UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

AFOSR-TR-96 0349

REPORT DOCUMENTATION PA[GE]

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S)
6a. NAME OF PERFORMING ORGANIZATION University of Connecticut 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION AFOSR/NM - Dr. Neal Glassman
6c. ADDRESS (City, State and ZIP Code) Box U-157 260 Glenbrook Road Storrs, CT 06269-2157	7b. ADDRESS (City, State and ZIP Code) Dept. of the Air Force Air Force Office of Scientific Research Bldg. 410, Bolling Air Force Base Washington, DC 20332-6448
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFPSR/NM 8b. OFFICE SYMBOL (If applicable) NM	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F49620-93-1-0399
8c. ADDRESS (City, State and ZIP Code) Department of the Air Force Air Force Office of Scientific Research Bldg. 410, Bolling Air Force Base Washington, DC 20332-6448	10. SOURCE OF FUNDING NOS. PROGRAM ELEMENT NO.: 61103D PROJECT NO.: 3484 TASK NO.: 2S WORK UNIT NO.:
11. TITLE (Include Security Classification) Optimization for Trajector Estimation (Unclassified)	

12. PERSONAL AUTHOR(S)
Robert Popp, Yaakov Bar-Shalom and Krishna Pattipati

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Yr., Mo., Day)	15. PAGE COUNT
final	FROM 93/06/01 TO 96/05/31	96/06/14	12

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB. GR.	Mathematical Optimization Target Tracking Estimation

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This work has dealt with the development of practical advanced multisensor-multitarget tracking algorithms for optimal processing of the information obtained from various remote sensing devices (e.g., radar) for tracking targets. The processing consists of filtering, i.e., state estimation of sensor data, across time, coupled with the fusion, i.e., association of that data, across sensors, wherein the main goals are overcoming the inherent limitations of real-world sensors (accuracy and reliability) due to noise and clutter while not sacrificing computationa performance.

19960726 077

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS ☐	21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a. NAME OF RESPONSIBLE INDIVIDUAL ~~Yaakov Bar-Shalom~~ Neal Glassman	22b. TELEPHONE NUMBER (Include Area Code) ~~860-486-4823~~ 202-767-5026
22c. OFFICE SYMBOL NM	

DD FORM 1473, 83 APR EDITION OF 1 JAN 73 IS OBSOLETE.

SECURITY CLASSIFICATION OF THIS PAGE

FINAL REPORT

AASERT Grant F49620-93-1-0399
01 June 1993 — 31 May 1996

Optimization Techniques for Efficient Ballistic Trajectory Estimation

Robert Popp, Yaakov Bar-Shalom, and Krishna Pattipati
Department of Electrical & Systems Engineering
University of Connecticut, U-157
Storrs, CT 06269-3157
Phone: (860) 486-4823
Fax: (860) 486-3789
email: ybs@ee.uconn.edu

June 14, 1996

1 Objectives

This work has dealt with the development of practical advanced multisensor-multitarget tracking algorithms for optimal processing of the information obtained from various remote sensing devices (e.g., radar) for tracking targets. The processing consists of filtering, i.e., *state estimation* of sensor data, across time, coupled with the fusion, i.e., *association* of that data, across sensors, wherein the main goals are overcoming the inherent limitations of real-world sensors (accuracy and reliability) due to noise and clutter while not sacrificing computational performance.

The state estimator used in our studies is the Interacting Multiple Model (IMM) estimator, the most versatile self-adjusting variable-bandwidth filter available. This filter, developed under earlier AFOSR grandparent grant sponsorship, has seen numerous applications [2, 3]. The data association algorithm considered is a 2D (two-dimensional) assignment algorithm — bipartite matching. This algorithm, called "modified auction", was used in earlier works in our development of a relaxation-based multidimensional (S-D, $S \geq 3$) assignment algorithm [1, 3].

The following are some of the specific objectives of the present research:

1. Association and fusion of measurements from multiple, asynchronous heterogeneous sensors based on discrete mathematical optimization techniques (assignment) for real-world practical problems.

2. Development of a tracking algorithm for multitarget problems wherein the state estimation and data association (assignment) problems are integrated.

3. Development of robust, scalable parallelizations of the tracking algorithm, for both shared-memory and distributed-memory computing architectures, adaptable to diverse multitarget scenarios, maintaining a high level of computational performance independent of the number of modules used in the state estimator.

4. Integration of the tracking algorithm into the Rome Laboratory (RL) Multi-Integration Software Tool (MIST) testbed, an in-house tool used to analyze and evaluate tracking algorithms against various data sets and sensor technologies.

5. Application of matching the assignment algorithms developed for the tracking problems to scheduling and resource allocation problems in command and control as well as manufacturing systems.

6. Development of suboptimal scheduling algorithms using mathematically-based techniques well-rooted in "classical" scheduling theory, i.e., Lagrangian relaxation, dynamic programming, and precedence-constrained network optimization on directed graphs.

2 Status of Effort

1. Developed a multitarget tracking algorithm for the tracking of multiple airborne targets, termed IMM-2D in the sequel, wherein the state estimation problem, via an IMM state estimator, is coupled with the data association problem, via an optimization-based 2D assignment problem.

2. Developed parallelizations of IMM-2D for both shared-memory (namely, on 2- and 4-node SPARCstation 20s, 4-node SPARCserver 1000, and a 12-node SPARCcenter 2000) and distributed-memory (i.e., 32- and 321- node Intel Paragon supercomputer) computing architectures.

3. Developed a robust, "coarse-grained", shared-memory parallel algorithm for IMM-2D, adaptable to diverse multitarget scenarios, maintaining a high level of computational performance independent of the number of filter modules used in the state estimator.

4. Developed a scalable, SPMD distributed-memory parallel algorithm for IMM-2D using relatively simple task allocation algorithms that minimize load imbalances and offer great promise in practice.

5. A sequential IMM-2D tracking/assignment algorithm has been successfully integrated into the RL MIST testbed, while the (shared- and distributed-memory) parallel IMM-2D algorithms are presently being integrated into MIST.

6. Developed several suboptimal, Lagrangian relaxation-based algorithms for the scheduling of parallelizable tasks, where the overall completion time for tasks (make-span) and the weighted sum of task completion times (weighted response) are minimized.

3 Accomplishments and New Findings

1. For a multitarget tracking algorithm where an IMM state estimator is coupled with a 2D data association (assignment) algorithm (IMM-2D), on shared-memory multiprocessor systems, a "coarse-grained" (dynamic) parallelization *across* the IMMs is superior, in terms of computational performance, to a "fine-grained" (static) parallelization *within* the IMMs.

2. For multitarget problems, the computational performance of the dynamic parallelization of IMM-2D, unlike a static parallelization, is *independent* of the number of filter modules used in the state estimator; the former realizes marginal performance only when an unrealistic number of filter modules are used in the state estimator.

3. When IMM-2D is configured with four filter modules in the IMM, the static parallelization has worse execution time than sequential time.

4. In addition to having the potential of realizing superlinear speedups, the dynamic parallelization of IMM-2D scales to larger multiprocessor systems and is robust, i.e., it adapts to diverse multitarget

scenarios maintaining the same level of efficiency given any one of numerous factors influencing the problem size.

5. The *interface* to the data association (optimization) problem is the major computational bottleneck, namely, computing the rather numerous gating and IMM state estimates, covariance calculations, and likelihood function evaluations (used as cost coefficients in the assignment problem), rather than the solution of the assignment problem itself.

6. Using a measurement database based on two FAA air traffic control radars, courtesy of Rome Laboratory, a SPMD distributed-memory parallelization of IMM-2D on a 32-node Intel Paragon supercomputer realizes near linear speedups using relatively simple task allocation algorithms that minimize load imbalances.

4 Personnel Supported

- Robert Popp, Jim Monte

5 Publications and Abstracts

5.1 Published

1. R. Popp, K. Pattipati, and R. Gassner, "Multitarget Tracking Parallelization for High-Performance Computing Architectures," *4th Annual SCS Symp. on High-Performance Computing*, New Orleans, LA, April 1996.

Abstract

The problem under investigation is the detection and estimation of potentially hundreds of maneuvering airborne targets in "real-time", using measurements from one or more active sensors. In this paper, we present a robust, parallel algorithm to track the states of such targets for the distributed-memory "message passing" class of parallel computing architectures. Specifically, we embed the Interacting Multiple Model (IMM) estimator into a two-dimensional (2 D) assignment-based algorithm to handle both the state estimation and measurement-to-track data association problems. Based on the supervisor/worker paradigm, two parallelizations are described, differing in the track distribution strategies employed. In particular, we show that a "static" mod p (p processors) uniform distribution scheme does not reflect the computational load across the processor set, and, hence, degraded performance ensues because of load imbalances. However, a distribution scheme based on a *track age* metric, which is an indicator of the expected workload, indirectly improves performance by "dynamically" balancing the workload across the processor set via a reallocation of tracks. We develop and demonstrate the parallelizations using a measurement database based on two FAA air traffic control radars and a 32-node Intel Paragon high-performance computer (HPC).

2. R. Popp, K. Pattipati, and R. Gassner, "Heuristic Task Assignment Algorithms Applied to Multisensor-Multitarget Tracking," *SPIE Conf. on Signal & Data Processing of Small Targets*, Orlando, FL, April 1996.

Abstract

In this paper, we are concerned with the problem of assigning track tasks, with uncertain processing costs and negligible communication costs, across a set of homogeneous processors within a distributed computing system to minimize workload imbalances. Since the task processing cost is uncertain at the time of task assignment, we propose several fast heuristic solutions that are extensible, incur very little overhead, and typically react well to changes in the state of the workload. The primary differences between the task assignment algorithms proposed are: (i) the definition of a task assignment cost as a function of past, present, and predicted workload distributions, (ii) whether or not information sharing concerning the state of the workload occurs among processors, and (iii) if workload state information is shared, the reactiveness of the algorithm to such information (i.e., high-pass, moderate, low-pass information filtering). We show, in the context of a multisensor-multitarget tracking problem, that using the heuristic task assignment algorithms proposed can yield excellent results and offer great promise in practice.

3. R. Popp, K. Pattipati, Y. Bar-Shalom, and M. Yeddanapudi, "The Parallelization of a Large-Scale IMM-based Multitarget Tracking Algorithm," *SPIE Conf. on Signal & Data Processing of Small Targets*, San Diego, CA, July 1995.

Abstract

The Interacting Multiple Model (IMM) estimator has been shown to be superior, in terms of tracking accuracy, to a well-tuned Kalman filter when applied to tracking maneuvering targets. However, because of the increasing number of filter modules necessary to cover the possible target maneuvers, the IMM estimator also imposes an additional computational burden. Hence, in an effort to design a real-time IMM-based multitarget tracking algorithm that is *independent* of the number of modules used in the IMM estimator, we propose a "coarse-grained" (dynamic) parallel implementation that is superior, in terms of computational performance, to previous "fine-grained" (static) parallelizations of the IMM estimator. In addition to having the potential of realizing superlinear speedups, the proposed implementation scales to larger multiprocessor systems and is robust. We demonstrate the performance results both analytically and using a measurement database from two FAA air traffic control radars.

4. R. Popp, "The Parallelization of a Multitarget Tracking Algorithm for a Class of High-Performance Computing (HPC) Architectures," BRC/CSE-TR-95-13, also, *AFOSR Final Report*, Culver City, CA, September 1995.

Abstract

In recent years, it has been shown to be both viable and economically feasible to use large-scale parallel and distributed high-performance computing (HPC) architectures when solving compute-bound type problems. Within the aerospace community, with the availability of such powerful and affordable parallel processing systems, presently there has been increased interest in the parallelization of "traditionally" computationally-intensive algorithms utilized for multitarget tracking. Multitarget tracking algorithms based on an Interacting Multiple Model (IMM) estimator have been shown to be very effective when applied to tracking maneuvering targets. However, the IMM estimator imposes a computational burden in terms of both space and time complexity, since more than one filter module is used to calculate state estimates, covariances, and likelihood functions.

Hence, in an effort to improve its performance computationally, the primary focus of this report is to present the results of our investigation on the parallelization of an IMM-based multitarget tracking algorithm for a distributed memory, message passing based class of high-performance computers. In particular, using a measurement database based on two FAA air traffic control radars, courtesy of Rome Laboratory, we utilize an Intel Paragon model XP/E HPC as our system platform to develop and evaluate the parallel tracker. We show that good computational performance (in terms of speedup and efficiency) results when utilizing a parallelization technique modeling the supervisor/worker approach and distributing the track data from the multitarget problem randomly and uniformly via the *modulo* scheme. However, since the number of association costs computed, not the size of the local track list, is directly correlated to the nodal computation time, intelligent load balancing schemes are necessary because, in an increasing number of processors, the nodal computation time becomes imbalanced resulting in degraded performance for the multitarget tracker.

5. R. Popp, "Multisensor-Multitarget Data Fusion Using an S-Dimensional Sliding Window Assignment Algorithm," BRC/CSE-TR-94-14, also, *AFOSR Final Report*, Culver City, CA, September 1994.

Abstract

The purpose of this paper is to demonstrate the use of the "sliding" multiscan windows approach in solving the multidimensional assignment problem for both track initiation and track extension. As part of AFOSR funded research and the current research effort for Rome Laboratory, the University of Connecticut has developed several data association and estimation algorithms for actual tracking

problems of interest to the Air Force. One such algorithm developed by Pattipati and Deb is the near-optimal, highly recursive, polynomial time S-dimensional assignment algorithm. In this work, as part of AFOSR's Summer Research Program, we implemented and tested (via an actual non-stressing multisensor-multitarget data set) a multiscan sliding-windows tracker utilizing the S-dimensional assignment algorithm as our data association component. The overriding objective was to demonstrate the use of a (sliding) multiscan window approach in addressing the problem of efficient and reliable data association, over time, for track initiation and extension.

5.2 Accepted for Publication

1. R. Popp, K. Pattipati, Y. Bar-Shalom, and M. Yeddanapudi, "Parallelization of a Multiple Model Multitarget Tracking Algorithm with Superlinear Speedups," *IEEE Trans. Aerospace & Electronic Systems*, (forthcoming 1996).

Abstract

The Interacting Multiple Model (IMM) estimator has been shown to be very effective when applied to air traffic surveillance problems. However, because of the additional filter modules necessary to cover the possible target maneuvers, the IMM estimator also imposes an increasing computational burden. Hence, in an effort to design a real-time multiple model multitarget tracking algorithm that is *independent* of the number of modules used in the state estimator, we propose a "coarse-grained" (dynamic) parallelization that is superior, in terms of computational performance, to a "fine-grained" (static) parallelization of the state estimator, while not sacrificing tracking accuracy. In addition to having the potential of realizing superlinear speedups, the proposed parallelization scales to larger multiprocessor systems and is robust, i.e., it adapts to diverse multitarget scenarios maintaining the same level of efficiency given any one of numerous factors influencing the problem size. We develop and demonstrate the dynamic parallelization on a shared-memory MIMD multiprocessor for a civilian air traffic surveillance problem using a measurement database based on two FAA air traffic control radars.

2. R. Popp, K. Pattipati, Y. Bar-Shalom, and R. Gassner, "Multitarget Tracking Algorithm Parallelization for Distributed-Memory Computing Systems," *5th IEEE Int. Symp. on High-Performance Distributed Computing*, Syracuse, NY, (forthcoming August 1996).

Abstract

In this paper we present a robust scalable parallelization of a multitarget tracking algorithm developed for air traffic surveillance. We couple the state estimation and data association problems by

embedding an Interacting Multiple Model (IMM) state estimator into an optimization-based assignment framework. A SPMD distributed-memory parallelization is described, wherein the interface to the optimization problem, namely, computing the rather numerous gating and IMM state estimates, covariance calculations, and likelihood function evaluations (used as cost coefficients in the assignment problem), is parallelized. We describe several heuristic algorithms developed for the inherent task allocation problem, wherein the problem is one of assigning track tasks, having uncertain processing costs and negligible communication costs, across a set of homogeneous processors to minimize workload imbalances. Using a measurement database based on two FAA air traffic control radars, courtesy of Rome Laboratory, we show that near linear speedups are obtainable on a 32-node Intel Paragon supercomputer using simple task allocation algorithms.

5.3 Pending Publications

1. R. Popp, K. Pattipati, Y. Bar-Shalom, and R. Ammar, "Parallel Design of a Multitarget Tracking Algorithm for a Shared-Memory Multiprocessor System," *IEEE Trans. Parallel & Distributed Systems*, (1995).

Abstract

Multitarget tracking algorithms based on an Interacting Multiple Model (IMM) estimator have been shown to be very effective when applied to tracking maneuvering targets. However, the IMM estimator imposes a computational burden in terms of both space and time complexity, since more than one filter module is used to calculate state estimates, covariances, and likelihood functions. Indeed, the processing time may become impractical with an increase in the number of filter modules. Hence, in an effort to improve its computational performance, in this paper we present the results of our investigation and evaluation of various parallelizations of an IMM-based multitarget tracking algorithm on a general-purpose shared-memory, MIMD multiprocessor system. In particular, we propose a "coarse-grained" (multithreaded) parallelization *across* the numerous IMMs found in a multitarget tracking problem that is robust, scalable, and demonstrates superior computational performance to previously proposed "fine-grained" parallelizations *within* the IMM. We show analytically that the coarse-grained parallelization has better execution time than a fine-grained parallelization for any number of filter modules used in the IMM estimator. Furthermore, using a measurement database from two FAA air traffic control radars, we empirically show that the coarse-grained parallelization can realize superlinear speedups *independent* of the number of modules used in the IMM estimator, whereas the performance of the fine-grained parallelization is *dependent* on the number of modules used. In fact, when using fewer modules (four or less), the execution time of the fine-grained parallelization is greater than sequential time!

2. R. Popp, K. Pattipati, and Y. Bar-Shalom, "Multitarget Tracking Algorithm Parallelizations for Air Traffic Surveillance," *Annals of Operations Research: Special Issue on Parallel Optimization*, (1996).

Abstract

In this paper, we present robust scalable parallelizations of a multitarget tracking algorithm with application to air traffic surveillance. In an effort to develop practical algorithms for tracking maneuvering airborne targets, we solve the state estimation and data association problems by embedding an Interacting Multiple Model (IMM) state estimator into the assignment (i.e., weighted bipartite matching) framework. However, contrary to conventional wisdom, the motivation behind the parallelizations is the *interface* to the optimization problem being the major computational bottleneck, namely, computing the rather numerous gating and IMM state estimates, covariance calculations, and likelihood function evaluations (used as cost coefficients in the assignment problem), rather than the solution of the assignment problem itself! We describe a coarse-grained (dynamic) parallelization *across* the state estimators for general-purpose shared-memory MIMD multiprocessors, wherein superlinear speedups are obtainable *independent* of the number of filter modules used in the state estimator. In addition, a SPMD distributed-memory parallelization for an Intel Paragon supercomputer is described; near linear speedups are realized using relatively simple task allocation algorithms that minimize load imbalances. Using a measurement database based on two FAA air traffic control radars, we demonstrate the performance of our parallel algorithms and show that they offer great promise in practice.

3. J. Monte and K. Pattipati, "Scheduling Parallelizable Tasks to Minimize Make-Span and Weighted Response," *ORSA Journal on Computing*, (1996).

Abstract

This paper presents a generalization to classical scheduling theory by removing the restriction that only one processor can work on a given task at a particular time. Instead, it is assumed that each task can be allocated any number of identical processors from one to the maximum number available, with each task's completion time being a function of the number of processors allocated. Tasks may be started at any time, but once started, a task must not have its processor allocation altered or be preempted. Two objective functions are considered: minimizing the overall completion time for the tasks (make-span) and minimizing a weighted sum of the task completion times (weighted response). Both are considered subject to a constraint on the total number of processors available. Sub-optimal algorithms are developed for both of these NP-hard problems using Lagrangian relaxation, and their performances are analyzed through simulations.

4. J. Monte and K. Pattipati, "A Shelf-Based Relaxation Algorithm to Schedule Parallelizable Tasks," *IEEE Trans. Parallel & Distributed Systems*, (1996).

Abstract

This paper generalizes classical scheduling theory by removing the restriction that only a single processor can work on a given task at a particular time. Instead, it is assumed that each task can be allocated any number of identical processors from one to the maximum number available, and that a task's completion time is a non-increasing function of the number of processors allocated. Once started, a task must run to completion without altering the number of processors given to it. Furthermore, a task can start only when no task is currently executing. The objective considered is minimization of overall completion time (make-span) for the tasks subject to the constraint of a limited number of available processors. To approximately solve this problem, an algorithm based on Lagrangian relaxation is developed, and its performance is analyzed through extensive simulations. Approximate duality gaps range from 0 to about 60% on average and are strongly a function of problem type and size.

6 Interactions and Transitions

We successfully completed the proposed work together with AF Rome Laboratory people, in particular, R. Gassner, V. Vannicola, M. Linderman, and Gerald Bright. Several joint papers coauthored with R. Gassner were presented at the 1995 SPIE, 1996 SPIE, SCS 1996 HPC, and IEEE 1996 HPDC-5 conferences. A Ph.D. candidate, R. Popp, successfully completed a second summer at RL through the AFOSR Summer Research Program, working with Gassner and others on parallelization of assignment and tracking algorithms on high performance computers and their integration with the RL MIST testbed. We also transferred our assignment-based parallel tracking algorithms to the RL testbed, where Northrop Grumman is presently using the assignment algorithm transferred to RL in their Joint STARS (E-8) program.

7 New Discoveries, Inventions, or Patent Disclosures

- It was discovered that, contrary to conventional wisdom, the *interface* between state estimation and data association (the cost calculations for the optimization) is the main computational bottleneck, and this was the motivation behind the parallelizations.

8 Honors and Awards

- NASA Summer Research Fellowship awards (R. Popp — 1993,94,95)
- Visiting Research Scientist awards as part of AFOSR's Summer Research Program (R. Popp — 1994,95)

References

[1] Y. Bar-Shalom (editor), **Multitarget-Multisensor Tracking: Advanced Applications**, Artech House, 1990.

[2] Y. Bar-Shalom and X. R. Li, **Estimation and Tracking: Principles, Techniques, and Software**, Artech House, 1993.

[3] Y. Bar-Shalom and X. R. Li, **Multitarget-Multisensor Tracking: Principles and Techniques**, YBS Publishing, 1995.