

First-Principle Approach to Functionally Decomposing the JDL Fusion Model: Emphasis on Soft Target Data

Richard T. Antony
SAIC

488 Long Mountain Rd. Washington, VA, USA
22747
rantony@ieee.org

Joseph A. Karakowski
US Army RDEC

Ft. Monmouth, NJ, USA 07703-5000
joseph.a.karakowski@us.army.mil

Abstract - *The paper summarizes work to date directed at defining a service-based functional decomposition of the fusion process. The resulting architecture accommodates (1) traditional sensor data, as well as human-generated input, (2) streaming and non-streaming data, and (3) the fusion of both physical and non-physical entities. Fifteen base level fusion services are identified then utilized to construct a service-based decomposition of JDL fusion model levels 0 - 2. Concepts, such as clustering, link analysis and database mining, that have traditionally been only loosely associated with the fusion process, are shown to play key roles within this framework. A pedigree summary metric is defined that characterizes the informational distance between individual fused products and source data.*

Keywords: JDL fusion model, all-source fusion, service oriented architecture, soft targets, unstructured natural language, similarity, functional decomposition, fusion principles, fusion metrics, pedigree

1. Fusion first principles

Numerous papers and books have been published that address various aspects of the data fusion problem. Some works have focused on mathematical foundations of data fusion [1,2,3]. Others have concentrated on fusion techniques and practices [4,5]. Still others have attempted to develop a more holistic view of the fusion process [6].

In toy world fusion applications, (1) entities are readily distinguishable, (2) uncertainty is well defined (or non-existent), and (3) the environment has a trivial (or non-existent) effect on entity behavior. None of these simplifications apply in realistic applications. Due to various forms of uncertainty encountered in applications of interest and the non-trivial impact of environment on behavior, we treat information fusion from a first-principle's perspective as the process of *reasoning about similar and dissimilar entities*. Entities can be *physical* (individuals, vehicles, equipment), *non-physical* (events,

organizations), or *abstract* (financial transactions, "state of unrest").

Fusing *similar* entities involves associating the "same" individual, vehicle, aircraft, or organization, while fusing *dissimilar* entities involves (1) operations among entities within the same class (two individuals) or (2) entities from different entity classes (e.g., an individual and an organization). While *similarity* relies on an *identity relation*, *dissimilarity* involves a (non-trivial) *derived relation*. Due to the many forms of uncertainty associated with most fusion applications, multi-hypothesis reasoning provides a mechanism to defer decision-making until additional supporting or detracting information becomes available. Truth maintenance (either manual or semi-automated) seeks to resolve conflicts that arise within the fusion product space.

In general, entities can have a large number of relevant dimensions, but virtually all possess both *temporal*, as well as *spatial* dimensions. In this paper, we expose these critical dimensions then treat the triple (*entity, location, time*) as the core fusion dimensions. While *entities* are normally discrete and possess discrete attributes, *space* and *time* are continuous dimensions of potentially large extent. By focusing on these three core dimensions, information fusion can be characterized as the process of reasoning about discrete entities embedded in a 4-D space-time continuum. This viewpoint, in turn, leads directly to the definition of the eight canonical fusion forms shown in Table 1. To better align with conventional terminology, *location & time comparisons* will be referred as being either *near* (similar) or *not near* (dissimilar). If either or both space and time are unknown (or inadequately specified), *time* and *space* comparisons are assigned to the category *not near*.

A given fusion operation can be formally characterized by the number of different entities participating in the fusion operation and defined as the *fusion operation order*. In this paper, we restrict our discussion to fusion processes involving just two operands (i.e., $A \times B$). Thus, the fusion of two *similar entities* is a fusion operation order-1 process, while the fusion of two

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUL 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE First-Principle Approach to Functionally Decomposing the JDL Fusion Model: Emphasis on Soft Target Data				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army RDEC,Ft. Monmouth,NJ,07703-5000				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES 11th International Conference on Information Fusion, June 30 ? July 3, 2008, Cologne, Germany.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

dissimilar entities is a fusion operation order-2 process. In the next two sections, we demonstrate that each of the eight canonical fusion forms defined in Table 1 plays a critical, well-defined role within the overall analysis process.

1.1 Fusion operation order-1 forms

Form 1 represents traditional single entity multi-source *space-time association*. Radar tracking is the most familiar streaming media example of single source, similar entity fusion. Fusing an ELINT and an EO report from similar locations and at approximately the same time is an example of multi-source, similar entity fusion form 1. Linking a text report with an ATM transaction that occurred near in both time and location represents a more contemporary example.

Form 2 represents single entity fusion based on location, independent of time. The *absence of location change* for a single entity (e.g., no change detection in imagery) has often been used to identify stationary equipment. When the measurement interval is relatively long, however, the entity being observed may only be quasi-stationary (e.g., a vehicle could have left and returned before the subsequent observation). In a more contemporary application, Form 2 supports the identification of locations associated with an entity.

Because physical objects cannot be at widely differing places at the same time, **Form 3** represents a constraint that can be used to detect inconsistent hypotheses, as well as deliberate deception. For non-physical entities, Form 3 can be used to trigger higher level reasoning processes.

Form 4 represents opportunistic, sparse (non-kinematics based) entity *tracking* where track updates are *not near* each other in either time or space. Such “tracks” may evolve over relatively long periods of time (e.g., hours, days, or months). In traditional kinematics-based

tracking applications (Fusion Form 1), tracks tend to “break” when the updates fail to be close enough in space-time to “connect the dots.” Assembling a set of tracklets believed to be associated with a single entity represents a special case of this fusion form.

1.2 Fusion operation order-2 forms

Form 5 represents *multiple entity coincidence* (two entities discovered to be near in space and time). A traditional example would be (1) optical detection of several side-by-side vehicles, (2) an electronic report on a radar system, and (3) a communication report associated with a command vehicle. In a more contemporary application, a report on **Individual A** might be associated with **Individual B** who placed a cell phone call from approximately the same location and at approximately the same time.

Form 6 associates two dissimilar entities based on *location independent of time*. Different individuals who visit the same safe house or mosque at different times can potentially be linked due to their associations with these locations (e.g., buildings, user specified key locations, hot spots). Because Form 6 is essentially time-independent, exploitation inherently involves *data mining-like operations* rather than traditional bottom-up “streaming” sensor data. As with “forensic” track development (Form 4), the temporal “distance” between two states or events might not be a key attribute even though activity state *ordering* might be.

By contrast, **Form 7** represents dissimilar entity linking based on temporal coincidence, but no spatial coincidence. To establish such relationships requires the exploitation of other problem dimensions, such as phone records, on-line chat, or financial transactions.

Finally, **Form 8** represents the case where *no apparent spatial or temporal link* exists between two dissimilar entities. Entity association thus relies solely on

Table 1. Eight canonical fusion forms

Form	Entity	Location	Time	Type of Analysis	Interpretation
1	Similar	<i>Near</i>	<i>Near</i>	Fuse one or more information sources from a single entity near in location and near in time	Traditional space-time correlation/tracking
2	Similar	<i>Near</i>	<i>Not near</i>	Fuse multiple sources associated with a single entity (near in location & not near in time)	Stationary object
3	Similar	<i>Not near</i>	<i>Near</i>	Infeasible condition	Potential constraint
4	Similar	<i>Not near</i>	<i>Not near</i>	Associate single entity observations that do not satisfy form 1 & 2 conditions	Non-kinematics space-time tracking
5	Dissimilar	<i>Near</i>	<i>Near</i>	Associate co-located objects	Traditional space-time correlation
6	Dissimilar	<i>Near</i>	<i>Not near</i>	Link different entities at the same location at different times	Location <i>alone</i> might link objects
7	Dissimilar	<i>Not near</i>	<i>Near</i>	Link different entities by temporal coincidence	Coordinated activities (cell phone call)
8	Dissimilar	<i>Not near</i>	<i>Not near</i>	Multiple entities at different locations & times	Non spatiotemporal associations

other entity dimensions, or on more complex entity relationships. Such links can be either *explicit* or *implicit*. Explicit links include associations among different entities derived from records showing joint ownership of property, known familial relationships, and organizational memberships. Implicit links imply more subtle associations including both indirect and purposely hidden relationships (e.g., money laundering operations). Traditional *link analysis* ideally supports Fusion Forms 5-8.

Table 1 leads to a number of important insights:

1. The first four fusion forms represent level-1 (same/similar-entity; fusion operation order-1), while the latter four forms represent level-2 (dissimilar-entity) processes.
2. The eight canonical forms represent base-level fusion services that apply across all entity classes, be they physical, non-physical, or abstract. While the *interpretation* of the results vary depending on the entity class(es) involved, the fusion forms themselves are application-independent.
3. The fusion forms are *complete* in the sense that they represent all possible combinations of the three principal fusion dimensions: *entity*, *location*, and *time*.
4. Fusion Forms 1, 2 & 5 are strongly associated with conventional hard-target fusion applications. The remaining fusion forms are required by applications that involve human-generated information. In the case of kinematics-based tracking, Kalman filter-based trackers begin to fail when input data transitions from being *near* (Fusion Form 1) to *not near* (Fusion Form 4) along the 4-D space/time dimension.
5. Because Fusion Form 8 relies on *other than spatiotemporal* relationships, it possesses the most open-ended data mining requirements.

1.3 Supplemental fusion services

In addition to the eight base-level fusion services defined in Table 1, Figure 1 depicts seven additional services that support the overall fusion process. The Filtering Service is straightforward and we focus on the six remaining services.

Message Extraction & Normalization Service prepares source data for fusion processing. Traditional sensors employ well-defined formats, normalization, and registration protocol. Non-traditional input (semi-structured and unstructured natural language) requires language translation, extraction, parsing, co-reference & ambiguity resolution, as well as other forms of analysis to “normalize” source data so that it can be properly handled by a semi-automated or fully automated fusion system. For natural language input, the message normalization service must produce well-structured facts that describe, at a minimum, *who*, *what*, *to whom*, *when* & *where*.

Location Normalization Service translates deterministic, probabilistic, and semantic *location descriptions* into a common representational framework that provides a uniform location representation (grid coordinate, polygon, *near* Mosque A) [7].

Temporal Normalization Service translates deterministic, probabilistic, and semantic *temporal descriptions* into a common framework that supports fusion across all data sources (military time/date representation, *yesterday*, *late afternoon on the 3rd of May*).

Location Support Service provides search, set operation, and spatial correlation services across the normalized spatial representations (deterministic, probabilistic, and fuzzy). In addition, this service supports the identification of *hotspots*, maintains all *user-defined* and *derived* key locations, and supports generation of alerts when certain entities/entity classes are discovered in or near *locations of interest*.

Temporal Support Service provides search, set operation, correlation, and other forms of temporal analysis across normalized temporal representations (deterministic, probabilistic, and fuzzy).

Context Support Service provides the interface to all supporting databases, including spatially organized data layers (GIS, human terrain knowledge, political districts, user defined regions of interest) that provide input to the fusion process or that constrain the interpretation of both source data, as well as fused products. Because context knowledge can have high dimensionality the service may need to be tailored to focus on relevant problem dimensions.




Figure 1. Base-level fusion services associated with first three levels of the JDL fusion model.

2. Source alignment and normalization

Source alignment and normalization has traditionally focused on registering or otherwise aligning hard-target sensor data. Introducing natural language input has led not only to required extensions to the existing fusion framework, but to a renewed focus on level-0 data preparation requirements. While semantic “fuzziness” promotes communication and understanding among humans, it challenges traditional machine-based reasoning approaches. To be effective, the extraction process must capture the *essence* of the *meaning* of potentially complex, interrelated, context-sensitive text strings. To facilitate automated processing, these extracted products must be converted into a simple, yet robust data set that preserves the essential meaning of message sets.

To illustrate just one aspect of the technical challenge involved, consider the following example:

Person A wired money to **Person B’s account** on June 1.
who verb what to whom where when

The goal of meaning extraction is to produce a distilled, unambiguous product that captures all the basic facts from an arbitrarily complex text string. Even in this relatively simple example, several possible interpretations exist.

Attempting to achieve *automated* all-source

fusion effectively demands that traditional (hard) sensor input and text-based (soft) information sources be normalized without creating separate “stove pipes” for these distinctly different classes of source data. Employing a uniform level-0 input syntax offers perhaps the simplest means of meeting this objective. Because natural language input will generally be more “complex” than the systematic output associated with most sensor systems, it seems appropriate to map conventional sensor data into the syntax selected to handle extracted natural language input. For instance, radar detections can be mapped rather straightforwardly into the above syntax as follows:

Unknown object detected at (X, Y, Z) at time t1
Who/what verb where when

In addition to the challenge of “normalizing” natural language input so the fusion process can operate uniformly across all information sources, introducing human-generated information alters the traditional “sequential processing” view of the fusion process. Unlike conventional sensors that tend to capture “primitive” data regarding entities, humans can provide information that directly (1) refines entity descriptions (level-1 “facts”), (2) explicitly links individuals to organizations (level-2 “facts”), and (3) reveals intent of individuals or groups (level-3 “facts”). Thus, while traditional fusion approaches

rely on the accumulation of low-level facts that must be fused or otherwise manipulated to produce higher level of abstraction products, “out of the box” human-generated messages potentially provide input applicable to fusion levels 1-3. Despite these additional considerations, the objective of the fusion process remains unchanged, that of developing a maximally consistent understanding of a diverse set of data.

3. Tracking and cross-entity association

The objective of level-1 fusion is development of a consistent understanding of primitive entities, be they individuals or various kinds of equipment. Level-2 fusion plays several roles, including (1) cross-entity association, (2) characterization of relationships among entities, and (3) characterization of composite entity behavior. Figure 2 presents an abstract view of three entity classes (*Events*, *Organizations*, *Individuals*) linked by specific *relations*. Just as a time-ordered sequence of primitive entity states can be viewed as the *track* of that entity (i.e., primitive entity behavior over time), a time-ordered sequence of *event* states can be treated as the “track” of a composite entity (e.g., organization behavior over time). The actual organization “behind” an event might initially be unknown or even rather nebulous (e.g., concerned citizens against global warming). Relations permit Order-2 fusion operations between two dissimilar entities (either from the same or different entity class). In general, *relations* can be unidirectional or bidirectional, one-to-many or many-to-one.




Figure 2. Three Level-1/2 entity classes and associated implied/derived relations.

In level-1, information from a variety of sources is assembled and correlated in an on-going process that seeks to better characterize (primitive) entities of interest. In traditional ISR applications, level-1 generally focuses on kinematics-based tracking. In applications involving soft targets, motion-based tracking may represent a relatively minor component of a process that seeks to capture (relatively static) facts about individual entities in terms of their biometrics, associations, training, long-term behavior, as well as more dynamic information, such as

current activity and present location. Tracking cell phone calls, electronic financial transactions, meeting attendance, purchases, chat sessions, and Internet activity represents a potentially important sequence of activity states even though few of these states are associated with physical movement. Thus, conventional tracking is just one component of entity behavior (*activity “tracking” over time*).

Entity tracks, essentially independent of entity class, can be exploited to deduce patterns and identify expected, as well as anomalous behavior. While complete understanding of entity behavior, prediction of future behavior, and intent estimation are highly desirable objectives, achieving these objectives involve many challenges. Despite the complexity of these objectives, viewing an entity track (observed or inferred behavior) within the appropriate context permits model-based and/or case-based categorization of behavior.

4. Fusion model decomposition

We next offer a service oriented functional decomposition of levels 0 through 2 of the JDL fusion model based on the three entity classes (individuals, events, and organizations) depicted in Figure 3. The decomposition tacitly assumes that events are associated with appropriate organizations just as reports on individual (or vehicle) are associated with the appropriate track. Thus, while the track of an individual is the sequence of time-ordered reports associated with that individual, an organization track is the time-ordered sequence of events/activity states associated with that organization.

Table 2 identifies three level-0 services, two level-1 services, and seven level-2 services. The fusion level-0 processes were defined in the Supplemental Services (Section 1.3). Fusion level-1.1 performs primitive entity tracking (Fusion Forms 1-4), while level-1.2 serves as the “situation-understanding” component of the level-1 process. Fusion level-2.1 links *dissimilar entities* within the *same entity class* (e.g., Individual A x Individual B – Fusion Forms 5-8). The level-2.2 fusion component links *different entities* from *different (observable) entity classes* (e.g., Individual A x Event A). Level-2.3 links observable entities with abstract (non-observable) entities (e.g., Event A x Organization A). Level-2.4 links individuals with organizations (Individual A x Organization A). Level 2.5 tracks organization behavior over time (e.g., Organization A event track x event B). Level-2.6 links two *dissimilar entities* from within the *same (aggregated) entity class* (e.g., Organization A x Organization B). Finally, Level-2.7 provides “situation understanding” with respect to composite entity behavior (Organization A event track x Behavior model).

Table 2: Functional decomposition of first three levels of the JDL fusion model.

Level	Fusion composition arguments	Entity class	Entity	Product	Required processing
0.1				Message translation	Extraction into generic fact database
0.2				Location normalization	Normalized deterministic, probabilistic, fuzzy spatial representations
0.3				Temporal normalization	Normalized deterministic, probabilistic, fuzzy temporal representations
1.1	Individual A x Individual A	Same	Similar	Individual A track	FF1-4
1.2	Individual A track x Behavior model	Same	Similar/dissimilar	Individual A behavior class	Model match Behavior categorization
2.1	Individual A track x Individual B track	Same	Dissimilar	Entity relation graph	FF5-8
2.2	Individual A track x Event A	Different /observable	Dissimilar	Link individuals to events	FF5-8
2.3	Event A x Organization A	Different/ observable-non observable	Dissimilar	Link events to organizations	FF5-8
2.4	Individual A x Organization A	Different /non-observable entity	Dissimilar	Organization membership	Link Individual A via <i>Event history</i> to Organization A
2.5	Organization A event track x Event B	Same	Similar	Organization A track	FF1-4
2.6	Organization A event track x Organization B event track	Same	Dissimilar	Organization relationships	FF5-8
2.7	Organization A event track x Behavior model	Same	Similar/dissimilar	Organization behavior class	Model match Behavior categorization




Figure 2. Expanded functional-level view of fusion levels 1 & 2.

From Table 2, the base level fusion forms defined in Table 1 are seen to directly support fusion levels-2.1, 2.2, 2.4, and 2.5. Note that fusion level-2.3 “breaks the pattern” in terms of its composition form. Rather than being associated with a *correlation/matching* operation, this form involves a *mapping or assignment* between an observable entity (an individual) and a non-observable entity (organization). Figure 2 offers a block diagram view of the level 1 & 2 services.

5. Extensions and generalizations

This section discusses extensions relevant to cascaded fusion operations. Recall that *fusion operation order* was defined as the number of *dissimilar* entities combined during a fusion operation. We now assign an *order* metric to the *product* of each such operation. Although there exist other possibilities, we employ a conservative rule that “favors” source data:

Fusion product order = Min (*fusion product order* of the operands) + *fusion operation order*

where source data (data that has undergone no (explicit) fusion processing) is considered fusion product order-0 information. Thus, while *fusion operation order* characterizes the complexity of a given fusion operation, *fusion product order* characterizes the “informational distance” of a fused product from source data. Based on these definitions, the track of a primitive entity (e.g., an individual) created by combining two *fusion product order*-0 data sets produces a fusion product order-1 result (the *minimum* of the fusion product orders of the two operands = 0; fusion operation order = 1, i.e., similar entities). On the other hand, the fusion of two dissimilar fusion product order-0 entities generates a fusion product order-2 result (the minimum of the two operands = 0; fusion operation order = 2, i.e., dissimilar entities).

To illustrate the practical utility of these definitions, suppose two prior fused products are to be fused. Further, suppose the first operand has fusion *operation order*-4 and fusion *product order*-3 and the second operand has fusion *operation order*-1 and fusion *product order*-2. The resultant fused product has a *fusion product order* of 4 (i.e., Min (3, 2) + 2). Alternatively, suppose a natural language message explicitly states a level-2 “fact” representing information that is comparable to this derived order-4 fused product. While source data is by definition an order-0 product, the order-4 fused product was derived through multiple stages of reasoning (correlation and inference). To compare these two products, we define the *fusion product order distance* (the difference in the *fusion product orders* of two products). In the above example, the

two “competing” level-2 products possess a fusion product order distance of 4 (fusion product order-4 v. fusion product order-0).

Fusion of these two independently derived products results in an order-1 fusion product (i.e., Min (4,0) +1 to account for *similar entity* fusion). On the other hand, suppose the two products lead to an inconsistent conclusion. Rather than “combining” the products, we would likely want to select the product with the higher confidence level. (Although confidence assignment and aggregation has not been discussed in this paper, they represent important aspects of information fusion.) Unless the confidence in the derived fusion product is quite high, information “closest” to source data (i.e., lower fusion product order information) might be preferred for the simple reason that uncertainty tends to increase through multiple stages of reasoning.

Assigning fusion *operation* and *product orders* to fused products offers several benefits: First, both measures are independent of the class and/or complexity of the entities being fused. Second, the *product order* metric provides a systematic means of maintaining the pedigree-relevant “informational distance” of a given product relative to source data. Third, process and product orders provide a foundation for exploiting set theoretic properties such as inheritance, equivalence, and transitivity. Inheritance allows higher fusion product order products to be viewed as functional compositions of lower level of abstraction products. Equivalence properties apply to all fusion operation order-1 operations. Transitivity (associated with fusion operation orders > 2) facilitates discovery and higher order inference.

6. Specialized fusion products

In addition to being the fundamental building blocks of a service-oriented fusion architecture, the base level fusion services support the development of specialized applications, such as *monitoring* and *automated alert generation*. Suppose an analyst formulates a rule for detecting *candidate* high value individuals (HVIs) as a weighted linear combination of the following four pieces of information: (1) number of explicit links to known HVIs, (2) Number of times candidate was known to be at a user-defined key location, (3) number of times candidate was at the same location at the same time as any HVI, (4) number of times a candidate was at the same location (but not the same time) as any HVI. In this case, four base level fusion services (message extraction & normalization service, Fusion Form 4, Fusion Form 5, and Fusion Form 6) directly provide the information required to implement this rule.

Subtler links between individuals can be investigated by exploiting *transitivity*. For example, suppose a potential

indirect relationship exists between a candidate individual and a known HVI through another individual. Fusion level 2.1 identifies all links (explicitly stated, as well as FF5-8) between all individuals of interest. If any individual who is linked to the *candidate* has links to an HVI, the candidate could conceivably be linked to the HVI through this intermediate individual. The *belief* associated with such conjectures depends on the number and character of those links. For example, a Fusion Form 5 product (location/time association) generally represents stronger evidence of an association that does Fusion Form 6 (location only association).

7. Summary and future directions

The paper presented a generic approach to fusion that permits human-generated, soft target information to be seamlessly integrated with existing hard-sensor, hard-target information fusion approaches. Eight canonical fusion forms, as well as seven base-level supplemental fusion services were defined. A serviced-oriented framework for implementation of the first three levels of the JDL fusion model was proposed. The concepts of *fusion operation* and *fusion product order* were introduced. Fusion operation order provides a description of the complexity of an individual fusion process while fusion product order measures the informational distance of a fused product from source data.

Based on the principles presented in this paper, a prototype fusion system (Context and Fusion Support System) has been developed. Earlier papers present selected level-1 and level-2 output from CFS2 for an unclassified 200 message set [8,9]. Future research will be directed at extending the level of abstraction of the current fusion process decomposition, continued formalization of the principles underlying all-source fusion, as well as investigation of fusion operation order-3 and higher processes. As the theory matures, these new capabilities will be added to CFS2 so that the utility of these extensions can be systematically evaluated.

REFERENCES

- [1] I.R. Goodman, Ronald P. Mahler, and Hung T. Nguyen, *Mathematics of Data Fusion*, Springer, 1997.
- [2] David L. Hall and Sonya A. H. McMullen, *Mathematical Techniques in Multisensor Data Fusion*, 2005.
- [3] Ronald P. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Artech House, 2007.
- [4] Edward L. Waltz and James Llinas, *Multisensor Data Fusion*, Artech House, 1990.
- [5] David L. Hall and James Llinas, *Handbook of Multisensor Data Fusion*, CRC, 2001.

[6] Richard T. Antony, *Principles of Data Fusion Automation*, Artech House, 1995.

[7] Richard T. Antony and Joseph A. Karakowski, Towards Greater Consciousness in Data Fusion Systems, National Symposium on Sensor and Data Fusion, Session SA04 pp. 1-23, June 2007.

[8] Richard T. Antony and Joseph A. Karakowski, Fusion of HUMINT & Conventional Multi-Source Data, National Symposium on Sensor and Data Fusion, Session SC04 pp. 1-16, June 2007.

[9] Richard T. Antony and Joseph A. Karakowski, Service-Based Extensions to the JDL Fusion Model, SPIE Defense and Security Conference March 2008.