Blighter®Track
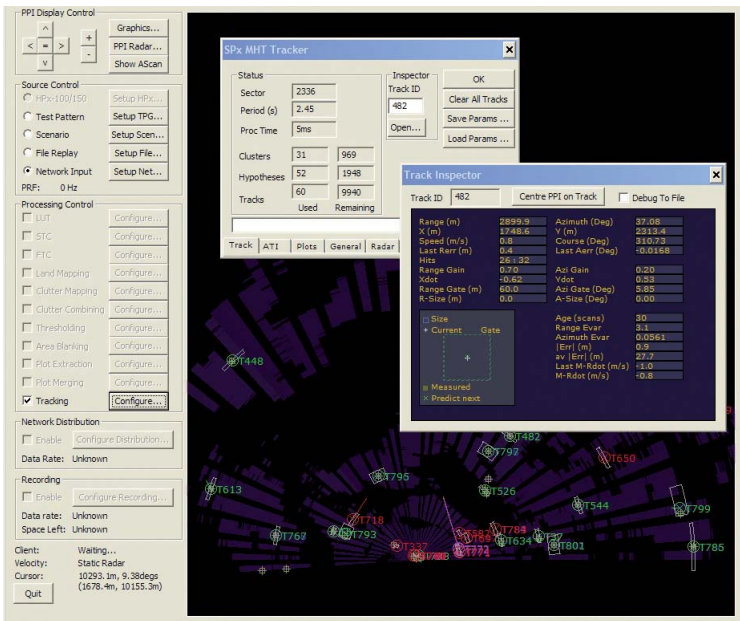

- Software based radar target tracker

- Adds full Track-While-Scan (TWS) capability to any Blighter radar system

- Fully parameterised and configurable

- Automatic Track Initiation (ATI) using built-in M from N detector

- Supports multiple hypotheses

- Multiple track filtering modes

- Dynamically calculated gating function

BlighterTrack is Blighter Surveillance Systems' fully configurable, software based radar target tracker for use with any Blighter ground surveillance radar. BlighterTrack offers automatic track initiation and has full support for multiple hypotheses. It incorporates multiple track filtering modes and gates are calculated dynamically.

BlighterTrack interfaces to the output data stream from the Blighter radar's built-in plot extractor and correlates from scan-to-scan to output positional and motion updates. The BlighterTrack software is highly configurable and may be used to identify targets according to defined rules.

When plot data is correlated from scan-to-scan, BlighterTrack uses multiple hypotheses to support ambiguous interpretations of the radar plot stream. The filter uses position, size and historical measurements to correlate existing tracks with new data, providing updated positions and dynamics, as well as a confidence estimate. The behaviour of historical track data is analysed to help interpretation and provide a first level classification capability.

Track Creation
The tracker maintains an active track database. The contents of the database are updated with new plot data sent by the Blighter radar. New tracks are added to the database automatically. The automatic track creation occurs when plots entered into the

database are seen to be uncorrelated, or ungated, with any existing known target. A new preliminary track is created and is updated with future detections until confidence is established that the track is likely to be a target of interest.

The time a track is held in the preliminary stage is a programmable option and needs to be set to balance the speed of detection with the likelihood of a false alarm. In a low clutter environment, where extracted plots are likely derived from real targets, the acquisition time may be as short as 2 detections. For noisy situations, where the plot extractor is reporting false detections, the integration time in the preliminary stage may be extended. The criterion for promoting a track from the preliminary to full track status is defined through either a binary integrator or an M from N detector, where a minimum of M valid detections are required in N observations

Track Correlation
Established tracks are updated using new measurements provided by the data extraction processing. The first stage of that is the association

at the frontier of security ™

Blighter®
Surveillance Systems

problem, by which a measurement is associated with the most likely track. In a simple situation, a true target will give rise to a single plot that can be directly associated with the expected target position. In the general case, there may be ambiguity as to which measurement relates to which track.

A distance measure is computed from the expected position of a track to the position of measurements that are within the track gate. Within BlighterTrack, the gating function uses a dynamically computed area around the expected position of a track. This area is computed to reflect the likely error in track position derived from the target's possible movement and the measurement noise.

The BlighterTrack gating function uses the best estimates of target dynamics and measurement noise models to compute an accurate gate area at each new measurement. This helps to ensure that association only considers plots that are likely to derive from the true target and that when the association is made a measure of statistical distance can be used to weight each candidate association.

Multi Hypothesis Tracking
BlighterTrack supports multiple hypotheses, which means that decisions to associate tracks with measurements can be deferred until additional information becomes available. If the tracker isn't sure whether a track should be associated with plot p1 or plot p2, for example, it can create two hypotheses. This allows both possibilities to be propagated to the next stage of processing. The next measurement might make it clear which of p1 or p2 was the right choice.
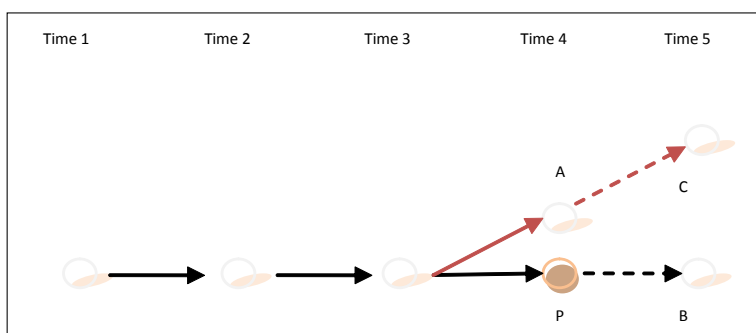
The job of a tracker is to interpret radar observations to distinguish real targets from noise, and to construct models to describe the motion of the true targets. The tracker is provided with data, typically in the form of plots, derived from the processing of the radar video. These plots are connected regions of radar video that satisfy some rules of position, amplitude, size and signal strength.

Unfortunately, measurements from the radar are imperfect. There will be noise from the measurement process, clutter from the environment and unpredicted manoeuvres of the targets of interest. This means that the tracker will be presented with noisy and possibly multiple measurements from the target of interest. The tracker's responsibility is to provide the best interpretation of the data using assumed or calculated statistics for the noise and the likelihood of change.

In the single hypothesis situation, the tracker is forced to make the best interpretation it can of the available data at each update. For some updates, where there is a clear interpretation of the measurement, the best interpretation may be obvious and the single hypothesis offers a satisfactory solution. Problems arise, however, if the interpretation of the measurements is not obvious. In this case it may be desirable to defer a decision until the next update when additional information will help to decide on the correct interpretation. The ability to simultaneously consider multiple interpretations of the system is the key to the multi hypothesis tracker.

In the diagram to the left, a target is moving from left to right. At each update time (Time = 1..4) the position of the target is shown. The target moves in a predictable way at Time = 1, 2 and 3. The tracker will estimate the target's speed and heading so that the predicated position at Time = 4 is shown as P. However, suppose that at Time = 4 the actual position observed by the measurement is A. The question now arises as to whether the observation of the target at A is due to a manoeuvre by the target (so that the expected position at Time = 5 will be C), or whether it is due to a measurement error (so that the expected position at Time = 5 will be B).

The single hypothesis tracker has a difficult choice at Time = 4. It has to make a single interpretation of the results, so it has to decide on the single best explanation for the measurement at A. If it assumes that the measurement is the result of a manoeuvre it will expect the target to be around C at time 5. But if the measurement is really an error, the tracker may not correctly detect the next measurement at B. The ability of a tracker to detect a measurement around the expected position is defined by a search gate. This is finite size geometrical shape around the expected position of the target. The size of the search gate is a compromise between minimising false detections (suggesting a small gate) and accommodating uncertainty in the estimate and allowing a target to manoeuvre (suggesting a larger gate size).

In the face of the uncertainty in making a decision at Time = 4, the single hypothesis tracker may try to "hedge its bets" by selecting one of the choices, but enlarging its search gate to accommodate the next update's measurement if the alternative explanation turns out to be correct. This might work, but enlarging the search gate in this way is a reflection of the uncertainty in the interpretation at Time = 4. The enlarged search gate makes the tracker prone to detect additional measurements that may further complicate the hypothesis. Furthermore, for the tracker to be able to respond to either the Time = 5 measurement being B or C, it will probably need to increase its filter gains, further enhancing the likelihood that the track will be corrupted by noisy measurements.

In contrast to the above, the multi-hypothesis tracker solves the problem at Time = 4 by considering both explanations simultaneously. It propagates two hypotheses to Time = 5. One hypothesis says that the next expected measurement is at B, with the measurement at Time = 4 being due to a measurement error. The other hypothesis believes the measurement and expects the detection at Time = 5 to be at C. At Time = 5, the discovery of the target at either B or C confirms one hypothesis and refutes the other. The refuted one is deleted and the confirmed one continues on to the next update.

Track Filter

For each hypothesis, the tracker updates the current estimated position with the new measurement. If the measurement were known to be completely accurate, the update process would believe the measurement and the new estimate would be exactly the measured value. For various reasons, the measurement is inaccurate so the update process must take a weighted combination of the expected position and the measured position. This is the track filtering. BlighterTrack offers a number of track filtering modes. The simplest mode uses fixed gains in the components of the measurement, which can be successful for tracking applications where

the target is clearly identified and relatively clutter free. Additionally, BlighterTrack also provides a dynamic gain filter, which automatically adjusts the filter gains to provide good filtering in the steady state, whilst retaining the ability to track through a manoeuvre.

The filter works by computing a dynamic filter gain, K, based on estimated system noise and measurement noise models. The system noise is used to model uncertainty in the known dynamics of the target, including its ability to manoeuvre. As system noise increases, or equivalently as measurement noise decreases, the filter places more weight on the measurement – the filter gains increase. As system noise decreases or as measurement noise increases, the filter gains decreases causing less emphasis to be placed on the new measurement. The filter gains are continually changing and provide, under certain assumptions of the noise characteristics and linearity, an optimal estimation of the true target position. As part of the update process, the filter also provides a convenient measure of the estimation. This provides a useful confidence assessment of the estimation.

Track and Plot Reporting

BlighterTrack outputs plot and track data onto a standard Ethernet network for delivery to local or remote clients for data fusion or display. The time stamped reports are delivered with low latency, and may include both the filtered and measured components of the track's state vector.

Client-side software libraries are provided to receive the network reports and make them available as data structures for client processing or display.