Conclusion

In conclusion, the author emphasizes the following points:

- The expert system development methodology requires an extensive set of sample problems. It encourages system developers to test their ideas against these problems early in the development life cycle.
- In the ATC setting, this requires a simulated traffic environment, possibly augmented by real-time human inputs in areas difficult to simulate.
- The current simulation is adequate for initial expert system development. Some additions will be needed to demonstrate all aspects of the first training problem.
- Continued development of an air traffic controller expert system will require continued development of the simulation to add to its functionality and to improve the fidelity with which it simulates those functions.

Discussion

George E. Swetnam, Mitre Corporation Does your simulation do anything to help present the conflicts to the controller? What you have done essentially is to replicate the information that is available to him presently on the traffic display. Has any thought been given toward showing him the conflicts in some other form that will help him grasp what the expert system is doing?

David A. Spencer Not particularly. The nice thing about this is that it is a simulator, and you can stop it. If somebody wants to analyze the situation we can stop it at some particular point and look at it, for as long as we want. We do not have to instantly present complex graphics so that a real-time decision can be made.

FAA Comment If I can add to that. Right now we are trying to see if we can get the system to work, but you are exactly right. There are major human factors issues, not only in presenting the conflict but presenting the resolution. These kinds of issues are very crudely understood.

David A. Spencer We are basically working with the problem-solving aspect of air traffic control, defining the problems, defining the solutions, and discovering how to generate reasonable solutions. How to present these to the controller and how he is to use them are complex questions that we are not equipped to handle at this time.

Robert H. Brown, NASA Johnson Space Center
Two questions. How many rules do you have, and are you bothered with garbage collection?

David A. Spencer The second one is easier. No, because we just use a large virtual memory.

Robert H. Brown, NASA Johnson Space Center
You are using the Symbolics machine?

David A. Spencer We are using the Symbolics 3670.

Robert H. Brown, NASA Johnson Space Center
You can run it long enough?

David A. Spencer You can run it long enough, yes. It can run through an entire simulation without garbage collecting.

Robert H. Brown, NASA Johnson Space Center
How long is that?

David A. Spencer An hour of real-time, 15 minutes if you run it in fast time mode.

Robert H. Brown, NASA Johnson Space Center
How many rules?

David A. Spencer We have not really gotten into the expert system. The expert system portion was shown dotted for a reason. At the time the display was made it did not exist. At this point in time it is a shell. The interfaces have been put in. We are now implementing some rules for finding path intersections, lines and line intersection points, that sort of thing. It is at that level at this point. We do not have rules that actually implement air traffic control.

Robert H. Brown, NASA Johnson Space Center
Do you have an estimate of how many rules?

David A. Spencer There is some feel for that. Several people, some here, in fact, have demonstrated that a relatively small number of rules can handle surprisingly complex cases, on the order of 20 to 30 rules.

Curtis A. Shively, Mitre Corporation We have 100 rules.

EXPERT SYSTEM FOR DOPPLER WEATHER RADAR INTERPRETATION

Steven D. Campbell
Massachusetts Institute of Technology
Lincoln Laboratory

Overview

This presentation concerns an expert system being developed for Doppler weather radar interpretation. The objective is to use artificial intelligence (AI) techniques to interpret weather radar displays to recognize wind shear hazards. The reason for developing an automatic recognition capability is that terminal Doppler weather radars will be placed at many airports to detect these hazards and it is not cost effective to put expert radar meteorologists at each of these locations. Thus, an automatic recognition capability is desired which can place a warning on the air traffic controller's screen so that these hazards can be avoided.

The approach being taken is to capture the expertise of a radar meteorologist in recognizing these hazards. Radar meteorologists exist who are very good at picking out microbursts from Doppler radar displays. The goal of this project is to understand what their expertise is, and to try to build it into a computer program so that it can be replicated at many sites.

This presentation will discuss expert systems briefly, summarize the characteristics of wind shear hazards and Doppler radar, outline the design of the weather interpretation system, and finally present some initial results.

Rule-Based Expert Systems

Rule-based expert systems have been built to perform expert level tasks in a number of different domains: medical diagnosis, VAX computer configuration, geological data analysis, and a number of other areas. These systems consist of a set of production rules in the form of condition - action (IF-THEN or antecedent-consequent) pairs, and these encode the heuristic knowledge of the system. There is a working memory that contains known facts about the situation, and an inference engine that matches those facts against the condition (IF) part of the rules. When the condition part of a rule is satisfied, then the action part is carried out.

Wind Shear Hazards

Wind shear is a change in wind velocity over some distance. When this change in velocity is large over a small distance, a wind shear hazard results. There are two kinds of wind shear hazards of primary interest here - microbursts and gust fronts. Microbursts are known to have caused crashes at New Orleans and John F. Kennedy Airports, and possibly the recent event at Dallas International Airport. Gust fronts are less hazardous than microbursts, but have a major impact on runway operations at airports because they are associated with wind shifts. The ability to anticipate such wind shifts and assess what kind of impact they are going to have on an airport's operations is very important.

Figure 1. Aircraft encounter with a microburst.

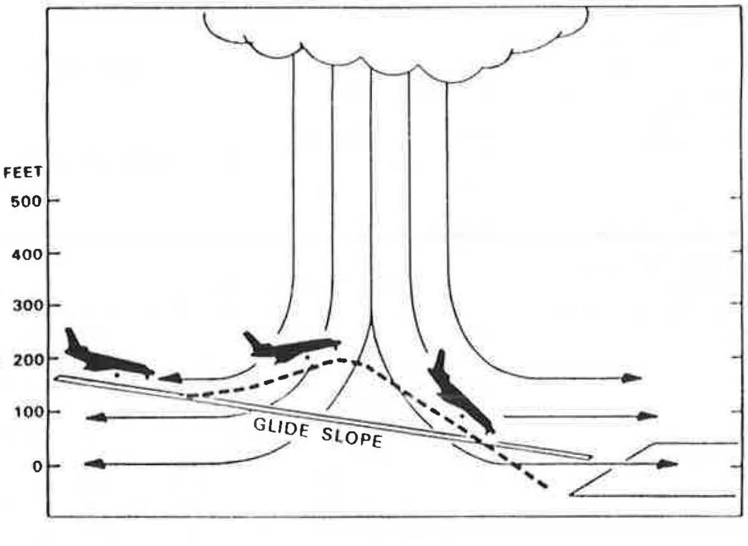

Microbursts

In a microburst there is a very strong downdraft which spreads out at the surface, and this can pose a problem for aircraft. For example, in Figure 1 the aircraft is on the glide slope and gets into a region of strong head wind as it enters the microburst. This causes the aircraft to gain lift, hence go above the glide slope. The pilot typically will try to correct for this by putting the nose down and/or reducing thrust. The aircraft then gets into a downdraft and starts losing altitude. Finally, it gets into a tailwind and loses airspeed and lift. This can cause the plane to crash short of the runway. Similar effects can also cause a crash on takeoff, as happened in the New Orleans crash.

Microbursts are very short-lived events, typically lasting on the order of 5 to 10 minutes (Figure 2). They begin in the upper atmosphere and descend to the surface where they spread out. About five minutes after their onset at the surface the most severe winds occur; they then dissipate and are usually over in ten minutes. The spatial scale of these events is on the order of 4 kilometers, at least for the initial outflow.

A model of at least one type of microburst has been proposed by Fujita of the University of Illinois and is shown in Figure 3. It involves an inflow or convergence of winds at upper altitudes, a downflow which may be rotating and then a surface outflow. This is the type of meteorological model that the knowledge base of our system needs to capture.

Doppler Weather Radar

The relevant characteristics of Doppler weather radar will now be summarized. Primary products are reflectivity, which measures rainfall rate, radial velocity, which measures the component of the wind velocity along the radar beam, and spectrum width, which is an indication of turbulence. There are also some derived products. One of them is radial shear, which is the derivative of the velocity taken along the beam and is an indication of inflow or outflow. Azimuthal shear is the derivative of the radial velocity, but taken in the azimuthal or cross-beam direction. Azimuthal shear can be used to detect rotation.

This data is collected in a set of constant elevation angle scans which are called tilts (Figure 4). A volume scan consists of several tilts which start at low elevation, then step up to successively higher elevation angles. For the terminal Doppler weather radar environment, a volume scan takes about two minutes.

Figure 2. Microburst life cycle.

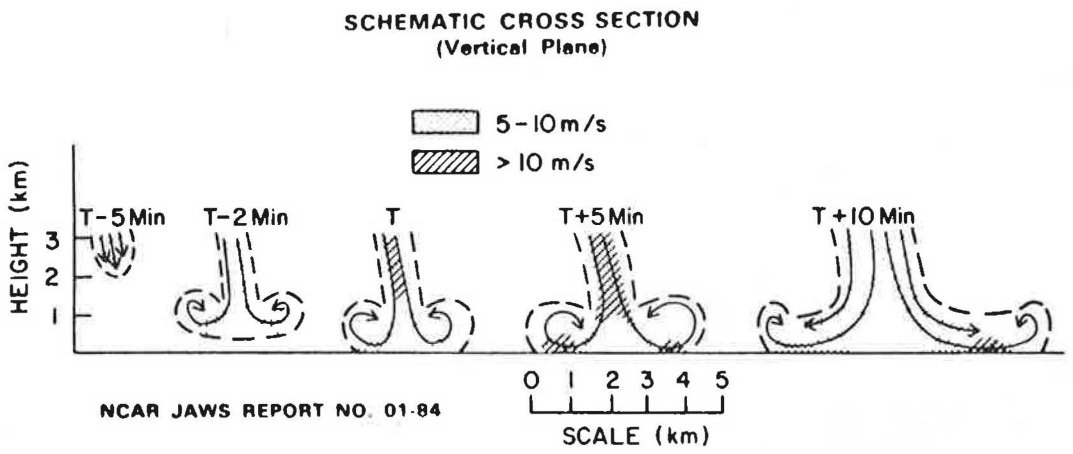

Figure 3. Model of a surface microburst.

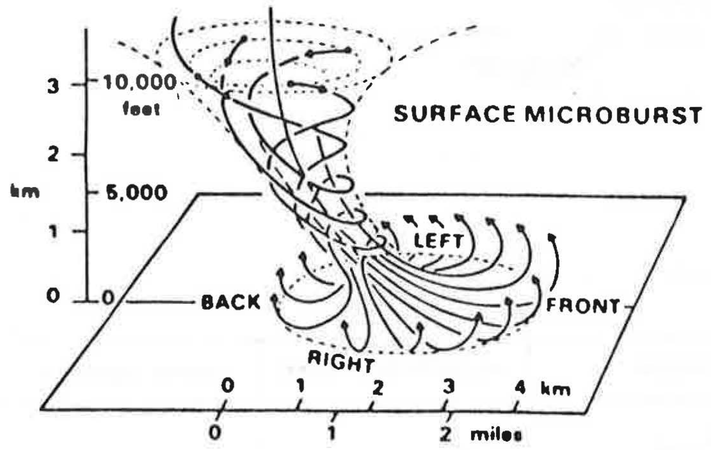

Figure 4. Single radar volume scan: A collection of tilts at various elevations.

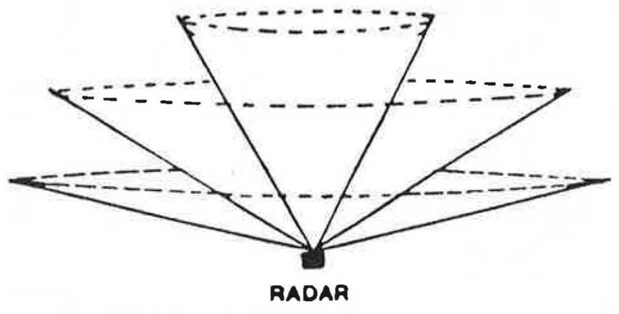

Figure 5. Primary signature of a microburst.

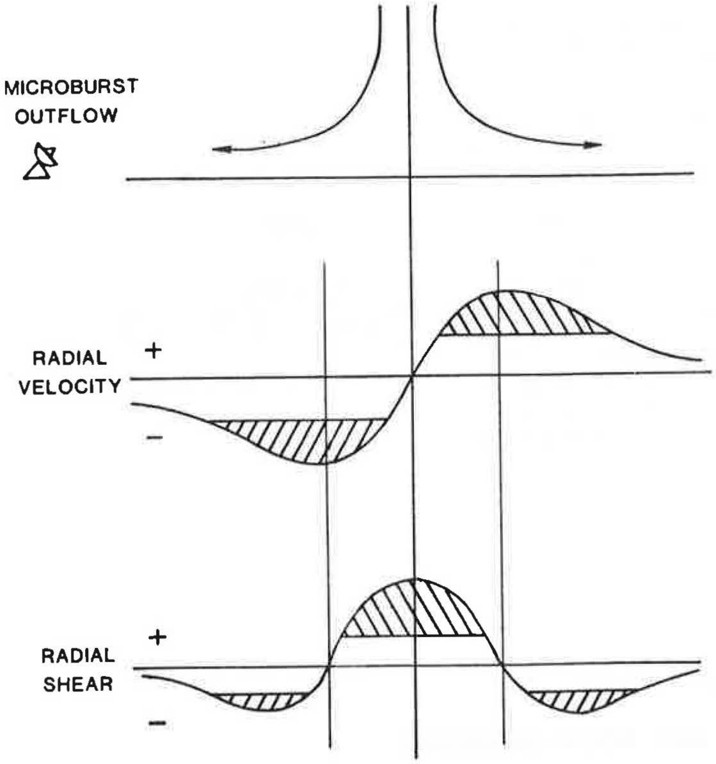

Figure 6. Summary of single-Doppler radar signatures.

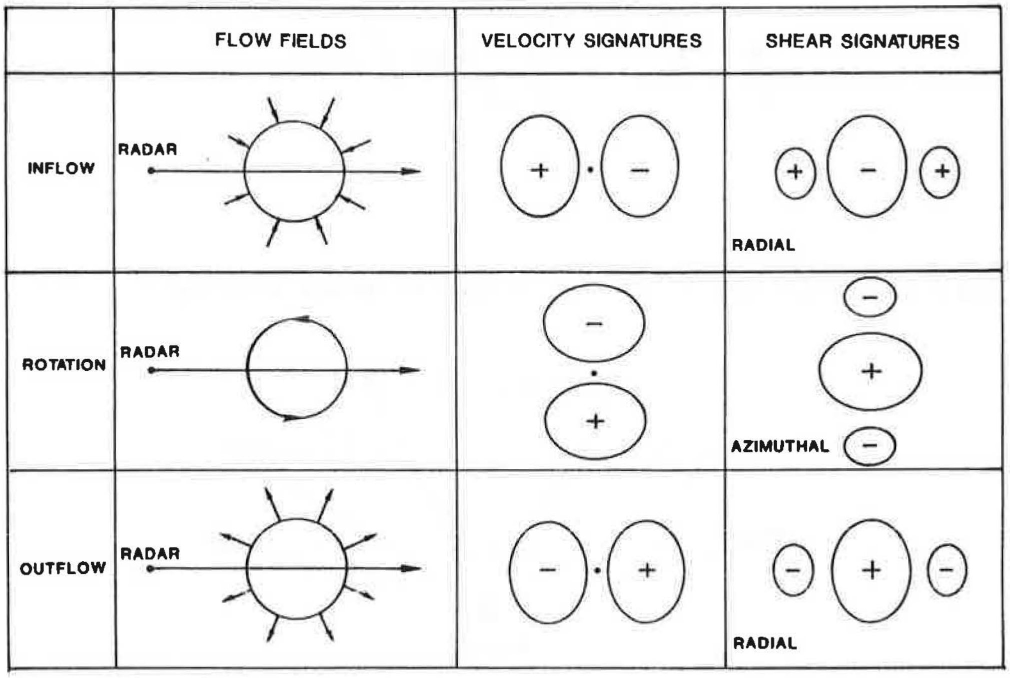

Wind Shear Signatures

The primary signature of a microburst is the surface divergence or outflow. This outflow creates a flow toward the radar (negative radial velocity) on the side nearer the radar and a flow away from the radar further out (Figure 5). If the derivative of radial velocity is taken along the radial, the resulting radial shear is first negative, then positive and then negative again. This region of positive radial shear is an indication of a strong divergence area or region of outflow.

A similar sort of analysis for inflow (Figure 6) shows first a positive velocity region and then a negative velocity region, resulting in a negative radial shear region. For rotations, the result is a velocity couplet which is oriented at right angles to the radar beam. This is where the azimuthal shear is used. As the radar beam moves clockwise, it goes from a region of negative velocity to a region of positive velocity. This gives positive azimuthal shear.

What is wanted is to tie the model of a microburst to the radar observables (Figure 7). The model shown here for a surface microburst has surface divergence, middle-level rotation and an upper level convergence. The signature for a surface divergence could be a velocity couplet or it could be a radial shear or it could be both. Similarly for rotation it could be velocity couplet or positive azimuthal shear region. For convergence it could be a negative radial shear region or a velocity couplet.

System Design

The first generation system now being developed is intended to recognize a limited set of weather hazards, microbursts and gust fronts, in non-real time. The approach is to couple a rule-based expert system to a powerful image processing package, since this particular application has a very high visual processing component.

The basic system design is shown in Figure 8. There are two components, an observer component and an expert component. It is as if there was an expert meteorologist in one room who cannot see the radar displays, and a naive observer in another who can see the displays but does not have meteorological expertise. The expert system sends queries to the observer asking about the radar data and about the features that have been extracted from it by the feature processing system. The responses by the observer are processed by the expert system, and a symbolic representation of the radar data is built up in the working memory. The expert system then operates on that symbolic representation using the production rules to recognize wind shear hazards.

The basic processing flow of the system is shown in Figure 9. First a radar data base is built up. This data base is a structured data base consisting of a set of volume scans. Each volume scan is composed of tilts, each tilt is composed of a set of radar products (reflectivity, velocity and so forth, plus derived products) and each product consists of a set of Cartesian resampled pixels.

Figure 7. Model of Denver microburst.

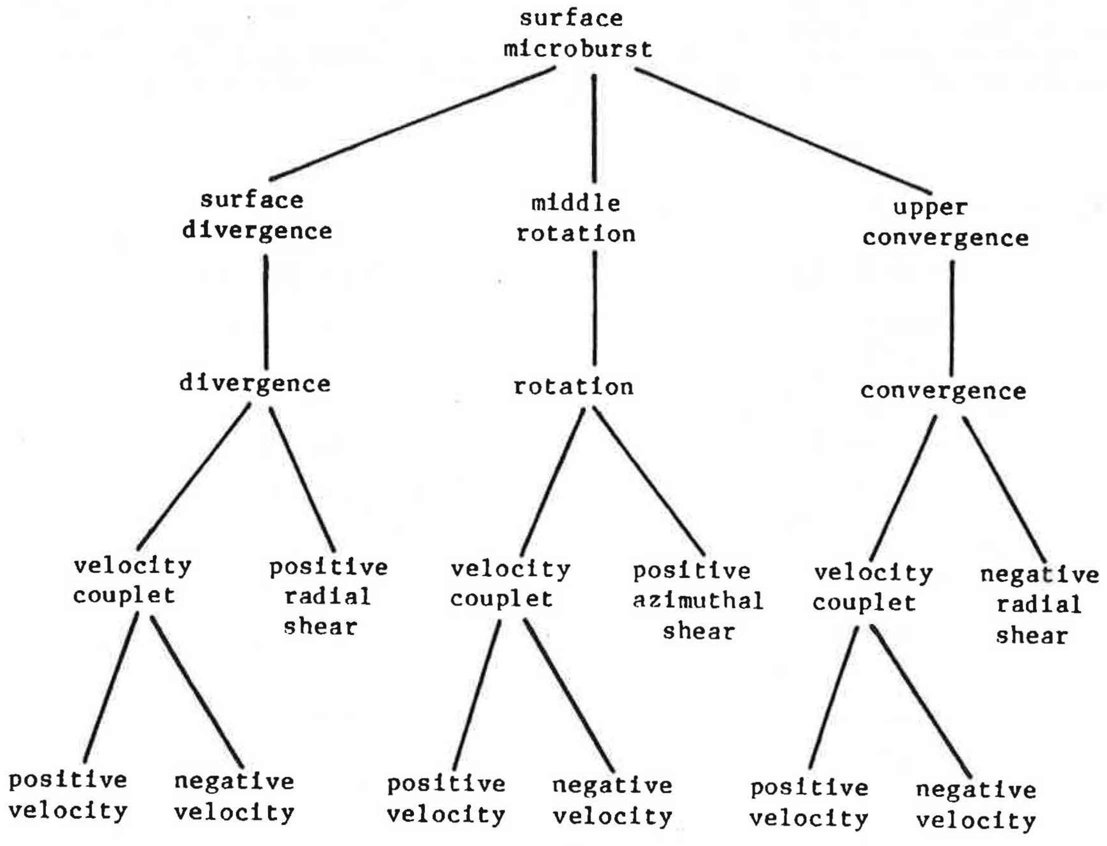

Figure 8. WX1 system design.

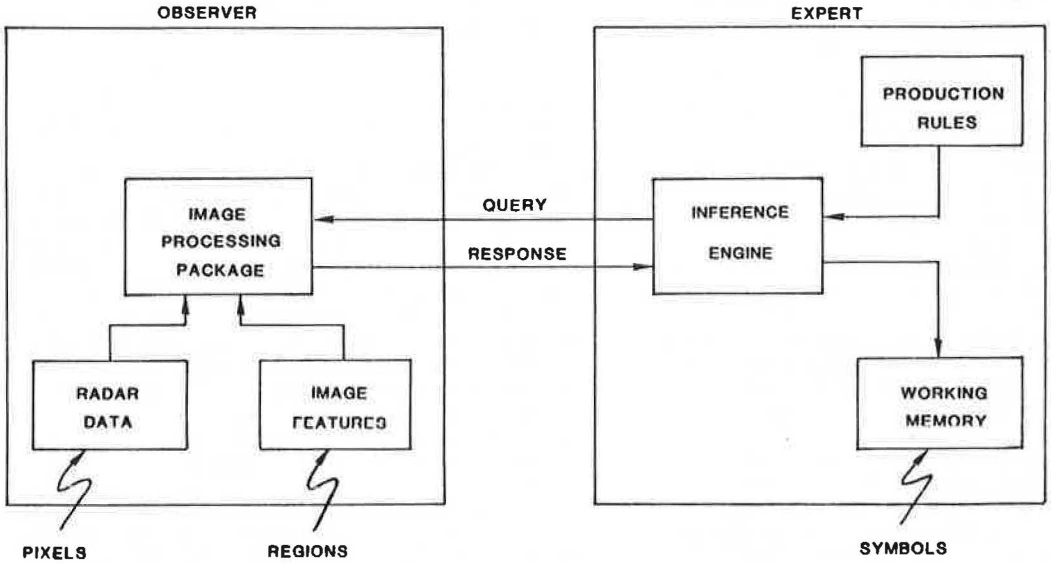

The next step is a feature extraction process. First, the pixels are thresholded to form classes, such as positive velocity, negative velocity, positive azimuthal shear and so forth. Next, connected regions are found and labeled as features. Then the program starts building up what we call a feature data base in which features are assembled into more complex features. For example, putting a positive velocity feature and a negative velocity feature together to form a velocity couplet creates a tilt feature. Tilt features are combined to create volume features such as microbursts. An example of a data set feature would be a microburst that is recognized over several successive volume scans.

Figure 10 shows an example of a feature as it is represented in an object-oriented programming system on the Lisp machine. This is a connected region of a particular class, in this case a positive velocity feature. This feature has instance variables attached to it, such as the feature label, class, number of pixels, etc. Also computed is the bounding box, which is defined by the maximum and minimum X and Y values of the object. This is very useful in expediting processing. There are also many methods (message handlers) which are attached to features. Messages can be sent to a particular feature to ask it, for example, to return its centroid, its size, its length, and so forth.

Figure 9. Basic processing flow of WX1 system design.

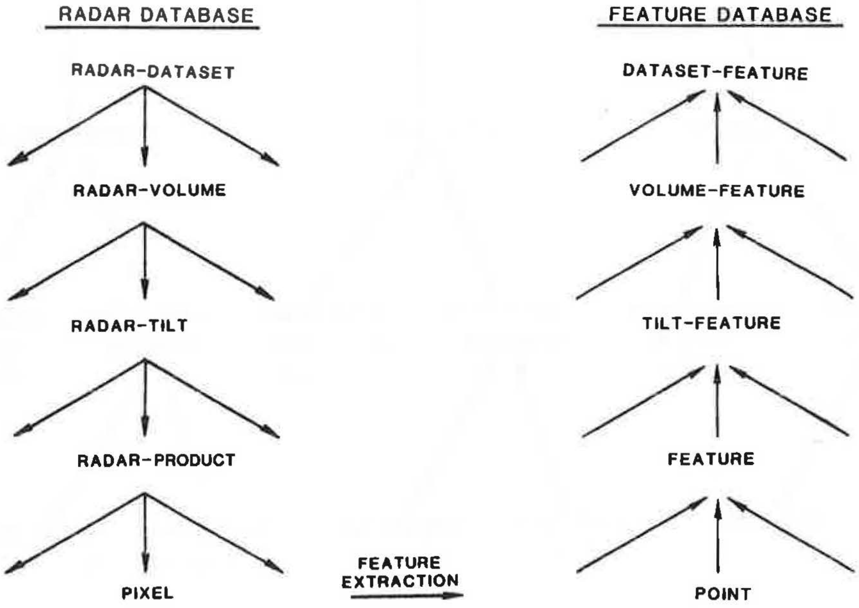

Figure 10. Example of a feature represented on the LISP machine.

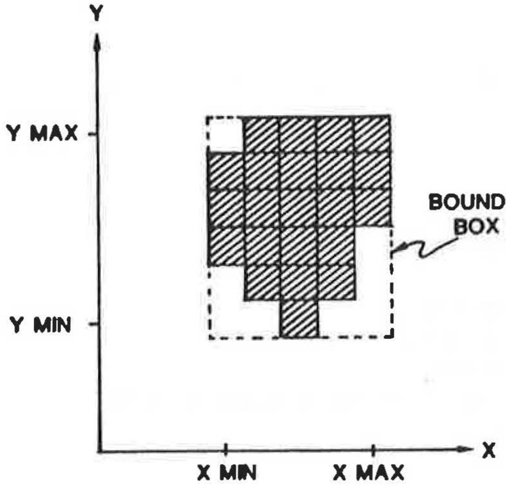

```
#<FEATURE 1>, an object of flavor FEATURE,
 has instance variable values:

 RADAR-FIELD: #<RADAR-FIELD VELOCITY 1>
 FEATURE-NUMBER: 1
 FEATURE-SIZE:    30
 FEATURE-CLASS:   :POSITIVE-VELOCITY
 XMIN: 162
 XMAX: 166
 YMIN: 130
 YMAX: 135
```

The linkage between the radar feature data base and the symbolic representation in the production system's working memory will now be described. Suppose there is an input field, say a velocity field V_1, and the feature extraction process has been performed on it (Figure 11). Suppose that positive and negative velocity features F1 and F2 are extracted from V_1. These features are represented as facts in working memory. Feature F1 is a positive velocity feature. Feature F2 is a negative velocity feature. The expert system does not know any of the details of a feature, just that it is a region of a certain class. But it can ask questions about that particular feature. Basically the expert system knows there is a blob out there, and it can ask questions about the blob.

Figure 12 shows an English language representation of rules for processing such a velocity field. After completing feature extraction, the system does an evaluation of the raw features and decides which ones are likely to be of interest (see second rule in Figure 12). Call these candidate features.

Now suppose F1 and F2 satisfy the conditions to become candidate features, a positive velocity feature and a negative velocity feature. The third rule in Figure 12 will match those features, with the variable FP matching F1 and the variable FN matching F2. The rule now tests that the distance between the two is less than 4 kilometers, and that the difference in velocity between the two is greater than 10 meters per second. If the tests are satisfied then a velocity couplet fact is created by the action part of the rule.

The creation of that velocity couplet fact triggers another rule (the fourth rule in Figure 12). The rule says that if a velocity couplet has been found and the orientation is appropriate, then label it as a divergence signature. The next rule asks

Figure 11. An expert system's radar feature data base as represented in working memory.

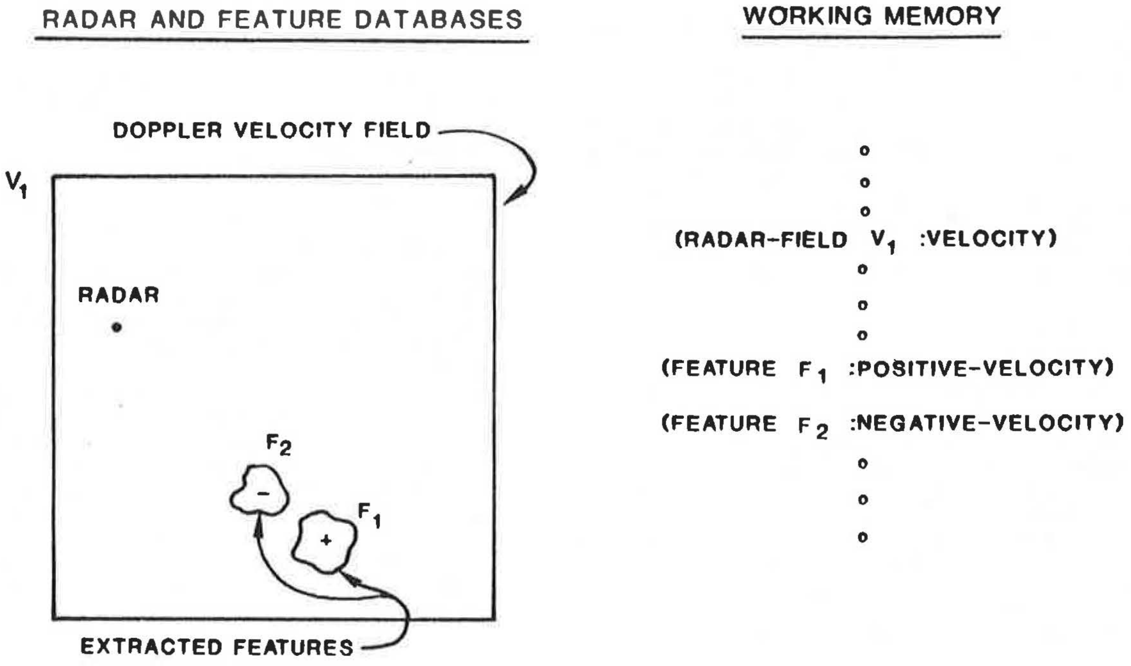

Figure 12. An English language representation of rules for processing a working memory's velocity field.

```
( If    VF is a velocity field
  Then  extract regions from VF
        for each region, create unevaluated velocity feature )

( If   FV is an unevaluated velocity feature
       size of FV is less than 4.0 km
       shape of FV is compact and not elongated
  Then change FV to a candidate velocity feature )

( If   FP is a candidate positive-velocity feature
       FN is a candidate negative-velocity feature
       distance from FP to FN is less than 4.0 km
       difference in velocity between FP and FN is greater than 10 m/s
  Then create a velocity-couplet fact from FP and FN )

( If   VC is a velocity-couplet
       orientation of VC is less than 45 degrees w.r.t. radar beam
  Then create a divergence-signature fact from VC )

( If   DS is divergence-signature
       altitude of DS is less than 1.0 km
  Then create surface-divergence fact from DS )

( If   SD is a surface-divergence
       MR is a middle-level-downdraft
       UC is a upper-level-convergence
       overlap exists between SD and MR
       overlap exists between MR and UC
  Then create surface-microburst fact from SD, MR and UC )
```

whether the divergence signature has an appropriate altitude. If it is at the surface, then a surface divergence fact is created. Finally, this line of reasoning is put together with other lines of reasoning (not shown) in a rule which says that if surface divergence, middle-level downdraft, and upper-level convergence are present, and the appropriate overlap occurs between these features, then create a surface microburst fact from those lower level features. In this way low-level features are built up into high-level features.

Initial Results

With that introduction to the methods being used, some initial results will now be presented. The original radar data for these examples has been obtained from a number of sources. One is the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. Two projects were done there, the JAWS project in 1982 and the CLAWS project in 1984. These were primarily projects to gather data on microbursts. Gust front data was obtained from the National Severe Storms Laboratory (NSSL) in Norman, Oklahoma. Finally, the MIT Lincoln Laboratory has a terminal Doppler radar program with a test radar in Memphis, Tennessee.

Microburst Example

This example shows the system's performance on one set of microburst data taken near Denver. Figure 13 shows the result after the system has performed feature extraction on the radar data [the raw radar data has not been shown due to technical difficulty in reproducing the color images]. It produces a set of candidate features which are shown in Figure 13. In the low elevation tilt, Tilt 1, it finds two big radial shear features which are likely microburst outflows. It has put a box around each. It also finds one velocity couplet region which is appropriately oriented to be a surface outflow.

In Tilt 2 it does not find the velocity couplet for the rotation because it is so asymmetric, but it does find an azimuthal shear region or region of rotation. The system relies on the overlap of the surface divergence features with the middle altitude rotation feature, and declares that a microburst exists at that location (Figure 14).

Figure 15 shows the result of processing some NCAR data where a sequence of seven volume scans was available. On four of those volume scans the system found a microburst. There were actually two microbursts. One microburst was recognized on two successive volume scans, and the system correctly

Figure 13. Candidate microburst features detected in Denver radar data.

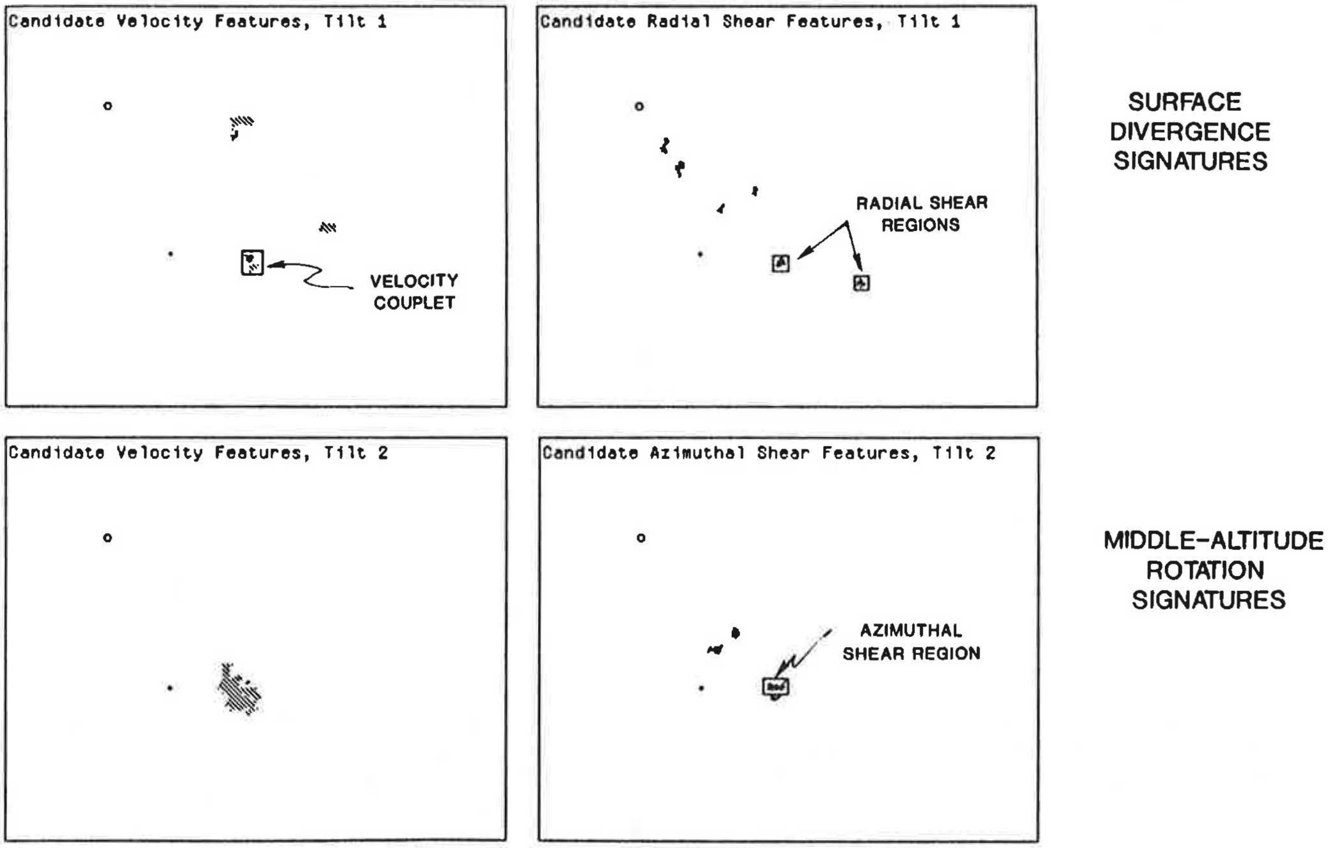

Figure 14. Recognized surface microburst features in Denver radar data.

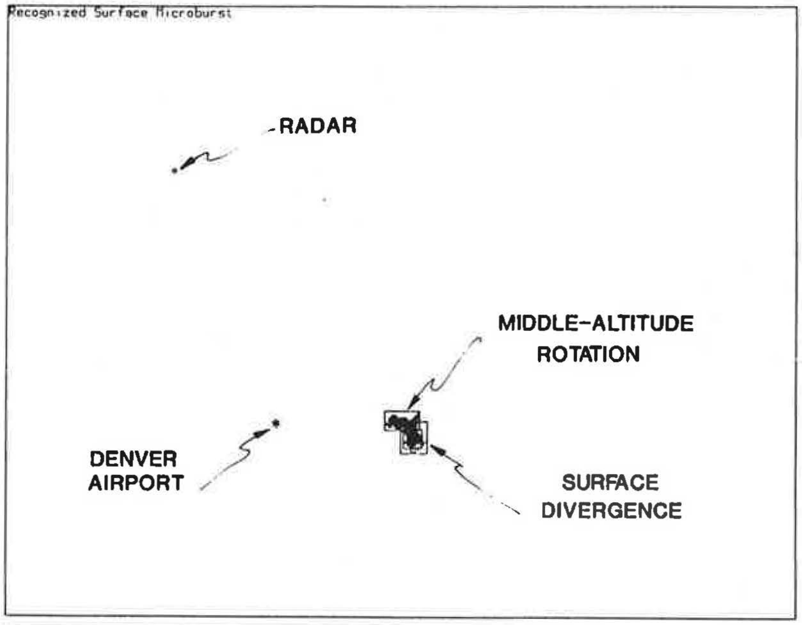

Figure 15. Additional microbursts detected in Denver radar data.

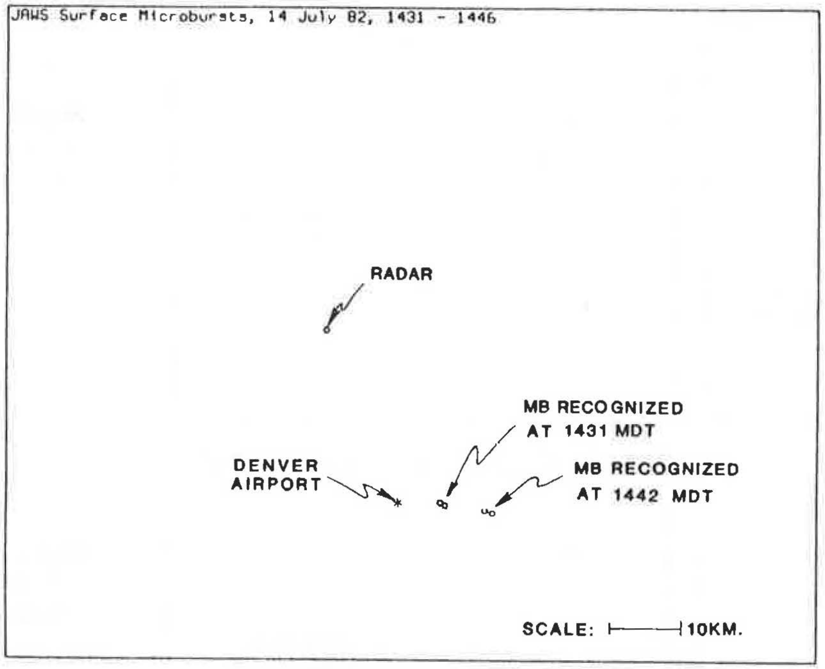

put them together as the same microburst. Subsequently it found another microburst on two other volume scans. These results have been confirmed from published reports.

Gust Front Recognition

Figure 16 is a diagram showing the origin of gust fronts. In a convective storm there is a downdraft and a region of cold air flows out of the storm. This cold air outflow running into a warm air inflow creates a region of turbulence called a gust front.

Figure 17 shows this more clearly. At the top of the figure is shown the cold air outflow colliding with the warm air inflow. A line of convergent airflow results, in contrast to the outflow caused by a microburst, and there is therefore a region of negative radial shear. The rule system looks for this region of convergence. It expects to find long, thin regions of negative radial shear.

One of the aspects of the gust front recognition system is that it is able to put together features that may become fragmented during the feature extraction process. In the hypothetical example in Figure 18 a long thin region of shear was found in Tilt 1. On the next tilt up, two separate such regions are found. The system is capable enough to go back, knowing that there was an overlapping line of convergence found on another tilt, and merge these two pieces together into one shear line.

Gust Front Tracking

Figure 19 shows the results of processing some data from the National Severe Storms Laboratory. The raw radar data has again been omitted due to the technical difficulty of reproducing color images. A squall is propagating to the east, and there is a shear line due to its gust front. About 15 minutes later it has moved to the east some distance. Four volume scans of data were available, and this figure shows a schematic representation of the results. The system was programmed to draw a line for each shear line found and to mark the centroid of that line. There is a clear propagation of the gust front to the east, and for this it is possible to estimate the propagation speed. Work is now going on to track gust fronts and predict where they will be in successive volume scans. This will be very useful in assisting air traffic control in anticipating wind shifts.

Summary

To summarize, a first generation system to interpret Doppler weather radar data is under development. It employs an expert system, coupled with a powerful image processing capability, and rules are being developed for detecting low altitude wind shear hazards. At the moment the microburst algorithm runs about six times slower than real time, and the gust front algorithm runs about four times slower than real time. The image processing calculations are the primary limitation on processing speed.

Discussion

Question Regarding the signature that you described before on the microburst, would that be the same for all over the country with the different microbursts?

Steven D Campbell No, they do not have the same structure. It turns out that microbursts vary quite a bit from place to place. Memphis, Tennessee microbursts always have a lot of rain. Microbursts in Denver usually have very little rain. When we started out I thought it was very simple -- convergence, rotation, divergence. We are finding out that it is not that simple. We just keep having to add more and more rules, and more and more ways

Figure 16. Aircraft encounter with a gust front.

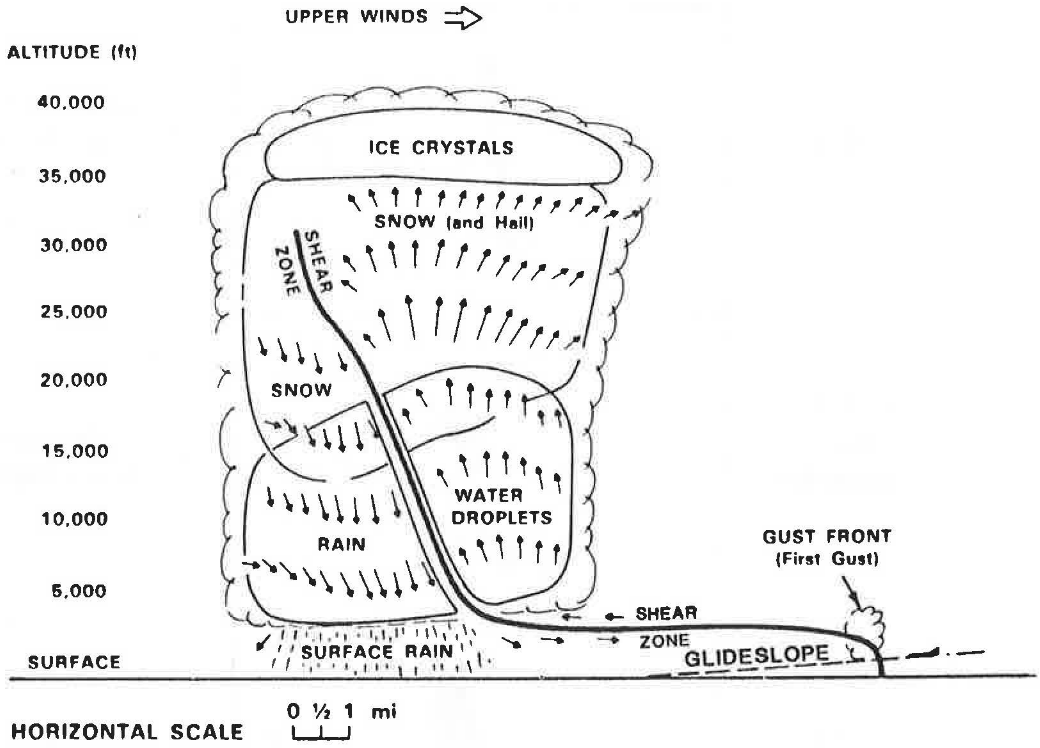

Figure 17. Gust front convergence signature.

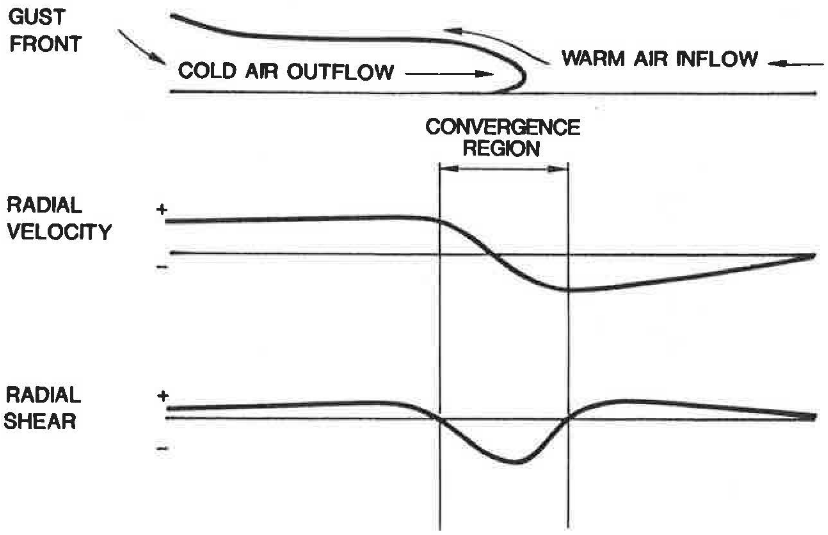

Figure 18. Merging fragmented shear features.

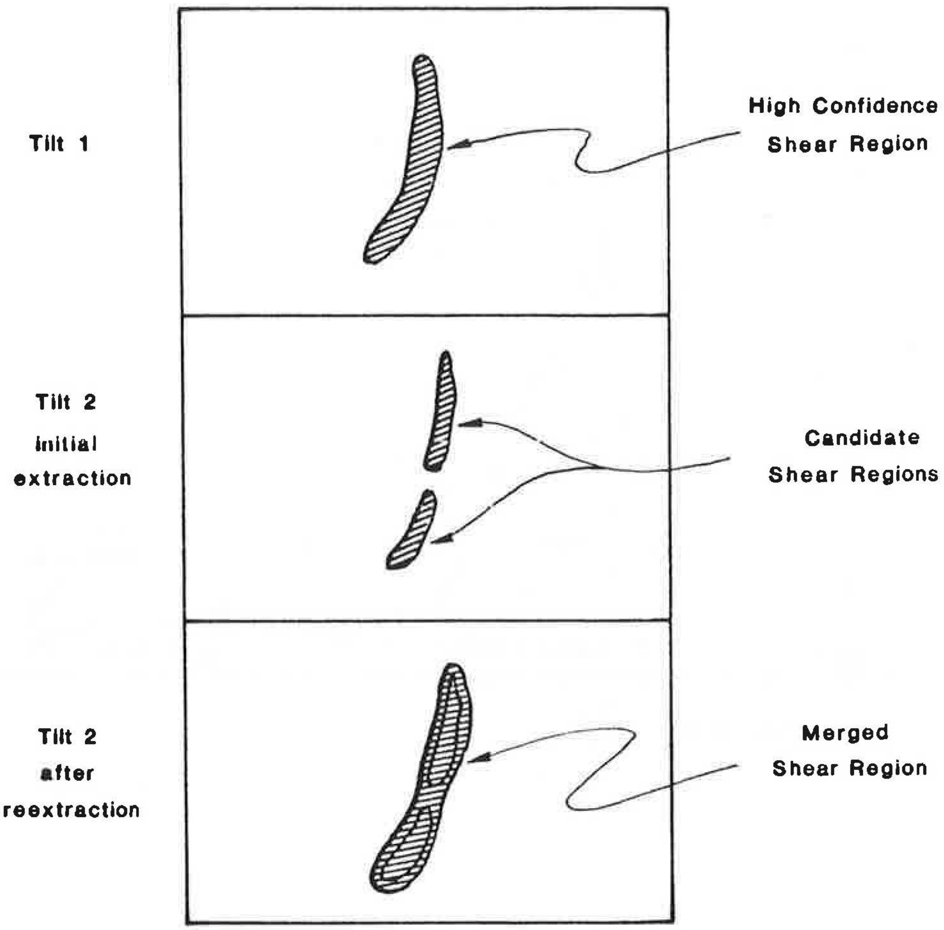

Figure 19. Gust front propagation.

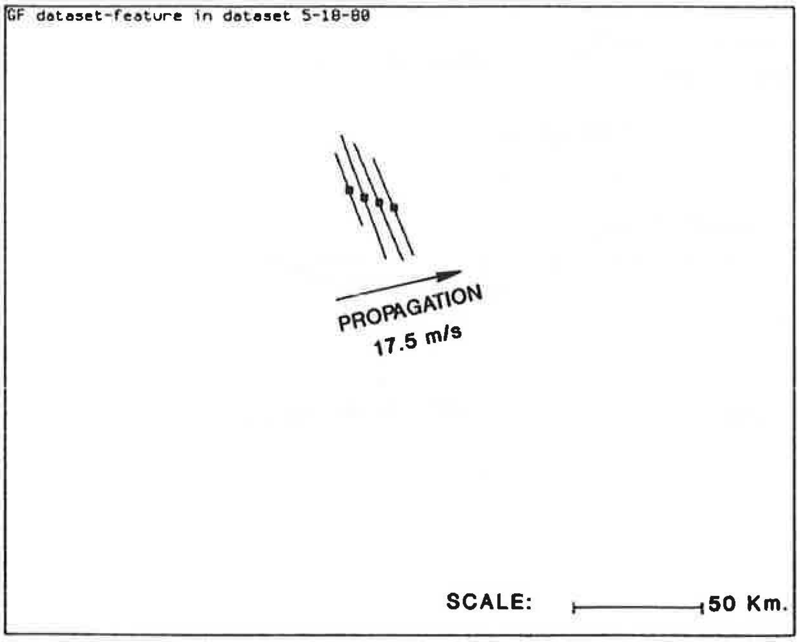

of detecting different types of microbursts. Even in Denver, for example, there are a number of different kinds of microbursts. Some have a lot of rain. Some have almost none. Some have rotation, and some do not, and so forth.

Question The one that you described, a Memphis microburst. Which one would that be related to?

Steven D. Campbell That one had very heavy rain. The microburst described before, the basis for the microburst model described in the presentation, and the example in Figures 13 to 15, was the dry microburst, with almost no rain at the surface.

Question How can the radar pick it up if there is no rainfall?

Steven D. Campbell You don't need to have surface rain to pick it up. You just need water particles in the air. Those particles are not necessarily enough to cause rain at the surface. There just has to be some water, or perhaps other airborne particles.

Question Since there are so many different kinds of microbursts, how would you be able to detect all types or would the rule-based system be designed just for that particular area of the country?

Steven D. Campbell There might be some changes for different parts of the country, but basically the answer is that you add rules to take care of various different types of situations. The model that was discussed can be augmented, particularly in the downdraft area. It could be cyclonic rotation, counterclockwise. It could be anti-cyclonic rotation, clockwise. It could be a strong region of reflectivity. It could be many things. The system could just use the upper level convergence. The thing about the rule-based system is that we can add to our existing body of knowledge to cover more cases. We start off with some fairly simple assumptions, find out where these do not work, and add some more rules.

Question Wouldn't that slow down your real time interpretation of it?

Steven D. Campbell Yes, but it turns out that we are not being limited by the expert system, but by the image processing.

Question That is the question I was going to ask. As you consider more possibilities you extend the rule base. If there are more possible candidates being developed, will there be more false alarms?

Steven D. Campbell We have not had any false alarms so far. This system is very conservative, and we have not had a problem with false alarms. I think it is because of our knowledge-based approach. I think that false alarms are going to be a real problem for some of the strictly algorithmic approaches that are currently being developed.