Figure 10: Mixed convective and stratiform case. In the left panel you can see the leading convection (the storm is moving eastwards) and the trailing stratiform. The right panel displays the vertical section along line (1-2), showing the convective region to the right and the trailing stratiform to the left.

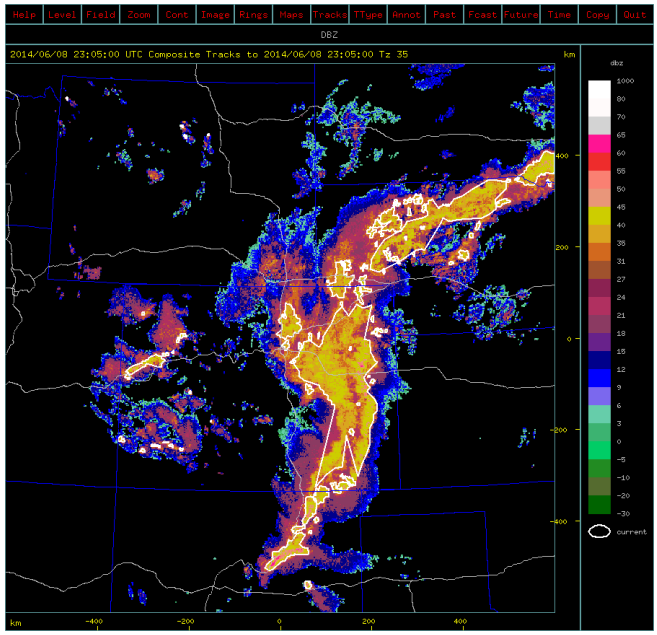

Figure 11: At a threshold of 35 dBZ, TITAN combines the convective and stratiform regions into one large storm. This is not very useful for nowcasting purposes.

In order to improve on the identification and tracking, we need to separate the convective and stratiform areas. We tested the method described by Steiner et. al (1995) on this case. The Steiner method computes the background mean reflectivity for a circle of 11 km radius around each grid point. The *reflectivity difference* is then computed as the reflectivity at a point minus the background value. Based on this difference, a circle of convection may or may not be associated with that point, with the size of that circle being a function of the background reflectivity.

Figure 12 below shows the reflectivity difference in the left panel and the convection/stratiform partition on the right. We may not have tuned this method optimally, but in this case it seems to identify some stratiform areas as convective. The red ellipse highlights the stratiform area from figure 10, and it is identified as convective in this case.

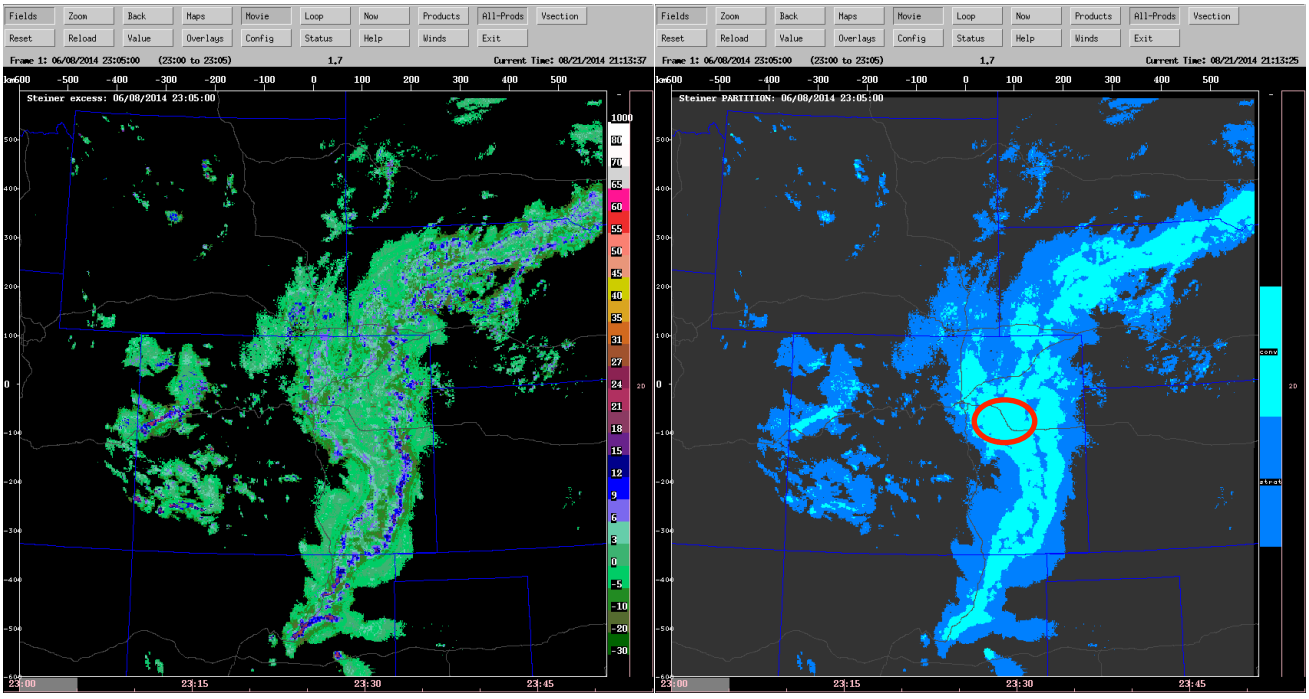

Figure 12: applying the Steiner method. Left: the 'reflectivity difference'.
Right: the convection/stratiform partition, with convection in cyan and stratiform in blue. The red ellipse shows a stratiform

As an alternative, we tried a somewhat modified method that computes a 'reflectivity texture' for a radius of 7km around a grid point, where the texture is defined as:

$$texture = \sqrt{sdev(\mathrm{dbz}^2)}$$

Borrowing ideas from Steiner, a point in the grid is considered to be convective if either (a) the reflectivity exceeds 53 dBZ or (b) the texture exceeds 15 dBZ. If a point is flagged as convective, all points within 5 km of that point are also flagged.

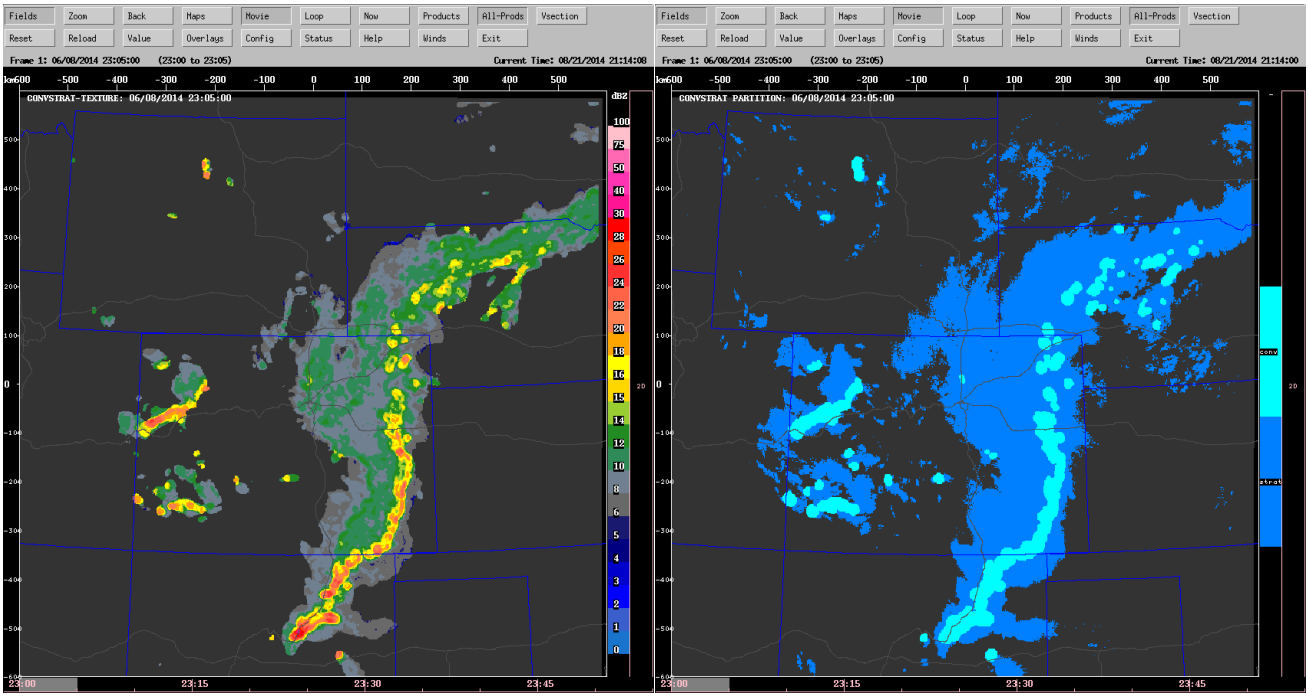

Figure 13: modified method. Left: mean reflectivity texture. Right: convective areas in cyan, stratiform in blue.

Figure 13 shows the reflectivity texture on the left and the convective/stratiform partition on the right. This method provides TITAN with the data required for tracking convective regions. Figure 14 shows the comparison of having TITAN track all regions vs. the convective regions only.

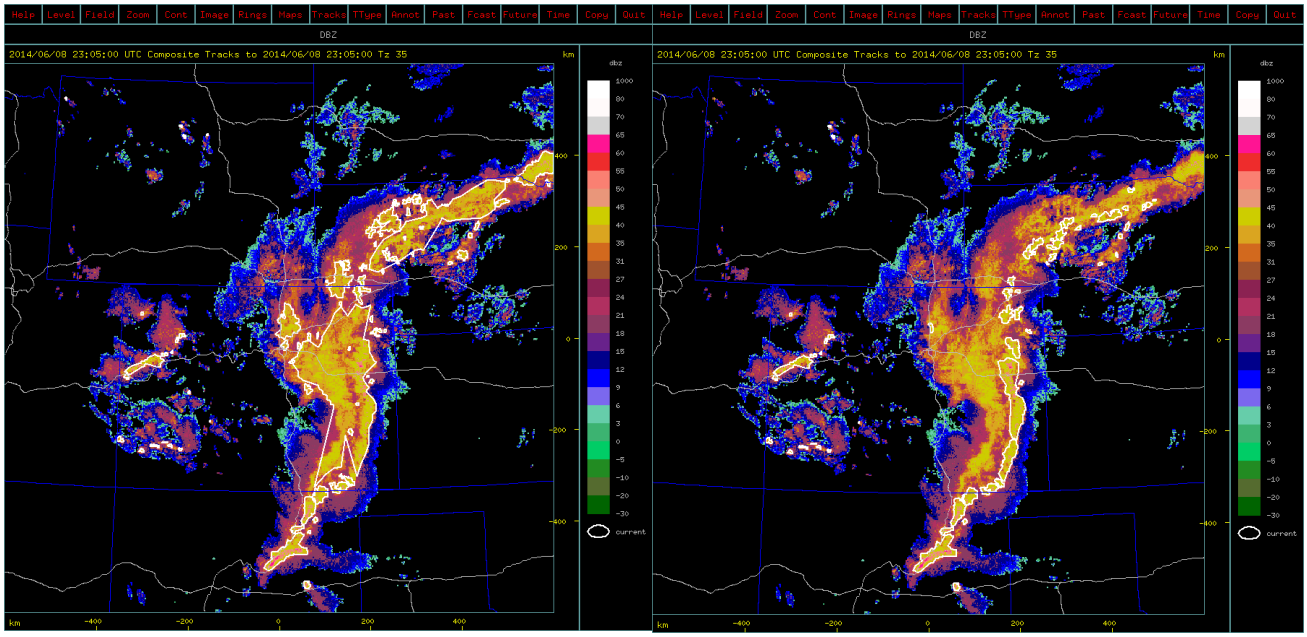

Figure 14: Left: TITAN storm identification on all regions (repeat of figure 11).
Right: TITAN storm identification for convective regions only.

Figures 15 and 16 below show a case from a single radar (near Sydney Australia) with a large stratiform region close to the radar. Without partitioning for convection, the storms identified by TITAN are amorphous and somewhat meaningless for forecasting purposes. With convective identification added, the result is much improved.

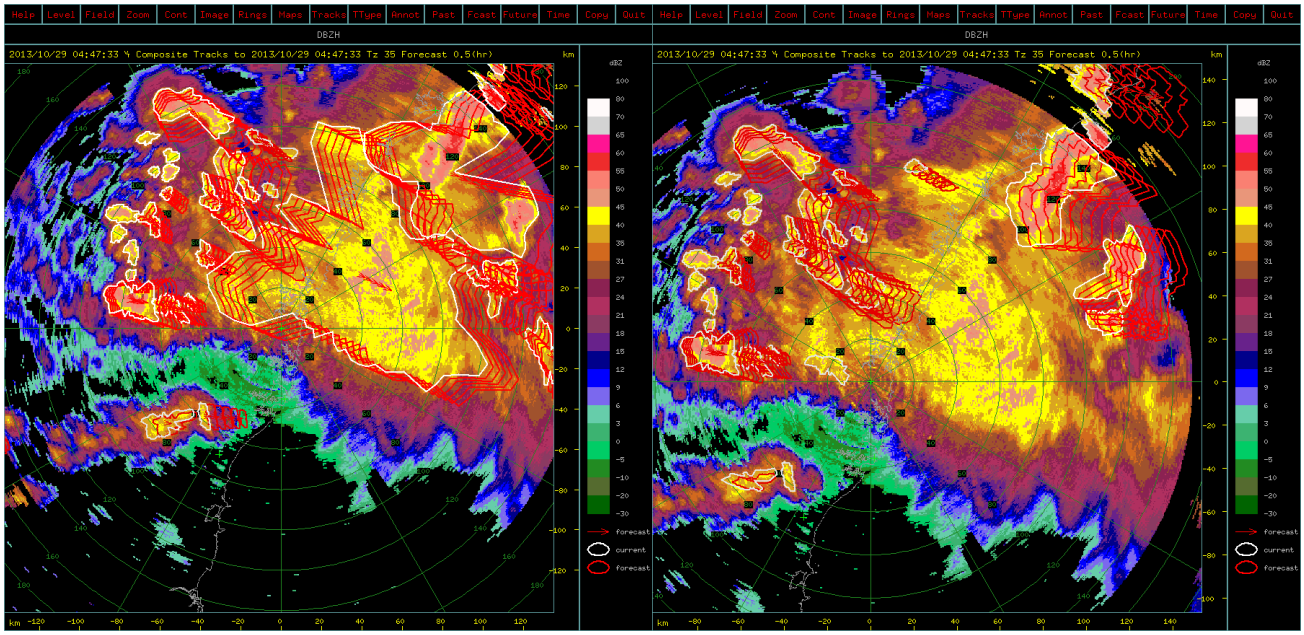

Figure 15: case for a single radar with stratiform close to the radar.
Left: using all regions. Right: using convective regions only.

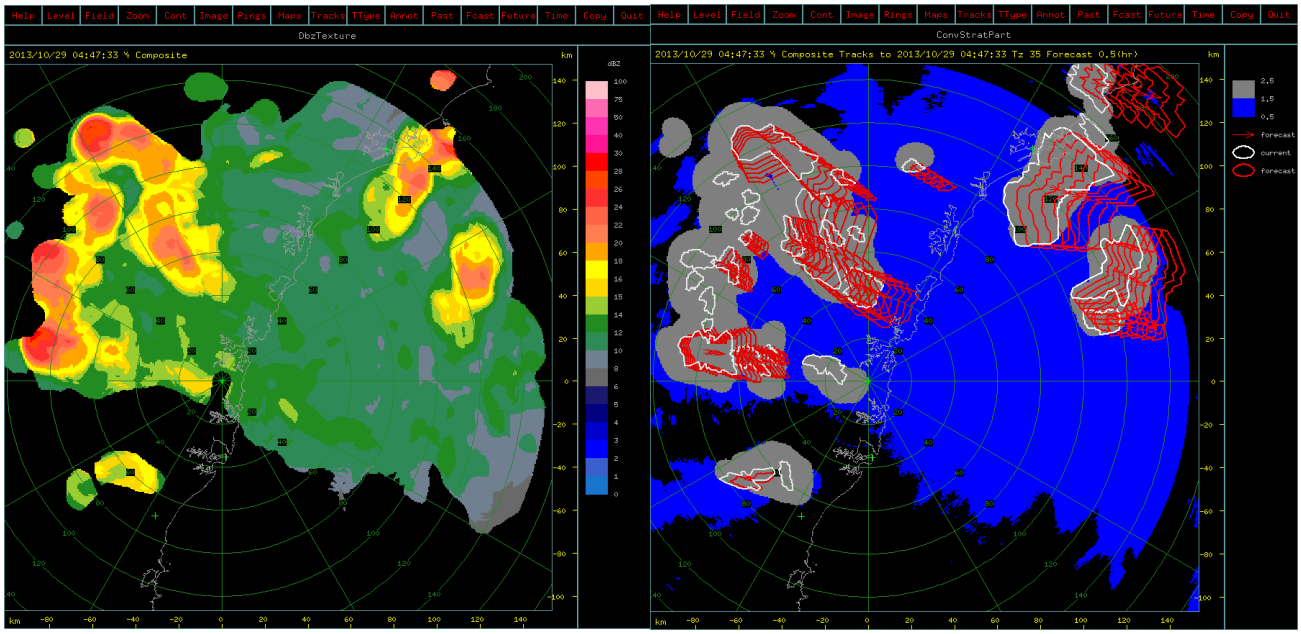

Figure 16: For single radar case from figure 15.
Left: reflectivity texture. Right: convective partition

4 Spatial filtering for longer lead-time nowcasts

4.1 Previous work

For nowcasts using extrapolation based on radar data, a 2-hour forecast is desirable but difficult to achieve. For these longer lead-time forecasts, we need to match the spatial scale of the phenomenon with the desired lead time. There is no point in making a 2-hour forecast for small-scale convection – we need to concentrate on larger systems such as fronts, lines and MCSs.

Seed (2003) showed that the correlation time calculated over the lifetime of an event varies as a function of spatial scale. Figure 17 gives some guidance about the scale of the phenomenon applicable to different forecast lead times.

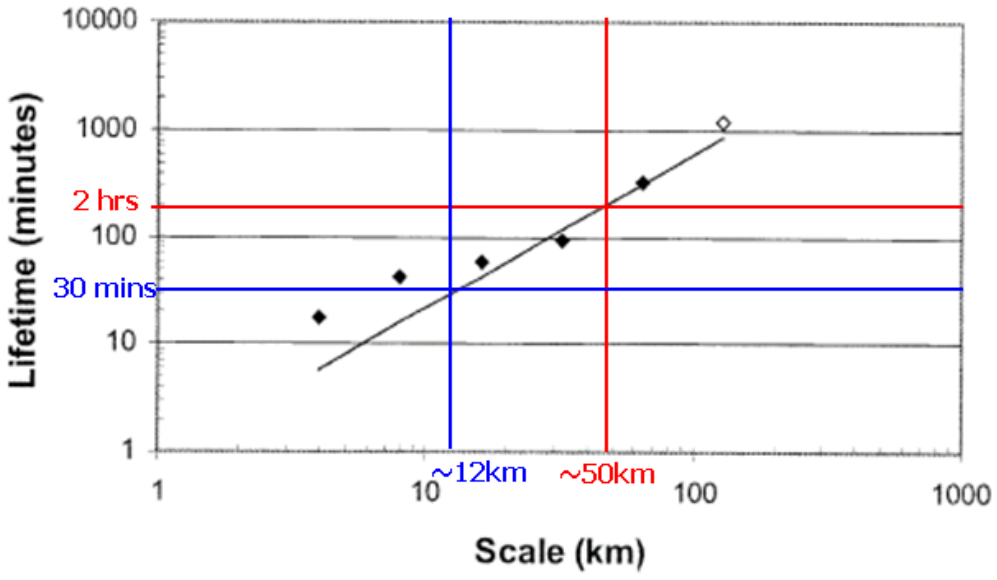

Figure 17: From Seed (2003) - median correlation time as function of spatial scale.
(Note: colored annotation added, not in original)

Similarly Germann et. al (2006) studied the spatial scale of events in terms of the event lifetimes when viewed from a Lagrangain perspective – see figure 18 below.

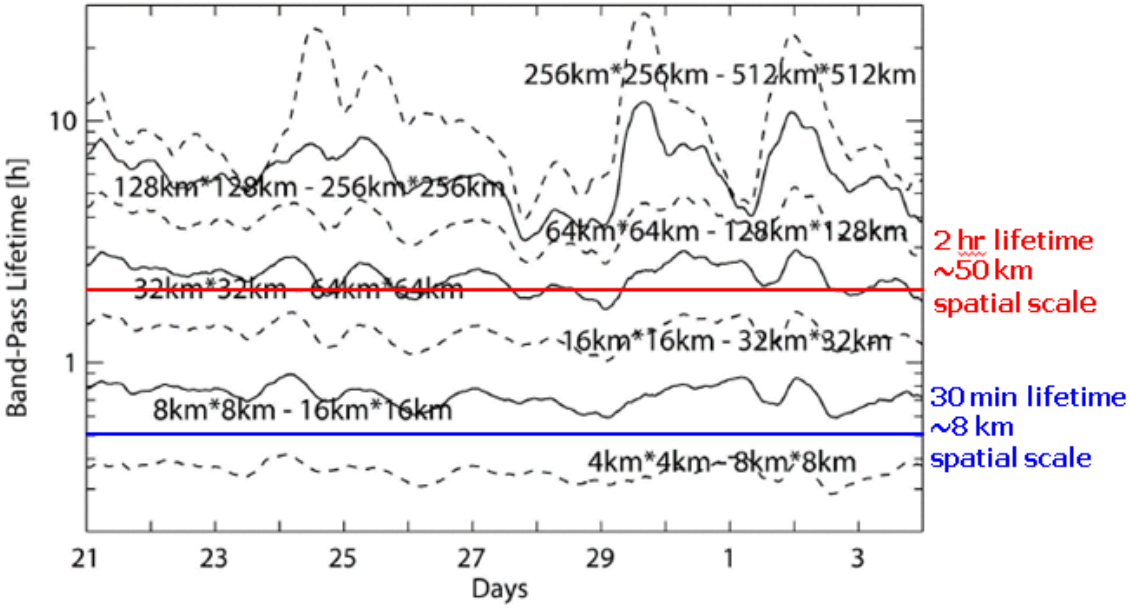

Figure 18 (from Germann 2006): Evolution of bandpass lifetime for Lagrangian persistence at scales ranging from 4 to 512 km, during a 13-day period.
Note: colored annotation added - not in original.

The red and blue annotations were added to these diagrams. From these we learn that, to make a 2-hour forecast, we need to extrapolate phenomena that have a spatial scale of around 50 km.

4.2 Filtering to match the desired 50 km spatial scale

One approach to identifying phenomena on the correct scale is to apply an appropriate spatial filter. Seed (2003) uses a 2D FFT-based filter to decompose a scene into components at various spatial scales. Figure 19 below shows our case of the squall line in Colorado, along with the 2D spectrum of that scene.

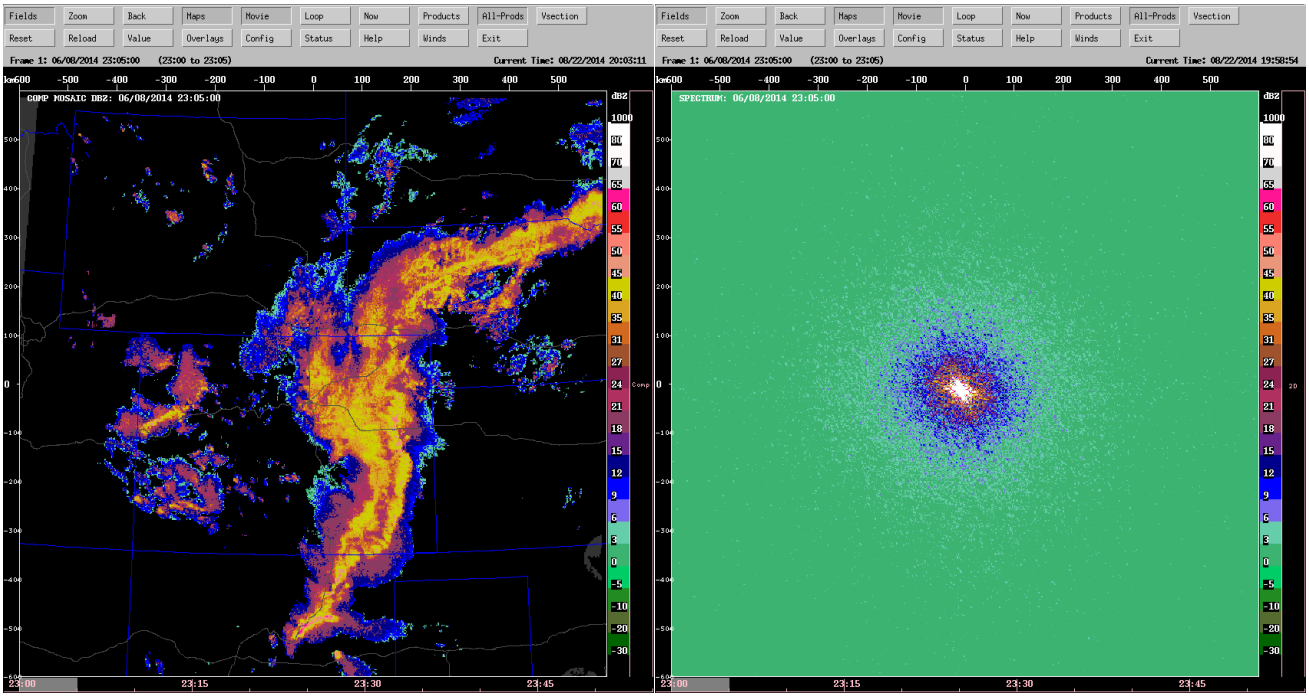

Figure 19: Left: 1200km x 1200 km reflectivity scene, at 2 km resolution.
Right: FFT-computed 2D spectrum for this data set.

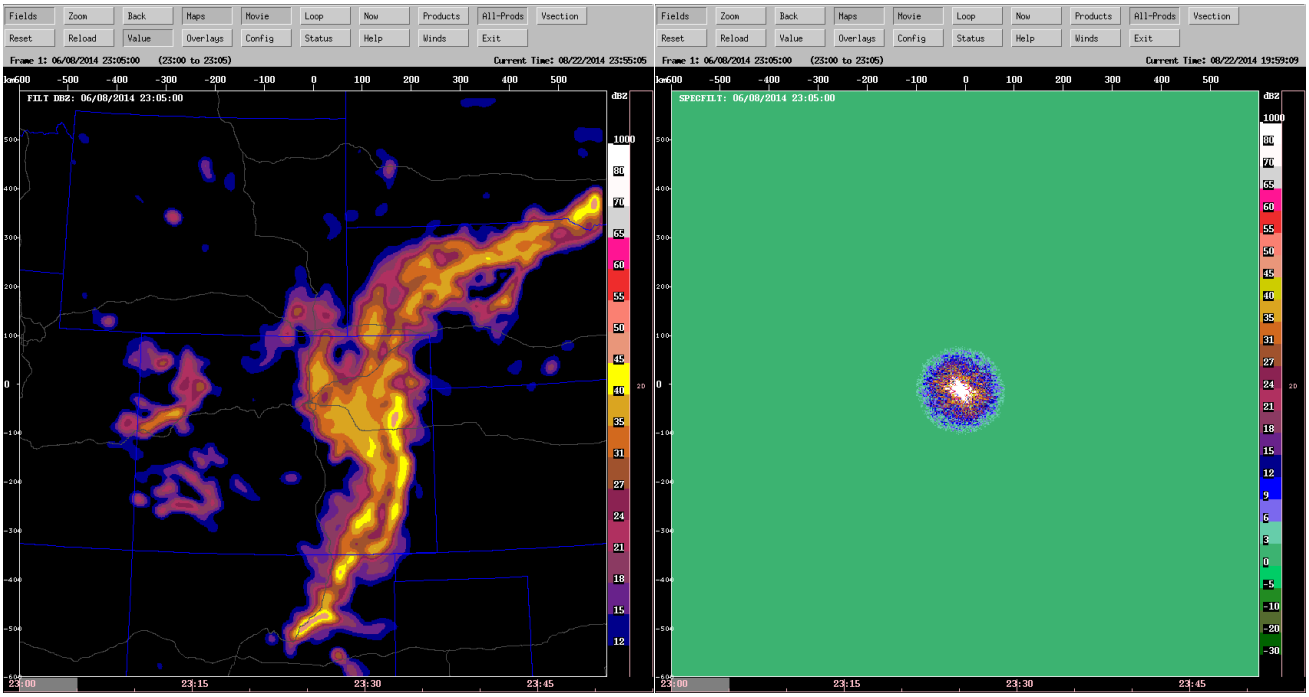

Figure 20: Left: reflectivity filtered to remove features below 50km in scale.
Right: 2D spectrum with features smaller than 50km removed.

Figure 20 shows how the reflectivity scene is altered by applying a filter that removes the smaller scale features and leaves only those above 50 km in scale. This smoothed result does highlight the larger scale features. However, it also includes the stratiform regions, which is not ideal for making forecasts of large-scale convective lines.

Going one step further, we can apply the spatial spectral filter to data containing only the convective regions. Figure 21 shows the result of doing so.

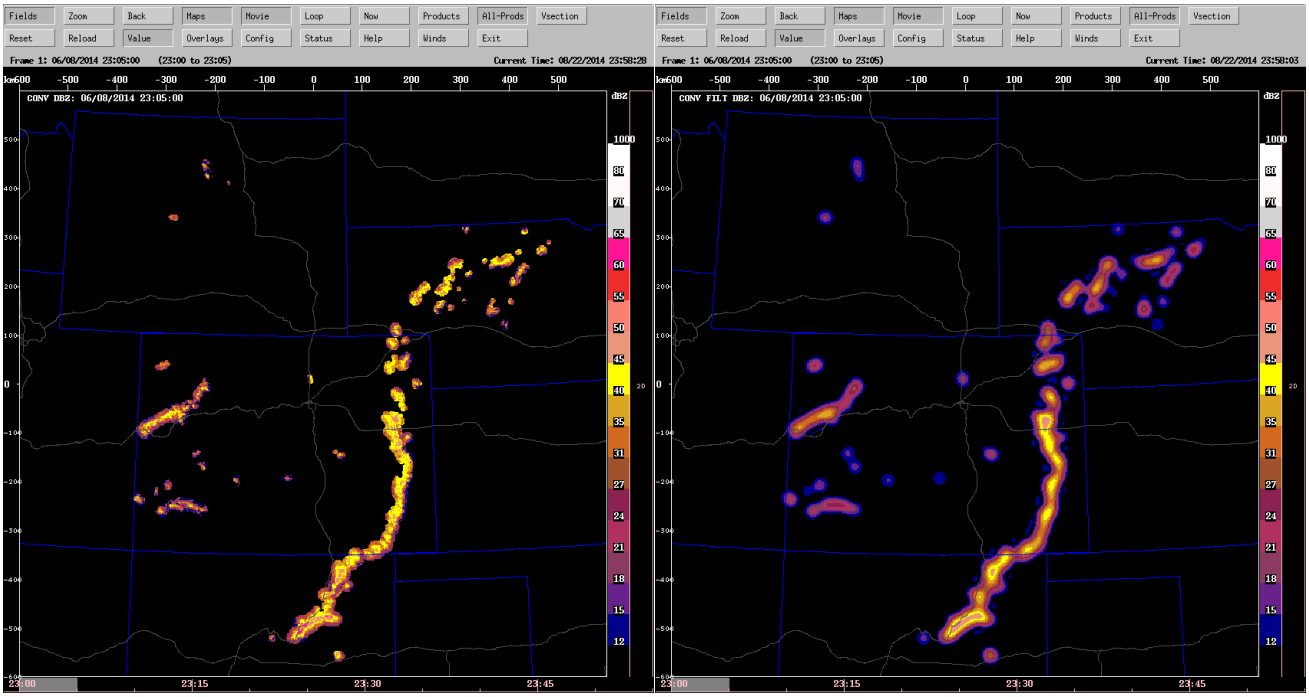

Figure 21: Left – reflectivity for convective regions only. Right: after application of the50km-scale spatial spectral filter.

Having filtered the convective areas to preserve the larger features of 50km scale and above, we can then run TITAN on this data to make longer-lead time forecasts. Figure 22 shows the TITAN identification polygons for this case, comparing the result for 'all scales' with '50km+ scale'. Note that in the spatially-smoothed cases, the value at the edges of the features is somewhat decreased, so we tracked the features using a 25 dBZ threshold instead of the normal 35 dBZ.

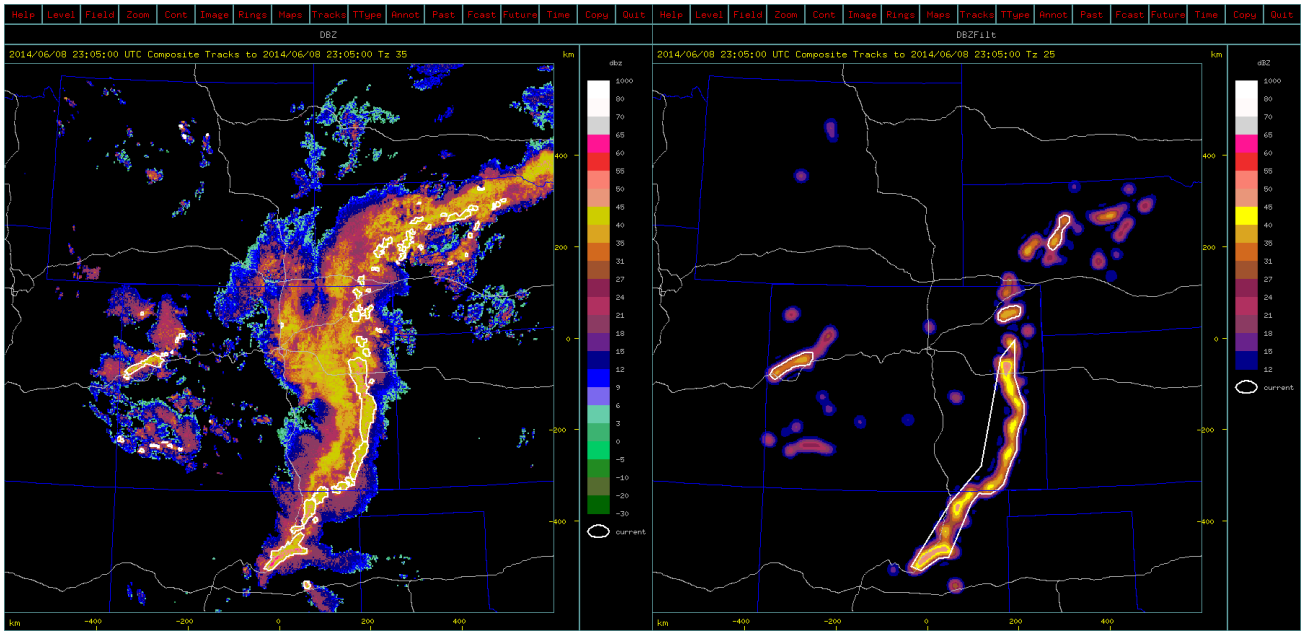

Figure 22: Left: storm identification for all convection including the small scale
Right: storm identification on large-scale features only

Figure 23 shows the 2-hour forecast made at 23:05 UTC 2014/06/08 for the larger scale features (left panel) and the verification 2 hours later at 01:05 2014/06/09. Of course we cannot claim success based on a single case. However, this result does demonstrate the possibility of making 2-hour forecasts on large-scale features, given the correct level of spatial filtering. Further work is required for validation and scoring of the forecasts.

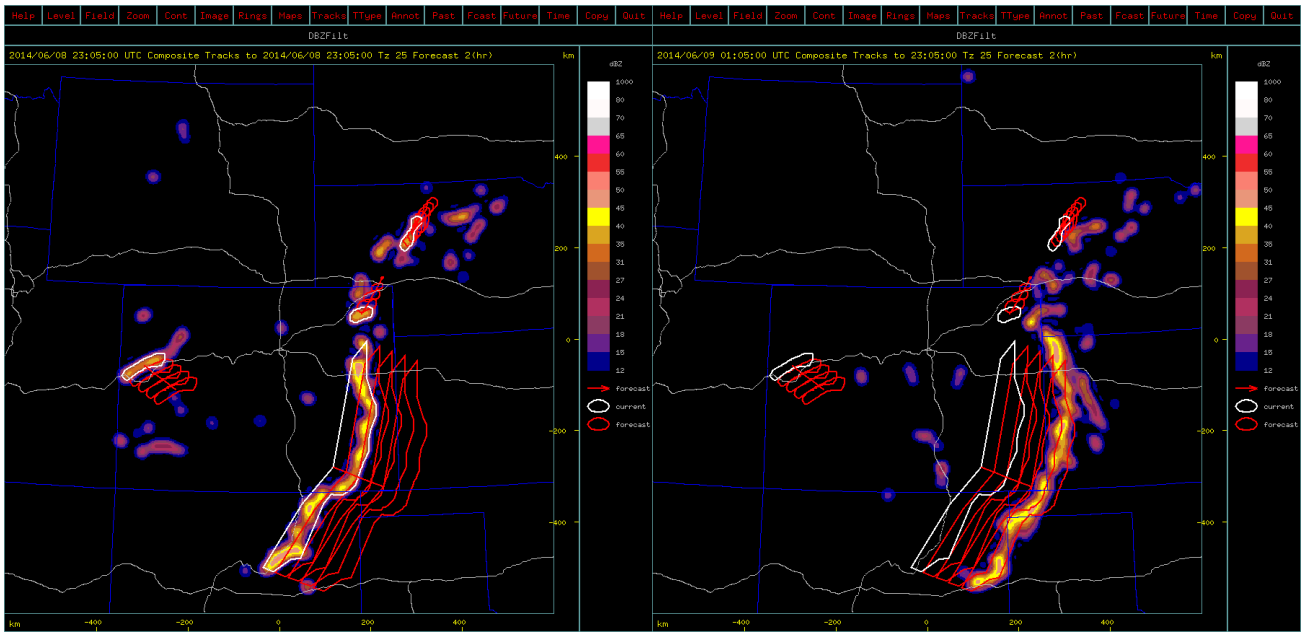

Figure 23: Left: 2-hour TITAN forecast(4 x 30 min polygons)
Right: verification of forecast using data from 2 hours later

5 Improved tracking using Optical Flow

Generally the TITAN tracking algorithm works well and produces correct tracks. However, there are challenging cases that do lead to mistakes. These generally occur when one or more of the following are true:

- storm motion speeds are high
- there is significant time between scans (10 mins or more)

- storm tracks have a short lifetime, or have a short history at the time of observation.

All of these can lead to a lack of overlap between storms from one scan to the next, so that the primary tracking step fails. This requires the tracking algorithm to rely more heavily on the secondary optimization step to match storms from time 1 to time 2. This latter method is the more error-prone of the two methods.

Figure 24 shows an example of such a situation.

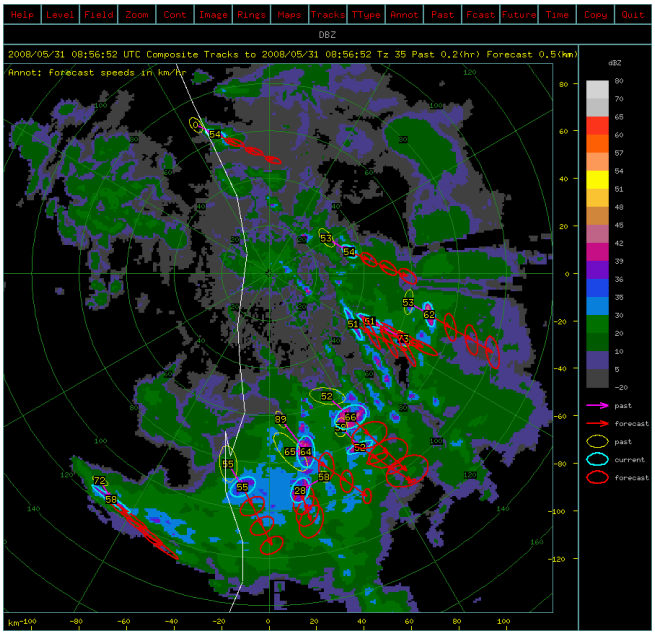

Figure 24: Example of radar scanning at 10 minute intervals, with fast moving storms. This can lead to problems with correct tracking because of lack of overlap.

One approach to solving this problem is to use a 'field tracker' such as Optical Flow to determine the vector field independently of TITAN. These vectors can then be used for storms that do not have significant history, to improve their forecasts. Figure 25 below shows the application of Optical Flow to compute field vectors.

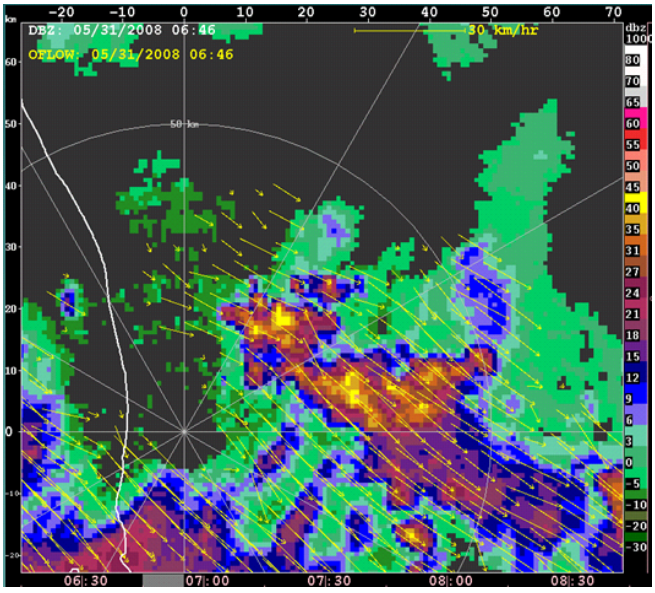

Figure 25: Using a field tracked such as Optical Flow allows us to estimate the 'background' movement of the echoes

Figure 26 shows how these vectors, when applied to short-lived storms, can significantly improve the forecasts associated with those storms.

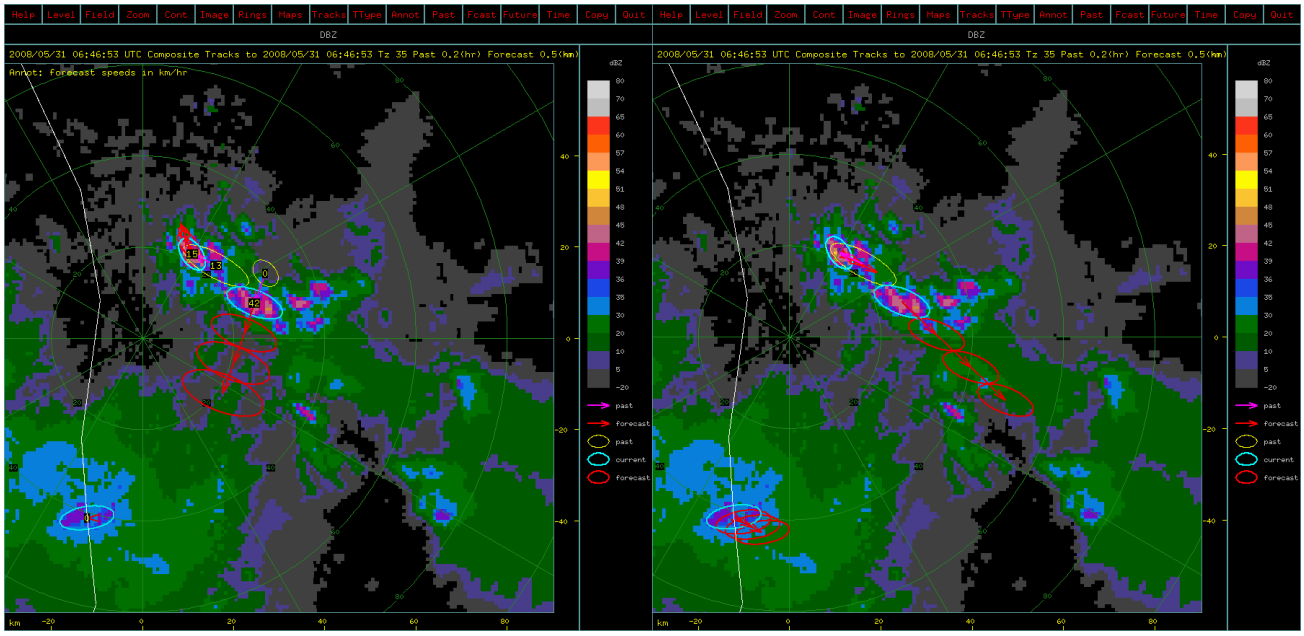

Figure 26:Left – errors in tracking, none of the storms has a correct forecast. Right: forecasts improved through the use of the Optical Flow vectors.

6 Conclusions

The TITAN storm tracking algorithm has been in use now for over 20 years, and has remained largely unchanged until recently. We introduce three new techniques that can help to improve the utility of TITAN. All of these show promise. Further work is required to validate their usefulness.

Acknowledgments

The National Center for Atmospheric Research is sponsored by the National Science Foundation. Any opinions, findings and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Dixon M. and G. Wiener, 1993: TITAN: Thunderstorm Identification, Tracking, Analysis, and Nowcasting—A Radar-based Methodology. *Journal of Atmospheric and Oceanic Technology, Vol 10, 6, 785 - 797.*

Germann U., I. Zawadzki and B. Turner, 2006: Predictability of Precipitation from Continental Radar Images. Part IV: Limits to Prediction. *Journal of the Atmospheric Sciences, Vol 63, 2092 – 2108.*

Seed A., 2003: A Dynamic and Spatial Scaling Approach to Advection Modeling. *Journal of Applied Meteorology, Vol 42, 381 - 388.*

Steiner M., R. A. Houze and S. E. Yuter, 1995: Climatological Characterization of Three-Dimensional Storm Structure from Operational Radar and Rain Gauge Data. *Journal of Applied Meteorology, Vol 34, 1978 – 2007.*