Toward a Living Web of Plans[1]

Brian Kettler[1], Rachel Hingst[1], Gary Edwards[1], and Richard Metzger[2]

Lockheed Martin Advanced Technology Laboratories
4301 N. Fairfax Drive, Ste. 500, Arlington, VA 22203
{brian.p.kettler, rachel.g.hingst, gary.r.edwards}@lmco.com

[2]**Air Force Research Lab**
Information Directorate (AFRL/RIS)
525 Brooks Road, Rome, NY
richard.metzger@us.af.mil

More challenging and dynamic future battlespace environments such Anti-Access/Area Denial (A2AD) require more operationally agile, higher ops tempo processes for planning and execution of synchronized operations beyond today's fairly rigid planning processes (e.g., the Joint Air Tasking Cycle for air operations planning). The modern battlespace increasingly includes multiple military domains (air, land, cyber, etc.), other governmental and nongovernmental organizations, and coalition partner countries. Driven by increasing threats to command and control (C2) infrastructure, new concepts of operations such as distributed control will further fragment the planning enterprise as operational-level planning shifts from a large, physically centralized center (e.g., a Coalition Air Operations Center) to more peer-to-peer planning among multiple, dispersed units. To better integrate diverse plans in various stages of planning/execution across these participants, commanders and planners face challenges in understanding and influencing this more interconnected, distributed planning enterprise; maintaining unity of effort across semi-independently authored plans; and detecting conflicts, gaps, inefficiencies, and opportunities across these moving parts.

Our novel Living Web of Plans vision will help commanders achieve a greater unity of effort out of diverse, decentralized planning through new automation that proactively discovers plans and plan fragments from distributed repositories and nontraditional open sources; links plans by identifying cross-plan relationships; and exploits the resulting dynamic, living "web of plans" to proactively identify synergies, conflicts, and provide planners and commander with visibility into and over an increasingly distributed planning enterprise. This paper describes the vision, and technology enablers towards a Living Web of Plans.

1 Introduction

More challenging and dynamic future battlespace environments such Anti-Access/Area Denial (A2AD) require more agile, higher ops tempo processes for planning and execution of synchronized operations beyond today's fairly rigid planning processes (e.g., the Joint Air Tasking Cycle for air operations planning). The modern battlespace increasingly includes multiple military domains (air, land, cyber, etc.), other governmental and nongovernmental organizations, and coalition partner countries. Driven by increasing threats to command and control (C2)

[1] This work was sponsored by the Air Force Research Laboratory (Contract FA8750-13-C-0113). That work benefited greatly from the operational and technical insights of Mr. Dan Fayette, Mr. Robert McHale, Mr. Kurt Lachevet, and Mr. Richard Fedors at AFRL.

infrastructure, new concepts of operations such as distributed control will further fragment the planning enterprise as *operational*-level planning shifts from a large, physically centralized center (e.g., a Coalition Air Operations Center) to more peer-to-peer planning among multiple, dispersed units (Hostage and Broadwell, 2014). Increasing operational agility[2] and "greater interoperability, transparency, and dynamic command and control will facilitate effective integration and teaming" with performance-optimized teams consisting of partnerships with joint, coalition, and interagency members (US Air Force, 2015).

To better integrate diverse plans in various stages of planning/execution across these participants, commanders and planners face challenges in understanding and influencing this more interconnected, distributed planning enterprise. How does a commander maintain unity of effort from this collection of semi-independently authored plans? How can options be combined into hybrid options earlier in the planning process? How can commanders and planners rapidly detect and resolve conflicts, gaps, inefficiencies, and opportunities across these moving parts?

Our novel **Living Web of Plans (LWOP)** vision will help commanders achieve a greater unity of effort out of diverse, decentralized planning through new automation that proactively discovers, links, and exploits "living" plans. Plans will be found by crawling distributed plan repositories (including less traditional sources such as chat); ingested from diverse formats using text analytics into lightweight, semantically grounded representations; linked to one another in a "web" by identifying cross-plan relationships and dependencies through heuristic reasoning; and exploited by tools such as a Google-like plan search, plan "mash ups" and other visual analytics. These will help commanders maintain "plan awareness" across the ecosystem of plans and proactively discover potential conflicts and opportunities across plans as they are developed. By exploiting nontraditional sources of plan information (e.g. open source intel such as social media), planners can also have extended visibility into the plans of external agencies (e.g. IGOs and NGOs), some of which may be averse to information sharing. A dynamic Web of Plans can be the means by which "pickup teams" of planners are assembled to collaborate on overlapping objectives and areas of interest and by which commanders have "topsight" over planning enterprises. For this vision to be adopted, one cannot assume every planning organization will adopt a common suite of planning tools. "Come as you are" planning must be supported, dealing with artifacts in today's common planning tools such as Microsoft Office, and respecting organizational spans of control.

Towards this vision of more agile, robust, and decentralized (yet coordinated) planning and execution, we have developed a preliminary CONOPS, technology roadmap, and several proof-of-concept tools, drawing upon lessons learned from prior planning automation efforts and exploiting advances in web technologies (including semantic and social web); text analytics, machine reasoning and learning; associate systems, and context-aware computing.

Section 2 presents additional details on the LWOP vision. Section 3 elaborates a concept of operations and illustrates this with an operational scenario. Section 4 presents some candidate technologies toward realizing this vision.

2 The Living Web of Plans Vision

Military planning is more appropriately viewed as "collection of different activities jointly aimed at producing a set of coordinated plans to achieve high-level mission objectives" (Patel et

[2] "Operational agility is the ability to rapidly generate —and shift among —multiple solutions for a given challenge" (US Air Force, 2015).

al., 2010). LWOP embraces this view of an "ecosystem" of plans, semi-independently[3] developed and executed by distributed planners employing heterogeneous planning processes, methods, and tools. LWOP attempts to help planners and commanders bring some unity of purpose and action across these disparate, yet ultimately intertwined activities by recognizing the relationships and interdependencies across plans as they are developed. The Living Web of Plans (LWOP) enables agile, integrated C2 for tomorrow's operational environment by discovering, linking, and exploiting a distributed, web of plans. Some **key challenges of the emerging operational environment** particularly relevant to the LWOP vision include:

- **Increasing proliferation of plans** driven by the presence of multiple organizations including DOD services, other USG (the Intelligence Community, Dept. of State), non-USG (coalition partners, IGOs, NGOs) across multiple functional components (e.g., air, space, cyber) and multiple echelons. Not all of these plans will fall under the command of a single JTF or as part of a single operation. Having multiple planning organizations and plans is driven by the goals of division of labor/expertise, maintaining agility (e.g., planning can proceed asynchronously and plan modifications can be better localized), distinct spans of control, security concerns, etc. Given the breadth of organizations, plans will necessarily be varied and in various stages of asynchronous refinement and execution at any given time, with communication between planning teams often being ad hoc or regimented to specific points in the overall planning processes. 2008)
- **Many different kinds of plans and planning tools** to support the variety of planning organizations, their missions, their planning processes, etc. These include more structured planning tools; semi-structured planning tools and general-purpose tools such as MS Office. This tool diversity is preferable to a one-size-fits-all, monolithic tool, and it is unlikely organizations will want to give up many of their current tools, which may be tailored to their unique planning tasks.
- **With the proliferation of heterogeneous plans, it becomes a challenge for decentralized human or automated planners to find all the plans that may be currently relevant to them.** Some of these plans they will know about: e.g., a strike planner knows there are potentially relevant cyber and space plans that will need to be coordinated with, but the details of such plans may not be fully available yet (the "*known* unknowns"). Some of these plans may be compartmentalized for reasons of security. But there may be a relevant coalition or NGO plan that the planner is completely unaware of the existence of (the "*unknown* unknowns").
- **Commanders will require increased operational agility**. We postulate they will need to be able to form ad hoc, "dream" teams of the best available planners that transcend organizational boundaries to work on specific objectives, including new objectives that are warranted by changes in the battlespace or mission. These teams – each rapidly and proactively working parts of the overall (re)planning problem – will need to interact frequently to share partial plans and other relevant information, collaborating via synchronous and asynchronous mechanisms. Thus dynamic workflows, unlike today's more rigid planning cycles (e.g., the AOC's Air Tasking Order production) and more, ad hoc coordination process (e.g., liaison officers to transmit plan rationale (Allen et al., 2008)), will be required to achieve higher levels of agility. These levels extend from

[3] Various organizations respond to the commander's guidance by independently authoring their parts of the plan. These parts must be later reconciled, deconflicted, etc. Some cross-team interactions happen during this process, facilitated by fixed points in the planning process (e.g., briefings to the commander), liaison officers, etc.

localized adaptation of plans and processes with greater levels of human autonomy to more agile, cross-component collaboration and ultimately to adapt the underlying C2 processes and supporting architectures (systems, comms, etc.) on the fly.

- **Multiple plans will need to be orchestrated to achieve a commander's overall objectives.** Once these moving parts are discovered and developed, a commander must ensure that they can be assembled to achieve his objectives, minimizing any conflicts, gaps, and risks and maximizing synergies, agility, and efficiency (e.g., effort, resource utilization, etc.). As these pieces are coming together, commanders need to get a *holistic* "view" of how they are coming together; where potential issues and opportunities are; the readiness of the composite plan; and how that plan is tracking with the developing battlespace situation, in which the adversary "gets a vote". They also need to see information about the **planning enterprise**: how subordinate planning staffs are spending their limited time: e.g., are they focused on the right problems? Are some problems being ignored? We call this comprehensive plan awareness "**Topsight**".[4] Topsight is the planning analogue of situation awareness.

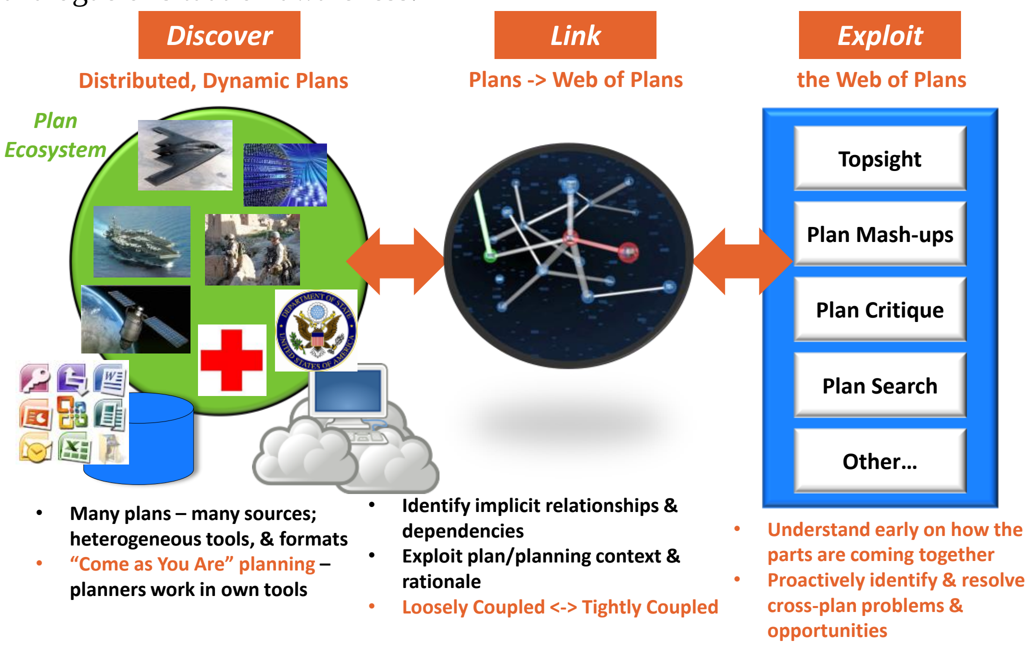

Figure 1. LWOP Vision At-A-Glance

As shown in **Figure 1** the **Living Web of Plans tools and services will:**

- **<u>Discover</u> Plans** (Section 4.3)**:** find plans and plan fragments by discovering or accessing artifacts as they are created by human planners and automated planning tools. These artifacts will be found by "crawling" various repositories including the web, intranets, clouds, etc. Plans may also be submitted to LWOP by planners. Once found, plans will be analyzed, indexed and then exploited by LWOP plan linking tools and, ultimately, by human (or automated) planning tools.
- **<u>Link</u> plans into a dynamic Living Web of Plans** (Section 4.4): once plans have been indexed, they will be linked to the extent possible by applying automated reasoning

[4] "Topsight" was popularized by the DARPA Command Post of the Future program, which developed a next-generation set of information sharing and collaboration tools.

techniques to discover linkages between plans, such as plan interdependencies including temporal, spatial, functional, organizational, and other dependencies. Plan rationale and context will aid inferring these linkages. Linkages are possible at varying levels of granularity ranging from "Plan A is-related-to Plan B" to specific dependencies among tasks or missions across both plans (e.g., a plan for destroying an IADS system may be a prerequisite to a plan involving air strike tasks). We term two plans with more fine-grained linkages as being "*tightly coupled*". LWOP can exploit such linkages when available but can also work with more "*loosely coupled*" plans. The resulting structure of found and interlinked plans is termed a "**living web of plans**".

- **Exploit the Living Web of Plans** (Section 4.5): provide users with comprehensive plan awareness across the planning enterprise ("Topsight"), visualizations of plans and linkages ("plan mash-ups"); analysis/critique of individual and composite plans including detecting and helping to resolve gaps or conflicts between plans; detection and exploitation cross-plan opportunities and synergies; search capability to find relevant plans; plan sharing and other collaboration capabilities.

In general, LWOP brings the "web" to planning to help human planners and commanders understand and coordinate over multiple plans. By "web" we mean a network or graph that links plans – where those links are explicitly specified by planners or inferred by LWOP tools – and "Web" as in information integrated by web-based technologies such as those found on the Worldwide Web. **Table 1** shows capabilities enabled by these technologies (discussed in more detail in Section 4) for both the Web (e.g., the Worldwide Web or its cousins) and the LWOP.

Table 1. "Web" capabilities in the Worldwide Web and envisioned Living Web of Plans

Capability	e.g., The Worldwide Web (today)	The Living Web of Plans (tomorrow)
Distributed Resources	Distributed *web resources* (pages, docs, apps) on the Internet/WWW	Distributed *plans/planning resources* (docs, apps) on Intranets/Clouds/etc.
Discover Distributed Resources	Discovery through web crawling & search (e.g., Google). Search documents by keyword. Ranked search results based on relevance (PageRank, etc.).	Discovery through web crawling/plan submission & search. Search documents by rich content, metadata, context, and plans. Search results ranked via plan importance, timeliness, proximity, etc.
Understand & Index Discovered Resources	Exploit human-generated, simple metadata tags (author, date). Text keyword indexing (fast, simple). Docs are human-readable.	Exploit human-generated metadata tags (author, date, etc). Plans are human-readable, though stylized. Text analytics to generate machine-readable tags (geospatial, temporal, objective, battlespace entity, etc).
Link Resources	Web resources linked explicitly by hyperlinks or implicitly by keyword content.	Explicit and implicit links between plans. Latter extracted via text analytics and machine reasoning.
Machine-understandability (Semantic Web)	Some web resources (e.g., DBpedia & other Linked Open Data resources) have metadata/content described using ontologies.	Some plans have metadata/content described using ontologies.

Collaboration (Social Media)	Social networking, collaborative authoring (e.g., Wikipedia) and sharing (Facebook, Twitter, You Tube, etc.)	Customizable plan mash-ups and other tools supporting synchronous and asynchronous collaboration

Key technical challenges are addressed in more detail in Section 4

3 An Illustrative Scenario

Currently, operational planning is by necessity typically a sequential process. Each of the military services has developed their own methodologies, and in many cases supporting tools, to meet their operational requirements. LWOP will provide Discovery, Linking, and Exploitation of emerging plans across components, revolutionizing the current hierarchical and linear planning process.

From an Air and Space Operations Center (AOC) planner's perspective, for example, LWOP will offer high impact results, without adding additional work, to enhance the current planning process. While AOC planners work to develop friendly courses of action to support commander's objectives, LWOP will continually watch for new planning activity and plan artifacts. LWOP supports "Come as you are" planning, because decentralized plan artifacts being developed will likely be in many formats (Word, Excel, ISPAN, etc.). The Web of Plans is augmented as new plans are found. LWOP can search not only traditional plan sources, but also plans from less traditional organizations (like non-governmental organizations) in a variety of forms – such as social media. Based on the context of the planner, such as target folders being developed or information being accessed, LWOP begins linking relevant elements of other planning activities.

In order to better illustrate the benefits of the Web of Plans approach, we'll start with a short humanitarian assistance sample scenario. Our scenario begins in October 2020 on the fictional island of Pacifica (**Figure 2**) which is comprised of 3 countries. The country of Califon has become increasingly aggressive, threatening to take land belonging to their neighbor Nevidah. Nevidah is an ally of the U.S., but the U.S. does not have a large military presence deployed there. The final country on the island, The Confederation of Washorgon States, has historically remained neutral.

Figure 2. Map of Fictional Island of Pacifica

As island tensions increase, a massive storm strikes, causing catastrophic loss of life and damage in all 3 countries. USPACOM sets up a Joint Task Force (JTF) to initiate humanitarian assistance and security missions. Simultaneously, operators at USSOCOM and USSTRATCOM have been planning operations to intervene in the escalating violence between Califon and Nevadah. Early during plan development 608th planners, using the LWOP tools, are able to rapidly take the plans they have developed using their preferred tools (ISPAN, MS Excel, etc.) and view any potential linkages or dependencies with other plans, being developed in parallel by, for example, USSOCOM, and in-theater Joint Task Force Pacifica.

As new or updated plans are found, LWOP agents begin to link plans to one another, inferring potential plan dependencies and other relationships between plans that may not be explicitly stated in the plans themselves. From an announcement in social media, LWOP finds a reference to a mobile camp that is due to be established in the town of Caliente in several days by Doctors without Borders – a non-governmental organization. LWOP flags the NGO plan for attention by the 608th planner, who views both plans in a mash-up display. Simultaneously, LWOP notifies the U.S. State Department, 624th and other relevant in-theater planners of this potential conflict.

Though these planners are in different time zones, through shared plan representations including plan mash-ups they are able to collaborate asynchronously, essentially crowd-sourcing a revised unified plan within their own ops tempos. LWOP also discovers that a coalition special forces unit has requested ISR support via chat for updated imagery of Caliente, which could indicate intent to conduct operations in the town. This is both a potential conflict and a potential opportunity to have coalition ground forces physically secure dangerous chemicals. For reasons of security, the details of this plan might not be fully visible.

In addition to enabling planners to discover, analyze, and view the impact of plans they may not be aware of, the Living Web of Plans will benefit *commanders* by allowing them to view the dynamic web of plans to understand how plans fit together to achieve key objectives. This ability – or Topsight – enables early, proactive detection of problems such as gaps in plans, conflicts between plans, missed opportunities, and inappropriately allocated resources. Topsight is more than a visualization, it provides the commander with a clear understanding of current blue-force operations AND evolving near and long term plans. Despite a complicated and evolving A2/AD environment, multiple objectives are simultaneously achieved through LWOP enabled technologies.

4 Technology Enablers

This section describes relevant technologies for realizing the Living Web of Plans (LWOP) vision, described in Sections 1-2 and illustrated in Section 3.

4.1 LWOP and Planning Approaches

Prior work in AI planning over the past couple of decades has yielded great progress in representation and automated reasoning techniques. With the exception of automated resource scheduling techniques, the transition of this work into operational planning applications for military users has been slower, despite relatively large government Science and Technology (S&T) investments with programs such as the ARPA-Rome Labs Planning Initiative (ARPI), DARPA Joint Forces Air Component Command (JFACC), DARPA Joint Air/Ground Operations: Unified, Adaptive Replanning (JAGUAR) program and other, more recent efforts. Many (semi-)automated planning tools rely on one or more of the following: (1) formal and deep plan representations; (2) a (logically) centralized, integrated plan (represented in such a representation); (3) standardization on the tools used to create, update, and modify these plans; and (4) deep domain knowledge for decomposing, resourcing, sequencing and deconflicting objectives, tasks, and missions. These approaches have often been unrealistic for real-world planning, especially given the trends in the operational environment as described previously. The LWOP approach has some key differences from this prior work, as summarized in **Table 2**.

Table 2. LWOP in the Context of Prior Planning Work

Prior Planning Tools (1970s-2000s)	LWOP Approach
One plan, known	Many (distributed) plans, some known & some discovered
Fully-automated plan generation	Semi-automated plan critique
Standardized plan authoring tool	Multiple plan authoring tools (including MS Office, etc.), formats – "*Come as You Are Planning*"
Plans must be fully understood by system	Can work with plans not fully understood by system
Unified, deep plan representations	Lightweight representations of varying depth
Little or no exploitation of plan metadata or context	Exploits plans' metadata or context
Deep domain knowledge from SMEs required	Deep domain knowledge *not* required for shallow indexing/linking, but can be exploited if available

Besides a better understanding of the current/future operational needs and the limitations of prior automated planning methods, LWOP is poised to exploit (relatively) recent innovations in web technologies such as scalable indexing methods; cloud-based tools and document storage for plan artifacts; text analytics advances; lightweight ontologies (in RDF/OWL) for metadata/data proven in the semantic web/linked open data communities; the expansion of context-aware computing outside simple, location-based views of context; and scalable big date analytics.

4.2 LWOP Functional Architecture

Figure 3 shows a high level, *functional* system architecture (module view) for the LWOP system, a collection of modules: tools, services, and agents. These modules will be implemented as a set of runtime components, which may have multiple, distributed instances. For example, there may be multiple Plan Crawlers for various clouds or networks. There will be multiple LWOP associates, each associated with a human planner or possibly a plan organization.

The details of modules/components are described in the sections below. In general, the Plan Discoverer's set of Plan Crawlers regularly crawls accessible intranets, clouds, document repositories (e.g., MS SharePoint sites), etc. (Web Resources & Planning Tools). As plans are found and recognized (Plan Detector), they are parsed (via text analytics, etc.), mapped into a common plan representation (CPR[5]), and indexed (Plan Indexer), using its knowledge base (PD KB) of known mappings, planning-related terminology, etc. Plan/planner context is stored where available. It is obtained from user activity and plan metadata (Context Capture). The Plan Indexer also leverages information about known battlespace entities from dynamic feeds (Situation Feeds). A plan's content (in a CPR) and metadata is stored persistently (Plan Index). The Plan Index effectively stores the Living Web of Plans, while maintaining pointers to the distributed web resources where plans were originally found.

[5] CPR is generally used in this document to a "core plan representation". This is used generically. When we are referring to the DARPA Core Plan Representation (CPR) specifically, we will preface it with "DARPA".

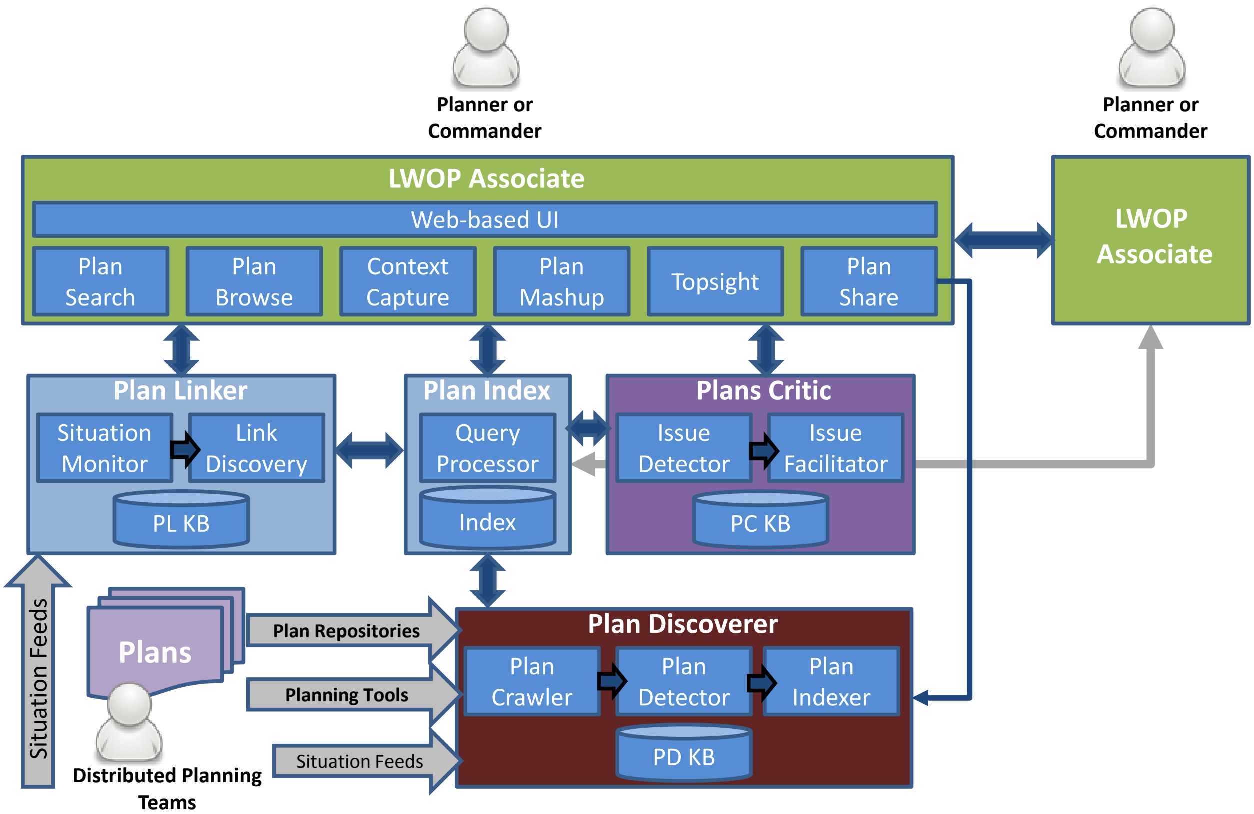

Figure 3. LWOP Functional Architecture (Module View)

The Plan Linker tries to make connections between newly discovered plans (or newly discovered updates to known plans) in the Plan Index using a variety of Link Discovery (heuristic) reasoning techniques and a supporting knowledge base (PL KB). These techniques also leverage information about known battlespace entities from dynamic feeds (Situation Monitor & Feeds). Links discovered, including explicitly specified links between plans and implicit links inferred, are added to the Living Web of Plans in the Plan Index.

The Plans Critic is one component that exploits the Living Web of Plans as maintained in the Plan Index. It attempts to find identify potential conflicts and opportunities between plans by analyzing the content and linkages of plans (Issues Detector), flagging conflicts and opportunities to users (via LWOP Associates), and aiding their resolution (Issue Facilitator). This exploits general and domain knowledge stored in the PC KB.

End users – e.g., planners and commanders – are able to search (Plan Search), browse (Plan Browse); view combinations of plans (Plan Mash-ups); view planning status and gaps/issues between plans (Topsight); and share plans from their desktop or private repositories (Plan Share). The latter results in a plan being submitted directly to the Plan Indexer. This set of capabilities is embedded in an LWOP Associate, a personal, customizable associate system that maintains both an awareness of a user's context (task, mission, etc.) (Context Capture) and is able to tailor its functionality accordingly. There will be many such LWOP associates. Users interact with associates (and all other LWOP functionality) through a web-based user interface. LWOP associates can communicate with one another.

Figure 4 summarizes the key technologies in the LWOP vision, to be described in the sections that follow.

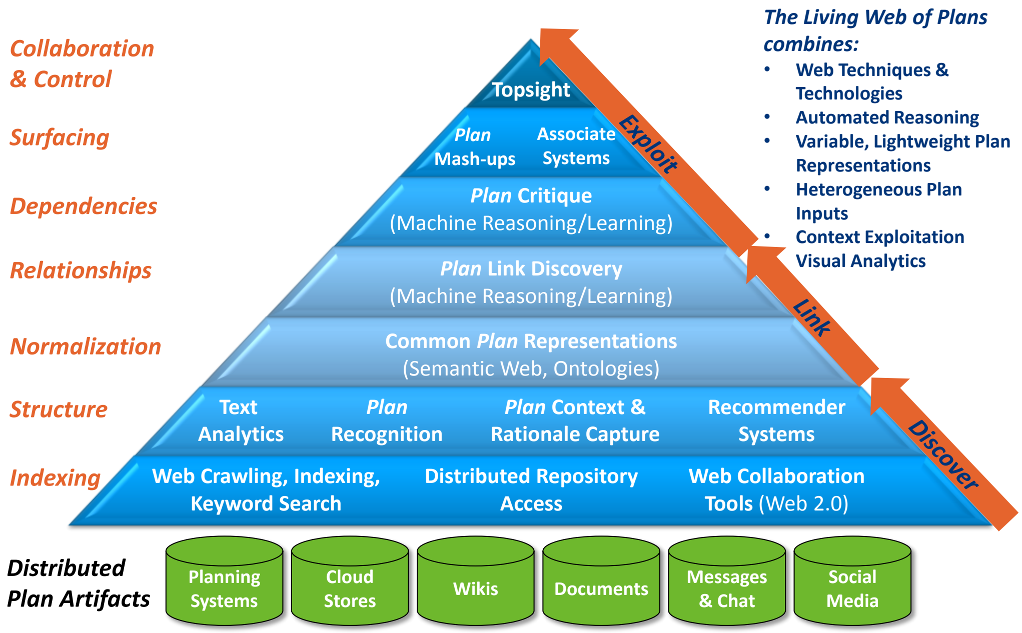

Figure 4. Key LWOP Technologies

4.3 Plan Discovery for the LWOP [*Plan Discover*]

As described in Section 4.2 and illustrated in Section 3, the Plan Discoverer components must find, identify, understand (to some extent), and index distributed plans.

4.3.1 Finding Plans [*Plan Crawler*]

The first task is identifying repositories of potential plans including intranets, clouds, document stores, wikis, and web services that provide access to planning systems and tools. Generally, location and access methods for many of these repositories may be known from prior planning efforts or organizational relationships. But if plan discovery is to take place in the broader Intranet or Worldwide Web – as illustrated in the scenario in Section 3 in which a plan fragment is discovered in social media – the discovery must range more widely. Similarly, discovery might be involve searching things like operator chat, as with the scenario's example of the discovery of the SOF plan. Search would leverage knowledge such as known lists of battlespace entities, missions, resources, and other keywords, as well as geospatial regions (named and unnamed by coordinates) and timeframes.

These repositories must then be crawled (or, more generally, accessed) by LWOP plan crawler agents or otherwise accessed, subject to security, privacy, and other access restrictions (e.g., prohibition of crawlers/spiders on parts of the web). These security issues, while very important, are beyond the scope of this paper. One potential approach would be to expose only the metadata (or part of the metadata) of plans with restricted access. This would let LWOP and its users know about the existence of a plan and whom to contact for further information.

If these repositories contain information besides plans, then plan artifacts must be identified as such. Given the LWOP vision of "come as you are planning", plan identification becomes a challenge as there may be many (known and unknown) plan formats, etc. These could range from highly structured documents (e.g., such as an ATO) or content from a database to less structured content in MS Office documents. Keywords used above to find plans can be used to help identify them, as well as supervised learning techniques trained on a variety of plan examples using content and metadata features.

4.3.2 Exploiting Plan Context

Context is something that can ideally be exploited by the LWOP in building linkages among related plans and recommending relevant plans to users. This has the potential to help when plans are incompletely specified. For example, consider two human planners each viewing the same electronic target folder from a cloud repository. The fact they are looking at the same target might mean their resulting plans may be relevant. This could be determined even if their plans don't yet reference the target. Plan context could include the planner's identity, organization, job function, interests (based on job or recent activity), and current tasking with regards to a specific operation.

Context can also include rationale for plan elements (e.g., specific objectives, tasks, missions, and actions in a plan). This rationale can range from a model of the battlespace situation, the key elements or "levers" therein, and how those are broadly addresses by elements of the operation – its operational design (Schmitt, 2008) – to an understanding of why a specific action is in a plan. Howe et al., (2011) talks about planning-relevant context as including those aspects of a situation relevant to a task/action and the communication showing it. They describe five key dimensions or element types of context including people, events/activities, settings, objects, and communication. They also described a Context Identification Framework that captures these elements and some experiments about how an explicit representation of context can facilitate planning by improving communication in coalition operations.

Context elements need to be represented. There are several frameworks proposed for context in general and plan context in particular (e.g., (Howe et al., 2011)). There is work on representation of plan context – defined as constraints, assumptions, and rationale - using both free text and controlled natural language (see below) (Mott et al., 2010). Rationale may also include traces of human or machine reasoning that help explain why, for example, a particular action is included in a plan. The CPM plan representation, described in Section 4.3.4, provides a means for the representation of plan rationale.

An alternative or addition to tools that implicitly capture user context via observation, desktop or mobile sensors, etc. is to ask the user. One of the concepts of an "associate system" is that a user spends time providing some context to the associate, similar to the way an executive might provide context at the start of the day to their admin assistant about their interests, tasks, priorities, etc. for the day. That effort is ideally repaid by the assistant's (or machine associate) the ability to deliver more relevant information, take appropriate action on behalf of the "executive", etc.

4.3.3 Indexing Plans [*Plan Indexer*]

Once plans – and ideally metadata and plan context – are found they must be indexed for subsequent exploitation, much as Google indexes web pages (and documents) found by web crawlers to enable a search engine to rapidly find web pages relevant to a user's query. In LWOP, plans that are found will be indexed by their content and metadata to support later "queries" by LWOP tools. In the Google search engine case, web pages are indexed by their text content (e.g., an inverted index of keywords) and metadata. A user's query is a set of keywords and the results returned are a ranked list of web pages (hits). In the LWOP case, plans are indexed by their planning elements – plan and situation elements (e.g., entities, actions, geospatial and temporal information – rather simply their keywords.

Determining plan indices requires extracting plan elements from a plan. This in turn may require a variety of techniques depending on the format and representation used in the plan artifact (e.g., a document, web page, etc.) containing the plan. More structured plans from planning tools/services or databases will be easier to parse given the regularities of their content (representation). Plans from semi-structured documents such as Air Tasking Order will involve some natural language parsing, but message formats, for example, are more structured than free text. In the scenario (Section 3), the parsing of the ISR collection require in the chat window between a SOF operator and the ISRD in the AOC involves semi-structured text. An example of this in the scenario (Section 3) is the parsing of social media content such as a tweet or blog to find fragments of a relevant plan (e.g., the Doctors Without Borders posting that a medical treatment camp is going to be set up).

Extracted information can be very important to understanding the context of a plan – who produced it, how recently, etc. This metadata might help infer that a particular plan is supporting another plan due to the organizational relationships among the authoring organizations. For example, social media metadata has proven useful to validate the accuracy of information in blogs and tweets purporting to be factual. A tweet might claim that soldiers are seen in a particular village but that tweet could be suspect if based on geolocation it was determined that the tweet came from a location far from the village.

For semi-structured and unstructured text, a variety of natural language processing techniques can be leveraged for these include named entity recognition, entity extraction, event extraction, and more sophisticated NLP based on statistical and/or linguistic methods. In some cases, one must deal with differences in temporal or spatial granularity in entities, events, or their properties that have been extracted. For example, location might be a country, provenance, city, or even street depending on the specificity of the text. Gazetteer-type information can be helpful in resolving place names. Temporal aspects are more challenging. Ling and Weld (2010) describe some recent work to resolve temporal references in text to specific intervals.

4.3.4 Representing Plans [*Plan Index*]

The content of plans needs to be made machine-understandable for purposes of indexing plans and establishing links between them. Over the past two decades, there have been several attempts at specifying a core or common plan representation (CPR) for military plans. This is in parallel with AI research in domain-independent plan representations (e.g., I-N-OVA (Tate, 2000), including representations of planning domain knowledge (e.g., PDDL). Military plan representations have included the Core Plan Representation (CPR) developed for DARPA in the mid 1990's (Pease & Carrico, 1997). More recently, the DARPA JAGUAR program's plan representation was derived from the DARPA CPR and later automated using the UCore[6] and Air Ops COI (AO COI) (Allen et al., 2009) representations. AFRL's Cornerstone effort has assembled a Unified Plan Representation (UPR) in the Integrated C2 domain from several higher level representations (e.g., UCore) with several domain-specific ones for air, space, and cyber (Pioch et al., 2012; Lebling et al., 2013).

Other recent work includes the Collaborative Planning Model (CPM)[7], explicitly designed to support human planners using multiple tools (e.g., (Patel et al., 2013)). CPM's stated goal is to

[6] http://www.ise.gov/universal-core-ucore. UCore may be superseded by the National Information Exchange Model (NIEM), https://www.niem.gov/Pages/default.aspx.

[7] This work, part of an International Technology Alliance (ITA) program to support coalition operations, was funded by the US Army Research Laboratory, the UK Ministry of Defence, etc.

develop a plan representation language that supports both human understanding and synthetic reasoning agents that automated the more tedious and repetitive tasks that human planners do not want to do (e.g. constraint checking, etc.). This is similar to goal of LWOP, and both visions include dynamic teams continually sharing partial plans in context and collaborating over those plans (Patel et al., 2013; Mott et al., 2010). In fact the CPM work claims to make plans "alive" by digitizing them to facilitate their dynamic use, modification, dissemination, and reuse. This is contrast to today's planning environment which is not fully digitized and can involve exchange of more static artifacts (e.g., documents. Such artifacts are difficult to integrate, modify, and, thus, keep them "alive" and relevant in highly dynamic environments.

A key distinguishing feature of CPM is its sharing of knowledge about the problem-solving process itself including rationale, assumptions, constraints and other information that can provide context to human (and machine) planners trying to understand a plan from another planning team and potential dependencies between plans (Patel et al., 2013). Rationale often includes recurring patterns, which may be discovered for later detection (Mott et al., 2010). In CPM, this contextual information can be specified using either free text or controlled natural language (CNL). A CNL is a controlled vocabulary with a more natural syntax for human input that can be more easily mapped into machine-understandable representations by reducing ambiguity and complexity. For example, Controlled English Common Logic maps from Controlled English (CE), a CNL and subset of English, to Common Logic, which is a formal language supporting machine-reasoning.[8]

In general, an issue of how to best bridge the "semantics" of representations by human planners and automated planning tools remains a topic of active research. A related issue has emerged in recent work on the development of machine-facilitated, human-developed crowdsourced plans (Talamadupula et al., 2013). In such applications, a crowd (which be experts and/or more general users) is enlisted to asynchronous collaborate (via a web site) to develop a plan for a particular objective subject to constraints. Like LWOP, this aims to provide machine-facilitated, collaborative development of plans, primarily by humans, but applying automation to help with the more tedious tasks such as plan critique (constraint checking, etc.).

Plan representations (schemas, taxonomies, ontologies, etc.) differ in several dimensions: (1) level of representation, (2) support for planning and/or execution; and (3) extensibility, semantics, and other properties of the underlying representation language (e.g., UML, XML, OWL/RDF, etc.). Many of these representations actually complement one another. Which representation is "best" is largely a function of how the representation will be used and one size does *not* typically fit all.

Our LWOP approach does say less is (usually) more when it comes to representations. The proposed approach does not require fine-grained representations of plan (and situation) elements (e.g., for automated planning), because we are generally working with more coarse linkages – often inferred – between the "big rocks" of plans (e.g., linking higher level objectives). Rather than design and dictate a universal, comprehensive, "one size fits all" plan representation, the Web of Plans can exploit multiple, lightweight domain representations with a thin layer of "bridging" semantics to capture relationships, dependencies and co-references. Work by the Semantic Web (and Linked Open Data) community provides a model (and supporting languages/tools[9]) for this approach with the mantra "a little semantics goes a long way."[10] This approach is well suited for

[8] http://www.jfsowa.com/clce/specs.htm
[9] e.g., the Resource Description Framework (RDF) language and Web Ontology Language (OWL)
[10] From Sir Tim Berners-Lee, the inventor of the Worldwide Web and the Semantic Web.

this problem because we are seeking to identify and assess relationships amongst multi-domain plan elements and situational elements, as opposed to constructing or repairing the plans themselves.

4.4 Plan Linking for the LWOP [*Plan Linker*]

As described in Section 4.2 and illustrated in Section 3, the Plan Linker components attempt to link plans from the discovery phase into the Living Web of Plans

4.4.1 Link Discovery Methods [*Link Discovery*]

The crux of the Web of Plans is in the linkages that form the web. These linkages will enable a range of exploitation capabilities described in Section 4.5. As **Figure 5** shows, there are a variety of potential linkages among plan and situation elements. An important lesson we have learned from the DARPA JAGUAR program and other related programs is that attempting to model the complete set of possible inter-relationships involved in military operations is daunting, if not unachievable. The LWOP approach does not require the complete modeling of all potential relationships to achieve benefits. These ideally can be inferred (as described below) instead of requiring humans to specify them through existing or new planning tools. If linkages are specified, we will be able to leverage them. For example, a plan may contain an objective that directly references an objective in a higher (or lower) level plan. While these dependencies are more likely to be known to human operators, it is the "horizontal" linkages across (sub)plans that may prove the most useful to discover and exploit. In contrast to automated planning techniques which require pervasive capture of linkages, our focus on "high value" linkages at higher levels of plans will mean fewer linkages need be captured and maintained.

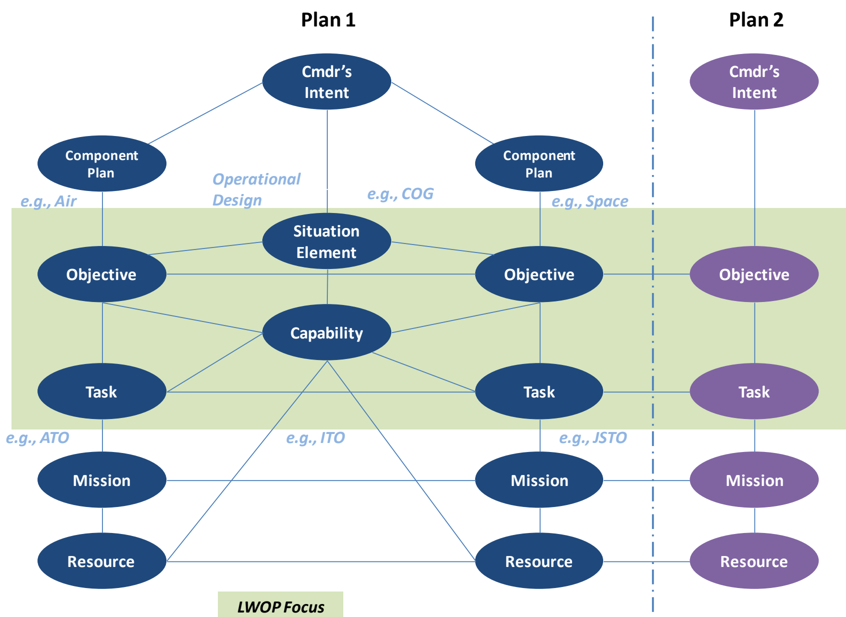

Figure 5. This shows two top-level plans (blue and purple), their elements, and potential linkages (explicitly stated, potentially inferred, etc.). The green shaded area is the area that LWOP where LWOP may have the most impact, as existing tools detect and resolve issues across plans at the lower mission/resources levels.

Linkages can reach into plans to connect plan elements – objectives, tasks, etc. – to recognize plans that are related or interdependent. This is especially useful to help promote awareness across top-level plans that might not normally exist (e.g., the blue and purple plans in **Figure 5**. Such plans, however, might share resources. Recognizing potential interdependencies among plans early, when the plans are in formation at higher levels, means potentially earlier detection of conflicts. Other linkages will connect plan elements to situation elements (e.g., centers of gravity, etc.).

Plan linkages can range from the general (two plans are "related" – e.g., Plan A supports Plan B) to the more specific ones. A useful "granularity" of linkages may be to connect subplans (or plan "legs") to one another. As shown in **Figure 5**, we believe focusing on subplans at the objective/task level is where the sweet spot may lay. At this level, one might connect an air objective to "Interdict threats to ground operations" with a specific area(s) of the battlespace (e.g., kill boxes) with supporting strike and ISR (air and space tasks). Once established, linkages can help identify key dependencies or opportunities: e.g., the addition of SEAD objectives/subplans pertaining to a an emergent threat (e.g., a new IADS component) in one plan might trigger a Cyber planning cell to evaluate potential alternative non-kinetic plans (e.g., addressing the IADs component via a cyber-attack, etc.).

While the web of plans can link plans together at *all* levels from strategic and operational objectives to tactical tasks and missions, linkages at higher levels provide the ability to determine where plan adaptation needs to be done through greater "visibility" into plan interdependencies (including how individual plans support the commander's intent). This enables more surgical ("local") replanning versus "global" plan regeneration and thus results in more effective and efficient handling of a broader range of battlespace changes.

Linkages are not necessarily binary. For example, an element from the commander's intent may be linked to a battlespace situation element and to a plan element or capability. This ternary linkage could capture part of the operational design. A focus on linking capabilities (vs. specific resources) to objectives will also help maintain flexibility during the planning process as plans are developed in parallel and specific resources may not be "bound" to capabilities until later in the process. Linkages may also be "hard" or "soft", representing inflexible or flexible dependencies, respectively. Linkages may have associated metadata: e.g., whether the linkages represent a hard constraint from a stated dependency, a soft constraint, or an inferred dependency, etc.

A key requirement is the "discovery" of candidate linkages as a prerequisite to their exploitation. Some of these linkages will be explicitly stated: e.g., one plan references another plan or another plan's plan element, such as an objective, task, or mission. Explicit relationships might also be stated in plan rationale, as with the CPM work described previously. To the extent that this rationale can be machine-processed, candidate linkages can be mined. More often, linkages need to be inferred since they are implicit. These inferences need not be guaranteed to be correct if there is a mechanism for later human validation of them.

To infer implicit plan linkages among plans, a variety of heuristic techniques to infer the presence of and ideally the types[11] of plan linkages such as the following – *not* meant to be an exhaustive list.

- **Spatio-temporal Correlation**: two plan's plan elements are in spatial or temporal proximity. For example, this could be two objectives associated with the same campaign phase or location; two resources in close proximity at some point in time (e.g., a strike

[11] The type would be one of the properties defined in a CPR that links plan elements. For example, one task/mission could "support" another task/mission. "Support" may be more specific: e.g., "provide fuel for", "provide ISR for", etc.

aircraft and tanker aircraft); etc. This can be done by spatial or temporal distance metrics. LM ATL has developed and applied algorithms for this on past efforts (e.g., (Allen and Dibona, 2004). In this work, plans in a DARPA's Core Plan Representation processed to find candidate relationships among tasks in an Air Tasking Order by categorizing temporal, spatial, and role relationships then applying domain-specific knowledge to heuristically link tasks that might be related.

- **Functional Association**: using domain knowledge, one might know, for example, that a strike mission typically has several associated missions (ISR, SEAD, etc.). This would require capturing of such domain knowledge in machine-understandable ontologies, leveraging the existing work in that area. Existing AI planning models could also be exploited. For example, Hierarchical Task Network (HTN) models for a domain, normally used to generate plans, could be used to "parse" a plan to recover its internal task decomposition structure.
- **Named Entity Correlation:** two plans' plan elements might reference the same named entity (e.g., BE number, location), etc. This can be done using open source text analytics for entity extraction and probabilistic entity (co-reference) resolution techniques. The common reference could be plan entities, as well as situation entities. For example, two plans' plan elements may reference the same high-level objectives.
- **Organizational, Doctrinal and Process Associations**: knowing the organizations behind two plans might give license to infer a particular relationship between those plans. This requires additional organizational and process knowledge. Similarly, being able to connect plan/situation elements to commander's guidance and corresponding doctrine (e.g., via text analysis of those) may also help with plan understanding and inference of linkages.
- **User Context Inference**: by knowing and observing a particular user, one may make inferences. For example, a planner who has two plans open simultaneously on his desktop might be cause to infer some relationship between those plans.

LWOP Associates could be used to present candidate linkages to users to validate LWOP Associates (Section 4.5.2). This could be done at link discovery time, with the benefit of providing feedback useful to improve, for example, machine learning techniques for link discovery. The tradeoff would be user effort to wade through a list of link candidates and, if many were false positives, potentially a loss of trust in the link discovery process. Another, non-mutually exclusive option, would be to mark such linkages as provisional and then take that into account when presenting candidate conflicts and opportunities to the user (Section 4.5.1). Note as plans change, linkages no longer valid will need to be retracted and new ones found.

Link Discovery will require some domain knowledge, though likely much less than required by more automated planning approaches. The Plan Linking Knowledge Base (PL KB) will need to contain knowledge for use by techniques such as those described above. A CPR will provide a taxonomy of plan elements, linkage types, etc. These will be the "vocabulary" of the PL KB. The PL KB will include knowledge of organizations and their planning products and the general relationships among these; knowledge about which plan elements might be related to which other plan elements (e.g., an airstrike may have a supporting SEAD and ISR missions); etc. More general reference knowledge such as situation entities for the current battlespace will be needed. This knowledge will be fairly dynamic and will require tools, processes, and SME support for its ongoing validation and maintenance. Some of this can be automated (e.g., entity extraction).

4.5 LWOP Exploitation [*LWOP Associate*, etc.]

4.5.1 Plans Critique [*Plans Critic*]

As illustrated in the scenario, the LWOP can be used to discover potential gaps, conflicts, and opportunities between plans. This discovery is likely to be mixed-initiative, with automated tools flagging some issues, subject to human validation. For example, the proximity of plan elements from two plans in the same area at the same time does not necessarily constitute a problem or conflict, especially if these plans are designed to be complementary (e.g., an air mission providing close air support to a ground unit). But the presence of a future NGO operation (e.g., the Doctors Without Borders camp) in proximity to the target of an air strike could be cause for concern. Differentiating these automatically generally requires a fair bit of domain knowledge, which in turn raises questions of how such knowledge is acquired, validated, operationalized (represented and applied), and maintained over time. Furthermore, prior work on plan critique may be less applicable given that LWOP will typically only have parts of plans represented in its index (see Section 4.3.3) and those plans may themselves be incomplete as they may be in early stages of refinement.

The **Issue Analyzer** takes plans in the LWOP and interplan linkages, relevant Situation Information (via the **Situation Monitor**), and applies knowledge in the **Plans Critic Knowledge Base (PCKB)** to flag potential conflicts, opportunities, or other issues between plans (e.g., Conflicts might include friendly fire or collateral damage potential; potential violation of ROEs; resource contention; single points of failure or other risks; etc. Opportunities might include reduction in duplication of effort, resource utilization, risk, etc. Because of the combinatorics involved (number of plans, plan elements, kinds of conflicts or opportunities, etc.), some kind of prefiltering or prioritization might be necessary. For example, LWOP is probably not well-suited to finding optimal global resource allocations across an operation's constituent plans. Rather it may be better at finding more localized opportunities or conflicts, perhaps using temporal/spatial constraints to bound the search. Applicable techniques for recognizing plan issues include: rule-based reasoning (where rules capture heuristic information about potential conflicts), case-based reasoning (i.e., finding similarities to prior plans with identified/validated issues), etc. Here some of the prior work on plan critique may be applicable if adapted to work across multiple, incomplete plans. By "incomplete" we include both unfinished plans and plans for which LWOP has an incomplete representation of their existing content.

The **Issue Facilitator** takes candidate issues flagged (and perhaps rank-ordered) and either automatically takes action to resolve them, or, more likely, presents issues to planning team(s) for collaborative resolution, ideally with suggested resolution strategies (which could also be linked to specific issue types or identified by machine reasoning). Since LWOP is primarily addressing cross-plan issues, these modifications may extend outside the scope of a single plan. LWOP collaboration tools such as Plan Mash-ups and Topsight LWOP Associate

An LWOP Associate is an associate system that manages the interaction of a user such as a planner or commander with LWOP capabilities. We use the term "Associate" for intelligent, taskable, personalized assistant software system that complements a specific human user through a rich model of that user's needs, tasks, preferences, etc. used to act to varying degrees of autonomy on behalf of the user. There are multiple LWOP associates, each ideally affiliated with a human planner (or other user) or, at a minimum, a planning team.

LWOP Associates provide a web-based user interface for most user interaction with LWOP including the following capabilities:

- **Plan Search**: allows a user to flexibly query the Plan Index with keyword(s), facets (faceted search), more structured queries, or their current plan (query-by-example). Plan Search returns a ranked list of relevant answers, which may be plan elements or entire plans. Plan Search may use elements of a user's context to refine the query and/or rank the results. Plans may be ranked by relevance based on similarity to a user's context, time/spatial proximity to a user's plan, organizational relationships with user's, etc.
- **Plan Browse**: allows a user to browse the elements in a plan in a variety of views (e.g., map-based, timeline, etc.). These views may include views in a plan's native authoring tool.
- **Plan Share**: allows a user to submit a plan for indexing by the Plan Indexer. This allows a new plan to quickly be "published" versus waiting for it to be discovered later (e.g., by a crawler). A user can submit a plan with metadata and other annotations. The plan can be a draft/partial plan.
- **Alerts**: displays relevant alerts to a user including plan issues (conflicts, opportunities, etc.), new plans of interest recently found, notifications from other LWOP Associates (e.g., those of frequent collaborators), etc. An example of these are shown in **Figure 6.**
- **Context Capture:** this captures information and actions from a user through user input/dialog and observation, respectively (Section 4.3). This information is used to tailor an LWOP Associate's behavior and that of other LWOP tools. A user will be able to view/edit LWOP's model of their context.

4.5.2 Visualization [*Plan Mash-up, Topsight, etc.*]

Users can exploit the LWOP through a variety of interactive visualizations. These views could support asynchronous and synchronous collaboration. Just as web mash-ups combine data from multiple-web sites (e.g., showing for example the locations of Starbucks stores superimposed on a map of subway stops), a Plan Mash-up can present a view of multiple plans via a map, timeline, or other visualization. A user could select plans from this view or drag plans from the Plan Search, Plan Browse, or other elements of an LWOP Associate's web-based UI**.** Plan mash-ups can be interactive with users able to add, delete, edit, and annotate1 LWOP information – such as interplan linkages.

As described above, the Topsight capability provides commanders and their staffs with comprehensive plan awareness, analogous to comprehensive situational awareness (e.g., through various kinds of common operating picture (COP) applications). Topsight provides awareness across the planning enterprise, which includes how Living Web of Plans relates to the dynamic battlespace, commanders' objectives, and planning-related activities of distributed planning teams across many organizations. Topsight is not just visualizations but also includes the analytics that drive those visualizations and the dynamic data feeds that the analytics operate on.

The commander wants to see how the various parts (constituent plans) are coming together (or not) to satisfy his objectives. This applies to commanders at various echelons. Based on the plan mash-up views, they can provide additional information of interest to the commander including plan's status (development, execution, etc.), metadata (who's working the plan), level of effort (how much planning effort or physical resources are associated with the plan), etc. The Topsight views show the commander's objectives, status of plans, the locus of planning effort, plan critic alerts, etc. A notional view is shown in **Figure 6**.

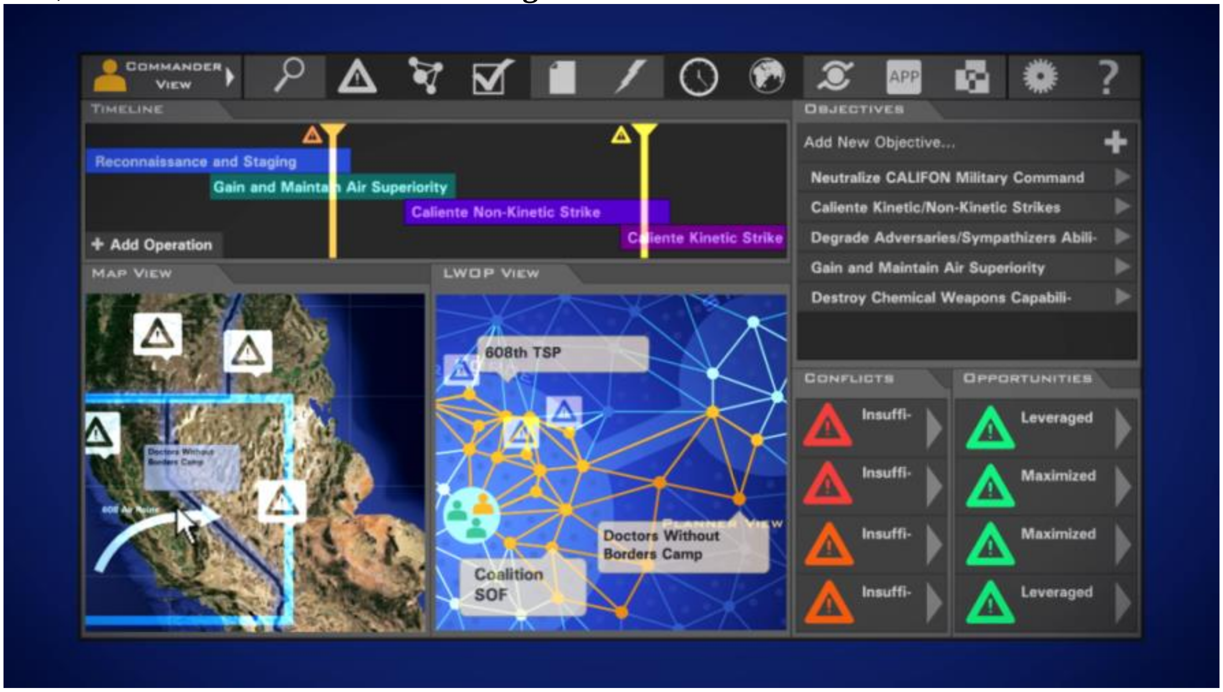

Figure 6. Notional Topsight View for Living Web of Plans

5 Conclusion

Global demographic, economic, technological, and military trends indicate an increasingly competitive and turbulent future. Commanders, backed by innovative command and control systems, will have to plan and execute operations in highly contested environments to win complex, high tempo conflicts. They must successfully integrate and orchestrate plans across all war fighting domains—air, space, cyber, land, and maritime. This requires Living Plans that preserve operational agility in the face of constant change via new planning paradigms and dynamic planning teams.

The Living Web of Plans will make multi-domain synchronization of operations more agile, flexible, and efficient by proactively discovering, linking, and exploiting existing and developing plans. LWOP lets planners use their own tools, yet share evolving plans early in their development. This provides commanders – via Topsight – an understanding of how these plans are coming together to meet their dynamic objectives. LWOP gives human planners greater awareness of relevant plans, fostering the proactive detection of cross-plan conflicts and opportunities. This early awareness enables changes to plans to be made early in the process, when it is less expensive to do so in terms of planning effort expended and perturbations made across multiple plans.

LWOP presents a new take on familiar, though increasingly difficult, planning challenges brought about by more complex operational environments (e.g., A2AD), an increasing number and diversity of players, and greater competition for more limited resources. By eschewing monolithic plans and mandated planning tools, LWOP avoids some of the pitfalls of prior efforts to automate the planning process. It works with plans as it finds them, extracting as much insight as it can from these diverse representations. It leverages plan rationale and context to aid in plan understanding and linking.

Realizing LWOP presents research challenges, exacerbated by the scale and dynamism of the planning enterprise. Recent advances in web technologies, big data stores, machine-understandable plan representations, and text analytics have made the Living Web of Plans feasible.

6 References

Allen, J., and Dibona, P., 2004. "Plan Understanding: Inferring Implicit Dependencies from Explicit Elements in Multi-Agent Plan Representations", 2004 International Multi-Conference in Computer Science & Computer Engineering, Las Vegas, NV, June 2004.

Allen, M.D., Macheret, C., and Malloy, M.A. 2009. C2 Core and UCore Message Design Capstone. MITRE Technical Report, MTR090537.

Allen, J.A., Mott, D., Bahrami, A., Yuan, J., Giammanco, C., and Patel, J. 2008. A Framework for Supporting Human Military Planning. *2nd Annual Conference of the International Technology Alliance*.

DiBona, P., Belov, N., Pawlowski, A., 2006. "Plan Driven Fusion: Shaping the Situational Awareness Process using Empirical Plan Data", *Proceedings of the 9th International Conference on Information Fusion*, July 2006.

Hostage III, G.M. and Broadwell, Jr., L.R. 2014. "Resilient Command and Control: The Need for Distributed Control," *Joint Force Quarterly*, no. 74, pp. 38-43.

Howe, S., Bond, R., Poteet, S.R., Xue, P., and Kao, Anne. 2011. Minimising Cross-Cultural Miscommunication during Coalition Operations. *5th Annual Conference of the International Technology Alliance.*

Lebling, D., Sexton III, W.A., Hunter, D., Macbeth, W., and Liu, W. 2012. Cornerstone: Final Technical Report. Air Force Research Laboratory Information Directorate TR. AFRL-RI-RS-TR-2013-116, May 2013.

Ling, X. and Weld, D.S. Temporal Information Extraction. *Proceedings of the 25th National Conference on Artificial* Intelligence, 2010.

Mott, D., Giammanco, C., Dorneich, M.C., and Patel, J. 2010. Hybrid Rationale and Controlled Natural Language for Shared Understanding. *6th Annual Knowledge Systems for Coalition Operations, Conference*.

Patel, J., Dorneich, M.C., Mott, D., Bahrami, A., and Giammanco, C. 2010. Making Plans Alive. *6th Annual Knowledge Systems for Coalition Operations, Conference*.

Patel, J., Dorneich, M.C., Mott, D., Bahrami, A., and Giammanco, C. 2013. Improving Coalition Planning by Making Plans Alive. IEEE Intelligent Systems, January/February 2013, pp. 1541-1672.

Pease, A. and T. Carrico, 1997. The JTF ATD Core Plan Representation. AAAI Technical Report SS-97-06. http://www.aaai.org/Papers/Symposia/Spring/1997/SS-97-06/SS97-06-012.pdf

Pioch, N. et al., 2012. Cornerstone: Foundational Models and Services for Integrated Battle Planning. *Proceedings of the 17th International Command and Control Research and Technology Symposium (ICCRTS).*

Schmitt, J.F. 2008. A Systemic Concept for Operational Design. Quantico, Virginia: Marine Corps Warfighting Laboratory. http://www.mcwl.usmc.mil/concepts/home.cfm

Smith, D. *Planning as an Iterative Process*. 2012. Proceedings of the American Assoc. for Artificial Intelligence (AAAI) Conference.

Talamadupula, K.; Kambhampati, S.; Hu, Y.; Nguyen, T.; and Zhuo, H.H. 2013. *Herding the Crowd*: Automated Planning for Crowdsourced Planning. In the *First* International *Conference on Human Computation (HCOMP),* 2013.

Tate, A. 2000. <I-N-OVA> and <I-N-CA> - Representing Plans and other Synthesised Artifacts as Sets of Constraints, AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems.

US Air Force, 2014. Air Force Future Operating Concept: A View of the Air Force in 2035. September, 2015.