1. Introduction

Cyber Ontology (CybOnt) is an AI system that produces indication and warnings of cyber attack behavior hypotheses to support cyber situational understanding (SU). CybOnt leverages established data fusion architectures and algorithms to generate the hypotheses, each with a mathematically principled likelihood ratio. The likelihood ratio is critical for SU so hypotheses can be filtered, thresholded, sorted, and prioritized on commander's Common Operational Picture (COP). CybOnt is architected using the Joint Directors of Laboratories (JDL) fusion levels documented and discussed in hundreds of technical publications e.g., [1, 2, 3, 4, 5].

This white paper provides an overview of the CybOnt fusion algorithms, screen captures of the User Interface (UI), a description of the demonstrations that have been conducted at APG to date, and a brief plan ahead.

Table 1. Fusion Levels Applied to Cyber

Level	Applied to Cyber
0	extracts features, computes features, and receives observations from cyber sensors such as Intrusion Detection Systems (IDS), Intrusion Prevention Systems (IPS), network devices, Host Based Security System (HBSS), and Continuous Diagnostics and Monitoring (CDM).
1	develops hypotheses and associated likelihood ratios for Attack Pattern Steps, Attack Patterns, Attacker Types, and Objective Types. Correlates hypotheses and merges beliefs with other level 1 fusion processes.
2	develops hypotheses and associated likelihood ratios for Spatio-temporal Group Associations, Mission Attack Associations, and Critical Capability Attack Associations. Reinforces or deconflicts with other level 2 fusion nodes.
3	develops hypotheses and associated likelihood ratios for Objective Types to Actual Vulnerabilities, Tactics Techniques, and Procedures (TTP) Correlation, and Attacker Types to Candidate Attackers. Reinforces or deconflicts with other level 3 fusion nodes.

2. CybOnt Data Fusion Algorithms Overview

Current cyber threat detection systems are rule based. A probabilistic cyber ontology is a better real-world model because understanding of the real-world is uncertain, e.g., due to measurement errors and gaps, attacker obscuration and deception. In addition much work over the past decade has shown benefits to conducting data fusion over ontologically structured data [e.g., 6, 7, 8]. Data fusion involves multiple distributed nodes and distributed data fusion algorithms require an unambiguous ontology so that algorithms and operators can work independently but in coordination via the ontology, i.e., interoperably.

A probabilistic system allows operators and analysts to adjust the probability-of-false-alarm (pFA) to probability-of-detection (pD) ratio to the level that supports their operational need, e.g., for timeliness, operator workload, and completeness.

2.1 Notational Conventions

A key to successfully developing data fusion and AI algorithms, particularly ones that work with a formal ontology, is compact, well understood, and consistent notional conventions. In CybOnt the following notational conventions were developed and will be used in this white paper:

1. Because CybOnt is ontology-based fusion, it is essential to distinguish Type (T-Box) from Individual (A-Box). Bold font denotes Types while non-bold denotes spatio-temporal extents,

also known as Individuals. For example, $\mathrm{OF}_u \in \mathbf{OF}_v$ says an Individual is an typeInstance of element of a Type.

2. Mathematical functions are italicized while variables are not. For example, a probability value from the knowledge base is ${}^{KB}\mathbf{P}\left(\mathbf{OF_u}\middle|\mathbf{APS_j}\right)$ while one that is computed is $P\left(\mathrm{OF}_u\middle|\mathrm{AP}_i\right)$.
3. Knowledge Base (KB) data is preceded with a KB superscript, e.g., ${}^{KB}\mathbf{P}\left(\mathbf{OF_u}\middle|\mathbf{APS_j}\right)$.
4. Lower case Greek letters are used for ephemeral indexes, superscripts, and subscripts.

2.2 Mathematical Fundamentals

Fundamentally, inference employs some type of inversion to infer causes from effects in the following manner:

$$P\left(\{\text{cause}\}_i\middle|\{\text{effects}\}\right)=\frac{P\left(\{\text{effects}\}\middle|\{\text{cause}\}_i\right)\mathrm{P}_0\left(\{\text{cause}\}_i\right)}{\sum_{\text{j=all possible cause sets}} P\left(\{\text{effects}\}\middle|\{\text{cause}\}_j\right)\mathrm{P}_0\left(\{\text{cause}\}_j\right)}$$

where:

$\{\text{effects}\}$ is a set of effects for which the cause(s) are sought

$\{\text{cause}\}_i$ is a set of possible causes of $\{\text{effects}\}$

$\mathrm{P}_0\left(\{\text{cause}\}_i\right)$ is the "prior" probabity

(1)

See, for example, [9]. There are some problems with this intuitively obvious formulation. Most importantly, P_0 says a probability can be known without any evidence, an arguable proposition. If accepted, a common way to handle lack of prior evidence - ignorance - is to assign equal probability to all events (flat priors). However, it is easy to prove this results in non-equal odds [10] - a contradiction since non-equal odds implies knowledge. Also, pragmatically the denominator, which is equivalent to $\mathrm{P}_0\left(\{\text{effects}\}\right)$, may be uncomputable or unknowable. CybOnt's solution is the fiducial likelihood ratio:

$$\Lambda\left(\{\text{cause}\}_i\middle|\{\text{effects}\}\right)\triangleq\frac{L\left(\{\text{cause}\}_i\middle|\{\text{effects}\}\right)}{L\left(\neg\{\text{cause}\}_i\middle|\{\text{effects}\}\right)}\triangleq\frac{P\left(\{\text{effects}\}\middle|\{\text{cause}\}_i\right)}{P\left(\{\text{effects}\}\middle|\neg\{\text{cause}\}_i\right)} \qquad (2)$$

The likelihood ratio is a sufficient statistic and is invariant under transforms. Even better, for operators and analysts, Λ is intuitively easy to understand: it says how much more likely the positive hypothesis is compared to the null hypothesis. For example, a Λ = 2 says the evidence is twice as strong for the positive as null hypothesis.

2.3 Level 0 Fusion – Probabilitic Auto-Adaptive Sensor Interfaces

Level 0 fusion is performed by COTS and GOTS hardware and software intrusion detection, network monitoring, endpoint device monitoring, virus and malware detectors and others – generally called sensors -- for observations and feature extraction. The range of types of outputs from sensors are called sensor events $\mathbf{SE}_v^t$ where the superscript t is for the type of sensors and v is an index into the types of events that sensor can generate. These need to be mapped to the CybOnt Observations and Features $\mathbf{OF}_u$ ontology where the index u is for each of the **OF** types.

Because the types of sensors and their message formats are many across sites and time as new technologies come along and the number of event types are in the thousands to tens of thousands, manual maping of the $\mathbf{SE}_v^t$ to $\mathbf{OF}_u$ is impractical and imprecise. To solve this problem, CybOnt computes adaptation parameters for each $\mathbf{SE}_v^t$ that probabilistically relate the $\mathbf{SE}_v^t$ to the ontology $\mathbf{OF}_u$'s. The current CybOnt version uses metadata describing each $\mathbf{SE}_v^t$ and $\mathbf{OF}_u$ to compute:

$$\Lambda\left(\mathbf{OF}_u \left| \mathbf{SE}_v^t \right.\right) = \frac{L\left(\mathbf{OF}_u \left| \mathbf{SE}_v^t \right.\right)}{L\left(\neg\mathbf{OF}_u \left| \mathbf{SE}_v^t \right.\right)} = \frac{smr\left(\mathrm{H}_1\right)}{smr\left(\mathrm{H}_0\right)} \qquad (3)$$

where the semantic mass ratios (smr) are computed using a cyber term frequency database and the analog of probability mass. These adaptation parameters are computed offline and maintained in CybOnt's knowledge base.

2.4 Level 1 Fusion – Λ's for Attack Pattern Steps

CybOnt's level 1 fusion uses a KB derived from the Common Attack Patterns Enumeration and Classification (CAPEC™) database [11]. CAPEC currently has 523 attack patterns with three phases – Explore, Experiment, and Exploit – and steps within the phases. Level 1 fusion association takes Observations and Features (OF) and infers the likelihood ratio of hypotheses Attack Pattern Step (APS) of type $\mathbf{APS}_j$ by a source asset towards a destination asset, an algorithm called OF2APS. It is analogous to track data association in kinetic data fusion in which contact reports are associated-to or used to form tracks. As in track data association, the OF may come from one sensor source (time series association) or multiple sensors (multi-sensor and time series association). An early but representative reference for this type of data fusion is [12].

Depending on the APS process model, the Λ depends not just on current observations and features, but also past ones. If there is a prior APS hypotheses for the destination in the fusion file, CybOnt will update the prior hypothesis maintained, either providing more or less support. (If the Λ in the fusion file falls below a Drop threshold, it is removed from the fusion file.) The recursive updating is patterned after the Kalman filter, i.e., a Maximum Likelihood Estimation (MLE) along the lines of [13, 14, 15]. Unlike some forms of the recursive Bayesian belief updater, the Kalman formulation for MLE overcomes initial estimate values, often quickly if the gain is high due to low measurement error and large process (or plant) noise. The measurement update is then fed into an α ß filter – a simplified Kalman - to perform the state estimate.

2.5 Level 2 Fusion – Associating Objects and Events into Patterns

CybOnt's architecture includes the following types of Level 2 fusion:

a. Attack Pattern Step to Attack Pattern (APS2AP) Stitching. This function threads APS to AP using the CAPEC ontology with CybOnt state transition probabilities and transition time statistics added. The APS hypotheses can be from any sources and do not have to be the some ones from step to step so coordinated attacks such as Distributed Denial of Service (DDoS) hypotheses can be formulated.

b. Group Association. Analogous to group tracking in kinetic data fusion, it is source assets involved in the same Attack Pattern within the Attack Pattern's likely spatio-temporal extent.
c. Mission Association. Mission association is similar to group association except that the members are not necessarily spatially or temporally correlated but all members are mission aligned, that is, their destinations belong to a mission group.
d. Tactics, Techniques, and Procedures (TTP) Stitching. TTP association is similar to APS2AP stitching but at a higher and multi-source level, stitching together larger and multi-warfare steps (e.g., kinetic, cyber, EW) into TTP models.

APS2AP stitching threads APS to AP. In the ontology, an Attack Pattern is just a set of temporal parts (steps), represented in the ontology as TemporalWholePartType (TWPT), plus an arrangement in before-after patterns, represented in the ontology as BeforeAfterType (BAT). An AP may be conducted with optional or non-detected APS, as shown in the example of a three step attack in **Figure 1**.

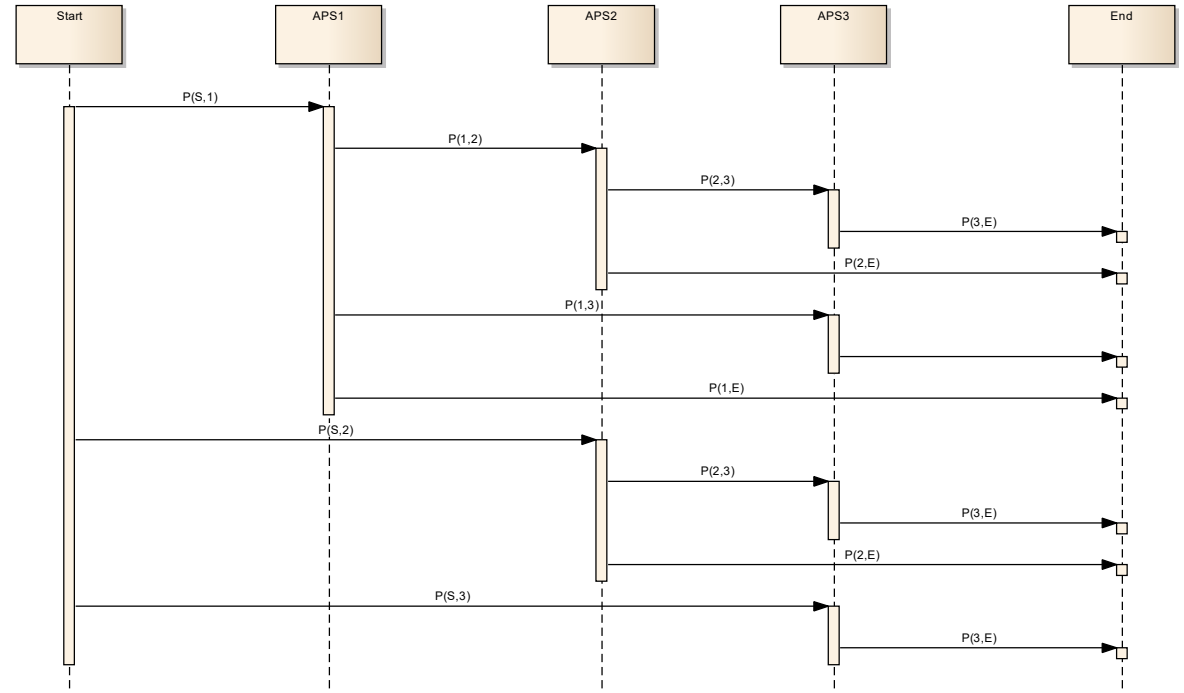

Figure 1. APS Sequences Notional Example

The likelihood ratio value at time t for the hypothesis of Attack Pattern $\mathbf{AP}_i$ against Destination Asset $A_d(x)$ by a set of Source Assets $\{A_s(y)\}$ is:

$$\Lambda\left(\mathrm{AP_i}\left(\mathrm{A_d},\{\mathrm{A_s}\},\mathrm{t}\right) \;\middle|\; \left\{\overrightarrow{\mathrm{SE_v^t}}\left(\mathrm{A_d},\{\mathrm{A_s}\},\tau\right)\right\}_{\tau=0}^{\mathrm{t}} \right)$$

where:

$\mathrm{AP_i}\left(\mathrm{A_d},\{\mathrm{A_s}\},\mathrm{t}\right)$ is the hypothesis for Attack Pattern type $\mathbf{AP_i}$ against destination asset $\mathrm{A_d}$ by a set $\{\mathrm{A_s}\}$ of source asset(s) *(4)*

$\left\{\overrightarrow{\mathrm{SE_v^t}}\left(\mathrm{A_d},\{\mathrm{A_s}\},\tau\right)\right\}_{\tau=0}^{\mathrm{t}}$ are Sensor Events relevant to $\mathrm{AP_i}$, meaning

$$\exists\Lambda\left(\mathrm{APS_j}\left(\mathrm{A_d},\mathrm{A_{s_i}},\mathrm{t}\right) \middle| \left\{\overrightarrow{\mathrm{SE_v^t}}\left(\mathrm{A_d},\mathrm{A_{s_i}},\tau\right)\right\}_{\tau=0}^{\mathrm{t}} \right) \neq 0 \text{ for any } \mathbf{APS_j} \in \mathbf{AP_i} \text{ and } \mathrm{A_{s_i}} \in \{\mathrm{A_s}\}$$

t is current time, i.e., time this estimate is being made

CybOnt solves this equation with a series of conditionalizations and prior computed APS Λ values. The powerful feature of this algorithm is that it works against any set of attackers for which there are APS hypotheses so that coordinated attacks can be detected.

3. User Interface (UI)

The focus of CybOnt has been on ontology and fusion algorithms but a UI essential. Currently CybOnt has an offline UI for the knowledge base (T-Box) and an online one for the SU display. The KB display is of superSubtype, wholePart, and temporalWholePart, hierarchies so cyber Subject Matter Experts (SME) can review and edit the T-Box ontology. There are two types currently in use, a hierarchy viewer/editor and a triples viewer/editor. The triples viewer/editor allows viewing and editing of ontology T-Box RDF/OWL triples where the predicate expresses overlap and beforeAfter relationships.

The online (A-Box) UI is a graph database tool called Node Edge 7 (NE7) wherein the nodes are blue, red, and unknown assets, attack patterns, and attack pattern steps and the edges are the hypothesis links. A screen capture from a CybOnt demo is shown in **Figure 2**. An explanation of what is being displayed is provided in the figure caption notes.

4. Demonstrations

There have been six incremental demonstrations of CybOnt thus far:

1. Phase I demo. Small set of Wireshark alerts with small matrix of OF2APS and a COTS UI called NodeXL for display. OF2APS-based Λ demonstrated
2. Phase II Increment 1 - KB loaded into ontology; use of ontology tools to view hypothesis results. Scaled OF2APS "matrix" to all CAPEC™, using viewable/editable DBMS and hierarchy and triple tools
3. Phase II Increment 2 - Initial use of DISA Joint Communications Systems Simulator (JCSS) with partial CAPEC™, APS2AP Level 2 fusion. Demonstrated counter-intuitive but correct case of AP Λ >> sum(APS Λ) due to what can be thought of as a jigsaw puzzle effect. Once so many of the APS filled in for an AP, that hypothesis vastly outweighed all others.
4. Phase II Increment 3. - JCSS-synthetic scenario two sets of all CAPECs, one set all 3σ off (abnormal) and the other to 0σ (benign). In Tactical Cloud Reference Implementation (TCRI) with graph DB, NE7. Employed Accumulo, Rya, Storm, and other tactical cloud technologies.

5. Phase II Increment 4 – Open source PCAP from National CyberWatch Mid-Atlantic Collegiate Cyber Defense Competition (MACCDC) [16]. PCAP file processed through SNORT to generate alerts. Mission impact info. Looking at alert graph leads to Source-Pattern-Destination (SPD) graph showing source experimenting with several assets.
6. Phase II Increment 5 – Open source of 3 days of 26,000 SNORT alerts. Probabilistic mapping to ontology OF using the adaptive interface. Alerts show possible compromise of multiple assets.

5. Candidate Next Steps:

a. CybOnt Λ and Deep Learning Synergistic Architectures. It may be beneficial to incorporate the advances in Deep Learning (DL) in a hybrid architecture. Some possibilities are:
    1) Bootstrapping DL. CybOnt is used to develop training sets for DL and the DL real-time learning is used to adjust CybOnt's knowledge base.
    2) Augmenting Feature Extractor. CybOnt and DL both receive the sensor inputs and CybOnt's hypotheses and Λ become additional inputs to the DL.
    3) CybOnt used as rapid data triage. CybOnt processes some subset of the SE or OF, reducing their dimensionality for the DL.
    4) Validator and Explainer. CybOnt is used to validate the DL has not be spoofed (e.g., adversarial AI) or come into an environment it does not know. At the same time, CybOnt explains DL results it agrees with using pedigree or provenance embedded in CybOnt's fusion file and displayable in graph UI.

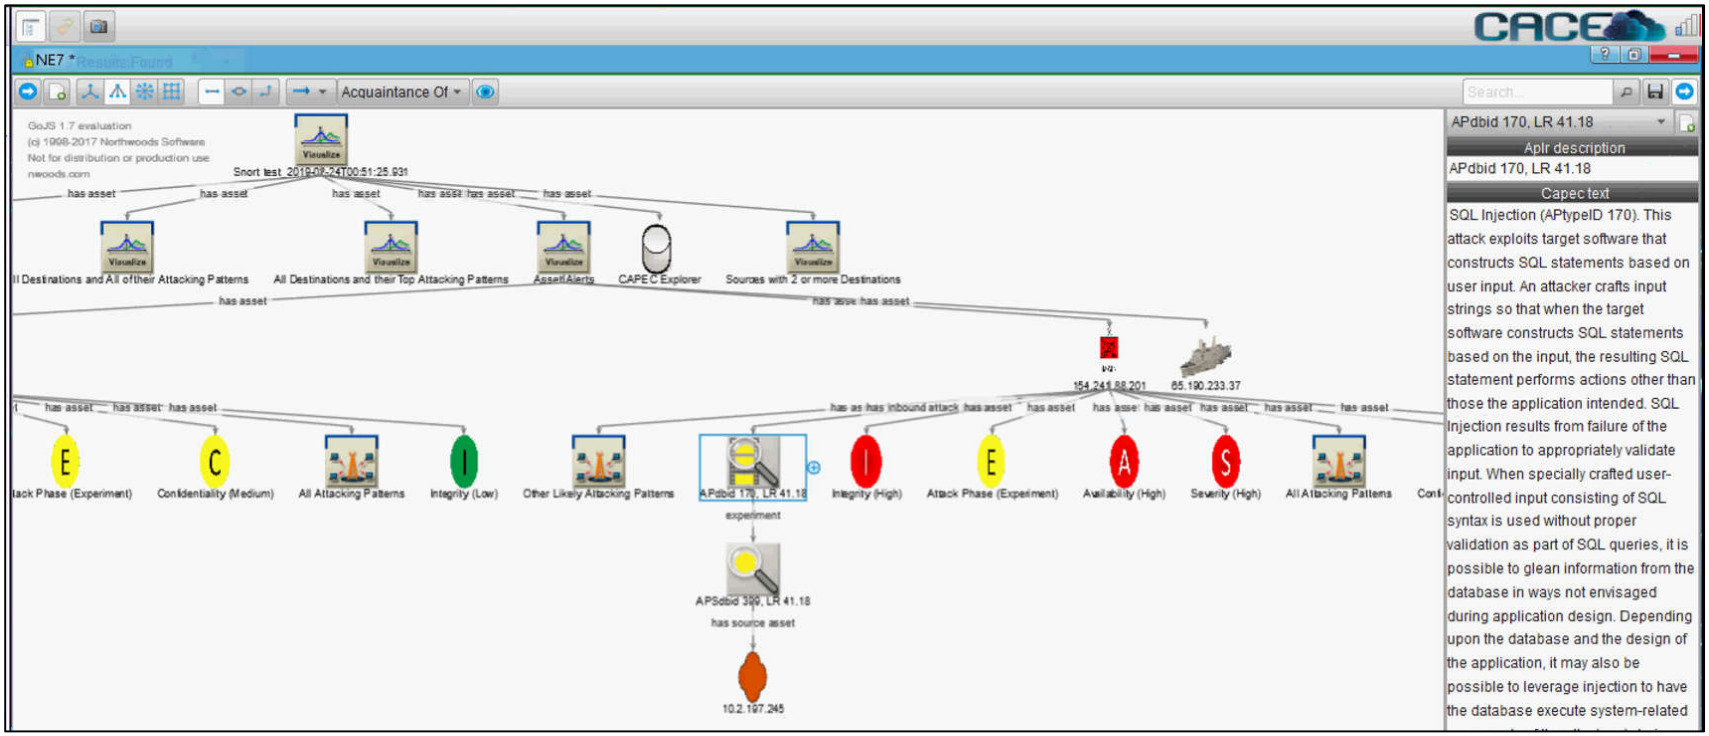

Figure 2. Online (A-Box) UI Screencap

Displayed in the first node expansion level are several graphs generated from the fusion file that allow the operator/analyst to explore source-centric or destination-centric. There is also a graph for high Λ hypotheses called “Asset Alerts”. The subgraph that is expanded is for alerts which are Λ thresholded at whatever level is required by the operator in the current situation. The alert being expanded is for the Marine Air Ground Task Force (MAGTF) Headquarters (HQ). It shows color-coded alerts for Confidentiality, Integrity, and Availability (CIA) as well as Severity and Attack Phase (Explore, Experiment, and Exploit). The AP causing the alert is shown in the middle and the APS and sources are expanded, in this case only one each. AP description information is shown in the node amplification frame to the right. .

6. References

[1] Office of Naval Technology, "Functional Description of the Data Fusion Process", Data Fusion Development Strategy, Office of Naval Technology, November, (1991)

[2] Steinberg, A. N., Bowman, C. L., White, F. E., "Revisions to the JDL Data Fusion Model", http://www.dtic.mil/dtic/tr/fulltext/u2/a391479.pdf

[3] David L. Hall, Martin Liggins II (Editor), James Llinas (Editor), Handbook of Multisensor Data Fusion: Theory and Practice, Second Edition, CRC Press, (2012)

[4] Yaakov Bar-Shalom, X. Rong Li, Thiagalingam Kirubarajan, Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software, Wiley, (2001)

[5] Lawrence A. Klein, Sensor and Data Fusion Concepts and Applications, SPIE Press, 1999

[6] Waltz, Edward; Ontologies and Data Fusion; Second Annual CMIF Workshop on Critical Issues in Information Fusion; Beaver Hollow, NY; October 2003

[7] Boury-Brisset, Anne-Claire; Ontology-based Approach for Information Fusion; Proceedings of the International Sensor and Information Fusion conference; ISIF; 1993.

[8] McDaniel, D.M., Regian, J.W., and Schaefer, G., "Ontology Based Fusion for E-2D", in *Proceedings of the National Symposium on Sensor and Data Fusion*, Military Sensing Information Analysis Center (SENSAIC), 2005

[9] Pearl, Judea; Probabilistic Reasoning in Intelligent Systems: Patterns of Plausible Inference; 1988

[10] Tore Schweder, Nils Lid Hjort; Confidence, Likelihood, Probability: Statistical Inference with Confidence Distributions; Cambridge University Press; 2016

[11] http://capec.mitre.org/

[12] Yaakov Bar-Shalom, Thomas E. Fortmann; Tracking and Data Association; Academic Press; 1988

[13] The Analytic Sciences Corporation (Author), Arthur Gelb (Editor); Applied Optimal Estimation; The MIT Press), 1974

[14] David Lee Hall; Mathematical Techniques in Multisensor Data Fusion; Artech House; 2004

[15] E. E. Holmes, "Kalman filtering for maximum likelihood estimation given corrupted observations"

[16] https://www.netresec.com/?page=MACCDC