

Ontological Support for Living Plan Specification, Execution and Evaluation

Erik Thomsen
Charles River
Analytics
Cambridge, MA
ethomsen@cra.com

Frederick Reed
Charles River
Analytics
Cambridge, MA
freed@cra.com

William Duncan
NCOR
Buffalo, NY
wdduncan@gmail.com

Tatiana Malyuta
New York City College of
Technology, NY, NY
tmalyuta@thedata-science.org

Barry Smith
NCOR
Buffalo, NY
phsmith@buffalo.edu

Abstract—Maintaining systems of military plans is critical for military effectiveness, but is also challenging. Plans will become obsolete as the world diverges from the assumptions on which they rest. If too many ad hoc changes are made to intermeshed plans, the ensemble may no longer lead to well-synchronized and coordinated operations, resulting in the system of plans becoming itself incoherent. We describe in what follows an Adaptive Planning process that we are developing on behalf of the Air Force Research Laboratory (Rome) with the goal of addressing problems of these sorts through cyclical collaborative plan review and maintenance. The interactions of world state, blue force status and associated plans are too complex for manual adaptive processes, and computer-aided plan review and maintenance is thus indispensable. We argue that appropriate semantic technology can 1) provide richer representation of plan-related data and semantics, 2) allow for flexible, non-disruptive, agile, scalable, and coordinated changes in plans, and 3) support more intelligent analytical querying of plan-related data.

Keywords—adaptive planning; outcomes assessment; ontology

I. THE NEED FOR ADAPTIVE PLANNING

“No plan survives first contact with the enemy” (Clausewitz, *On War*). Real world uncertainties all but guarantee that even the most carefully developed plan will not be carried out exactly as intended. The military response, as in the business domain, has been to increase the speed and agility of planning and execution [1-4]. On the strategic level, the transition from the Joint Operation Planning and Execution System (JOPES) to an Adaptive Planning and Execution (APEX) system exemplifies this trend. In addition to speeding up the deliberate planning and review cycle, these efforts seek to increase the number of planned options and contingencies.

According to the Adaptive Roadmap II, signed by the Secretary of Defense in March 2008, the ultimate goal is to provide plans that are “maintained continuously within a collaborative environment” to reflect any changes that impact any significant aspects of a plan. Such plans will together form something the Adaptive Roadmap calls a “living plan.” Plans may need to be adjusted to maintain their relevance based on changes in the world (e.g., weather, location of enemy troops, troop readiness, air assets). Additionally, they may need to be adjusted in order to maintain their coherence within a system of plans, such as when the goals of supporting or supported plans change.

II. THE IDEA OF THE LIVING PLAN

In the current state of military planning – as encapsulated in Joint Doctrine (JP 5.0) – a distinction is drawn between deliberate planning and crisis action planning. Deliberate planning is supply driven. Plans are static information objects created as the outputs of a deliberative, rule-governed process, and stored in a repository until needed. They may be created years ahead of actual use, or they may never be used at all. Crisis actions plans are demand driven: something happened and we need an urgent response; because the response should involve a degree of organized action, planning is needed. Crisis action planning is a response to the uncertainty involved in our knowledge of real-world states. But even deliberate planning rests on an institutional acknowledgement of our inability to accurately predict the future, in that Doctrine allows the making of ad hoc resource requests which deviate from the deliberate plan as specified. Sometimes, on first contact with the enemy, deliberate plans break and workarounds are needed. Regardless of the quality of the prior deliberation that went into the deliberate plan, the need for such corrective actions as a result of the unanticipated interactions between blue forces and the world make for suboptimal procedures.

The goal of the living plan is to remove this ‘breaks because it would not bend’ feature of the deliberative plan by minimizing the distinction between deliberate planning and crisis action planning through a new type of planning process that is marked by constant update in light of updates in our real-world knowledge. The idea is to embed into the very fabric of plan representation our uncertainties about the world, so that the activity of planning is transformed from one of the creation of plans as outputs to a process of continuous plan development. The living plan itself becomes a probabilistic, branching information artifact – a representation of the moment-to-moment intentions not merely of single platoon commanders but of the military as a whole. It incorporates at each phase representations of multiple alternative courses of action which are continuously changing in light of actual and projected states of the world, adjacent plans, supporting and supported plans.

III. ADAPTIVE PLANNING REQUIREMENTS

We believe that any computational approach to supporting the Secretary of Defense’s goal for living plans must meet six critical requirements.

First, it must be able to represent all the types of entities and relationships, knowledge about which is important to maintaining a living plan. This requires a *highly expressive representational capability* to capture, manage, and reason over plans, plan elements (e.g., goals, available assets, weather, battle terrain), and their relations within a system of plans.

Second, any approach must be able to detect meaningful changes that impact plan relevance and coherence. This requires *effective monitoring and sensitivity analysis* to identify in a reliable and scalable way those changes which are of significance to the system of plans [5,6]. Recognition of the significant changes must then trigger processes that maintain the relevance and coherence of this system at multiple levels and across plan elements.

Third, any approach requires *coordinated adjustment processes*, which are needed to fulfill the second requirement (above). Such processes must be able to run independently, be applicable (when necessary) to real-time conditions, and be capable of harmonizing with other large-scale plan adjustments.

Fourth, any such approach requires *automated information extraction and routing* because maintaining realistic plans requires more information processing than can be achieved through manual methods alone.

Fifth, whether in support of human planners, warfighters during mission execution, operations assessment staff, or automated systems performing the same tasks, any approach needs to *support analytical queries against the ensemble of plan-related data*. Since plan-related data is very heterogeneous, this amounts to applying a unified structured query front end to structured and unstructured data on the backend.

Sixth, joint warfighters at all levels of command will need to collaboratively plan and execute in conjunction with semi-automated adaptive planning systems. Therefore, any approach for providing living plans must support *extensible and versatile*

interactive applications that can deal with the sorts of diverse but integrated user environments required for living plans.

Relative to the six requirements described above for supporting the Secretary of Defense’s goal for living plans, our overall approach is based on the idea that semantic representation of data by means of ontologies, combined with probabilistic classifiers operating in a transactional environment, will allow the needed representation, monitoring, analysis, sharing and querying of information at distinct levels of granularity and detail and across distinct applications. The system will be required, for example, to display a JFACC’s view of ATO mission plans, a squadron Commander’s view of the day’s mission plan, and STRATCOM’s view of a Theater. As in other domains, the semantic approach is designed to reduce information siloes, and enable effective tailoring of knowledge and information to different needs. It is designed also in such a way as to allow incremental improvements over time, as shortcomings in the framework uncovered at any given stage are rectified in subsequent stages.

In what follows we focus on the first and fifth requirements described above: for rich representations of data and semantics, and for the capacity to use such representations in mounting queries against plan-related data.

As regards the former, we describe the coverage domain of our proposed Plan Ontology (see Figure 2) in terms of how we: (a) model plans in terms of cyclical phase-specific attributes; (b) embed metrics that relate plans to world conditions; and (c) embed meta-metrics that use the metrics under (b) to create an incremental plan and plan-execution improvement process across the whole system. On each level multiple families of related terms will be required, including definitions and axioms specifying the relations between them.

As regards the latter, we describe how queries are passed through parts of the system in order to illustrate some of the semantic relations that need to be computed in order to support analytically useful queries over living plan data.

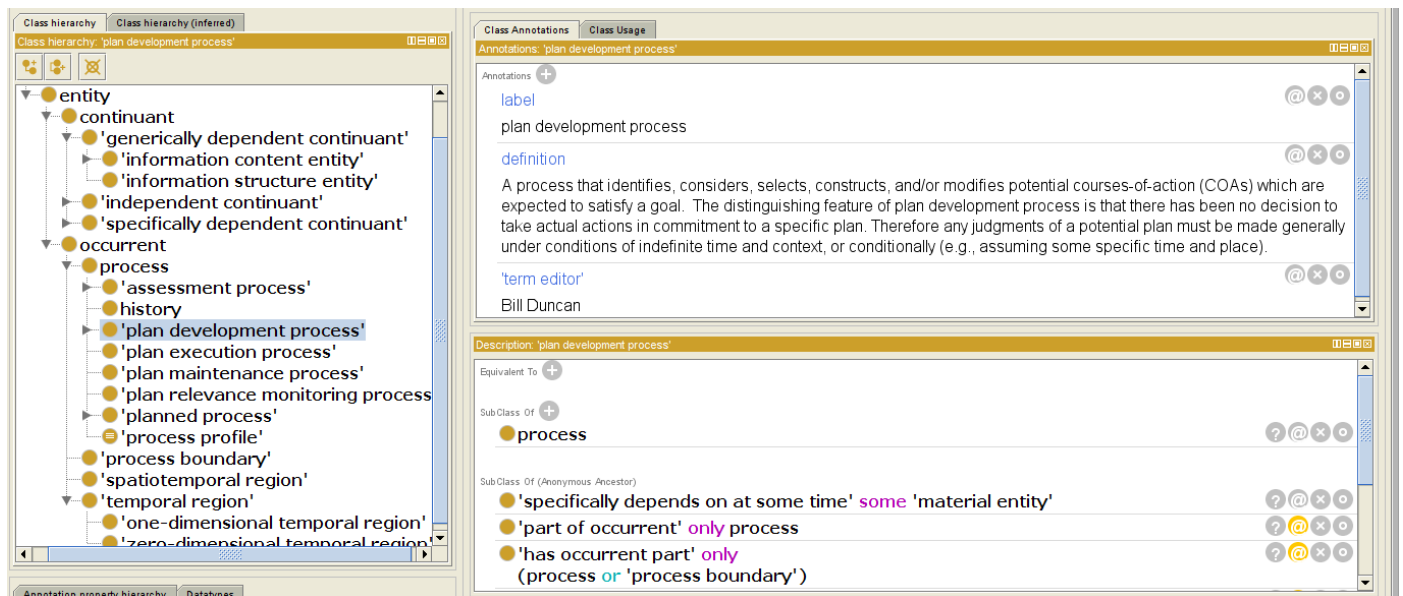


Figure 1: Fragment of the draft Plan Ontology at <http://ncor.buffalo.edu/plan-ontology>

IV. REPRESENTING PLANS IN RECURRING PHASES

To better understand and support the notion of continuous, living plans, we require a view of planning that is more abstract than is traditionally employed. The simplistic notion of ‘the plan’ created prior to ‘the execution’ is at odds with our view of planning as a dynamic, continuous, iterative process that not only adapts to the effects of planned actions, but also adapts the process of planning itself in ways designed to achieve more satisfactory outcomes over time.

Our model focuses on three primary factors in the planning process:

1. different phases of the planning process (successive phases within a given course of planning processes),
2. types of judgments within each of those phases that enable effective planning, and
3. information, including metrics, on which these judgments are grounded.

On the traditional view, planning only happens periodically as a precursor to its execution. Here, in contrast, we view the total planning process computationally as forming a series of parallel, interacting courses or flows at a number of different levels. These processes unfold dynamically, with changes in any given course being communicated to parallel and hierarchically related courses wherever changes in the latter are required. The system is organized in such a way that updated versions of needed plans and subplans can be generated at any point in time.

Each parallel course is itself seen as being organized into a succession of three phases corresponding roughly to the first three phases of the well-known Plan Do Check Act (PDCA) cycle, and similar models. A difference is that the phases in our framework are viewed as continuous and intermeshed with each other rather than discreet. Especially the *Act* phase, where adaptive actions are taken, is distributed and continuous across the other phases.

- *development* – This phase consists of processes of identifying, considering, selecting, constructing, and/or modifying potential courses-of-action (COAs) that are expected to satisfy a goal. This includes the process of creating and maintaining potentially executable ‘plans sitting on the shelf’ in traditional, deliberate planning – referred to in our ontology as ‘plan specifications’. The distinguishing feature of this phase is that there has been no decision to take actual actions in conformity with and under commitment to any specific plan.
- *execution* – This phase involves processes of planning while acting according to a particular planned COA. Unlike random or spontaneous actions, such planned processes can be evaluated relative to the plan. For example, indicators can be used to judge whether the intermediate effects of planned actions are consistent with expectations. But, as the plan has not yet terminated, the net effect of all planned actions relative to the goal set forth in the plan cannot be judged. A key planning process in the execution phase is the making of a decision to terminate execution because

the goal has been achieved, or because the plan is no longer relevant or coherent, or is being executed unsatisfactorily.

- *post-execution* – This phase involves the post-execution processes of interpreting and judging an executed plan and its outcomes relative to expectations. In this process, all actions taken under commitment to the plan have been taken. Thus their net effect can be assessed relative to the specified goal. The primary purpose of the processes involved in this post-execution phase is to enhance future planning, for example by:
 - defining new goals;
 - clarifying existing goals;
 - improving effectiveness in achieving goals.

Associated with processes of each of the mentioned types are four basic planning-related judgments that enable reasoning aimed at leading to the creation and selection of better plans:

- *relevance* – How well does the current state of planning relate to actual or anticipated external world conditions, such as constraints, opportunities, planned outcomes, unplanned side-effects, etc.?
- *coherence* – How well do the processes of planning on-going in the current phases relate to other synergistic planning processes. In other words, are they in conflict or coherent with other friendly force, coalition, political, etc. planning?
- *planning-assessment* – How well were the processes in each phase of planning performed by the *planner*, from a single person to an organization?
- *meta-metric learning* – How well does the current set of metrics support the goal of evolutionary improvement of the entire planning process (and, as a consequence thereof, the entire process of creating and executing and evaluating plans)?

V. REPRESENTING RECURRING CLASSES OF METRICS IN SUPPORT OF CYCLICAL PLAN PHASES

In this section, we bring together the three factors of planning outlined above – phases, judgments, and metrics – to see how they merge to form a more complete picture of a continuous adaptive planning process. For each combination of planning phase and judgment we provide example metrics. These are provided here for illustrative purposes only, and especially as concerns plan execution our framework will draw on the extensive list of Measures of Effectiveness and Performance identified in salient doctrine for the tasks of the Universal Joint Task List, for example as described at: <http://www.dtic.mil/dtic/tr/fulltext/u2/a398683.pdf>

A. Plan Development Phase

1) Relevance

Example metrics informing the judgment whether a potential plan will be relevant to some anticipated world state:

- **Values and locations of relevant adversary assets** (a plan to invade a country to remove WMD stockpiles would be irrelevant if there were no stockpiles, or if existing stockpiles were unreachable in a timely manner)
- **Number of red operational defensive SAM sites** (a plan that did not either act to reduce this number, or account for blue attrition because of them, would not be relevant)
- **Number of blue re-fueling tankers available during a period** (a plan with more missions than could be supported for refueling would not be relevant)
- **Network of adversary command communications** (a plan that intends to cripple communications by taking out a central node is not relevant if the network is decentralized and/or has alternate paths)

2) Coherence

Example metrics informing the judgment of whether a potential plan will be coherent with other related planning:

- **Rates of attrition of shared assets** (a plan that over-optimistically assumes assets will remain available after another plan executes is not coherent)
- **Times of anticipated/actual actions that are signs of intentions** (a plan that assumes an element of surprise is not coherent with another plan that takes earlier actions that signal a shared or related intent)
- **Intentions of non-military planning in Area of Operations** (a military plan that depends on large-scale destruction of economic infrastructure, apparatus of civil authority, etc. is not coherent with a political plan that seeks to rapidly restore civil rest and order)

3) Planning Assessment

Example metrics informing the assessment of planning performance during plan development:

- **Time required to reach plan execution phase** (compared to predicted, needed, historical, and so on)
- **Number of substantially different COAs and embedded options considered** (based on the assumption that the larger the number of options the better is the understanding of the space of options)
- **Number of relevance and coherence metrics considered** (by some definition of *considered* and a procedure for *counting* separate metrics)
- **Length of review chain prior to approval by Commander** (includes first-pass and re-review cycles)

4) Meta-Metric Learning

Example meta-metrics describing how well the relevance, coherence and planning assessment metrics support plan development, and enable improvement of the metrics – and thus of the total planning process – over time. Meta-metric

learning often requires data over combinations of planning phases.

- **Inter-phase meta-metrics deriving from correlations between some earlier-phase metric with some later-phase metric relating to outcomes** (for example: if the number of options embedded in COAs has historically correlated positively with post-execution assessment metrics indicating greater satisfaction of plan goals, then it may become a more positive metric that is given greater weight in future plans)
- **Correlation between intra-phase metrics generally considered positive (or negative)** (for example: the number of COA options considered is itself to be viewed as a positive metric; but if this number goes up in such a way that the time required to bring a plan to execution goes up at the same time (which is considered negative), then this suggests an optimization is possible, or perhaps a different metric, such as measuring the difference in time between completing a plan and its estimated time of execution rather than total time)
- **Percent of relevance and coherence metrics with measures above a certain level of belief/confidence** (over time, the confidence in metrics should be driven up, for example the confidence in metrics of adversary state such as number of SAM sites should be actively improved with better sensors and analysis processes)
- **Number of corrections made to a metric** ('corrections' means: substantial changes in a metric which are made on the basis of evidence contradictory to the original estimate of what sort of metric would be needed; for example: contradictory evidence that the current WMD estimate, made by whatever process, is wrong leads to improving the process that led to this estimate).

B. Plan Execution

1) Relevance

Metrics informing the judgment whether an actual plan being executed remains relevant to actual conditions, such as constraints and opportunities:

- **Cloud height over intended target** (may violate constraint of target visibility)
- **Number/rate of adversary unit surrenders or other change in adversary offensive activity** (may indicate plan assumptions regarding adversary's will to fight are incorrect or not relevant)
- **Aggregate Measures of Performance (MOPs) for current actions** (low levels of mission performance may indicate that the pre-conditions and contexts for actual actions were not satisfactorily planned – for example low levels of destroy, degrade, deny, disrupt (4Ds) may indicate poor intelligence, weaponeering, etc.)

2) *Coherence*

Example metrics informing the judgment whether a plan remains coherent over time:

- **Changes in planned asset availability** committed by other plans (for example: there are assets which the plan assumes other plans do not require)
- **Success rate of synchronization points** (if plans have explicit specifications of COA relationships, defined-execution windows, handoffs, meetings, supporting events, and so on, then what is the rate at which these relations are successfully maintained?)

3) *Planning Assessment*

Example metrics informing the assessment of whether the plan is being executed satisfactorily:

- **Percent of scheduled missions flown on time** (assessing compliance with plan, not outcomes)
- **Rate COA modifications made per unit time** (a better specified plan might require a lower rate of modifications)
- **Aggregate time delays of actual execution for planned simultaneous actions** (for example in massing fires in planned combined air strike and artillery)
- **Time from a relevant change in world state to the appropriate change in COA** (for example: time from when the new target location information is obtained to time when a new mission tasking has been created that accounts for the new information)

4) *Meta-Metric Learning*

Meta-metrics describing how well the relevance, coherence and planning assessment metrics support plan execution, and enable improvement of the metrics:

- **Inter-phase metric correlation** (for example: low correlation between missions flown on time and post-execution MOE metrics may suggest that flight promptness is not as important as thought, perhaps because late flights were able to act on better, more recent information)
- **Intra-phase metric correlation** (for example: a negative correlation of rate of COA changes and aggregate time delays of planned simultaneous actions may suggest that allowing more frequent COA changes to constructively maintain coherence is beneficial, notwithstanding the expected disruptive effect of the changes; better metrics might distinguish COA changes by class of initiating event, such as new information, command decision, and so on; as the framework itself becomes more sophisticated in its reasoning power, more frequent COA changes will themselves become more easily accommodated by the planning system)

C. *Post Execution*

1) *Relevance*

These are metrics informing the judgment of the effects of the executed plan on world state, particularly relative to intended outcomes. In addition to the more typical post-operations assessment process, there are other ways to conceptualize post-execution relevance. For example: do the lessons drawn from assessment have relevance to the current or future world? Is the originally desired outcome – such as destroying (or building up) another actor’s offensive capability (for example arming the Taliban) of continued relevance? Or is it becoming less relevant, for example because they have changed sides?

- **Number of missiles landing in homeland** (this is said to have been the post-execution operations metric for the recent Gaza invasion)
- **Number of computer systems not patched for exploit X** (exploit X might have worked well on this occasion, but if the adversary has since learned about it and therefore patched the prior vulnerability, the simple assessment that it worked well previously is not particularly relevant for future planning)

2) *Coherence*

Metrics informing the judgment how the net outcome is coherent with other plans (in any phase)

- **Actual asset attrition** (for example: achieving the current plan objective with more or fewer bullets may not matter to the current plan, but it may harm/limit other planning. This is following the notion that Relevance is assessing the relation of the outcome to the *current world state*, so Coherence would be the relation between the outcome and *other plans*.)
- **Degree to which actual net outcome facilitates or limits COAs of future plans** (e.g., confident removal of WMD threat makes other plans easier to develop and execute)

3) *Planning Assessment*

Metrics informing the judgment of how well the post-execution planning process is performed:

- **The number of indicator metrics integrated into the overall goal assessment** (for example: if goal end-state is to influence future behavior, then more indirect present indicators would potentially lead to better inference of future behavior tendencies)
- **The fraction of actually executed missions for which a reliable measure of performance exists** (for how many missions do we have the metrics needed to assess mission performance? for any given mission, how many salient performance metrics are we actually capturing for that mission?)
- **The number of lessons-learned distributed** (clearly depends on how lessons and distribution are counted)

4) *Meta-Metric Learning*

Meta-metrics describing how well the relevance, coherence and planning assessment metrics support plan assessment, and enable improvement of the metrics:

- **Inter-phase correlation** (e.g., correlation of lessons-learned distributed and follow-on planning preparation metrics over time might suggest little relationship between the two. Perhaps the value of the lesson should be included in the metric, or independently, whether the lesson-learned changed any process)
- **Intra-phase correlation** (e.g., no correlation between asset attrition and assessment of satisfaction of goal state suggests that it might valuable to distinguish between “productive” and “unproductive” attrition)

VI. ONTOLOGY-DRIVEN QUERYING OF PLAN INFORMATION

The kinds of representations described above are necessary to support Living Plan requirements. But they are not sufficient. Without the query support to populate them, the representations are vacuous. Since the underlying living plan-related data requires the inference-based identification of objects with associated attribute and location information under conditions of uncertainty, ontology-driven query mechanisms will need to include probabilistic functions in addition to more traditional deductive ones.

Consider the following metric where we have underlined ontology terms to be used by the Living Plan framework:

the percentage of operational anti-aircraft missile sites by area-of-operations for some given plan specification.

Such a metric would be useful in determining the progress of an operational objective for example related to suppression of air defenses. Though seemingly straightforward, even this metric raises a number of interesting semantic challenges that need to be resolved by a query processor.

As stated, the metric is conditioned on a user’s specification of a plan. Given a plan, the metric represents the percentage of operational anti-aircraft missile sites by area-of-operations for the specified plan. The query processor thus needs to be able to ascertain area-of-operations associated with a given plan, something which could possibly vary over time.

A. Indirect identification of plans

Even the identification of the plan may be a non-trivial exercise. While in theory it may be possible to use a unique plan identifier to locate the desired plan, in practice the plan may be identified indirectly in a number of ways, such as:

- **Attributes:** Using combinations of attributes such as *plan phase* (development, execution or post-execution), Commander in charge of plan execution, approval date, and so on.
- **Containment:** Identifying related plans through relations of *containing* or *being contained within* other plans: the AOP (Air Operations Plan) is *contained within* a specified Joint Campaign Plan, or conversely, for a Campaign Plan that contains a specified AOP.
- **Assets:** By relating a plan to the assets associated with it during a given time frame, as when an AOP is tasking Squadron X in some given week.
- **Operational relation:** For example, one plan *precedes* or *succeeds* another as *pre-condition* or *sequel*. Or two plans relate to each by having mutually dependent executions.

ONTOLOGY NAME	VISIBILITY	TERMS	NOTES	REVIEWS	PROJECTS	UPLOADED
Agent Ontology (agent)	Public	986	0	0	0	11/21/2013
AIRS Basic Formal Ontology (airs-bfo)	Public	40	0	0	0	11/21/2013
AIRS Emotion Ontology (emo)	Public	76	0	0	0	11/21/2013
AIRS Mid-Level Ontology (amo)	Public	738	0	0	0	11/21/2013
Artifact Ontology (artifact)	Public	344	0	0	0	11/21/2013
Basic Formal Ontology (BFO)	Public	39	0	0	0	02/04/2013
Bioweapons Ontology (BWO)	Public	145	0	0	0	04/14/2013
Chemical entities of biological interest (CHEBI)	Public	39,433	0	0	0	04/22/2013
Emotion Ontology (MFOEM)	Public	855	0	0	0	04/22/2013
Environment Ontology (ENVO)	Public	1,557	0	0	0	02/02/2013
Event Ontology (event)	Public	409	0	0	0	11/21/2013
Extended Relation Ontology (ero)	Public	0	0	0	0	11/21/2013
Gazetteer (GAZ)	Public	0	0	0	0	02/02/2013
Geospatial Ontology (geo)	Public	302	0	0	0	11/21/2013
Human disease ontology (DOID)	Public	8,671	0	0	0	02/02/2013
Infectious Disease Ontology	Public	508	0	0	0	02/04/2013
Information Entity Ontology (info)	Public	88	0	0	0	11/21/2013

Figure 2: I2WD ontologies at <http://milportal.org>

One or more of these methods could be used in the query to identify the desired plan, requiring the query processor to apply additional knowledge of plan attributes and relations to properly parse the query to eventually locate the desired plan and its area-of-operations.

B. *Ontology-driven queries*

The complexity and dynamic nature of relationships between the plans and the involved information cannot be adequately represented in non-semantic technologies (for example in traditional databases). Moreover, direct traditional querying of such representations will be difficult to automate and maintain in the necessary flexible manner, and the results of such querying may not be capable of the needed rapid update to incorporate emerging important data. Our hypothesis, therefore, which draws on the work described in [7,8] is that a comprehensive and incrementally evolving set of Living Plan ontologies, drawing on the I2WD suite of ontologies (see Figure 2) can provide the needed nuanced representation of the plans, metrics, and of the semantics of the source data against which the querying is performed, while taking account of relationships between all of these components. Such an approach will lay a foundation for sophisticated querying and analytics enhanced by inference, and is designed above all, to enable agile changes to all components. Additionally, the ontology framework will have to include representations of complicating factors such as those described below and their relationships with the plans and metrics.

C. *Probabilistic ontological classifications*

One example complication concerns the identification of the location constraint for those sites that are to be considered because they lie within the area-of-operations. The problem turns on the fact that there may be sites physically outside this area that are identified as harboring capabilities that project into the area-of-operations. This may imply an ambiguity at the operational level. If the focus is on assessing the performance of missions to disrupt or destroy sites physically within the area-of-operations, then the metric should be interpreted in one way. If, on the other hand, the intent of the metric is to assess the security of aircraft within the area-of-operations, then the better interpretation may extend the focus to include sites that have an air defense capability that reaches into the area-of-operations from outside. In order to properly respond to a query based on the latter interpretation, the system would need to be able to infer such projection capabilities and perform spatial reasoning to find substantial intersections with the physical boundaries of the area-of-operations. Such capabilities may depend on the type of missiles available, requiring further information about specific missile capabilities and deployment.

Another potential complication is bias in the identification of individual sites for counting. For example, the adversary might expend additional effort to hide remaining operational sites rather than sites that may have already been degraded in some way. Conversely, missile firings from operational sites make them more difficult to hide. At the same time, own-forces may not expend as much effort in identification and counting of non-operational sites as those which still pose a

threat. In short, the process of identifying and counting sites may be substantially different according to whether they are operational or non-operational. To provide appropriate measures of confidence in the associated metrics, the query processor would have to know what sorts of biases to consider and their relative magnitudes in terms of attributes such as power projection capability, which will be defined in our ontology framework.

A likely more difficult counting complication would arise from semantic assembly of information regarding the very attribute of being *operational* as applied to sites. Whether a site is operational may be difficult to determine for multiple reasons. For example, if a site loses some part of its targeting capacity but retains ability to launch, then it is operational as a launch site, but without targeting it will pose little threat to modern aircraft. The state of the site may also be time-dependent; for example, a site that is partially degraded could be anticipated to be restored at some point in the future. Such expectations would depend on the nature of the degradation and the resources available to make repairs and restore operation. At any particular time, the query processor would have to combine operational state attributes based on reports from different times and with varying levels of confidence arising from uncertainty in expectations as to whether a site will remain operational.

Other complications might arise in classifying a site as functioning or not functioning as an ‘anti-aircraft missile site’. It is certainly possible that the raw intelligence information and sensor data on which counts are made will not directly and unambiguously classify a facility as an anti-aircraft missile site. Instead, there may be reports of a more specific nature (for example, that we are dealing with a specific type of missile capability) which through interaction with weapons ontology would be determined to qualify more generally as ‘anti-aircraft’. On the other hand, some reports may refer only to a ‘missile site’, which would then require further inference to determine if the site is likely to have a more specific type of anti-aircraft capability. Such inferences generally require the knowledge of type-subtype relations and the attributes on which such classifications are based. For example, information about a missile site supertype could be inferred to be also of the anti-aircraft missile site subtype through examination of other potentially known attributes, such as size and location of the site, imagery features, connectivity to other assets, and so on. Such information will be incorporated as probabilistic functions into our ontology framework.

D. *Missing, inconsistent and other invalid data*

Considering the fog of war, some information will at any given stage be incorrect, inconsistent, or missing. Barring independent evidence to the contrary, incorrect information, such as a site being reported as operational that is not, cannot be rectified. However, when there are multiple reports in conflict, it may be possible to reach a most likely conclusion. A query processor that maintains, or has access to, meta-information regarding the typical or historically-observed believability of reports from various sources can combine conflicting reports as weighted evidence to reach a most believable conclusion. The needed provenance-related

attributes, too, will be incorporated into our ontology framework.

A conflict in evidence may be due to understandable reasons, the simplest being that they were made at different times in relation to something that is changing, such as the state of a missile site. A more complex case would involve the ability of different sources to provide substantial evidence at different times or under different circumstances. For example, prior to actually observing an anti-aircraft missile site launch a weapon, a determination of its state of operation may be difficult to establish. An intact-looking site might be non-operational for reasons that are not directly observable, such as broken electronic or computer-based equipment. Under these circumstances, direct observation might provide credible evidence of non-operational status (the physical structure may be visibly degraded or destroyed), without being able to provide evidence of operational status. Intelligence reports from intercepted communications would be a better source of information under these circumstances, but only if they are to be believed as genuine and not intentional misinformation. Of course, direct observation of a successful missile launch at a later point in time would over-rule any prior assertions about the site's state, but only until contravening reports are later received indicating that its state may have changed, such as a battle damage assessment that it was successfully struck and destroyed at an even later point.

Such issues, related to reports of the changing state of a missile site, may be interpreted differently depending on the purpose of the associated metric. If the intent is to assess progress of given actions toward an operational objective of reducing the risk of operations in a given airspace, then the most important information is the conversion through those actions of known operational sites into non-operational sites. In that case, for example, it would be less important to know which sites were non-operational for other reasons prior to the start of the campaign. At the same time, the change in state of a particular site would presumably be the effect of some action taken, and such information would aid in the interpretation of the action reports. For example, if the site were observed to be launching missiles prior to a kinetic strike on the facility and no launches were observed after the strike, it would be reasonable to believe that the strike had its intended effect in rendering the site non-operational. On the other hand, if the metric is being used primarily to ascertain the relative risk of operations in that airspace, then the numbers of operational and non-operational sites prior to the campaign become important, as well as the previously-discussed issue of sites being restored to operation over time.

In addition to incorrect and conflicting information, the query processor must also deal with missing information. In some circumstances reports may be available only for certain time periods, or concerning certain types of information. For example we may have reports on site location without state of operation information, or only assertions of being operational but not of being non-operational. Such differences in missing information will add complexity to making a reliable estimate

of the ratio of operational to non-operational sites over a given area of interest.

VII. CONCLUSION

To support the Secretary of Defense's vision for Living Plans, we believe that plan-related ontologies need to be extended into two areas:

- A generic planning process ontology that is based on the Information Artifact Ontology and that takes into account the cyclical process of planning.
- Ontologies containing representations of each of the kinds of attributes and relations needed to identify desired plans according to relevant areas-of-operations, assets, capabilities, and so forth.

Additionally, the query processing component of any plan-related computational framework that converts potentially huge stores of plan-related expressions (data types, values, natural language expressions), into user-oriented actionable metrics needs to be aware not merely of the ontologies, but also of the needed types of deductive transformations and, as we showed above, of probabilistic classifications. Materialized query processing tools will rely on the principles set forth in [7, 8] which are being used to integrate diverse data in a variety of disciplines. The approach is designed to achieve integration in an agile, flexible and incremental way, and also to incorporate into our system the ontology content created for related purposes by our collaborators in different military domains and disciplines.

REFERENCES

- [1] Mintzberg, H. (1994). The fall and rise of strategic planning. *Harvard Business Review*, 72, 107-114.
- [2] Grant, R. M. (2003). Strategic planning in a turbulent environment: evidence from the oil majors. *Strategic Management Journal*, 24, 491-517.
- [3] Davis, P. K. (2012). Lessons from RAND's Work on Planning Under Uncertainty for National Security. DTIC Document.
- [4] Boukhtouta, A., Bedrouni, J., Berger, J., Bouak, F., and Gulati, S. (2004). *A Survey of Military Planning Systems. International Command and Control Research and Technology Symposium (ICCRTS)*.
- [5] Pollack, M. E. and McCarthy, C. (1999). Towards focused plan monitoring: A technique and an application to mobile robots, *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA99)*, 144-149.
- [6] Boutilier, C. (2000). Approximately optimal monitoring of plan preconditions. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 54-62.
- [7] Smith, B., Malyuta, T., Mandric, W.S., Fu, C., Parent, K., Patel, M. (2012). Horizontal Integration of Warfighter Intelligence Data: A Shared Semantic Resource for the Intelligence Community, *Proceedings of the Conference on Semantic Technology in Intelligence, Defense and Security (STIDS)*, George Mason University, Fairfax, VA, October 23-25 (CEUR 996), 112-119.
- [8] Salmen, D., Malyuta, T., Hansen, A., Cronen, S., Smith B. (2011). Integration of Intelligence Data through Semantic Enhancement, *Proceedings of the Conference on Semantic Technology in Intelligence, Defense and Security (STIDS)*, George Mason University, Fairfax, VA, November 16-17 (CEUR 808), 6-13.