

Generalized Information Representation and Compression Using Covariance Union

Ottmar Bochardt
IDAK Industries,
Canada
ottmar.bochardt@gmail.com

Jeffrey K. Uhlmann
Department of Computer Science
University of Missouri — Columbia
Columbia, MO 65211-2060, USA
uhlmannj@missouri.edu

Ryan Calhoun
Department of Computer Science
University of Missouri — Columbia
Columbia, MO 65211-2060, USA
rcalhoun@mizzou.edu

Simon J. Julier*
ITT AES / Naval Research Laboratory
4555 Overlook Avenue SW
Washington, DC 20375, USA

Abstract - *In this paper we consider the use of Covariance Union (CU) with multi-hypothesis techniques (MHT) and Gaussian Mixture Models (GMMs) to generalize the conventional mean and covariance representation of information. More specifically, we address the representation of multimodal information using multiple mean and covariance estimates. A significant challenge is to define a rigorous fusion algorithm that can bound the complexity of the filtering process. This requires a mechanism for subsuming subsets of modes into single modes so that the complexity of the representation satisfies a specified upper bound. We discuss how this can be accomplished using CU. The practical challenge is to develop efficient implementations of the CU algorithm. Because of the novelty of the CU algorithm, there are no existing real-time codes for use in real applications. In this paper we address this deficiency by considering a general-purpose implementation of the CU algorithm based on general nonlinear optimization techniques. Computational results are reported.*

Keywords: Covariance Intersection, Covariance Union, Data Fusion, Kalman Filter, Multimodal distributions.

1 Introduction

Level-1 information management has matured significantly over the last decade, with the development of rigorous algorithms robust to the effects of unmodeled correlations, and corrupt and/or spurious information, in the context of general distributed data fusion networks. Despite the dramatic theoretical and practical results in the Level-1 arena, very few inroads have been made into higher level information management applications. This is due in large measure to the discrepancy

between the relatively simple types of information encountered in low level tracking and control applications, and the much more varied and richer forms of information that must be processed in high level applications.

In this paper we explore a methodology for generalizing the unimodal information representation scheme used in Level-1 contexts to permit the representation of information that has a more complicated multimodal structure. This is accomplished using a set of unimodal state estimates to capture the multiplicity of possible states of the target of interest. The challenge is to bound the computational complexity issues that arise from this approach. In this paper we describe how a mechanism called Covariance Union (CU)[6, 3] can be applied to reduce the complexity of a multimodal representation to satisfy a fixed complexity budget, while rigorously guaranteeing information integrity.

The structure of the paper is as follows: Section 1 discusses the issue of information representation. Section 2 discusses the need for an information compression mechanism to bound the computational complexity of the fusion process. CU is shown to be a solution to this problem. Section 3 discusses computational issues that must be addressed in order for CU to be applied in practice. Practical algorithms for implementing CU are described. Section 5 provides experimental results demonstrating the application of CU. Section 6 discusses the results presented in the paper.

2 Information Representation

Determining how to represent information and uncertainty is a key first step that impacts all aspects of the data fusion problem. The representation must provide both an estimate of the state of the target or system of interest *and* its associated degree of error or uncertainty, and the uncertainty must be defined in a form that permits it to be empirically determined. There must be a rigorous algorithm for fusing information in

*The author is now with the Department of Computer Science, University College London. Email: S.Julier@cs.ucl.ac.uk

the representation, and the computational complexity of the representation and its associated fusion algorithms must be bounded for practical application.

By far the most widely used information representation is the mean and covariance form, where the mean vector defines the best estimate of the state of the target and the error covariance provides an upper bound on the expected squared error associated with the mean. For example, the measured position of an object in two dimensions can be represented as a vector \mathbf{a} consisting of the object's estimated mean position, e.g., $\mathbf{a} = [\mathbf{x}, \mathbf{y}]^T$, and an error covariance matrix \mathbf{A} that expresses the uncertainty associated with the estimated mean. If the error in the estimated mean vector is denoted as $\tilde{\mathbf{a}}$, then the error covariance matrix is an estimate of—is greater than or equal to—the expected squared error, $E[\tilde{\mathbf{a}}\tilde{\mathbf{a}}^T]$. The estimate is said to be consistent (or conservative) if and only if $\mathbf{A} \geq E[\tilde{\mathbf{a}}\tilde{\mathbf{a}}^T]$ or, equivalently, $\mathbf{A} - E[\tilde{\mathbf{a}}\tilde{\mathbf{a}}^T]$ is positive definite or semidefinite (i.e., has no negative eigenvalues). The full estimate of a target's state is given by the mean and covariance pair (\mathbf{a}, \mathbf{A}) .

Given two mean and covariance estimates (\mathbf{a}, \mathbf{A}) and (\mathbf{b}, \mathbf{B}) , the data fusion problem consists of determining a fused estimate (\mathbf{c}, \mathbf{C}) that is guaranteed to be consistent, and that summarizes the information in the two estimates with error (in terms of the size of \mathbf{C}) less than or equal to that of either estimate. If the two estimates are consistent and have a precisely known degree of correlation, the Kalman filter can be applied; otherwise, Covariance Intersection (CI) must be used. Both algorithms guarantee consistent results when used appropriately. The limitations of the mean and covariance representation of information can be found in a variety of practical contexts. For example, suppose a vehicle is being tracked along a road in an urban environment. Assuming that it travels at a speed that is average for the road, its future position can be predicted forward a short length of time reasonably accurately; however, if it encounters a T-junction that requires turning left or right, there are two distinct possible future positions. The future state can be represented with a single mean and covariance estimate, but doing so requires establishing a mean position at the junction with a covariance large enough to account for its position after a left *or* right turn. This produces a clearly unsatisfactory result in which the mean vector does not correspond to either of the possible states of the vehicle and consequently has a very large error covariance. Intuitively it seems clear that a better option would be to maintain information about the *two* possible future states rather than subsuming them into a single mean and covariance estimate.

Historically there have been two distinct approaches for representing “multimodal” information (e.g., as in the above example). One involves Multiple Hypothesis Tracking (MHT), which maintains multiple mean and covariance estimates corresponding to distinct possible states[1]. The other approach is to attempt to maintain a parameterization of the Probability Density Function (PDF) that defines the uncertainty dis-

tribution associated with the state of the target. In practice, PDF approximation methods typically only represent the significant modes of the distribution in terms of their means and covariances, thus making its representation all but identical to MHT. A key distinction is that the PDF-based approach treats the set of estimates as a union of Gaussian probability distributions. More specifically, the distribution is expressed as a Gaussian Mixture Model (GMM) of the form:

$$p(\mathbf{x}) = \sum_{i=1}^N p_i \mathcal{N}\{\mathbf{x}; \mu_i, \mathbf{P}_i\}. \quad (1)$$

The reason for adopting this form is that GMMs can conveniently approximate a wide class of PDFs and are identical in implementation to MHT. Unfortunately, representation is only one aspect of the overall information management problem. There also must be tractable algorithms for fusing information in a given representation.

The fusion of a set S of mean and covariance estimates—each defining a possible state of the target, only one of which is guaranteed to be consistent—with another set T can be accomplished under the MHT interpretation simply by forming the Cartesian product $S \times T$ and applying the appropriate fusion algorithm (Kalman or CI) to the pairs. Unfortunately, this yields a combined estimate that has $O(|S|*|T|)$, which implies that the complexity of the fused estimate exceeds that of the original estimates. This increasing complexity will tend to exhaust available resources and therefore must be mitigated.

3 Representation Compression

One of the most important features of the mean and covariance representation of information is its constant complexity. Specifically, the amount of information required to describe the state of the target does not increase as new information is incorporated. However, when the representation of state is generalized to maintain one mean and covariance estimate for each of the different modes, the update/fusion operation multiplies the number of modes. In order to manage the complexity of the representation, some form of representation compression must be applied.

In most MHT applications, the proliferation of hypotheses is managed by pruning the least likely ones according to some measure. A practical problem with pruning is that the likelihood measure typically includes many assumptions (e.g., PDF-related) that lead to more loss of correct hypotheses than is expected, and any loss of the hypothesis that corresponds to the true state of the target undermines the rigor of the entire information management framework. Therefore, pruning cannot be the primary mechanism for the limiting the representational complexity of our multimodal estimates.

If it is not possible to prune estimates (discard modes), then the only alternative is to somehow coalesce similar modes to stay within a fixed representa-

tional complexity budget. The key question is how to achieve this coalescence such that the integrity of the information is maintained. If it is assumed that one mode of an estimate corresponds to the true state of the target, and the others are spurious, then a mechanism called Covariance Union (CU) can be applied. For example, given n modes represented by estimates $(\mathbf{a}_1, \mathbf{A}_1) \dots (\mathbf{a}_n, \mathbf{A}_n)$, CU produces an estimate (\mathbf{u}, \mathbf{U}) that is guaranteed to be consistent as long one of the mode estimates $(\mathbf{a}_i, \mathbf{A}_i)$ is consistent. This is achieved by guaranteeing that the estimate (\mathbf{u}, \mathbf{U}) is consistent with respect to *each* of the estimates:

$$\begin{aligned} \mathbf{U} &\geq \mathbf{A}_1 + (\mathbf{u} - \mathbf{a}_1)(\mathbf{u} - \mathbf{a}_1)^T \\ \mathbf{U} &\geq \mathbf{A}_2 + (\mathbf{u} - \mathbf{a}_2)(\mathbf{u} - \mathbf{a}_2)^T \\ &\vdots \\ \mathbf{U} &\geq \mathbf{A}_n + (\mathbf{u} - \mathbf{a}_n)(\mathbf{u} - \mathbf{a}_n)^T \end{aligned} \quad (2)$$

where some measure of the size of \mathbf{U} , e.g., determinant, is minimized. The consistency of the CU estimate is assured for each of the n inequalities because the difference between the mean \mathbf{u} and \mathbf{a}_i is accounted for in the covariance \mathbf{U} by the addition of the square of that difference to the covariance \mathbf{A}_i .

Given a complexity budget of N modes, the fusion of two N -mode estimates will produce a new estimate with N^2 modes which must be reduced to N modes. This can be achieved by applying a clustering algorithm (e.g., standard k-means clustering based on a covariance-weighted distance measure such as Mahalanobis). Each of the N clusters can be combined into a single mean and covariance estimate using CU, and the rigor of the framework is guaranteed because one of the N estimates will be consistent as long as one of the original N^2 estimates was consistent.

This application of CU for mode reduction is appropriate for MHT-type applications. However, CU must be generalized to accommodate weights/probabilities associated with modes when the representation is interpreted to be a Gaussian mixture approximation of a multimodal probability distribution. This requires a generalization of the definition of consistency for multimodal estimates. We require that each probability p_i be greater than or equal to the actual probability that estimate/mode i corresponds to the true state of the target. The problem is that any small nonzero probability implies that the associated estimate may represent the true state of the target, so consistency requires it to have the same influence on the CU result as an estimate with a much higher probability. The only difference is that the final result can be interpreted as having an associated probability that is equal to $\min(1, \sum_i p_i)$, where the *min* function is required because the weights are assumed to be conservative and thus may sum to a value greater than unity. Thus, the MHT case is equivalent to having no probability estimates, which requires unity to be assumed for every mode.

4 Computational Methods

The Covariance Intersection optimization has simple linear constraints that are compatible with any generic constrained optimization package. In contrast, the Covariance Union optimization problem has matrix-valued semidefinite constraints which generally required specialized optimization techniques. For our experiments, however, we have applied generic optimization methods, which are discussed in this section.

The CU constraints contain a quadratic term, but a simple rearrangement converts each constraint into a linear matrix inequality (LMI) in \mathbf{u} and \mathbf{U} . Note that the k^{th} CU constraint in (2) can be rearranged as the Schur complement (see A.5.5 of [2]) of the entry '1' in:

$$\begin{bmatrix} (\mathbf{U} - \mathbf{A}_k) & (\mathbf{u} - \mathbf{a}_k) \\ (\mathbf{u} - \mathbf{a}_k)^T & 1 \end{bmatrix} \quad (3)$$

This means that the semidefinite constraints in (2) are equivalent to:

$$\begin{bmatrix} \mathbf{U} & \mathbf{u} \\ \mathbf{u}^T & 1 \end{bmatrix} \geq \begin{bmatrix} \mathbf{A}_k & \mathbf{a}_k \\ \mathbf{a}_k^T & 0 \end{bmatrix}, \quad k = 1 \dots m \quad (4)$$

The intersection of all of the constraints is represented as a larger block-diagonal inequality in which the diagonal elements are the LMIs in (4). This demonstrates that CU trace minimization is a standard semidefinite programming (SDP) problem[2], since trace is linear. There is no such SDP formulation for determinant minimization, so we vectorized the constraints by computing their eigenvalues. The resulting optimization is non-smooth so we used the *SolvOpt*[4] nonlinear non-smooth optimizer. *SolvOpt* is an implementation of Shor's r-algorithm[5]. An initial feasible solution is generated by setting \mathbf{u} to zero and summing the right-hand sides of the simplified constraints:

$$u_0 = 0 \quad (5)$$

$$U_0 = \sum_{k=1}^n (\mathbf{A}_k + \mathbf{a}_k \mathbf{a}_k^T) \quad (6)$$

We have developed several approximate solutions which are much faster while still preserving consistency. These solutions are suitable for real-time use and are also good starting points for iterative improvement. Most of them rely on separation of the \mathbf{u} and \mathbf{U} optimizations to achieve computational savings. If the \mathbf{u} vector is fixed at a specific value then the problem is considerably simplified: find a minimal \mathbf{U} such that $\mathbf{U} \geq \mathbf{F}_k$ where the \mathbf{F}_k are constant. This simpler problem yields closed-form solutions when there are only two estimates. For example, if determinant is the measure used then \mathbf{U} is computed directly via simultaneous diagonalization:

$$\mathbf{U} = (\mathbf{V}^T)^{-1} \max(\mathbf{V}^T \mathbf{F}_1 \mathbf{V}, \mathbf{V}^T \mathbf{F}_2 \mathbf{V}) \mathbf{V}^{-1} \quad (7)$$

where *max* is the component-wise maximum of two diagonal matrices. \mathbf{V} contains the generalized eigenvectors of \mathbf{F}_1 and \mathbf{F}_2 .

One such approximation is to assume that real-life applications produce estimates in which the optimal

mean \mathbf{u} is a convex combination of the input means. This constrains \mathbf{u} to a bounded region in \mathbf{R}^n . Indeed if there are only two estimates (\mathbf{a}, \mathbf{A}) and (\mathbf{b}, \mathbf{B}) then \mathbf{u} is constrained to the line segment between \mathbf{a} and \mathbf{b} . Let $\mathbf{c} = \mathbf{b} - \mathbf{a}$, $\mathbf{u} = \mathbf{a} + \omega\mathbf{c}$. The problem is reduced to:
Find a minimal \mathbf{U} such that:

$$\mathbf{U} \geq \mathbf{A} + \omega^2 \mathbf{c}\mathbf{c}^T \quad (8)$$

$$\mathbf{U} \geq \mathbf{B} + (1 - \omega)^2 \mathbf{c}\mathbf{c}^T \quad (9)$$

This is convex in ω and is easily solved via a simple one-dimensional search.

In our experiments we observed that convex-combination CU produces reasonably good approximations to the optimal values when applied to two estimates in low dimensions. However, its performance has not yet been fully characterized. We evaluated it using determinant minimization on pairs of estimates whose mean components and covariance eigenvalues were randomly chosen on the interval $(0, 1)$, and the dimensionality n varied from 2 to 20. For the two-dimensional data, the determinant of \mathbf{U} produced by convex-combination CU was on average only 4% larger than the optimal value. For $n = 20$ it was 20% larger. The overall increase was roughly proportional to \sqrt{n} .

Another fast real-time approximation is derived by noting that the optimal two-element CU update tends to produce a \mathbf{u} vector for which the two constraints are similar in size and shape:

$$\mathbf{A} + (\mathbf{u} - \mathbf{a})(\mathbf{u} - \mathbf{a})^T \approx \mathbf{B} + (\mathbf{u} - \mathbf{b})(\mathbf{u} - \mathbf{b})^T \quad (10)$$

This observation suggests a strategy in which \mathbf{u} is fixed at the point where the difference is minimized. Minimization of the Frobenius norm leads to a closed-form solution for \mathbf{u} . Let $\mathbf{c} = (\mathbf{a} - \mathbf{b})/2$. Then:

$$\mathbf{u} = \left(\mathbf{a} + \mathbf{b} + ((\mathbf{c}^T \mathbf{c}) \mathbf{I} + \mathbf{c}\mathbf{c}^T)^{-1} (\mathbf{A} - \mathbf{B}) \mathbf{c} \right) / 2 \quad (11)$$

We tested this solution with random data. It produces good estimates when the differences between the estimates' means is large compared to the differences between the estimates' covariance matrices.

Large problems with many estimates can be broken down into a set of smaller problems by recursively solving two estimates at a time. For example, if there are three estimates $(\mathbf{a}_1, \mathbf{A}_1)$, $(\mathbf{a}_2, \mathbf{A}_2)$, and $(\mathbf{a}_3, \mathbf{A}_3)$ they are separated into two smaller problems:

1. Compute $(\mathbf{u}_1, \mathbf{U}_1)$ as the union of $(\mathbf{a}_1, \mathbf{A}_1)$ and $(\mathbf{a}_2, \mathbf{A}_2)$.
2. Compute $(\mathbf{u}_2, \mathbf{U}_2)$ as the union of $(\mathbf{u}_1, \mathbf{U}_1)$ and $(\mathbf{a}_3, \mathbf{A}_3)$.
3. $(\mathbf{u}_2, \mathbf{U}_2)$ is the solution.

The main advantage of this approach is that two-element unions are solved quickly via convex combination CU using closed-form formulas such as equation (7). As discussed in subsection 4.1, this method has one serious drawback: it does not guarantee consistency. It does guarantee that the covariance matrices \mathbf{U}_k will never shrink and will most likely grow on every iteration. However, there is no guarantee that \mathbf{U}_{k+1}




Figure 1: *Example: original input estimates.*

will still be consistent with the earlier estimates when it is re-centered at a new mean \mathbf{u}_{k+1} . Previous experiments did not observe this effect due to the extra slack provided by the (sub-optimal) convex-combination formulation.

4.1 Consistency

Fast suboptimal CU methods are good approximations in certain simple cases, but do not guarantee consistency as the dimensionality and the number of modes grows large. Even the optimal method described in the next subsection does not maintain consistency with the original input when the union is applied pairwise recursively. This results directly from CU's ability to incorporate the mean shift of each estimate in the current union, and inability to remember the mean shift information from any previous union. As the sequence of pairwise unions grows longer, the center of the union grows farther away from the original estimates; eventually, some of those estimates may lie completely outside the union of which they are supposed to be members.

To illustrate this, consider a simple 2D example. Figure 1 shows a set of six input estimates, $(\mathbf{a}_k, \mathbf{A}_k)$; each covariance is represented by its 1σ ellipse with its mean centered at the $+$. As a baseline, Figure 2 superimposes the optimal union, given by

$$\mathbf{U} \geq \mathbf{A}_k - (\mathbf{u} - \mathbf{a}_k)(\mathbf{u} - \mathbf{a}_k)^T \quad (12)$$

which is consistent with all six inputs. This result can only be obtained when the union is calculated as a batch operation. Figure 3 shows the result of the same optimal method applied pairwise recursively. When plotted in this manner, the union ellipse should enclose at least all the input means; clearly, this union is not consistent with all of the inputs. This can be verified




Figure 2: *Optimal CU calculated as a batch operation. This result is consistent with all six inputs.*

numerically as well, since the calculated union violates the original constraints.

This behavior is true not only for a pairwise union of multiple modes, but also for any sequence of union operations, in which the result is used as input for another union. To correct this, we change the CU constraint from its optimal form, as in (12), to the general form,

$$\mathbf{U} \geq \frac{\mathbf{A}_k}{\alpha_k} - \frac{(\mathbf{u} - \mathbf{a}_k)(\mathbf{u} - \mathbf{a}_k)^T}{1 - \alpha_k} \quad (13)$$

by adding an α_k term for each input $(\mathbf{a}_k, \mathbf{A}_k)$, where α_k must be solved for in the range $(0, 1)$. This form guarantees consistency with the input, for any sequence of union operations, by finding the circumscribing ellipse, completely enclosing the input ellipses.

Continuing with the example, Figure 4 shows the general CU result when applied to the inputs as a single batch operation. Figure 5 shows the same result, but applied as a series of pairwise operations. While we introduce uncertainty into the system by inflating the covariance \mathbf{U} , we can now guarantee that the results of the current CU operation will maintain consistency when used as input to another CU operation.

4.2 Implementation

The *SolvOpt* package is able to find a minimizing vector x according to a cost function $f(x)$, which may be optionally constrained by some function $g(x)$. We choose x to be the n elements of \mathbf{u} plus the $\frac{n(n+1)}{2}$ elements of the upper triangle of \mathbf{U} ¹.

¹When finding the general form of CU, discussed in subsection 4.1, we include the α_k values for each of the m estimates in vector x as well




Figure 3: *Optimal CU method applied to input pairwise recursively. The mean shift from each estimate to the center of the union does not carry through successive pairwise operations, and the result is inconsistent.*

We minimize the determinant of the covariance, \mathbf{U} , subject to the constraint that

$$\mathbf{X}_k = \mathbf{U} - \mathbf{A}_k - (\mathbf{u} - \mathbf{a}_k)(\mathbf{u} - \mathbf{a}_k)^T \quad (14)$$

have non-negative eigenvalues, for all $k \in [1, \dots, m]$, where m is the number of estimates given.

To find $|\mathbf{U}|$, we perform an LU decomposition of matrix \mathbf{U} , to generate an upper triangular matrix \mathbf{W} and a lower triangular matrix \mathbf{L} , such that $\mathbf{LW} = \mathbf{U}$. \mathbf{L} and \mathbf{W} are given by

$$\mathbf{L}_{ii} = 1 \quad (15)$$

$$\mathbf{L}_{ij} = \frac{1}{\mathbf{W}_{ii}} \left(\mathbf{U}_{ij} - \sum_{k=1}^j \mathbf{L}_{ik} \mathbf{W}_{kj} \right); \quad i > j \quad (16)$$

$$\mathbf{W}_{ij} = \mathbf{U}_{ij} - \sum_{k=1}^i \mathbf{L}_{ik} \mathbf{W}_{kj} \quad (17)$$

Then $|\mathbf{U}| = \prod_{i=1}^n \mathbf{W}_{ii}$. The complexity cost of this operation is $O(n^3)$.

The single value *SolvOpt* uses to constrain the minimization must be nonpositive. Since we want to constrain the eigenvalues of (14) to be nonnegative for all $k \in [1, \dots, m]$, we simply find the most negative of all nk eigenvalues, λ_{min} , and return $-\lambda_{min}$ as the constraint.

To compute the n eigenvalues of each \mathbf{X}_k , we follow a two-step procedure:

1. Find the Hessenberg form of $\mathbf{H}_k = \text{Hess}(\mathbf{X}_k)$
2. Apply the QR transform to \mathbf{H}_k until the eigenvalues are isolated on the diagonal




Figure 4: General form of CU as a batch operation.

The Hessenberg form of a symmetric matrix is tridiagonal, which simplifies the actual eigenvalue calculations. This technique works because the original matrix \mathbf{X}_k and its Hessenberg form \mathbf{H}_k have the same eigenvalues.

The QR algorithm iterates on \mathbf{H}_k until it approaches the Schur normal form, which contains the eigenvalues on the diagonal.

Each QR decomposition of \mathbf{H}_k results in \mathbf{Q} , which is orthogonal, and \mathbf{R} , which is upper triangular, such that $\mathbf{QR} = \mathbf{H}_k$. The algorithm proceeds as follows

$$\mathbf{QR} = \mathbf{H}_{k,s} \quad (18)$$

$$\mathbf{H}_{k,s+1} = \mathbf{RQ} \quad (19)$$

for $s = 0, 1, 2, \dots$, until \mathbf{H}_k is in the Schur normal form.

As has been discussed, *SolvOpt* evaluates the cost and constraint function callbacks to minimize $|\mathbf{U}|$ over the $n + \frac{n(n+1)}{2}$ elements of \mathbf{u} and the triangle of \mathbf{U} . To merge m estimates, the cost function performs $O(n^3)$ operations, the constraints function $O(mn^3)$. The number of iterations which *SolvOpt* must perform varies widely, from 1500 to 15000, depending on the batch dimensions and also the input data values. In the next section we present results showing the overall computational cost of this approach.

5 Experimental Results

In this section we present experimental results for different implementations of the optimal² CU algorithm, using *SolvOpt*, written in both Matlab and C. We have timed the application of CU on sets of random data to explore actual execution times for various dimensions n , and modes m . The times listed in the

²There appears to be no predictable significant effect on running time when the α_k parameters are included for the general form.




Figure 5: General form of CU as a series of pairwise operations.

following tables were obtained on a single 1.5 GHz Pentium computer.

Avg. execution times for Matlab (in secs)

Dimensions	2 Modes	4 Modes	8 Modes	16 Modes
2	0.91	1.21	1.94	2.22
4	22.76	10.75	12.78	21.63
6	40.95	80.58	55.68	74.41
8	230.50	204.36	231.83	276.55

Average execution times for C (in seconds)

Dimensions	2 Modes	4 Modes	8 Modes	16 Modes
2	<0.01	<0.01	0.01	0.03
4	0.43	0.62	1.89	2.73
6	2.42	6.25	14.18	30.61
8	11.50	37.05	63.16	146.87

These results show that the generality of the *SolvOpt* algorithm incurs a significant computational cost that makes it impractical for most real-time applications when the dimensionality and number of nodes is high.

6 Discussion

In this paper we have examined the problem of representing multimodal information using MHT and GMMs. We have discussed the fusion of information represented in the form of multiple mean and covariance estimates corresponding to distinct possible states, or modes of a distribution, for a tracked target. We have discussed how the fusion operation results in a multiplicative increase in the complexity of the representation that will grow exponentially over time unless bounded by a mechanism that can compress the representation to a fixed number of modes. We have

described how Covariance Union can be used to coalesce modes while preserving the rigor of the information management framework. Experiments demonstrate the effectiveness of our approach.

The main result of this paper is our *SolvOpt*-based algorithm, with implementations in Matlab and C, for computing CU solutions. Experimental results corroborate the correctness of the algorithm, but they also show that *SolvOpt* is not practical for real-time applications. It is expected, however, that our experimental codes will prove useful in support of future research on applications of CU.

References

- [1] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, New York NY, USA, 1988.
- [2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, United Kingdom, 2006.
- [3] S. Julier and J. K. Uhlmann. Fusion of time delayed measurements with uncertain time delays. In *2005 American Control Conference*, Portland, OR, USA, 8 – 10 June 2005.
- [4] A. Kuntsevich and F. Kappel. Solvopt - the solver for local nonlinear optimization problems. *Technical Report*, <http://www.uni-graz.at/imawww/kuntsevich/solvopt/>, 1997.
- [5] N.Z. Shor. *Minimization Methods for Non-Differentiable Functions*, volume 3. Springer-Verlag, Berlin, Germany, 1985.
- [6] Jeffrey Uhlmann. Covariance consistency methods for fault-tolerant distributed data fusion. *Information Fusion*, 4(3):201–215, Sept 2003.