Defence Research and Development Canada
Recherche et développement pour la défense Canada

DRDC | RDDC

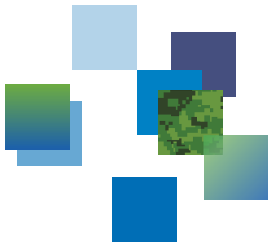

Identification of Non-Broadcasting Vessels

Basic Description for Non-Cooperatively Identifying Vessels with Combinatorial Statistics

D.E. Schaub
DRDC – Atlantic Research Centre

Defence Research and Development Canada

Scientific Report
DRDC-RDDC-2015-R286
December 2015

Canada

Identification of Non-Broadcasting Vessels

Basic Description for Non-Cooperatively Identifying Vessels with Combinatorial Statistics

D.E. Schaub
DRDC – Atlantic Research Centre

Defence Research and Development Canada
Scientific Report
DRDC-RDDC-2015-R286
December 2015

Abstract

Identification of targets has historically been conducted on an ad-hoc basis. Over the past two decades, the development of the recognized maritime picture—particularly in operations centres—has become heavily reliant on the continual intake of automatic identification system (AIS) messages. Arguably, situational awareness has improved considerably owing to the enormous increase in self-reports of vessels that would otherwise remain undetected or unidentified. While providing ready access to seemingly high-quality target identity information, AIS remains highly vulnerable to inadvertent or malicious disruption. Moreover, vessels posing threats are either unlikely to broadcast AIS messages at all or transmit false identification or location information. Under these circumstances, it becomes imperative to develop robust methods of identifying vessels using information besides self-reported identity. To this end, the present work qualitatively describes a combinatorial framework for non-cooperative identification, with the objective of motivating the development of a flexible system (both tactical and operational) that may be used to identify non-reporting (dark) vessels and confirm the stated identity of self-reporting vessels.

Significance for defence and security

This work describes a method of combining diverse information from multiple sensors to find the optimal vessel-identity-to-track assignments. The given approach may benefit the Royal Canadian Navy by enabling approximate identification of targets that would otherwise remain unidentified.

Résumé

Par le passé, l'identification des cibles était effectuée de façon ponctuelle. Au cours des deux dernières décennies, l'élaboration du tableau de la situation maritime—tout particulièrement dans les centres des opérations—dépendait largement de l'arrivée continue de messages du Système d'identification automatique (SIA). Sans doute, la connaissance de la situation s'est améliorée considérablement grâce à l'augmentation importante de l'autosignalement de navires qui autrement n'auraient pas été détectés ni identifiés. En fournissant un accès facile à des renseignements apparemment de grande valeur sur l'identification de cibles, le SIA demeure très vulnérable aux interruptions malveillantes ou causées par inadvertance. Par ailleurs, il est peu probable que les navires qui représentent des menaces transmettent des messages dans le SIA ou encore de fausses données sur leur identité ou leur emplacement. Dans ces circonstances, il devient impératif d'élaborer de solides méthodes d'identification des navires à l'aide de renseignements obtenus autrement que par autosignalement. À cette fin, le présent travail décrit qualitativement un cadre combinatoire relatif à l'identification non coopérative dans le but de favoriser la mise au point d'un système souple (sur les plans tactique et opérationnel) pouvant être utilisé pour déterminer l'identité des navires (mystérieux) qui ne l'ont pas signalée ou confirmer celle des navires qui l'ont fait.

Importance pour la défense et la sécurité

Le présent travail décrit une méthode pour combiner des renseignements provenant de divers capteurs, afin de déterminer les façons optimales de retracer les navires. La Marine royale canadienne pourrait profiter de l'approche énoncée, car celle-ci lui permettrait d'identifier de façon approximative des cibles qui autrement n'auraient pas pu l'être. est l'importance pour la défense et la sécurité.

Table of contents

List of figures

Acknowledgements

The author expresses gratitude to Anthony Isenor (Defence Research and Development – Atlantic Research Centre) for constructive comments and discussions during preparation of this report.

This page intentionally left blank.

1 Introduction

Identification of targets represents a fundamental cornerstone in situational awareness. Beyond discriminating between friend and foe in tactical scenarios, identification of vessels around the world—and particularly Canada's area of responsibility—is of significant concern to maritime operations centres. In contrast to the air domain, identification has historically been less comprehensive in maritime settings because of the vast number of ocean-going vessels and the limited availability of information. However, the last fifteen years has witnessed the recognized maritime picture (RMP) undergo a significant transformation brought about by the introduction of the automatic identification system (AIS) [1].

Originally developed as a collision avoidance system, the automatic identification system comprises shipborne transceivers that broadcast the vessel's identity and location (obtained through the global positioning system) and receive and display such broadcasts from other nearby vessels. Operations centres now experience large inflows of AIS data from a vast network of dedicated AIS receivers (coastal, airborne, and space (satellite) based). However, although AIS provides a seemingly exhaustive view of vessel traffic, its possesses several design weaknesses (see [2] for a partial survey) that undermine its fitness for surveillance applications. The three principal considerations are:

1. AIS depends on GPS, which is easily jammed, spoofed, or otherwise rendered inoperable. While countermeasures to GPS attacks have been developed, these are generally not deployed on civilian systems.

2. AIS is a *cooperative* sensing system. That is, its functioning requires the transceivers to be correctly configured, enabled, and maintained. It is trivial matter to disable an AIS responder or tamper with it so as to broadcast falsified position reports or identity. Geofeasibility checks—whereby an AIS message's reported location is compared to its estimated broadcast location—are presently under development [3] but offer no defence against identity falsification or the disabling of transponders. Also, many small vessels remain invisible to the AIS system, as operation of transponders are only mandated for vessels in excess of 300 gross tonnes.

3. It would require little effort to develop a radio broadcast system that could, in a certain geographic domain, overload the local AIS system with spurious messages that both displace legitimate broadcasts and saturate the system with 'ghost' vessels. In an information warfare context, the AIS system is extremely vulnerable to the injection of large quantities of malicious data that have the potential to disrupt surveillance activities that are predicated on the (implicit) veracity of AIS messages. It would not be infeasible to carry out such an attack over several domains simultaneously.

In view of the foregoing vulnerabilities and limitations inherent to AIS, and the fact that many vessels remain dark, it has become imperative to pursue robust and non-cooperative means of identifying targets. The related work of [4, 5] has culminated in the development of the SAR-AIS Association Systems (SAAS), which associates received AIS reports with space-based synthetic aperture radar (SAR) imagery obtained from the RADARSAT-2 constellation. In particular, this system allows the partitioning of detected vessels into to those exhibiting normal behaviour (broadcasting AIS messages bearing valid data) and those displaying suspicious characteristics (such as broadcasting dubious AIS messages or failing to transmit AIS messages at all), allowing resources to be directed to identifying the latter targets by other means.

Considering the difficulty—and frequently expense—in tasking assets to identify detected but non-broadcasting targets (such as those classified as dark by SAAS), it is strongly desirable to attempt to estimate the possible identities of detected but unidentified targets using the totality of available information. Developing such a capability requires the integration of information from various sensor / intelligence sources along with logical and mathematical analysis that yields an identity estimate for each detected target. To perform this task in a manner that is theoretically justified and rigorous is challenging, as identity can only be inferred indirectly from various observed target characteristics, and there are generally complex interrelationships between the identity assignments of different targets. Central to inference over multiple detected targets and possible identities lies a combinatorial mathematical structure that has remained (possibly due to its mathematical or computational complexity) unexplored in existing works, which have invariably considered identification as a collection of (suboptimal) single-target classification problems or have used identity information to assist track maintenance. This report will demonstrate that such a principled and mathematically-optimal combinatorial framework for non-cooperative identification is theoretically possible and should be practically realizable. Fundamentally, the described work seeks to establish a system whereby all dark targets are simultaneously assigned possible identities, representing a considerable departure from the conventional, piecemeal treatment of identity found in earlier efforts such as CASE-ATTI (Concept Analysis Simulation Environment for Automatic Target Tracking and Identification) [6, 7]. Ultimately, this work aims to ensure that decision makers can place their full confidence in the estimated identity of a dark target, and that the estimate is the best possible given the totality of available information.

The remainder of this report is organized as follows. In section 2, dark-target identification is first formulated as a problem in a branch of mathematics named combinatorics and subsequently contrasted to conventional target classification. It is shown that the essential complexity inherent to multitarget identification is in fact a standard, well-studied combinatorial sum. Section 3 provides an example of how the mathematical framework can be applied to resolve the identities of dark targets. In section 4, the necessary steps for implementing an automatic non-cooperative identi-

fication system in operations centres are proposed. Finally, section 5 briefly outlines future research in applying combinatorics to situational awareness and identification, and developing principled approximation methods that allow the implementation of real-time target identification. This report is intended for a general audience, and therefore descriptions of the mathematics are given only qualitative terms. Detailed mathematical derivations may be found in [8]. This work was carried out in the context of support to operations centres at DRDC, Atlantic Research Centre under the Maritime Information Warfare Project (01da).

2 Combinatorial Identification

2.1 Combinatorial Matchings

A combinatorial *matching* is a mapping that uniquely assigns each member of one group to a member in another. For example, assigning dance partners with an equal number of ladies and gentlemen may be formulated as a matching problem (Fig. 1). Obviously, more than one matching is possible when the membership of each group equals or exceeds two. Should there exist a means of computing a compatibility metric (a number that assigns a compatibility to each prospective couple—e.g. how well they are expected to dance together), then a matching can be found that maximizes aggregate compatibility over all of the resulting couples. Alternatively, matchings may be chosen randomly with a probability proportional to their compatibility (i.e. compatible matches are sampled more frequently than their incompatible counterparts). The principle difficulty with reasoning over combinatorial matches follows from the factorially large problem space. This is readily illustrated by considering a mere 100 ladies and 100 gentlemen; this arrangement allows for $100 \times 99 \times 98 \times \ldots \times 2 \times 1$ (more than 9×10^{157}) possible matchings, precluding the use of computational techniques that rely on exhaustive enumeration.

Combinatorial matchings also naturally arise in the context of target identification, where the task involves matching of vessels (detected and tracked by a suite of sensors) to a collection of known identities, such as those from a database. In this case, the 'compatibility' metric is the degree to which each target conforms to the known properties of a given identity. For example, the compatibility between a rapidly-manoeuvring target and a speedboat is high, while that of the same target to a sluggish bulk carrier is comparatively low. Compatibility comparisons similarly occur in the related field of target classification that seeks to assign observed targets to members of a predefined set of target classes (e.g. classifying an air track as a civilian or military aircraft). Taken to the extreme, each identity could be assigned its own class. However, while estimation of a given target's class is carried out using data solely from its own track, identification—by virtue of the fact that a given identity can only be at single track at a time—must be performed using all available data. Thus,

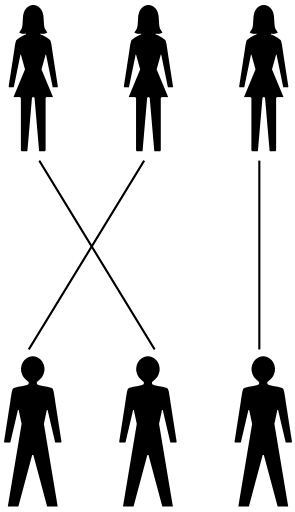

Figure 1: *The essential problem of combinatorial matching. There are* $3 \times 2 \times 1 = 6$ *possible arrangements of dance partners. One such matching is shown by the connecting lines.*

the characteristic feature of multitarget identification is that *target identification at one track is highly interdependent with target identification at all other tracks.*

The practical differences between identification and classification (specifically, the comparable case where each identity is given its own class) may be illustrated by considering an unknown track in the South Atlantic Ocean using the prior knowledge that the RMS *Queen Mary 2* is located at a track in the Indian Ocean. An optimal identification algorithm exploits the non-local information (that the RMS *Queen Mary 2* cannot be in the South Atlantic) and would report a zero probability for such an occurrence, i.e., $p(\text{QM2@SA}) = 0$, accordingly. On the other hand, conventional classification, using only information from the South-Atlantic track, disregards the non-local knowledge of the RMS *Queen Mary 2*'s whereabouts and would likely report a non-zero probability for the ship's presence (or the presence of its class) in the South Atlantic. In fact, $p(\text{QM2@SA})$ may be quite substantial should observations at the South-Atlantic track be vague or correspond to a vessel bearing features similar to those of the RMS *Queen Mary 2*.

As illustrated in Fig. 2, joint identification may be performed in two distinct forms, by computing either

1. Marginal identity-to-track probabilities (Fig. 2a). In this case (which is closely related to the aforementioned random sampling of matchings), each possible track / identity pair is assigned a probability. These statistics are suited to answer the questions, *'Which identities might track X be?'* and the converse,

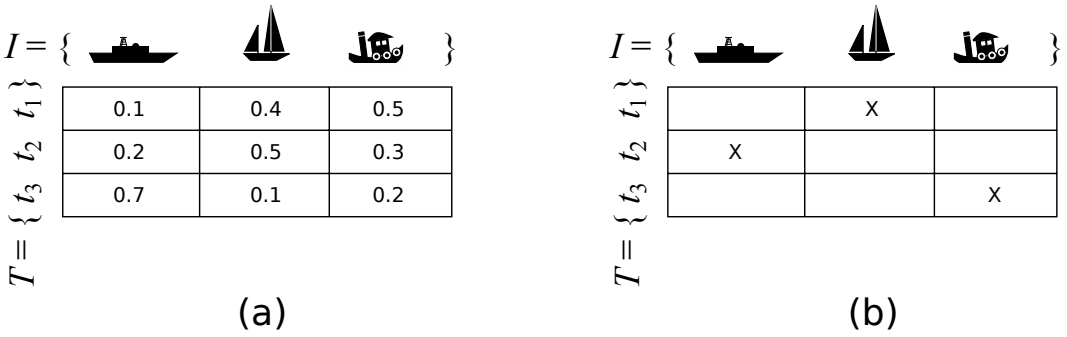

Figure 2: *Example of (a) soft and (b) hard matching. Matching can either be performed by (a) assigning marginal track (T)-to-identity (I) probabilities or (b) by finding a single global best match. In the case of soft matching the rows and columns must sum to one, reflecting, respectively, that at each track there must exist some identity and that each identity must reside at some track.*

'At which tracks might identity Y be located?'; or

2. Global maximum likelihood permutation (Fig. 2b). This statistic associates each vessel to a single identity and vice-versa (i.e., this statistic is the single, most probable matching between identities and observed targets), answering the question, *'What is the most likely spatial arrangement of identities?'*

In most practical scenarios, the former is usually of greater interest. For instance, for an unknown target that is approaching a sensitive asset, computing the identity probabilities would likely be preferable to finding most probable global matching. Furthermore, as illustrated in Fig. 2 the single most likely permutation (Fig. 2b) may actually correspond to individual track-to-identity assignments (Fig. 2a) that are themselves improbable, or at least not the most probable. Consequently, the remainder of this report concerns the computation of marginal identity-to-track probabilities.

2.2 Computation of Marginal Identity-to-Track Probabilities

Calculation of the marginal identity-to-track probabilities may be decomposed into two principal steps. The first involves performing Bayesian joint target tracking and classification in the (conventional) manner of [9], where, as described in section 2, each identity is assigned its own class. Provided that targets are reasonably well separated, the tracks may be processed individually. The output of this step comprises raw track-to-identity weights, shown in Fig. 4.

The second step involves identity deconfliction, wherein conflicting identity information between tracks (for example, where the previous instantiations of the classification algorithm assigned, with high confidence, a particular identity to more than

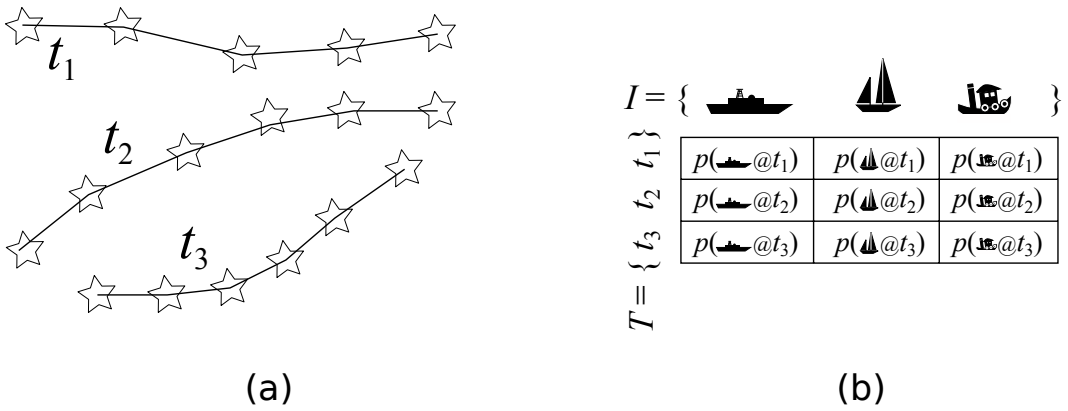

Figure 3: *Marginal Track-to-Identity Probabilities. Given a set of tracks (a), it is desired to find the assignment probabilities between tracks and identities (b).*

one track) is resolved. This requires computing the *matrix permanent* [10], a combinatorial sum that, in the present context, rebalances the table of Fig. 4b into a form that is logically consistent. This form requires that each row and column in Fig. 4b sum to one, so as to reflect the fact that, with a probability of 100%, each identity is located at some track and vice-versa (each track must correspond to some identity). It can be shown that the matrix permanent operation is both necessary and sufficient to ensure the Bayesian property of the identification system; therefore, the posterior probabilities of Fig. 4b are the best that are mathematically attainable through Bayesian statistical reasoning.

3 Illustration

A brief simulated example is provided to qualitatively demonstrate the utility of combinatorial identification. Five targets, each with distinct kinematic characteristics, are simulated to move in two dimensions (which may be taken, for example, as the surface of the ocean). The track-to-identity probabilities are shown as a function of time in Fig. 5. As expected, each track tends to converge on a particular identity as more observations are processed. However, even at the last timestep, no track is able to perfectly resolve the identity of its target, reflecting the residual uncertainty inherent to any system whose observations and assumptions about target motion are themselves uncertain. It should be noted that owing to the system's mathematical optimality, the probabilities at any given time are the most accurate than can be computed. A similar but more thorough and quantitative example (along with detailed simulation parameters) may be found in [8].

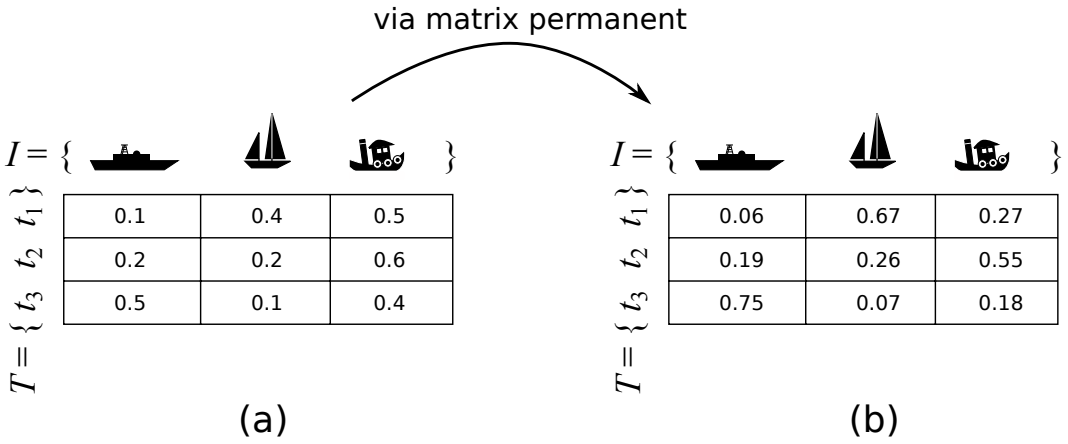

Figure 4: *The matrix permanent operation for reconciling conflicting information from different tracks. The track-to-identity probabilities generated by the local trackers (a) are refined by the matrix permanent into the globally optimal track-identity probabilities (b) using the totality of available information. For example, information from tracks 1 and 2 (specifically that they are very likely to be the sailboat and tugboat) leads to an increase of the probability that track 3 is the frigate. Note that only row sums equate to one in (a), while this also holds for column sums in (b).*

4 Identification System Requirements

Implementation of a complete identification system requires the development of several complementary components, listed here:

1. **Comprehensive Database of Ocean-Going Vessels.** *Required to perform conventional joint tracking and classification.*
   This database should contain as many vessels and their features as possible, including:
   (a) Unique, unchanging identifier (such as IMO) or in its place a class of ship (e.g. small fishing vessel)
   (b) Ship colors
   (c) Dimensions (beam, length, draft, etc.)
   (d) Radar cross section
   (e) Installed radar and communications systems
   (f) Signatures (infrared, electromagnetic acoustic, etc.) and management capabilities for military vessels
   (g) Engine parameters (maximum speed, acceleration, etc.)
   (h) Manoeuvrability

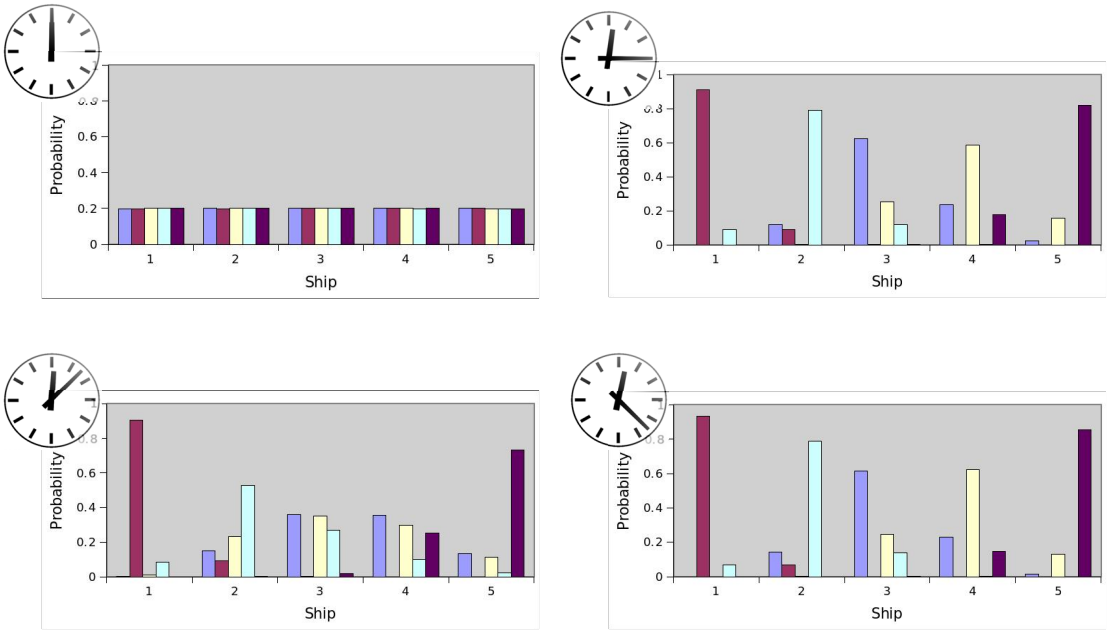

Figure 5: *Simulated identification example for five ships with distinct kinematic profiles. Observations (such as radar returns) are collected at five tracks. At time zero the absence of any observations results in complete identity-track ambiguity, as is evidenced by every identity being equally probable at each track. With the progression of time, each track tends to become associated with a only single identity, though some uncertainty remains.*

(i) Fuel capacity (or maximum distance that can be travelled between port visits)

(j) Ice class or ice capability

(k) Any other vessel characteristic that might be observed, even if indirectly

Clearly, it is unlikely that all of the above can be fully determined for each vessel, and it is therefore expected that the list of features will be incomplete. Frequent update of the database is imperative, and each feature entry should itself be a list comprising pairs of values and date ranges over which those values are valid (e.g. the color attribute for a given vessel would indicate that it was painted red between 01/01/1980 to 15/03/1991, and, since 16/03/1991 has been painted blue). Although it may appear that resources should be preferentially directed to acquiring further details on those vessels of greatest interest, such a strategy may not necessarily be optimal, as the identification of vessels of interest is deeply intertwined with being able to identify vessels of non-

interest[1]. Furthermore, considerable effort and specialized knowledge is likely required to convert physical attributes into observable kinematic characteristics (e.g. converting engine parameters and hull dimensions into manoeuvrability characteristics).

2. **Comprehensive Database of Sensors and their Error (Uncertainty) Profiles.** *Required to perform conventional joint tracking and classification.*
   Bayesian inference depends crucially on the correct characterization of error in reported data. For example, a report stating that '*the ship is at* 44.5° N*,* 62.1° W' is inadequate, as it cannot be fused with second report stating '*the ship is at* 43.5° N*,* 63.1° W'. Obviously these reports are different. Do we take an average—or a weighted average? What is the uncertainty in the fused result? Combinatorial estimation, while complex, is still just a form of Bayesian inference; uncertainty must be specified in the input data, and in turn, the system will provide uncertainty of the fused result.

3. **Comprehensive Database of Bathymetry, Coastal Profiles, Sea Ice, and Meteorology.** *Required to perform conventional joint tracking and classification.*
   This database assists the identification system in determining which vessels are able to traverse a given region of the ocean. In particular:
   (a) A region's local bathymetry data implies the maximum vessel draft
   (b) Sea ice information determines the minimum ice capabilities as well as maximum speed and manoeuvrability
   (c) Meteorology (sea state) may be used as a proxy to minimum seaworthiness

   While bathymetry and coastal data are expected to be static—the former requires continual (minor and predictable) adjustment with shifting tides—, sea ice and meteorological data is dynamic and must be frequently updated. Access speed of this database is crucial, as it is heavily queried by the core identification system.

4. **Tracking System that Retains Contacts and their Sensor Tags.** *Required to perform conventional joint tracking and classification.*
   In accordance with the discussion of item 2, the tracking system is required to retain all contacts associated with a given track. Moreover, the sensor tag, if available, should remain affixed to each contact, so that the measurement error can be readily obtained.

[1] Take, for example, a world where there exist only two ships and only one is of interest; by elementary logic, positively identifying one target as the vessel of non-interest automatically identifies the other target as being the vessel of interest. Where there are thousands of targets, the same essential logic applies, though at a scale that is much more complex.

5. **High-Performance Computer System.** *Required to compute (approximations to) the matrix permanent.* As discussed in section 2, the general case of the combinatorial deconfliction step is computationally intensive, and therefore, depending on accuracy requirements, substantial computational resources may be required. Fortunately, most algorithms that approximate the matrix permanent are easily and efficiently parallelized and are expected to scale well on cluster / cloud architectures.

4.1 Challenges

Implementation of the aforementioned system is faced with both technical and logistical challenges that arise from the volume (rate) of input data and its security classification(s), respectively. Considerations surrounding these issues are given below.

4.1.1 Classification Issues

As the identification system is expected to operate on classified data, special steps must be performed over the course of its implementation and operation. While the interconnection (networking) of computing systems of different classifications is possible—and indeed is not uncommon—, substantial investment of resources may be necessary to meet applicable security directives. More challenging, however, is the fact that unredacted information-fusion products usually inherit the highest classification of their constituent inputs. Thus, it is anticipated that the matrix permanent algorithm would necessarily operate at the system's highest classification, significantly limiting computational options for its implementation. In particular, commercial off-the-shelf (unclassified) cloud computing—with its desirable economies of scale—would be precluded, thereby increasing the overall system cost. One possible future solution would employ homomorphic encryption [11], which allows computation to be carried out on encrypted data such that the processing system remains agnostic to the data itself. Unfortunately, this approach is burdened with high computational overhead[2]. Moreover, homeomorphic encryption is not presently sanctioned for use in this manner.

4.1.2 Computational Issues

In wide-area surveillance, the identification system will likely encounter large numbers of dark vessels. The fastest known matrix permanent algorithm [12] (based on a Markov chain Monte Carlo formulation), is likely able to carry out near real-time

[2] At the time of this writing, research in this area is vigorous, and the computational overheads continue to be reduced significantly.

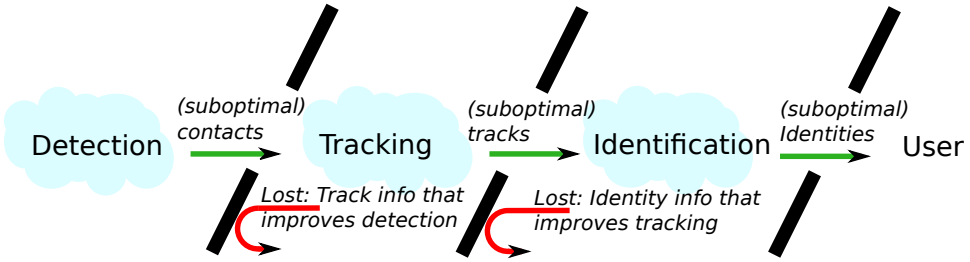

Figure 6: *Detection, tracking, and identification in the current setting. Calculations in each fusion layer are independent of those in adjacent layers. As shown, this approach is suboptimal.*

probabilistic identification of ∼100 targets on a typical computer cluster. This approach is particularly well suited for surveillance of the Arctic, where vessel traffic is relatively light. A faster method [13] (a perfect random sampler), is likely suitable for ∼1000 targets, and may be used when there are only vague feature observations at the individual tracks. Finally, there also exists a simple yet suboptimal method (based on primitive matrix scaling) which can process 10,000 targets or more [14]. It is likely that the various techniques can be combined in a manner that yields a satisfactory tradeoff between accuracy and computability. In particular, should the input to the matrix permanent step be decomposable into low-information (vague) and high-information (specific) components, it is conceivable that [12] and [13] could be used together. Whether such propitious conditions arise in practice remains an open question, and consequently, this area is ripe for further study using empirical data (i.e. the actual data produced by currently deployed sensors).

5 Future Directions

Where tracks have been previously formed by a tracking system, the system of identification described in the previous sections can be shown to be statistically optimal. On the other hand, should the unprocessed (raw) contacts be available, both tracking and identification can be improved by subsuming these tasks under a single mathematical framework. This may be seen by considering that, in conventional systems, the fusion layers of detection, tracking, and identification only propagate information in one direction (Fig. 6). In actual fact, the mathematics of the different fusion layers are deeply intertwined and would ideally be regarded as different facets to the same underlying estimation problem. At the present time, quasi-exact joint multitarget tracking and identification remains theoretical, as approximation schemes with boundable error (such as those that are available for approximating the matrix permanent) do not yet exist. Investigating extensions to the Markov chain Monte Carlo methods of [15, 16, 12] to include the lower fusion layers—if possible—may prove

fruitful.

6 Conclusion

The past two decades have seen maritime situational awareness improve dramatically with the introduction of AIS, which provides position and identification reports from vessels that would otherwise go undetected or unidentified. However, while AIS notionally provides vast quantities of high-quality target identity information, it is susceptible to inadvertent or malicious disruption. For these reasons, the development of a system of non-cooperative identification is imperative. To this end, the present work described a combinatorial framework for identifying non-reporting vessels (and confirming the identities of reporting vessels) using information obtained from non-cooperative sensors. In contrast to AIS, such an system can also be used to identify dark (non-reporting) targets. Implementations of the identification framework in various operational and tactical settings are forthcoming.

References

[1] Harre, I. (2000), AIS Adding New Quality to VTS Systems, *The Journal of Navigation*, 53(03), 527–539.

[2] Balduzzi, M., Alessandro, P., and Wilhoit, K. (2014), A Security Evaluation of AIS Automated Identification System, In *Proceedings of the 30th Annual Computer Secuirty Applications Conference*, pp. 436–445, ACM.

[3] Papi, F., Tarchi, D., Vespe, M., Oliveri, F., Borghese, F., Aulicino, G., and Vollero, A. (2014), Radiolocation and Tracking of Automatic Identification System Signals for Maritime Situational Awareness, *IET Radar, Sonar & Navigation*, 9(5), 568–580.

[4] Fitzpatrick, S., Gong, S., Salciccioli, M., Chen, W., Biron, K., and Vachon, P. (2014), SAR-AIS Association System Upgrade and Support: Technical Report to Close Out GEOINT Task 17, Contract W7714-091140/001/SV, (DRDC-RDDC-2014-C117) DRDC – Ottawa Research Centre.

[5] Vachon, P., Kabatoff, L., and Quinn, L. (2015), Operational Ship Detection in Canada using Radarsat, (DRDC-RDDC-2015-N091) DRDC – Ottawa Research Centre.

[6] Roy, J., Bosse, E., and Dion, D. (1995), CASE ATTI: (Concept Analysis and Simulation Environment for Automatic Target Tracking and Identification) An Algorithm-Level Testbed for Multi-Sensor Data Fusion, (DREV-9411) Defence Research Establishment Valcartier.

[7] Benaskeur, A., Yuen, S., and Triki, Z. (2003), Perfomance Evaluation within CASE_ATTI of MHT and JVC Association Algorithms for COMDAT TD, (DRDC-Valcartier TR-2003-287) DRDC Valcartier.

[8] Schaub, D. (2015), Joint Identification of Multiple Tracked Targets, *Submitted to the Journal of Advances in Information Fusion.*

[9] Ristic, B., Gordon, N., and Bessell, A. (2004), On Target Classification using Kinematic Data, *Information Fusion*, 5(1), 15–21.

[10] Brualdi, R. A. (2006), Combinatorial Matrix Classes, Vol. 13, Cambridge, United Kingdom: Cambridge University Press.

[11] Yi, X., Paulet, R., and Bertino, E. (2014), Homomorphic Encryption and Applications, Springer.

[12] Bezáková, I., Štefankovic, D., Vazirani, V. V., and Vigoda, E. (2008), Accelerating Simulated Annealing for the Permanent and Combinatorial Counting Problems, *SIAM Journal on Computing*, 37(5), 1429–1454.

[13] Huber, M. and Law, J. (2008), Fast Approximation of the Permanent for Very Dense Problems, In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 681–689, Society for Industrial and Applied Mathematics.

[14] Sinkhorn, R. (1964), A Relationship Between Arbitrary Positive Matrices and Doubly Stochastic Matrices, *The Annals of Mathematical Statistics*, 35(2), 876–879.

[15] Jerrum, M., Sinclair, A., and Vigoda, E. (2001), A Polynomial-Time Approximation Algorithm for the Permanent of a Matrix with Non-Negative Entries, In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, pp. 712–721, ACM.

[16] Jerrum, M., Sinclair, A., and Vigoda, E. (2004), A Polynomial-Time Approximation Algorithm for the Permanent of a Matrix with Nonnegative Entries, *Journal of the ACM (JACM)*, 51(4), 671–697.

DOCUMENT CONTROL DATA

(Security markings for the title, abstract and indexing annotation must be entered when the document is Classified or Protected.)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

DRDC – Atlantic Research Centre
PO Box 1012, Dartmouth NS B2Y 3Z7, Canada

2a. SECURITY MARKING (Overall security marking of the document, including supplemental markings if applicable.)

UNCLASSIFIED

2b. CONTROLLED GOODS

(NON-CONTROLLED GOODS)
DMC A
REVIEW: GCEC DECEMBER 2012

3. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Identification of Non-Broadcasting Vessels

4. AUTHORS (Last name, followed by initials – ranks, titles, etc. not to be used.)

Schaub, D. E.

5. DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF PAGES (Total containing information. Include Annexes, Appendices, etc.)	6b. NO. OF REFS (Total cited in document.)
December 2015	24	16

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Scientific Report

8. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

DRDC – Atlantic Research Centre
PO Box 1012, Dartmouth NS B2Y 3Z7, Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
01da	

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DRDC-RDDC-2015-R286	

11. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected.)

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

Identification of targets has historically been conducted on an ad-hoc basis. Over the past two decades, the development of the recognized maritime picture—particularly in operations centres—has become heavily reliant on the continual intake of automatic identification system (AIS) messages. Arguably, situational awareness has improved considerably owing to the enormous increase in self-reports of vessels that would otherwise remain undetected or unidentified. While providing ready access to seemingly high-quality target identity information, AIS remains highly vulnerable to inadvertent or malicious disruption. Moreover, vessels posing threats are either unlikely to broadcast AIS messages at all or transmit false identification or location information. Under these circumstances, it becomes imperative to develop robust methods of identifying vessels using information besides self-reported identity. To this end, the present work qualitatively describes a combinatorial framework for non-cooperative identification, with the objective of motivating the development of a flexible system (both tactical and operational) that may be used to identify non-reporting (dark) vessels and confirm the stated identity of self-reporting vessels.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

DRDC Scientific Report; Surveillance; Identification; Tracking; Bayesian; Combinatorial; Permanent; Markov Chain Monte Carlo; MCMC

DRDC | RDDC

SCIENCE, TECHNOLOGY AND KNOWLEDGE FOR CANADA'S DEFENCE AND SECURITY | SCIENCE, TECHNOLOGIE ET SAVOIR POUR LA DÉFENSE ET LA SÉCURITÉ DU CANADA

www.drdc-rddc.gc.ca