

Trajectory Optimization Estimator for Impulsive Data Association

Matthew Travers[†], Todd Murphey[†], and Lucy Pao[‡]

Abstract—A new approach to multi-target data association is presented. In this new approach, the problem of correctly associating measurements to the object from which they originate is viewed as an impulse optimization. In the systems considered, a single sensor attempts to measure a pre-specified object of interest. Interpolating the discrete set of measurements, it is possible to find a “measurement trajectory.” The action of the sensor randomly measuring a system target other than the object of interest is viewed as an impulse in the measurement trajectory. It is possible to correctly associate measurements by determining the times at which the impulses occur. The main contribution of the work presented in this paper is to provide a method for estimating the total number of impulses as well as the times at which they occur.

I. INTRODUCTION

Multi-target data association is an area of interest in many different fields [1], [4], [5], [8], [10], [11], [12]. The work presented in this paper provides a new approach, referred to as *impulsive data association* (IDA), to the multi-target problem that utilizes the maximum principle to efficiently search the space of possible solutions.

Figure 1 presents the basic structure upon which IDA is based. In Figure 1 there are two separate trajectories. The dotted trajectory represents the object of interest and the dashed trajectory some object nearby. There is a sensor that attempts to measure the trajectory of the object of interest at all times. The sensor makes a single measurement at each time a measurement is taken. The solid black line in Figure 1 represents the portions of the two trajectories that are measured by the sensor. Viewing the solid black line as a single trajectory, at time $t = 1$, an impulse of magnitude δ_1 occurs. This impulse models the sensor switching from measuring the trajectory of the object of interest to measuring the trajectory of some object nearby. At time $t = 2$, another impulse, of magnitude δ_2 , occurs. This impulse models the action of the sensor switching back to measuring the trajectory of the object of interest.

In the work presented in this paper, we assume that the total number of objects present in the system is known to be two. The assumption that the total number of objects is known a priori is similar to the assumption in single-scan Markov chain Monte Carlo data association (MCMCDA) [9]. The main difference between IDA and single scan MCMCDA is that no information about the second object, i.e., not the object of interest, other than its dynamics are




Fig. 1: Deterministic trajectories of two objects. The dotted line represents the trajectory of a pre-specified object of interest and the dashed line an object nearby. The solid line represents the portions of the two trajectories that a single sensor measures. Considering the solid black line as a single trajectory, an impulse of magnitude δ_1 occurs in the signal at time $t = 1$. A second impulse of magnitude δ_2 occurs at time $t = 2$.

assumed to be known. Thus, we have no initial distribution information with respect to the second object.

Optimizing over the times at which the impulses occur is the process through which the problem of correctly associating measurements to their objects of origin is solved. In previous work [13] it was assumed that the total number of impulses was known a priori. The main contribution of the work presented in this paper is to provide an algorithm for estimating the total number of impulses.

The total number of impulses are estimated using a trajectory optimization algorithm. The trajectory optimization algorithm is similar to a continuous-time analogue of the *probabilistic data association filter* (PDAF) [2]. The main theoretical difference between the PDAF and the trajectory optimization procedure presented in this work is that the entire measurement history is taken into account at the same time in the trajectory optimization, as opposed to recursively. As shown in Section IV, this distinction has important implications when the number of incorrect measurements received in a row is large.

The structure of this paper is as follows. In Section II the multi-target data association problem is analytically defined. Section III-A contains the main contribution of this paper, which is a trajectory optimization algorithm that provides a way to estimate the number of impulses that occur as well as estimating when. Section III-B is a short

[†]Department of Mechanical Engineering, Northwestern University, Evanston, Illinois 60208

[‡]Department of Electrical, Computer, and Energy Engineering, University of Colorado at Boulder, Boulder, Colorado 80309

summary of the second-order method derived in [13], where the technical details of the optimization over the impulse times are provided. In Section IV, IDA is compared to the probabilistic data association filter. Lastly, Section V contains conclusions and directions of future work.

II. PROBLEM DEFINITION

The systems of interest to the work presented in this paper contain two objects that share the same dynamics (which will in general be nonlinear). There is a single sensor that attempts to measure the position of a pre-specified object of interest at all times. At random times the sensor measures the other object. Neither the times at which the sensor measures the position of the second object nor the total number of times that it occurs are known. The goal of the data association problem in this work is to determine both.

Analytically, the systems considered are described by

$$\begin{aligned} \dot{x}^i &= f(x^i(t), t), & x^i(0) &= x_0^i, \text{ where } i = \{1, 2\} \\ z &= h(x^i(t), t) + \nu, \end{aligned} \quad (1)$$

where z is the measurement, and ν is a noise term. In general, none of the results presented in this work depend on the distribution from which the noise term ν originates. Note that the measurement in (1) always originates from a single one of the two targets in the system.

Let us assume that the object of interest has dynamics described by $\dot{x}^1 = f(x^1(t), t)$, $x^1(0) = x_0^1$. The goal of the data association problem for this work is to then determine the times at which measurements originate from the other object in the system.

III. IMPULSIVE DATA ASSOCIATION (IDA)

This section is split into two separate subsections. The first subsection contains the main contribution of this work, which is a trajectory optimization procedure that serves the purpose of estimating the total number of impulses in the system as well as estimating the times at which they occur. The second subsection contains the analytical description of the optimization over impulse times [13].

A. Trajectory Optimization

First consider the infinite dimensional space of curves \mathcal{L} , where $\xi \in \mathcal{L}$ such and $\xi = (\alpha, \mu)$. Next define the cost function

$$S(\eta) = \int_0^T \ell(s, x(s), U(s)) ds \quad (2)$$

where $\eta \in \mathcal{T}$ and $\mathcal{T} \subset \mathcal{L}$ is a trajectory manifold [7] defined such that $\eta = (x, U)$ satisfies $\dot{x} = F(x, U)$ subject to the initial condition $x(0) = x_0$.

The constraint on the trajectory manifold, $\dot{x} = F(x, U)$, defines a *projection operator* $P : \mathcal{L} \rightarrow \mathcal{T}$ [6]

$$\eta = P(\xi) : \begin{cases} \dot{x} = F(x, U), & x(0) = x_0 \\ U = \mu + K(\alpha - x) \end{cases} \quad (3)$$

where U is a diagonal matrix function that is linear in the state and

$$\dot{x} = F(x, U) = U(t)f(x, t). \quad (4)$$

The exact details of obtaining the control U in (3) are left to the references [6], [7], but it should be mentioned that the linear optimal control signal is found in a way very similar to the optimal control found in the *linear quadratic regulator* (LQR). Note that regardless of whether or not the dynamics (1) are linear or not, the overall trajectory optimization described in this section is nonlinear. This multiplicative structure is chosen so that the control is nominally equal to the identity when the reference signal does not experience an impulse.

Having defined the projection operator, let us return to the cost (2) and explicitly state the trajectory optimization problem. Since the ultimate goal of this optimization is to find curves $\eta \in \mathcal{T}$, let us first state the constrained optimization

$$\min_{\eta \in \mathcal{T}} S(\eta). \quad (5)$$

The purpose of defining the projection operator (3) was to rewrite the constrained optimization (5) as an unconstrained problem

$$\min_{\xi \in \mathcal{L}} S(P(\xi)), \quad (6)$$

since \mathcal{L} is an infinite dimensional manifold.

Steepest descent is the method chosen to solve the unconstrained optimization defined in (6). In order to apply the steepest descent algorithm, we must first find several curves: $\zeta = (\beta, \nu)$ such that $\zeta \in T_\xi \mathcal{L}$ and $\gamma = (z, v)$ such that $\gamma \in T_\eta \mathcal{T}$, where $T_\xi \mathcal{L}$ is the tangent space of \mathcal{L} at ξ and $T_\eta \mathcal{T}$ is the tangent space to \mathcal{T} at η . Equation (3) also defines another operator $DP : T_\xi \mathcal{L} \rightarrow T_\eta \mathcal{T}$ such that

$$\gamma = DP(\xi) \circ \zeta : \begin{cases} \dot{z} = Az + Bv, & z(0) = z_0 \\ v = \nu + K(\beta - z) \end{cases}$$

where the control v is linear in the state (as in (3) the process of determining K is again similar to the LQR problem). The rest of the technical details of the implementation of the optimization of (6) are left to [3], [6], [7].

The technical details of the trajectory optimization are left to the references [6], [7] so that the authors can focus on its implementation here. For example, consider the system where each of the two objects have dynamics described by

$$\begin{aligned} \dot{x} &= v(t) \cos(\theta(t)) \\ \dot{y} &= v(t) \sin(\theta(t)) \\ \dot{\theta} &= \omega(t), \end{aligned}$$

where the initial conditions between the two objects vary in the y -coordinate and $v(t)$ as well as $\omega(t)$ are constant. This difference in initial conditions leads to an impulse that also occurs in the y -coordinate. The measurement is assumed to be of the full state under the identity mapping. For this example, the $U(t)$ from (4) is defined to be

$$U(t) = \begin{pmatrix} u_1 & 0 & 0 \\ 0 & u(t) & 0 \\ 0 & 0 & u_3 \end{pmatrix}, \quad (7)$$

where $u_1 = u_3 = 1$ due to the fact that the impulse always occurs in the y -direction.

The scenario considered is illustrated at the link <http://www.youtube.com/watch?v=MPg7mSmRh2I>. This link shows a movie of two objects, assumed to be airplanes, moving through three periods on circular trajectories. The sensor is represented by the grey circle located at $(0, 0.5)$. When the sensor is correctly measuring the position of the object of interest, the line from the sensor to the planes is green. When the second object occludes the sensor's view of the object of interest, the line from the sensor to the planes turns red. We assume that there are three intervals over which the second object, i.e., not the object of interest is the object that is measured. The solid black line in Figure 2(a) shows the measurement signal produced by the sensor over a single period, where the dotted line represents the trajectory of the object of interest and the dashed line the trajectory of the second object. Note that the portion of the solid black measurement signal at the bottom of Figure 2(a), which roughly falls on top of the dashed trajectory, represents the portion of the measurement signal that originates from the second object. Figure 2(b) shows the control signal in the y -direction that results from applying the trajectory optimization. Due to the fact that there are three intervals over which incorrect measurements are received, we expect there to be six impulses modeling the switching between measuring the trajectories of the two objects (one impulse switching from the trajectory of the object of interest to the trajectory of the second object, and another to switch back). The impulses in the measurement signal show up as peaks in the control signal in the trajectory optimization. In order to detect these peaks automatically, we need to threshold the control signal $u(t)$.

In order to threshold the control signal, we need to make the assumption that the probability of detection, i.e., the probability that the correct measurement is received, is known. Knowing the total time over which measurements are received as well as the probability of detecting a correct measurement, it is possible to estimate the total time over which correct measurements are received.

In order to facilitate the thresholding procedure, it is helpful to first low-pass filter the control signal $u(t)$. The results of low-pass filtering the control signal shown in Figure 2(b) is shown in Figure 3(a). The maximum deviation in the filtered signal $u'(t)$ away from the nominal value $u'(t) = 1$ can be used to define upper and lower bounds for a window around $u'(t) = 1$. In the current example, the upper bound of this window is $u'(t) = 1.0778$ and the lower bound is $u'(t) = 0.9222$. The thresholding procedure is performed by evenly shrinking this window around until the percentage of the total measurement time enclosed by the window is equal to probability of association. Figure 3(b) shows the result of applying this procedure. The dashed lines in Figure 3(b) represent the boundaries of the window around $u'(t) = 1$ such that the percentage of the time enclosed by these bounds is equal to the probability of association. As shown in Figure 3(b), six peaks, corresponding to the six impulses in the measurement signal, are detected by this thresholding procedure.




Fig. 2: (a) Measurements for the two-plane example (solid line). The dotted line represents the trajectory of the object of interest and the dashed line the trajectory of some plane nearby. (b) The control signal that results from applying the trajectory optimization to the two-plane system when six total impulses occur.

A natural question is what happens to the thresholding described in this section when the amount of relative noise in the measurement signal is increased. Figure 4(a) shows the same basic picture as that in Figure 2(a), except that the initial conditions of the two airplanes are closer together. The dotted line again represents the trajectory of the object of interest, the dashed line the trajectory of the second object, and the solid line the measurement signal. The distribution from which noise is sampled is the same for the measurement signals shown in Figures 2(a) and 4(a). Decreasing the distance between the two airplanes increases the relative amount of noise present in the measurement signal. This increase can be seen by comparing the control signals that result from applying the trajectory optimization shown for the case where the airplanes are relatively far apart, Figure 2(b), and relatively close together, Figure 4(b). Note that the distance between the two airplanes in the example shown in Figure 4 is equal to twice the standard deviation of the




Fig. 3: (a) Filtered control signal and (b) thresholding applied to the filtered signal.

distribution from which the noise in the measurement signal is sampled (we are assuming that the standard deviation is equal in both the x - and y - directions).

Figure 5(a) shows the result of low-pass filtering the control signal shown in Figure 4(b). Figure 5(b) shows the result of applying the thresholding procedure described in this section to the filtered signal shown in Figure 5(a). Figure 5(b) shows that we are still able to identify six peaks corresponding to the six impulses in the measurement signal, but it is clear that the local minima in the filtered signal $u'(t)$ around the threshold values (represented by the dashed lines) are making this procedure less clear.

B. Optimization Over Impulse Times

As it was mentioned previously, the trajectory optimization procedure in Section III-A serves the purpose of providing an estimate of the total number of impulses as well as the times at which they occur in the multi-target systems that are the focus of this work. The estimates of the number of impulses and times at which they occur serve as inputs for a second optimization procedure.

The second optimization in the IDA algorithm is over the times at which the impulses occur. As it was mentioned earlier, the purpose of optimizing over the impulse times is that in doing so it is possible to determine the intervals over which measurements originate from an object other than the object of interest.

In this subsection, second-order adjoint equations for the second-order derivatives of the cost are provided. As it will be shown, the reason for providing the adjoint equations is that it drastically reduces the number of computations needed




Fig. 4: (a) Measurements for the two-plane example (solid line) where the relative distance between the two airplanes is relatively small when compared to the standard deviation of the distribution from which the noise in the measurement signal is sampled. The dotted line represents the trajectory of the object of interest and the dashed line the trajectory of some plane nearby. (b) The control signal that results from applying the trajectory optimization.

for calculating the derivatives when there are even a small number of impulses in the system.

Define the cost function

$$J(\cdot) = \int_{t_0}^{t_f} \ell(x(s), s) ds \quad (8)$$

where $\ell(x(s), s) = (x_d(s) - x(s))^T (x_d(s) - x(s))$.

The second derivative of the cost function $J(\cdot)$ with respect to the switching time τ_j where $\tau_i \geq \tau_j$ is

$$\begin{aligned} D_{\tau_j} D_{\tau_i} J(\cdot) \circ (\partial \tau_j, \partial \tau_i) = \\ D_1 \ell(x(\tau_i^-), \tau_i^-) \circ (D_{\tau_j} x_d(\tau_i^-) \circ \frac{\partial \tau_i}{\partial \tau_j} - D_{\tau_j} x(\tau_i^-) \circ \partial \tau_j) \\ - D_1 \ell(x(\tau_i^+), \tau_i^+) \circ (D_{\tau_j} x_d(\tau_i^+) \circ \frac{\partial \tau_i}{\partial \tau_j} - D_{\tau_j} x(\tau_i^+) \circ \partial \tau_j) \end{aligned} \quad (9)$$




Fig. 5: (a) Filtered control signal and (b) thresholding applied to the filtered signal when the distance between the two airplanes is small compared to the standard deviation from which the noise in the measurement signal is sampled.

$$-D_1 \ell(x(\tau_i), \tau_i) \circ X^i \partial \tau_i \delta_i^j + \psi(t_f, \tau_i) \circ X^{i,j} + \Omega(t_f, \tau_i) \circ (\Phi(\tau_i, \tau_j) \circ X^j, X^i) \quad (10)$$

where $\Omega(t, \tau) \circ (U, V) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is the bilinear operator found by integrating

$$\Omega(t, t) \circ (U, V) = 0_{n \times n} \quad (11a)$$

$$\begin{aligned} \frac{\partial}{\partial \tau} \Omega(t, \tau) \circ (U, V) &= -D_1^2 \ell(x(\tau), \tau) \circ (U, V) \\ &\quad - \psi(t, \tau) \circ D_1^2 f(x(\tau), \tau) \circ (U, V) \\ &\quad - \Omega(t, \tau) \circ (D_1 f(x(\tau), \tau) \circ U, V) \\ &\quad - \Omega(t, \tau) \circ (U, D_1 f(x(\tau), \tau) \circ V) \end{aligned} \quad (11b)$$

backwards over τ from t_f to τ_i and $X^{i,j}$ is an initial condition that results from taking the derivative of $D_{\tau_i} x(t) \circ \partial \tau_i$ with respect to τ_j . The proof of (10), (11a), and (11b) are again left to the reference [13].

Convergence results for the impulse optimization over six impulse times (Figure 2(b)) using a second-order method are shown in Figure 6. The horizontal axis in Figure 6 is the iteration number of the algorithm and the vertical axis the first derivative of the cost associated with the impulse times plotted on a log scale; quadratic convergence is indeed achieved.

IV. COMPARISON TO PDAF

In this section we compare the performance of a Kalman filter used in conjunction with IDA to the PDAF. The PDAF is a recursive method and is thus less computationally




Fig. 6: Second-order convergence results of applying impulse optimization to the two plane, six impulse system.

burdensome than IDA, which requires some explanation as to why this comparison is being made. Due to the fact that we are assuming no a priori knowledge of the second object in the system, and are considering only a single scan through the data, comparison to single scan MCMCDA [9] is not possible. While it is possible to compare to the multiple hypothesis tracker (MHT) [2], the computational complexity of MHT goes up as 2^n with the number of measurements. Adding measurements in the IDA algorithm amounts to integrating over a longer time period, and thus there is no exponential growth (with respect to the number of measurements received) in computational complexity.

Figure 7 shows a second scenario for the two airplane example. In Figure 7(a) the dotted line represents the trajectory of the object of interest, the dashed line the trajectory of a second object, and the solid line the portions of the two trajectories that are measured over a single period of the circular trajectories.

The solid black line in Figure 7(b) represents the measurement signal. The dashed line represents the result of applying the PDAF to the measurement signal, and the dotted line the result of applying a Kalman filter used in conjunction with IDA. The black point represents the final measurement originating from the object of interest before the impulse occurs, i.e., before the trajectory of the second object is measured.

Figure 8 shows an RMS error comparison between the PDAF (dashed line) and the Kalman filter (dotted line) used in conjunction with IDA. The horizontal axis represents the measurement number and the vertical axis the RMS error. The solid black point represents the last measurement originating from the object of interest before the second object begins to be measured. Figure 8 shows that when using the PDAF the RMS error is higher due to the addition of measurements from the second object (i.e., not the object of interest). The added error is due to the fact that we are not including the measurements determined to be from the second object into the updated state calculated using the Kalman filter in conjunction with IDA, we are thus partitioning the set of measurements. We are thus able to accomplish this partitioning without the added complexity inherent in MHT or multi-scan MCMCDA.




Fig. 7: Results of applying PDAF (dashed line) as well as IDA (dotted line) to a set of noisy measurements (solid line) over a single period of the counterclockwise-circular trajectories shown. The large point represents the last measurement from the object of interest before the first impulse is detected by IDA.

V. CONCLUSIONS AND FUTURE WORK

The focus of this paper is to introduce a method that provides an estimate of the number of impulses as well as the times at which they occur. This information serves as an input for an optimization procedure which optimizes over impulse times. This optimization is a data association procedure which determines the periods over which measurements not originating from an object of interest are received. It was shown that by treating every measurement received at the same time, i.e., by optimizing over the entire “measurement




Fig. 8: RMS error comparison between the Kalman filter used in conjunction with IDA (dotted line) and the PDAF (dashed line). The large point corresponds to the measurement at which the first impulse is detected by IDA. Note that the large point in Figure 7 corresponds to the same measurement as it does here.

trajectory”, an advantage in terms of RMS error could be achieved over recursive methods (PDAF).

The future direction of this work is to optimize over the impulse magnitudes as well as the impulse times. Optimizing over the impulse magnitudes, it is possible to eliminate any assumptions on the number of system targets. The resulting algorithm would perform multi-target data association with similar results to multi-scan MCMCDA [9].

REFERENCES

- [1] T. Bailey and H. Durrant-Whyte. Simultaneous Localization and Mapping (SLAM): Part II. *Robotics Automation Magazine, IEEE*, 13(3):108–117, Sep. 2006.
- [2] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, Inc., 1988.
- [3] T. Caldwell and T. D. Murphey. Switching Mode Generation and Optimal Estimation with Application to Skid-Steering. *Automatica*, 2010. In Press.
- [4] I. Cox. A Review of Statistical Data Association Techniques for Motion Correspondence. *International Journal of Computer Vision*, 10(1):53–66, 1993.
- [5] F. Dellaert, S. Seitz, C. Thorpe, and S. Thrun. EM, MCMC, and Chain Flipping for Structure from Motion with Unknown Correspondence. *Machine learning*, 50:45–71, 2003.
- [6] J. Hauser. A Projection Operator Approach to the Optimization of Trajectory Functionals. *IFAC World Congress*, 2002.
- [7] J. Hauser and D. G. Meyer. The Trajectory Manifold of a Nonlinear Control System. In *IEEE Conference on Decision and Control*, pages 1034–1039, 1998.
- [8] M. K. Kalandros, L. Trailović, L. Y. Pao and Y. Bar Shalom. Tutorial on Multisensor Management and Fusion Algorithms for Target Tracking. In *American Controls Conf. (ACC)*, 2004.
- [9] S. Oh, S. Russell, and S. Sastry. Markov Chain Monte Carlo Data Association for General Multiple-Target Tracking Problems. In *IEEE Conference on Decision and Control*, pages 735–742, 2004.
- [10] R. M. Powers and L. Y. Pao. Power and Robustness of a Track-Loss Detector Based on Kolmogorov-Smirnov Tests. In *American Controls Conf. (ACC)*, pages 3757 – 3764, 2006.
- [11] K. M. Reichard, E. C. Crow, and D. C. Swanson. Automated Situational Awareness Sensing for Homeland Defense. *SPIE System Diagnosis and Prognosis: Security and Condition Monitoring Issues Conference*, pages 64 – 71, 2003.
- [12] B. Schumitsch, S. Thrun, G. Bradski, and K. Olukotun. The Information Form Data Association Filter. In *Conference on Neural Information Processing Systems*, 2005.
- [13] M. Travers, T. Murphey, and L. Pao. Impulse Optimization for Data Association. In *IEEE Conference on Decision and Control*, 2010. found at: <http://robotics.mech.northwestern.edu/murphey/papers.html>.