

This article was published in 2002 shortly after the successful operational test and evaluation of the Cooperative Engagement Capability (CEC) in the Atlantic Fleet. CEC provides the U.S. Navy and Marine Corps sensor netting and enables integrated fire control. As of 2018, CEC is deployed on over 140 units (Navy ships, Navy aircraft, and land mobile Marine Corps units). CEC was conceived at APL during the 1970s and 1980s. During the 1980s and 1990s, APL prototyped key CEC innovations—namely its sensor netting algorithms, sensor netting software, and components of its Data Distribution System. Also during this time, APL worked closely with industry to productionize CEC and integrate it into Navy ships and aircraft. Today, as the CEC technical direction agent, the Laboratory continues to contribute to the success of CEC in new mission areas.

## Trade-Offs in Sensor Networking

William G. Bath



**A**ir defense dominance requires a precise and timely picture of the surrounding airspace that is built using inputs from sensors on many different ships and aircraft. This article examines the sensor networking process in terms of its three primary descriptors: (1) The robustness and capacity of the data distribution process used. (For example, what is the probability of a sent message never being received, and how much data can be sent?); (2) The data grouping approach used. (How are detections grouped in order to know that a particular detection by one unit and another detection by another unit do or do not correspond to the same object?); and (3) The data sharing approach used. (Which data are transmitted between units, and how are these data used to calculate network tracks?)

### INTRODUCTION


The Navy generally strives to establish air defense dominance in the regions in which it operates. A precondition for air dominance is having a precise and timely picture of the surrounding airspace. This “picture” is a digital representation (position and velocity) of every “piece of metal” in the sky—be that piece an airplane, a helicopter, or a missile. The air picture is essential to air dominance as it allows both the human operators and the computers under their control to sort through all these objects to identify any suspect or hostile targets.

The air picture is built using inputs from sensors on many different ships and aircraft. Sensor networking is the process of moving these data around and building the picture. A radar receives a reflected signal that is observed as a function of time delay—corresponding to target range. Many of these echoes are collected from different angles, producing a range-angle map of the surrounding airspace. Over time, these individual “looks”

are collected and combined to make an automatic track (Fig. 1).

This process is going on simultaneously on many ships and aircraft and at many different fixed and mobile land sites. Each is using its own sensors to attempt to create the best possible air picture. In an ideal world, each sensor could track every target continuously (e.g., from takeoff to landing, without interruption). However, in the difficult environment we face, the laws of physics do not permit this. Target fades, terrain blockage, and spurious signals from the natural environment or from countermeasures all conspire to make the situation not ideal. As a result, in general, no single ship, aircraft, or land unit can create a complete air defense picture.

Unfortunately, while advances in sensor and computing power greatly improved what could be done automatically on a single ship or aircraft, on the whole, the networking systems have not kept pace.



**Figure 1.** The tracking process examines anonymous automatic detections (b), collected over time (c), and produces tracks (d, e).

When these incomplete pictures are combined through today's tactical data links, considerable confusion results—including miscorrelations of data and ensuing dangerous misidentifications of targets. Thus there is now considerable discussion within the Navy and, indeed, throughout the DoD, as to how best to improve the networking of sensors. This article discusses some of the trade-offs inherent in different approaches to solving the problem.

## THE THEORY OF RADAR TRACKING

To understand the trade-offs in sensor networking it is necessary to delve into the basic theory of tracking. This theory is relatively well developed, having begun during World War II in the 1940s and being the subject of thousands of technical papers and dissertations since. Each sensor makes a digital detection when the received signal exceeds a threshold. The essence of the problem is that frequently each sensor's detections, such as those in Fig. 1, are anonymous. That is, to a first approximation, the detections look alike, regardless of what target they correspond to. The tracking process consists first of *association*—the process of grouping together detections that are believed to correspond to the same target—and then *filtering*—the process of calculating the track state (estimated position and velocity of the target)—and finally *prediction*—the process of extrapolating. These three processes interact recursively. First one groups together  $N$  detections and calculates a track state. Then one uses that track state as the basis for associating the  $N + 1$ st detection with the group. Because the detections are anonymous, the grouping is based largely on the measured positions of the detections.

Tracking works well when one knows in advance that the targets are sufficiently far apart from each other that detections from one target are not confused with those from another. When targets are close and detections do get confused, many serious tracking problems can result. For one thing, the tracks may be pulled off onto fictitious positions or assume inaccurate velocities. This then corrupts the association process on the next opportunity—potentially resulting in more erroneous associations. The outcomes of these types of errors fall into certain well-known categories. Track swaps occur when two tracks essentially swap their streams of detections so that track #1 assumes the position of track #2 and vice

versa (Fig. 1). This may seem innocuous, but it is a serious military problem. For example, if track #1 had been assessed as friendly/neutral and track #2 had been assessed as hostile, then track #2 could be engaged by shooting at it with surface-to-air or air-to-air missiles. If the measurement streams swap after the assessment of identity, then the surface-to-air or air-to-air missile will be fired at and guided to the wrong (friendly/neutral) target position. The result could be the killing of innocent or friendly people. In addition, since the missile will not be fired at a track tagged as friendly/neutral, the truly hostile target is not engaged. This can result indirectly in further damage or loss of life among friendly forces.

One could try to eliminate track swaps by making correlation criteria very strict, that is, by using an algorithm that will not group detections together unless they are very close to a common trajectory. Unfortunately, this results in a second type of error called a dual track. This is the erroneous representation of a single target by two tracks instead of one, which occurs when the design overcompensates for trying to prevent the track swap described above. If the correlation criteria are made unduly strict, then a single group of detections from a single target can be erroneously divided into two groups corresponding to two different tracks. Dual tracks produce less serious but still harmful effects in air battle management. One of the most important is that they tend to undermine the confidence of human operators in the air picture. (Operators often know from other sources that only one airplane is present, so why are the computers telling them there are two?) Another important harmful effect is dual engagements—launching two salvos of missiles when only one is needed.

W. G. BATH

There are many different variations on these errors, which will not be discussed here. Needless to say, the name of the game is to prevent mistakes in grouping the detections to form tracks. The tracking system optimizes this grouping process by calculating the most accurate track state possible and by making the best possible association decisions. A well-known theoretical framework exists for grouping measurements into tracks. Let the true target position as a function of time be given by the function  $\mathbf{x}(t)$ . A given sensor will make a measurement  $\mathbf{y}_k$  at time  $t_k$ :

$$\mathbf{y}_k = \mathbf{f}(\mathbf{x}_k) + \mathbf{n}_k, \quad (1)$$

where  $\mathbf{n}_k$  is an independent stream of measurement errors with covariance matrix  $\mathbf{R}$ , and  $\mathbf{f}(\cdot)$  is the function which determines the measurement coordinate frame. The filtering process that estimates the true target position and velocity is based on a classical Kalman filter formulation in which the target state is position and velocity and the target motion is modeled as linear. For example,

$$\mathbf{X}_k(\text{model}) = \begin{bmatrix} \text{target x position at time } t_k \\ \text{target x velocity at time } t_k \\ \text{target y position at time } t_k \\ \text{target y velocity at time } t_k \\ \text{target z position at time } t_k \\ \text{target z velocity at time } t_k \end{bmatrix}. \quad (2)$$

The modeled target state at time  $t_k$  evolves in time according to

$$\mathbf{X}_{k+1}(\text{model}) = \boldsymbol{\phi}_k \mathbf{X}_k(\text{model}) + \mathbf{w}_k. \quad (3)$$

The transition matrix  $\boldsymbol{\phi}_k$  represents linear motion. The vector  $\mathbf{w}_k$  represents maneuvers—deviations of the target from a straight line. Generally these are turns, changes in speed, or changes in altitude rate having any of an uncountable number of shapes. Examples are the brief portions of a sinusoid (constant rate turn or course), jinks or a sustained sinusoid (evasive target executing a weave), and a parabola (constant longitudinal acceleration typical of speed changes). The maximum sustained maneuver level or “g level” is determined by the density of the atmosphere as well as the shape, speed, and construction of the target. Generally a human pilot cannot withstand much more than several g’s, so this limits what a manned aircraft can do. Missiles can have higher limits before they become unstable or fail structurally. Most military systems are specified for some maximum g level based on these limits (e.g.,

the system must track targets performing no more than a specified number of g’s).

The process of calculating an estimate,  $\mathbf{X}(k|k)$ , of  $\mathbf{x}(t_k)$  given all measurements  $\{\mathbf{y}_1, \dots, \mathbf{y}_k\}$  is done recursively using  $\mathbf{X}(k|k-1)$ , the estimate of  $\mathbf{x}(t_k)$  given all measurements  $\{\mathbf{y}_1, \dots, \mathbf{y}_{k-1}\}$ :

$$\mathbf{X}(k|k) = \mathbf{X}(k|k-1) + \mathbf{K}_k \{\mathbf{y}_k - \mathbf{f}[\mathbf{X}(k|k-1)]\} \quad (4)$$

and


$$\mathbf{X}(k+1|k) = \boldsymbol{\phi}_k \mathbf{X}(k|k) + \boldsymbol{\gamma}_k, \quad (5)$$

where  $\mathbf{K}_k$  is the gain matrix, and the multiplication by  $\boldsymbol{\phi}_k$  has the effect of moving the estimate forward in time from  $t_{k-1}$  to  $t_k$ . The most general method for determining the gain matrix is the Kalman filter process<sup>1</sup> in which maneuvers are modeled as a random walk. However, other less general methods also exist for deterministic maneuvers.<sup>2,3</sup> Although shown here for simplicity in Cartesian coordinates, the update may also be done in the measurement coordinates.<sup>4</sup> The term  $\boldsymbol{\gamma}_k$  is a deterministic correction for any known acceleration (e.g., acceleration of gravity or Coriolis acceleration for ballistic objects). In general, the track state,  $\mathbf{x}(t_k)$ , and the estimates,  $\mathbf{X}(k|k)$  and  $\mathbf{X}(k|k-1)$ , are vectors. A six-dimensional vector is used for position and velocity in three dimensions, and a nine-dimensional vector is used for cases where instantaneous acceleration is also estimated. Equation 4 represents the update of the filtered track state with a new measurement. These appear on a graph as discontinuities, such as are seen in Fig. 2. Equation 5 represents the extrapolation of the track to the time of the next anticipated measurement. For the case where only position and velocity are being estimated, these are the straight lines in Fig. 2.

Since the measurements are anonymous, each one must be assessed and a decision made as to whether it belongs to each existing track. Both the measurement and the predicted track states have regions of uncertainty (ROUs) that contain the true target position with some statistical confidence, say 99%. A measurement is a candidate for grouping into a track when these ROUs overlap, as shown in Fig. 2. The measurement ROU is simply proportional to the measurement accuracy. The extrapolated track state ROU has a covariance term,  $\text{Cov}(k+1/k)$ , which includes the effect of the accuracies of all previous measurements, the gains used, and the degree of extrapolation. It also contains a lag term,  $\text{Lag}(k+1/k)$ , due to biases created by any potential maneuver:

$$\text{ROU}(k+1/k) = \text{Lag}(k+1/k) + 2.6 \text{Cov}(k+1/k).$$

To prevent the track swaps and duals described previously and in Fig. 2, these association decisions must



**Figure 2.** An expanded view of the tracking process shows the three main functions.

be correct. The best approach to making them correct is to keep this track ROU,  $ROU(k + 1/k)$ , as small as possible. This reduces the chance that measurements from a different target will be grouped into the track accidentally. Further, with the small ROU, duals are reduced because the designer is not tempted to use overly strict grouping criteria.

## ALTERNATIVES FOR SENSOR NETWORKING

Although sensor networking is a complex and multifaceted problem, there are three primary descriptors of a sensor networking system.

1. The robustness and capacity of the data distribution process used. (For example, what is the probability that a sent message will never be received, and how much data can be sent?)
2. The data grouping approach used. (How are detections grouped in order to know that a particular detection by one unit and another detection by another unit do or do not correspond to the same object?)
3. The data sharing approach used. (Which data are transmitted between units, and how are the data used to calculate network tracks?)

These descriptors and the fundamental trade-offs between different design approaches are detailed next.

### Data Distribution

In our personal dealings with computers we are used to relatively reliable distribution of data (e.g., over local area networks, the Internet, etc.). These sorts of data distribution generally have a communications media over established infrastructure—copper wires, fiber-optic cables, fixed point-to-point microwave relay, etc. Even cellular phones, while potentially of low quality from the handset to the cell antenna,

have a very reliable backbone infrastructure for interconnecting the cells. Data distribution between mobile platforms in a military theater (e.g., ship-to-ship, ship-to-air) is considerably different because no such established infrastructure exists. (Or if there is an infrastructure and it is still intact, it may belong to the enemy and be unavailable.) Denied this infrastructure, data distribution is accomplished by a network of radio communications. Units are often at the boundary of reliable communications from a fading or jamming perspective. Thus, robustness of data distribution is not guaranteed and becomes a strong driver of overall system performance.

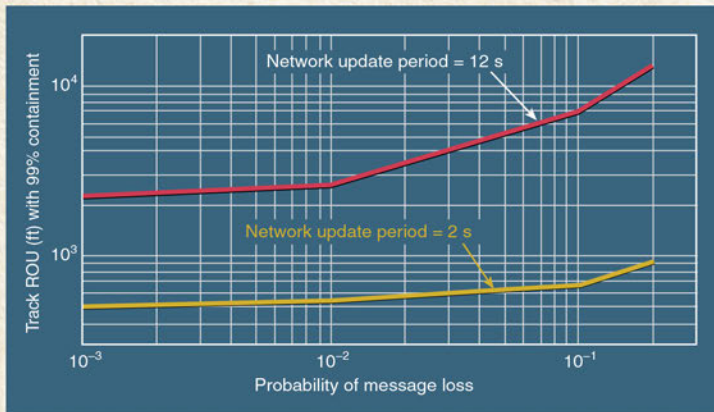
For simplicity, characterize the robustness of the data distribution by a single number—the probability of message loss. Not surprisingly, if one is looking for a certainty of 99% containment of the tracking errors, then one needs a probability of message loss somewhat less than 100% minus this number. Figure 3 shows the increase of the track ROU as the probability of message loss increases. Something in the .1 to 1% range is clearly needed if the radios are not to be the limiting factor in tracking quality.

Unfortunately, many of today's military data distribution systems are not this robust. Sensor networking systems generally require some sort of improvement to the existing data distribution process. APL has been instrumental in two such improvements: the development of robust high-frequency/ultra-high-frequency communications using Multi-frequency Link-11 and the development of robust microwave communications using the Cooperative Engagement Capability.

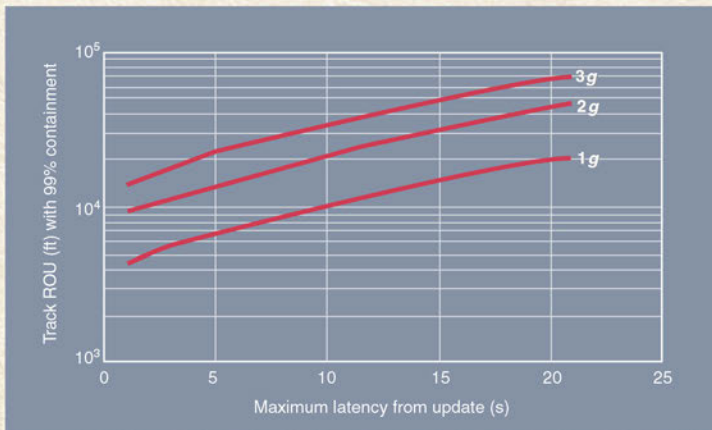
Again for simplicity, characterize the capacity of the data distribution by a single number—the update rate of tracks. In general, the amount of data that needs to be networked is on the order of (the number of tracks)  $\times$  (the number of bits in a message)  $\times$  (the rate at which the messages are sent per track).

The number of tracks is an independent variable determined by the environment. The number of bits in the message will be discussed later, but is likely on the order of a few hundred. If one views these two numbers as fixed, then the effects of data distribution capacity can be seen in the track update rate. Consider a sensor that makes measurements every 2 s. If one had the capacity to transmit some sort of message immediately following each update, then the maximum time any user would have to extrapolate the track would be the 2 s until the next measurement. If one can only update track data every 12 s, then the maximum extrapolation time would be 14 s. The sensitivity of track ROU to this extrapolation depends on the degree to which the target could be maneuvering (e.g., is there at least a 1% chance of this?). For example, in Fig. 4, we consider measurements that are relatively inaccurate—1200 ft. This would likely be a long-range angle case. Here the difference between a 2-s

W. G. BATH



**Figure 3.** The importance of data distribution robustness is shown. ROU increases sharply when the probability of message loss exceeds .01.



**Figure 4.** The importance of data distribution capacity is shown by ROU sensitivity to time latency (target max. maneuver = 1, 2, 3 g; probability of detection = .8).

and 14-s extrapolation is about 3 to 1. If one considers data that are more accurate to begin with, such as 30 ft, then this ratio grows to 5 to 1. Clearly, in either case, the capacity of the data distribution equipment greatly affects the accuracy of sensor networking.

### Data Grouping

Data grouping is key because, as previously discussed, the most serious sensor networking errors can result from mistaken grouping of the data. There are two basic data grouping approaches (Fig. 5):

1. *Measurement-to-track association* associates each measurement to the networked track potentially calculated using measurements from all sensors. Thus the entire stream of measurements (up to the present) is potentially available to calculate the track state used for the decision on the most recent measurement.
2. *Track-to-track association* associates each measurement to a single sensor track state calculated using only measurements from that sensor. The

single sensor track states are then grouped with each other to produce a netted track state.

The design decision as to which approach is better for grouping data depends on the sensors and targets involved. One case where measurement-to-track association is clearly better is when the sensors have a reduced probability of detection so there are potential gaps in the data stream or the data stream is sparse for a period of time. In these cases, a much more accurate track state can be calculated using multiple data streams than using only one, as multiple streams will tend to fill in the gaps in detection and restore a high, consistent data rate during periods of reduced probability of detection. Figure 6 illustrates the sensitivity to target fades by plotting track ROU versus probability of detection for single sensor tracking and multiple sensor tracking. When the probability of detection is much less than unity, the multisensor track is considerably more accurate since the probability of a significant outage of data is much reduced if two sources are available. With a more accurate track, tighter association criteria can be used for measurements. However, if the biases cannot be effectively removed, then there may be an advantage to associating to a single sensor track, which by definition is unbiased with respect to itself. If biases cannot be kept smaller than the ROU, then at the high probabilities of detection, one would prefer single sensor association followed by track-to-track association.

Another variant of this occurs when combining data from a high data rate sensor with those from a low data rate sensor. The netted track, because of its higher update rate, will provide a much more accurate basis for data association than the low update sensor alone. A common example of this is the combination of data from a low update rate airborne early warning (AEW)

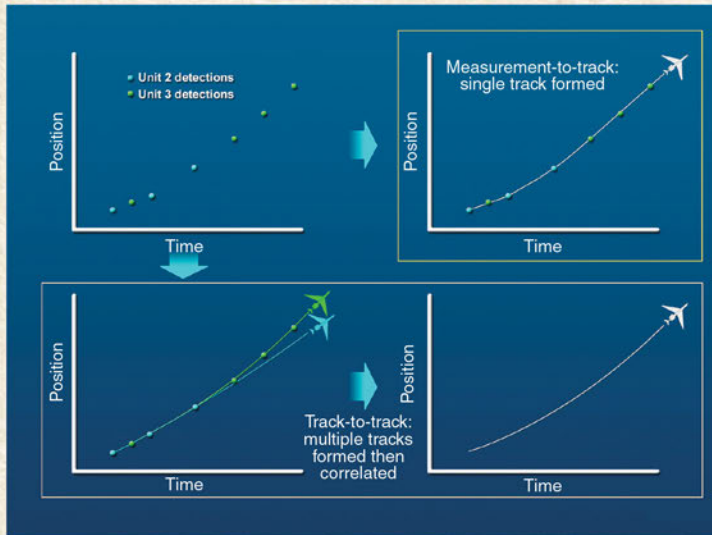


Figure 5. Two data grouping approaches.

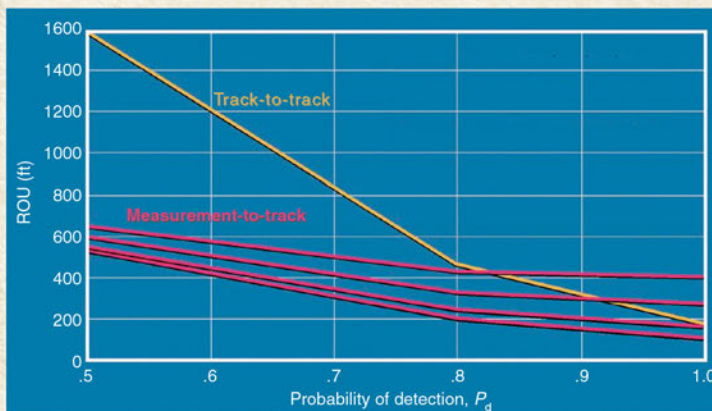


Figure 6. Comparison of measurement-to-track and track-to-track association. For fading targets ( $P_d < 1$ ), measurement-to-track is preferred. For large sensor biases and nonfading targets, track-to-track is preferred. (Second sensor biases = 0, 100, 200, and 300 ft, respectively, for measurement-to-track curves going from bottom to top.)

sensor with a higher update rate surface-based fire control radar. If one uses track-to-track association, then the AEW sensor is on its own for data association. Because the data rate is low, the ROU is relatively large and there can be significant chances that the track uncertainty will encompass a detection from another object. This potentially leads to track and identification swaps. If one uses measurement-to-track association, then the AEW sensor measurements are associated to a track that contains the high update rate measurements from the fire control sensor. As a result, the ROU is much smaller, and the likelihood of a track or identification swap can be reduced by an order of magnitude or more.

Finally, measurement-to-track association can actually increase the number of measurements available to the tracking process. In general, track formation requires a very low measurement false alarm rate, whereas tracking can proceed with a somewhat higher false alarm rate. Thus the

#### TRADE-OFFS IN SENSOR NETWORKING

detection threshold (Fig. 1) can be reduced at one sensor if another sensor has the target in track. Further, if the sensors can manage their energy, then additional energy can be used in the direction of known tracks. Both these factors can increase the number of measurements that a sensor is able to put into the network.

#### Data Sharing

There are two basic data sharing approaches, both based on the tracking theory already described:

1. *Measurement fusion* combines the measurement streams from all sensors to produce a single composite measurement stream. This composite measurement stream is then filtered using Eqs. 4 and 5.
2. *Track fusion* pre-filters the data before sending them. Thus Eqs. 4 and 5 are applied recursively to a sequence of single sensor measurements prior to transmission. The recipient then combines the tracks (for example a weighted combination<sup>5</sup> with the weights based on the relative size of the covariance  $P_k$ ). Track fusion has some interesting variants.

- Track selection is a very simple and data rate efficient logic which has been used in today's widely deployed networking systems (e.g., Link-11 and Link-16). A common decision is reached among all units as to which unit has the best single sensor track. Then only this track is transmitted. This selection is known as reporting responsibility. Each unit in the network (except for the one with reporting responsibility) then has two tracks to choose from—its local one and the "best" remote one.
- Tracklet fusion<sup>6</sup> pre-filters the data before sending them, then filters them again on reception. Thus Eqs. 4 and 5 are

W. G. BATH

applied recursively to a sequence of single sensor measurements prior to transmission, producing a stream of tracklets emanating from each site. The recipient then combines the streams of single sensor tracklets into a composite stream of tracklets and applies Eqs. 4 and 5 to the composite sequence of  $\{X(k|k)\}$ . To ensure that the input to the second filtering process is statistically independent, the first filtering is either over disjoint sets of measurements or is pre-whitened upon reception.

These approaches tend to be similar when the target is known with certainty to not be accelerating. However, when one is uncertain about the level of maneuver the target is currently performing, significant differences between measurement fusion and track (or tracklet) fusion exist. This is because as more sensors are combined through measurement fusion, the lags are significantly reduced. The reduction when more sensors are combined through track or tracklet fusion is considerably less. To illustrate this point, take a very simple case consisting of a single tracking filter. The filter gain is

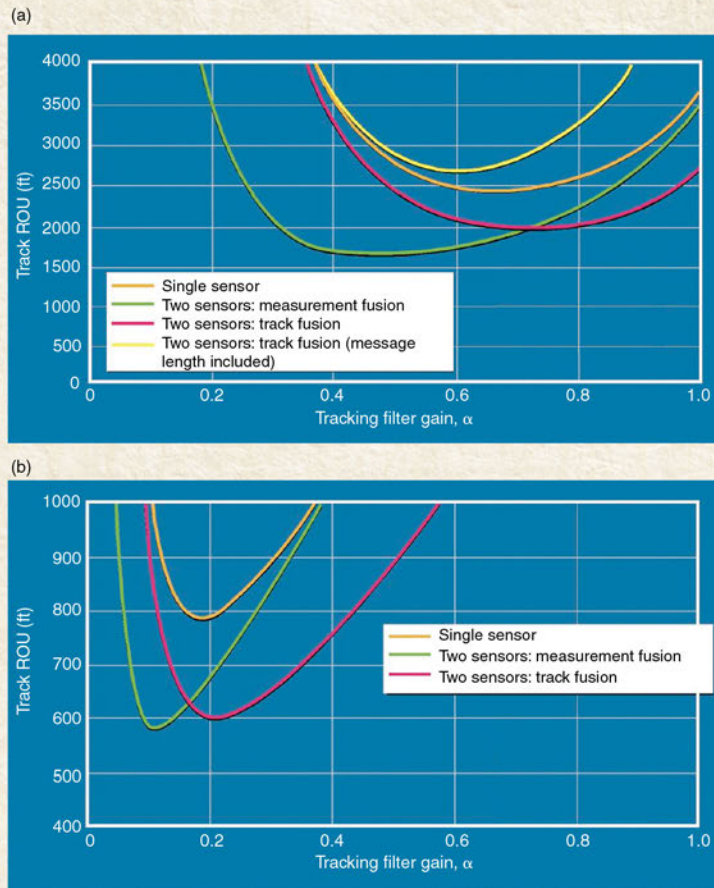
$$K_k = \begin{bmatrix} \alpha \\ \beta/T \end{bmatrix},$$

which is a constant, and the terms  $\alpha$  and  $\beta$  are the position and velocity tracking gains, respectively.

It is well known<sup>2</sup> that the steady-state ROU with 99% probability of containment is

$$ROU(k+1/k) = 2.6 \sqrt{\frac{2\alpha^2 + \alpha\beta + 2\beta}{\alpha(4 - \beta - 2\alpha)} R + \frac{\alpha T^2}{\beta}}.$$

Using this expression, it is possible to make simple comparisons between the accuracy of measurement fusion as opposed to track fusion (equal accuracy is one-dimensional). When the ROU is plotted as a function of position gain  $\alpha$ , it has the “bathtub” shape shown by the single sensor curve in Fig. 7a. The left-hand side of the



**Figure 7.** Comparison of data sharing approaches. (a) For air-breathing targets, measurement sharing produces the smallest ROU. (b) For exo-atmospheric motion, the ROUs are similar. (For parts a and b, sensor accuracy = 1200 ft, sensor update rate = 2 s, and probability of detection = 1.0; target max. maneuver = 2 g and 0.01 g in a and b, respectively.)

bathtub is dominated by the lag component, while the right-hand side is dominated by the covariance component. Since the gains (horizontal axis) are the designer’s choice, the single sensor ROU is the minimum of the bathtub curve.

Now consider the fusion of two sensors in a particular dimension. If one sensor has one-tenth the ROU of the other, then fusion is uninteresting because the more accurate sensor will dominate and essentially determine the result. At least in steady state, it is relatively easy to produce this dominance by any of the fusion methods. Of more interest is the case where the sensors are comparable in terms of accuracy and update rate, producing comparable ROUs. This case more clearly shows the difference in the fusion methods.

## TRADE-OFFS IN SENSOR NETWORKING

For example, when two identical sensors are combined by measurement sharing, the update rate is essentially doubled. This reduces the lag by a factor of 4, allowing a smaller gain to be selected (optimization more to the left of the bathtub) and reducing the covariance. The net result is the movement from the single sensor curve to the measurement sharing curve in Fig. 7a.

When two identical sensors are combined by track sharing, the update rate for each tracking process does not change, and so the lag does not change. However, the overall covariance is reduced, allowing a larger gain to be selected (optimization more to the right of the bathtub) and reducing the lag. The net result is the movement from the single sensor curve to the track sharing curve in Fig. 7a.

If any significant maneuver is possible (e.g., if the target is an aircraft or cruise missile), the factor of 4 in lag will have a more significant effect than the factor of 2 in covariance. Thus one can see that the measurement fusion curve achieves a significantly lower minimum than the track fusion curve. The number of sensors combined amplifies this difference. If maneuvers can be excluded (e.g., exo-atmospheric motion of an inert object such as a ballistic missile part with no thrusters), then the factor of 4 in lag does not have nearly the same effect, and measurement and track fusion become comparable for equal message lengths (e.g., Fig. 7b, where the acceleration is 0.01 g).

Tracklets will generally produce results in between those of measurements and tracks. If many measurements are combined into a tracklet, then tracklet fusion will compare with single sensor tracking as track fusion does. If only one or two measurements are used in a tracklet, then they will compare with single sensor tracking as measurement fusion does.


So far this discussion has centered on which data sharing approach is most accurate. Another important consideration is how well the approach characterizes the accuracy of the track. In many applications, such as the grouping of data, there is relatively little benefit to having a more accurate track unless one knows exactly how accurate it is in each dimension.

The importance of accurately representing the data covariance and lags is amplified by the physics of radar detection. With modern components, one can easily build a radar that can measure range to the target to an accuracy of 20 ft. However, in the microwave band, it is

impossible with a shipboard or aircraft antenna to measure angle to an equivalent precision. Thus the sensor ellipsoid of uncertainty (Fig. 8) is very narrow in one dimension (range) and much wider in the other two dimensions. When the viewing difference angle is significant, data fusion essentially triangulates the target position to the volume where the ellipsoids overlap. This can produce a factor of 10 or more reduction in ROU. However, knowing the result of this fusion requires a very precise knowledge of the shape and orientation of the ellipsoids. This is exactly the information contained in the covariances and lags.

An advantage of measurement fusion is that the recipient, using the measurements, can calculate the ROU. Thus it is only necessary to send the covariances of the measurements. This covariance matrix is diagonal, and so contains typically only a few non-zero numbers. In a track (or tracklet) fusion system, the recipient does not see the individual measurements and so needs additional data to calculate the ROU—the track noise covariance and the lags.

Table 1 compares the amount of data that must be sent for the simple example of three-dimensional sensors (range, bearing, and elevation). Clearly, update messages for track fusion will be significantly longer than for measurement fusion. As a result, the data must be sent less frequently. To reduce the track fusion message size, some approximations can be made. The lags and covariances can be combined and the resulting matrix block diagonalized. This results in a significant reduction in the fidelity of representing the ROU, but does permit the data to be sent more often. (Separate accounting for noise covariances and lags is important because they grow at different rates.) A very data



**Figure 8.** Measurement ellipsoids on uncertainty are very narrow in range dimension. Accurate triangulation requires precise characterization of ellipsoids.



W. G. BATH

**Table 1. How errors are represented in different data sharing approaches. Track selection and measurement fusion generally require the shortest messages.**

Networking approach	How errors are represented	What is sent
Measurement fusion	Full representation	6 numbers: 3 measured coordinates, 3 measurement covariances
Track or tracklet fusion	Full representation	33 numbers <sup>a</sup> : 6 filtered coordinates, 21 track covariances, 6 acceleration lags
Downsized track or tracklet fusion	ROU representation only; covariance matrix block diagonalized	15 numbers: 6 filtered coordinates, 9 "total" track covariances
Track selection	ROU representation only (heavily and logarithmically quantized)	Approx. 6 numbers: 6 filtered coordinates, 1 ROU index

<sup>a</sup> This is for a six-dimensional track state. For a nine-dimensional track state, 67 numbers are required.

distribution efficient variant is track selection. In this case, the ROU is heavily quantized into an ROU index in a logarithmic fashion. If only a single unit, the one with the highest ROU index, is allowed to report data, then networking is possible at very low data distribution rates.

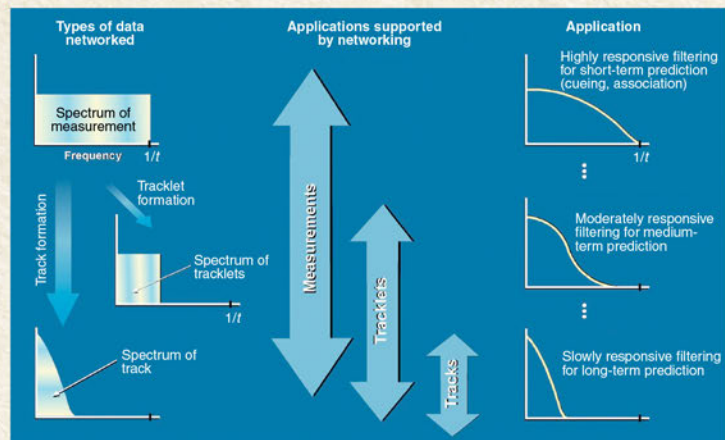
When one discounts track fusion for the additional length of the messages (a factor of 3 increase is assumed), the difference between measurement and track fusion becomes more pronounced (track sharing [message length included] curve in Fig. 7a).

Another consideration in selecting the data sharing approach is maintaining an open architecture for the user of the system. In general, target data are used for many different functions in a weapon system (e.g., situational awareness, data correlation, interceptor launch scheduling, handover of a target to an interceptor seeker, etc.). Each of these functions has different data filtering requirements. One of the main differences among these requirements is how heavily the data are smoothed (effectively, how small is  $\alpha$ ). Generally very heavy smoothing (small  $\alpha$ ) is used when predicting ahead for a long time (e.g., scheduling the launch time of a long-range interceptor). In this case one deliberately discards information by filtering out much of the higher-frequency components

of the measurement spectrum. Conversely, very light smoothing (large  $\alpha$ ) is used when predicting ahead for a short time (e.g., associating data at the next measurement). In this case one wants to preserve nearly all the spectral content of the measurements. The data sharing approach may restrict the types of functions that can be supported with the networked data (Fig. 9).

Measurement sharing provides the most open architecture because all functions can be supported (assuming that the original data are of sufficient quality). As shown in Fig. 9, the wide spectral content of the measurements allows various filtering processes to be used to achieve any desired spectral shape. Tracklets have a narrower spectrum because the higher-frequency components have been essentially averaged out. As long as the application needs a smaller spectrum (i.e., smoother data) than the tracklet itself, the application can be supported. However, applications requiring highly responsive filtering (large bandwidth) cannot be supported. This creates a somewhat undesirable coupling between the application design and network design, in that the application designer may have to come back to the network designer and request a different tracklet formation process.

Tracks are not easily refiltered for different applications, so a track-based system provides the least open architecture. Only a single bandwidth of track data is available, and this one form of data must be used for any



**Figure 9.** The type of data shared may restrict the applications for which they can be used. Measurement sharing provides the most open architecture.

## TRADE-OFFS IN SENSOR NETWORKING

application. This creates an even stronger coupling of the network designs, in that now different units with different applications may compete with each other to have the network tracking process designed to meet their needs.

## SUMMARY

The sensor networking process has been examined in terms of its three primary descriptors: (1) the robustness and capacity of the data distribution process, (2) the data grouping approach used, and (3) the data sharing approach used. The sensitivity of track accuracy to data distribution robustness and capacity has been examined. The measurement-to-track grouping approach has been shown to have advantages when single sensors have target detection fades, but it is more sensitive to sensor biases. The transmission of

measurements (vice tracks or tracklets) has been shown to be a more efficient use of data distribution bandwidth when targets are capable of significant maneuvers. The transmission of measurements also provides a more open architecture by giving the user more options for tailoring the track data to the application.

## REFERENCES

- <sup>1</sup>Kalman, R. E., and Bucy, R., "New Results in Linear Filtering and Prediction," *J. Basic Eng.* **83D**, 95–108 (1961).
- <sup>2</sup>Benedict, T. R., and Bordner, G. W., "Synthesis of an Optimal Set of Radar Track-While-Scan Smoothing Equations," *IRE Trans. Autom. Control* **AC-7**(4), 27–32 (1962).
- <sup>3</sup>Portmann, G. J., Moore, J. R., and Bath, W. G., "Separated Covariance Filtering," in *Proc. 1990 Int. Radar Conf.*, pp. 456–460 (7–10 May 1990).
- <sup>4</sup>Castella, F. R., "Multisensor, Multisite Tracking Filter," *IEE Proc. Radar Sonar and Navigation* **141**(2), 75–82 (Apr 1994).
- <sup>5</sup>Gelb, A., *Applied Optimal Estimation*, MIT Press (1974).
- <sup>6</sup>Blackman, S., and Popoli, R., *Design and Analysis of Modern Tracking Systems*, Artech (1999).

## THE AUTHOR



WILLIAM G. (JERRY) BATH received B.E.S., M.S.E., and Ph.D. degrees from The Johns Hopkins University in 1974, 1975, and 1980, respectively. Since joining the Laboratory in 1974, he has worked on the signal processing, automation, and integration of radar systems for Navy ships (AN/SYS-1, AN/SYS-2, MK-92 CORT, SPS-48C DDC) and the early networking of Navy radars using the Surface Gridlock System (SGS/AC). For 14 years, he led the engineering analysis group that developed the concepts, requirements, and algorithms for sensor netting in the Cooperative Engagement Capability. Since 1997, Dr. Bath has led the Sensor and Weapon Control Integration Branch of the Air Defense Systems Department. He has published over 30 papers on radar signal processing, tracking, and networking. He is a senior member of IEEE and the first recipient of the IEEE radar systems panel award for outstanding contribution to radar. His e-mail address is jerry.bath@jhuapl.edu.