AFRL-RI-RS-TR-2009-68
Final Technical Report
March 2009

COLLABORATION AND ANALYST SYSTEM EFFECTIVENESS (CASE) CONNECT

General Dynamics Advanced Information Systems

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

AFRL-RI-RS-TR-2009-68 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/	/s/
NANCY A. ROBERTS Work Unit Manager	JOSEPH CAMERA, Chief Information & Intelligence Exploitation Division Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
MAR 09

2. REPORT TYPE
Final

3. DATES COVERED *(From - To)*
Sep 07 – Dec 08

4. TITLE AND SUBTITLE

COLLABORATION AND ANALYST SYSTEM EFFECTIVENESS (CASE) CONNECT

5a. CONTRACT NUMBER
FA8750-07-C-0218

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)

Stephen J. Sickels

5d. PROJECT NUMBER
CASE

5e. TASK NUMBER
BA

5f. WORK UNIT NUMBER
71

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

General Dynamics Advanced Information Systems
10560 Arrowhead Drive
Fairfax, VA 22030-7305

8. PERFORMING ORGANIZATION REPORT NUMBER

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/RIED
525 Brooks Rd.
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TR-2009-68

12. DISTRIBUTION AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 88ABW-2009-0946

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The goal was to develop effective computer-based support for collaborative intelligence analysis, with a particular focus on facilitating tacit collaboration. Our approach can be characterized as a “Develop/Understand/Improve” cycle, where the “Understand” phase is the foundation of our work. We used statistical techniques, such as Latent Dirichlet Allocation & Relational Topic Models—to uncover “hidden” (latent) patterns and relationships that can be exploited to assess the analytic value of information objects and analyst-to-analyst connections. Two conceptual themes that are within the context of our overall “Develop/Understand/Improve” cycle: “Blending the Personal and the Collective” and “Infer, Rather than Enforce, Structure”, drive and are driven by our technology developments and the experimentation we undertake with our technology.

15. SUBJECT TERMS

Tacit collaboration, collaborative intelligence analysis, tags, web 2.0, structured argumentation, shared intelligence

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON** Nancy A. Roberts
a. REPORT U	**b. ABSTRACT** U	**c. THIS PAGE** U	UU	41	**19b. TELEPHONE NUMBER (*Include area code)*** N/A

Table of Contents

Table of Figures

1 Overview

This report provides a summary of the work accomplished by the General Dynamics Advanced Information Systems Team on AFRL Contract FA8750-07-C-0218, "CASE Connect," which began on September 28, 2007. The high-level goal for this effort was to develop effective computer-based support for collaborative intelligence analysis, with a particular focus on facilitating *tacit collaboration*.

Our overall research and development approach on this effort can be characterized as a "Develop / Understand / Improve" cycle, the "Understand" phase of which is the true foundation of our work. To gain the fullest understanding possible, we have *paired technologists and domain experts in experimentation*, which we believe is a truly transformative approach to developing effective computer-based support for intelligence analysis. (We describe this approach in more detail in Section 3, *"Understanding" Through Experimentation*.)

The "Understand" phase of our development cycle has focused on the following three dimensions of computer-based support to intelligence analysis—each of which are of critical importance:

> <u>Collaborative dimension</u>: includes issues such as consensus, diversity, avoiding "groupthink," negotiation, and workflow.
>
> <u>Analytic dimension</u>: includes various aspects of the intelligence analytical cycle, such as collection, collation, evaluation, and reasoning.
>
> <u>Technical dimension</u>: includes issues such as our tools' usability and utility, flexibility, and generalizability—as well as interoperability and support for Web Services-based and SOA-based integrations and "mashups."

Although we have developed a substantial body of *technical* accomplishments, we would like to highlight two *conceptual foundations* that are at the core of our CASE Connect work, which we have developed through our focus on *understanding* collaborative analytic dynamics:

> <u>Blending the Personal and the Collective</u>: Our tools are designed such that an analyst's individual work can "automatically" benefit others. This directly addresses the "zero sum" problem inherent in most collaborative approaches, in which individuals often perceive collaborative work as separate and distinct from their "normal" workflow (and thus as wasting valuable time and effort) and therefore tend to avoid it.
>
> <u>Infer, Rather than Enforce, Structure</u>: Structure is critical to analysis. And, in particular, a structure of *connections* (e.g., concepts to concepts, people to concepts, and people to people via concepts) is, arguably, a necessary ingredient in developing computer-based support to tacit collaboration (which is the CASE program's goal). However, any structural constraint built into a tool—such as enforcing the use of a predetermined taxonomy—may potentially limit the tool's adaptability in the face of new and previously un-envisioned situations and challenges. By comparison, for

example, our tag|Connect tool's open-ended flexibility provides tremendous benefits: users can respond instantly to new situations and requirements and can seize new opportunities to impose innovative order on the information space (such as, for example, the user-driven innovation on Intelink of employing NIPF—or National Intelligence Priorities Framework—identifiers as tags). A key research question driven by this tradeoff, then, is how to best *infer* structure from relatively unconstrained analytic artifacts.

Our approach is to use sophisticated statistical techniques, such as Latent Dirichlet Allocation (LDA)—and, more particularly, Relational Topic Models (RTMs)—to uncover "hidden" (latent) patterns and relationships that can be exploited, ultimately, to assess the analytic value of information objects and analyst-to-analyst connections.

As an aside, we would not claim that less structure is always better! Only that the benefits of a less structured approach are tremendous (in terms of user acceptance, and in terms of adaptability to new demands and the capability to create new opportunities), and that research into mitigating the potential downsides of a less structured approach with sophisticated statistical modeling may well lead to substantial and wide-ranging improvements to the analytic support capabilities that we can provide to the Intelligence Community. Also, it is important to emphasize that a less constrained data set can potentially provide a "richer" network of valuable interconnections (particularly across domains, areas of expertise, and user communities) than can be derived from data that is founded on a more structured approach.

Because we developed these conceptual themes within the context of our overall "Develop / Understand / Improve" cycle, they both drive and are driven by our technology developments and the experimentation we undertake with our technology.

2 Technologies

On the CASE Connect effort, the General Dynamics Team has continued the development of a several tools and technologies that were initiated on a previous contract: tag|Connect, Catalyst, and Context-Grounded Conversations. We have also undertaken work characterizing the underlying structure of the “tag space” using Latent Dirichlet Allocation (LDA)-based Relational Topic Models (RTMs).

We begin with a description of tag|Connect.

2.1 tag|Connect

2.1.1 Overview

tag|Connect, shown in *Figure 1*, is a flexible and easy-to-use Web application that allows analysts and others to organize Web-based resources using "tags" (or keywords) of their own choosing. Analysts can use tag|Connect to quickly assign multiple tags to Web-based resources and can then use an intuitive interface to browse or search those tags and tagged resources.

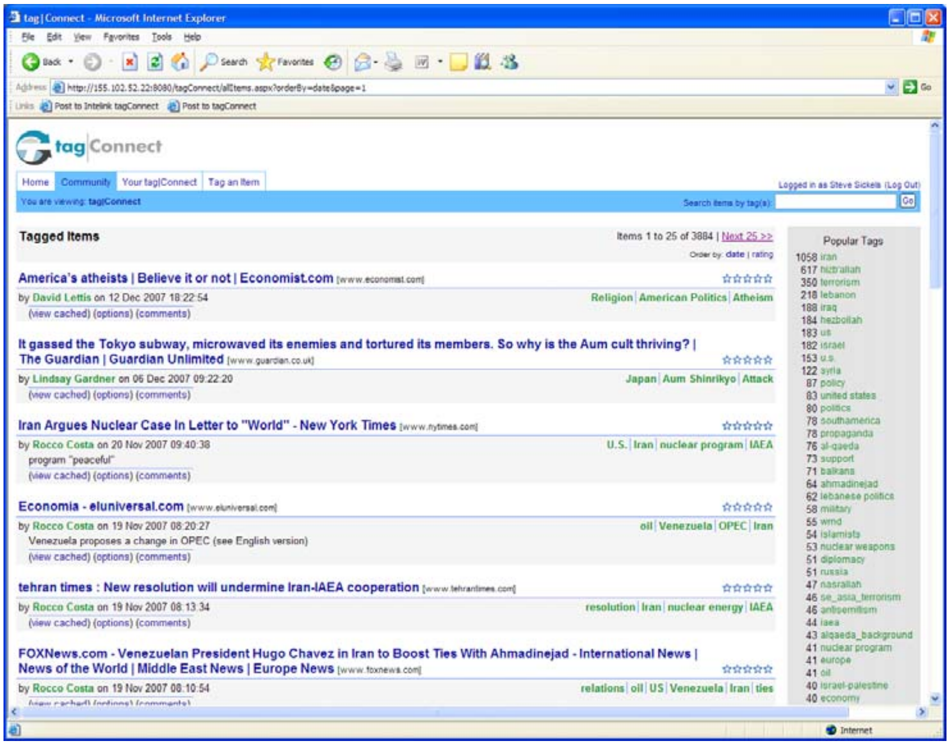
tag|Connect - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://155.102.52.22:8080/tagConnect/allItems.aspx?orderBy=date&page=1
Go
Links Post to Intelink tagConnect Post to tagConnect
tag|Connect
Home Community Your tag|Connect Tag an Item
Logged in as Steve Sickels (Log Out)
You are viewing: tag|Connect
Search items by tag(s): Go
Internet

Tagged Items

Items 1 to 25 of 3884 | Next 25 >>

Order by: date | rating

America's atheists | Believe it or not | Economist.com [www.economist.com] ☆☆☆☆☆
by David Lettis on 12 Dec 2007 18:22:54 Religion | American Politics | Atheism
(view cached) (options) (comments)

It gassed the Tokyo subway, microwaved its enemies and tortured its members. So why is the Aum cult thriving? | The Guardian | Guardian Unlimited [www.guardian.co.uk] ☆☆☆☆☆
by Lindsay Gardner on 06 Dec 2007 09:22:20 Japan | Aum Shinrikyo | Attack
(view cached) (options) (comments)

Iran Argues Nuclear Case In Letter to "World" - New York Times [www.nytimes.com] ☆☆☆☆☆
by Rocco Costa on 20 Nov 2007 09:40:38 U.S. | Iran | nuclear program | IAEA
program "peaceful"
(view cached) (options) (comments)

Economia - eluniversal.com [www.eluniversal.com] ☆☆☆☆☆
by Rocco Costa on 19 Nov 2007 08:20:27 oil | Venezuela | OPEC | Iran
Venezuela proposes a change in OPEC (see English version)
(view cached) (options) (comments)

tehran times : New resolution will undermine Iran-IAEA cooperation [www.tehrantimes.com] ☆☆☆☆☆
by Rocco Costa on 19 Nov 2007 08:13:34 resolution | Iran | nuclear energy | IAEA
(view cached) (options) (comments)

FOXNews.com - Venezuelan President Hugo Chavez in Iran to Boost Ties With Ahmadinejad - International News | News of the World | Middle East News | Europe News [www.foxnews.com] ☆☆☆☆☆
by Rocco Costa on 19 Nov 2007 08:10:54 relations | oil | US | Venezuela | Iran | ties
(view cached) (options) (comments)

Popular Tags

1058	iran
617	hizb'allah
350	terrorism
218	lebanon
188	iraq
184	hezbollah
183	us
182	israel
153	u.s.
122	syria
87	policy
83	united states
80	politics
78	southamerica
78	propaganda
76	al-qaeda
73	support
71	balkans
64	ahmadinejad
62	lebanese politics
58	military
55	wmd
54	islamists
53	nuclear weapons
51	diplomacy
51	russia
47	nasrallah
46	se_asia_terrorism
46	antisemitism
44	iaea
43	alqaeda_background
41	nuclear program
41	europe
41	oil
40	israel-palestine
40	economy

Figure 1. tag|Connect

Because tags are shared (unlike traditional browser bookmarks), tag|Connect also provides tremendous *collective* benefits by giving analysts the opportunity to view and leverage the work of others. For example, analysts can view other analysts' tag|Connect tags and tagged resources to find additional key resources relevant to their own work. And they can quickly understand which topics and analytic subjects are of broad or specific interest among peers and colleagues in the Intelligence Community.

Tag|Connect has a flexible user interface that allows analysts to quickly and easily leverage the "tag space" in a variety of ways. In general, it serves as both a personal and collective organizational tool. And it allows its users to seamlessly and easily pivot back and forth between these two uses. For example, an analyst can see all the items that she has tagged with both "Iraq" and "Shiite," (the personal organizational function) and can then easily “pivot” to see all items that *everyone* has tagged with those tags (the collective organizational function). From there, it’s easy for her to see all the other tags that have been applied by others to items tagged with “Iraq” and “Shiite,” or all people who have used those tags (sorted by the number of times the tags were used). By these “pivots,” an analyst at the CIA might discover, for example, valuable resources tagged by someone at the State Department, which she might not have found otherwise. Or she might find key related topics (via related tags). Or, she might discover other analysts (potentially at other agencies) with whom to collaborate. Any of which might lead to other forms of cross-agency collaboration.

Users can optionally add comments to tagged resources, which is especially useful when sharing with others.

Most recently, tag|Connect v0.9.7, released in September, 2008, includes the following new features:

- Tag clouds: Tag clouds depict, via the tags’ font sizes, the number of times that various tags have been applied across the tag|Connect user community. Users can easily toggle between a tag cloud view and a list view. (Tag clouds are available in a number of contexts in the tag|Connect UI). Tag clouds are a popular feature in a variety of tagging applications (including del.icio.us on the open Internet), and are a feature that the Intelink user community has expressed great interest in.
- Sort by alpha / count: Tag lists and tag clouds can be toggled between sort-by-alphabetical order and sort-by-count.
- Filters: Users can easily filter tag clouds and lists to display only those that have been used at least 1 (or 2, or 5) times. As with the other features listed above, these filters can easily be toggled on or off.

tag|Connect is a powerful tool for organizing resources and for discovering the resources and topics of interest to and in use by others. As such, it allows analysts to:

- Organize resources according to individual preferences, which optimizes recall and retrieval; and
- Leverage the collective organizational space, which provides diversity and breadth and relevant resources.

tag|Connect is currently deployed as a “Core Service” on the three Intelink networks: JWICS, SIPRNet, and Intelink-U. *Figure 2* shows tag|Connect on Intelink-U.

tag|Connect has also been deployed by the CIA’s Open Source Works (OSW), and is a component of the ODNI’s “A-Space” analytic environment. On the three Intelink networks, tag|Connect is currently utilized by thousands of users, with its popularity growing rapidly.

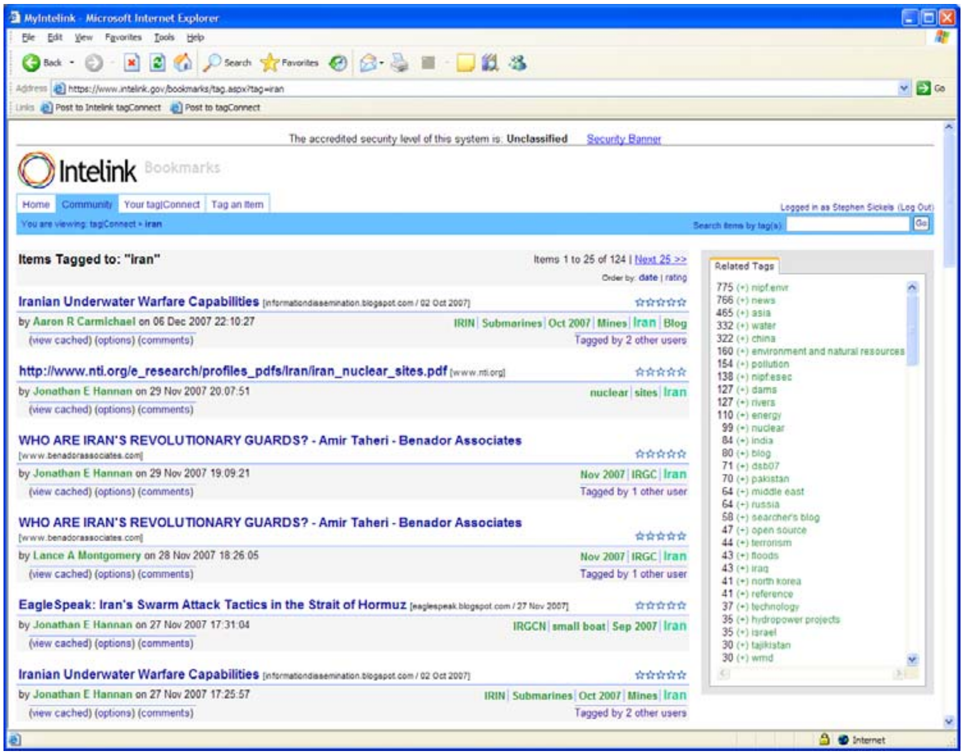
MyIntelink - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address https://www.intelink.gov/bookmarks/tag.aspx?tag=iran
Go
Links Post to Intelink tagConnect Post to tagConnect
Intelink Bookmarks
Home Community Your tag|Connect Tag an Item
Search items by tag(s): Go
Internet

The accredited security level of this system is: **Unclassified** Security Banner

Logged in as Stephen Sickels (Log Out)

You are viewing: tag|Connect » iran

Items Tagged to: "iran"

Items 1 to 25 of 124 | Next 25 >>

Order by: date | rating

Iranian Underwater Warfare Capabilities [informationdissemination.blogspot.com / 02 Oct 2007]
by Aaron R Carmichael on 06 Dec 2007 22:10:27 — IRIN | Submarines | Oct 2007 | Mines | Iran | Blog
(view cached) (options) (comments) — Tagged by 2 other users

http://www.nti.org/e_research/profiles_pdfs/Iran/iran_nuclear_sites.pdf [www.nti.org]
by Jonathan E Hannan on 29 Nov 2007 20:07:51 — nuclear | sites | Iran
(view cached) (options) (comments)

WHO ARE IRAN'S REVOLUTIONARY GUARDS? - Amir Taheri - Benador Associates [www.benadorassociates.com]
by Jonathan E Hannan on 29 Nov 2007 19:09:21 — Nov 2007 | IRGC | Iran
(view cached) (options) (comments) — Tagged by 1 other user

WHO ARE IRAN'S REVOLUTIONARY GUARDS? - Amir Taheri - Benador Associates [www.benadorassociates.com]
by Lance A Montgomery on 28 Nov 2007 18:26:05 — Nov 2007 | IRGC | Iran
(view cached) (options) (comments) — Tagged by 1 other user

EagleSpeak: Iran's Swarm Attack Tactics in the Strait of Hormuz [eaglespeak.blogspot.com / 27 Nov 2007]
by Jonathan E Hannan on 27 Nov 2007 17:31:04 — IRGCN | small boat | Sep 2007 | Iran
(view cached) (options) (comments)

Iranian Underwater Warfare Capabilities [informationdissemination.blogspot.com / 02 Oct 2007]
by Jonathan E Hannan on 27 Nov 2007 17:25:57 — IRIN | Submarines | Oct 2007 | Mines | Iran
(view cached) (options) (comments) — Tagged by 2 other users

Related Tags

775	(+) nipf.envr
766	(+) news
465	(+) asia
332	(+) water
322	(+) china
160	(+) environment and natural resources
154	(+) pollution
138	(+) nipf.esec
127	(+) dams
127	(+) rivers
110	(+) energy
99	(+) nuclear
84	(+) india
80	(+) blog
71	(+) dsb07
70	(+) pakistan
64	(+) middle east
64	(+) russia
58	(+) searcher's blog
47	(+) open source
44	(+) terrorism
43	(+) floods
43	(+) iraq
41	(+) north korea
41	(+) reference
37	(+) technology
35	(+) hydropower projects
35	(+) israel
30	(+) tajikistan
30	(+) wmd

Figure 2. tag|Connect on Intelink-U

2.1.2 Deployment Support

The General Dynamics Team has worked closely with the ODNI's Intelligence Community Enterprise Services (ICES) office, which is responsible for running and maintaining the Intelink networks, to ensure that tag|Connect meets ICES requirements and the requirements of the Intelink user community.

During our initial coordination meetings with ICES[1] in June and July of 2006, we jointly took the decision that the General Dynamics Team would be responsible for any "hardening" necessary to support a large-scale (world-wide) deployment as well as feature upgrades (and bug fixes) following the initial deployment.

On 2 October 2006, we delivered tag|Connect v0.5.0 to ICES. On 30 October 2006, ICES opened up tag|Connect for a limited beta test for small set of Intelink users. Then, on 5 February 2007, tag|Connect went "live" as a Core Service, available to all Intelink users on all three networks, worldwide.

[1] At that time ICES operated as the Intelink Management Office, or IMO.

The General Dynamics Team has also supported a deployment of tag|Connect to the CIA's Open Source Works (OSW), which necessitated a number of modifications to support their login/authentication approach, as well as to introduce new features as requested, such as RSS 2.0 feeds and better Unicode character set support. Our first delivery to OSW was v0.9.0 on 28 September 2007.

tag|Connect has also been incorporated into the ODNI's "A-Space" analytic environment, which went live in September, 2008.

Finally, on 23 September 2008 we provided an "as is" copy of tag|Connect to the Navy's SPAWAR, which is embracing Web 2.0 technologies, and which actively sought out the opportunity to deploy the tag|Connect social bookmarking application as one of their key Web 2.0 tools.

2.1.3 Web Services

The General Dynamics Team has architected tag|Connect to have a clean and well-defined separation between its underlying "engine" (and "business logic") and the Web-based user interface (UI) that exposes its functionality to users. Although the tag|Connect UI supports users in easily navigating and exploring the underlying "tag space," there is no reason why tags need to be exposed *only* through the tag|Connect UI. There is, in particular, great potential advantage in tags being exposed to and by *other* application UIs.

For example, the developers of any application that presents URL-based content to users could create additional UI functionality that would allow users to tag that content from *within that application*. Or perhaps to see what tags others have been applied previously. Although these enhancements would require that the application's developers implement new UI functionality, there would be no need for the developers to re-implement a tagging "engine" or database to manage the tag data. Instead these applications could access tag|Connect's underlying engine (which provides the capabilities to create and access tag data) via program-to-program calls to tag|Connect's Web Services API. An additional benefit in this approach—over and above saving the application's developers from having to re-implement a tagging engine—is that any tags applied to resources from within the new application's context are also "automatically" available from within the tag|Connect UI. And, conversely, resources listed or depicted in the new application can show tags previously applied from within tag|Connect as well.

To facilitate exposing tag data within other application contexts, we have developed an extensive Web Services API for tag|Connect. There are three different “levels” of calls that can be made against the API:

1. “Full” calls which are the most complex to construct, but are also the most powerful and flexible. These calls correspond roughly to those made programmatically from within tag|Connect against its underlying “engine.”
2. “Simple” calls which hide some of the complexity of the full calls at the cost of some flexibility and functionality.
3. “REST” calls, which are the easiest for tech-savvy non-programmers to use, and thus are the data sharing mechanism that enables many “mashups” on the Internet.

After we presented our “Full” and “Simple” SOAP-based Web Services APIs at the “tag|Connect Mashup Workshop” that we hosted in the fall of 2007 (and which was attended by people from across the Intelink user community), it became clear to us that the community was interested principally in our REST APIs, which we had then just begun to develop. Since that workshop, we have focused our Web Services development efforts on the REST API.

Currently, tag|Connect’s REST API is read only, but the essential infrastructure is in place to allow tags to be *applied* through the REST API (vs. only through the tag|Connect UI, as is currently the case). However, in order for this to be implemented, ICES will need to make some modifications on their end regarding how authentication is handled.

2.2 Catalyst

2.2.1 Overview

General Dynamics' Catalyst, shown in *Figure* 3, is a flexible application that allows analysts—individually and collaboratively—to represent and analyze a wide range of intelligence issues, problems, and challenges. Catalyst supports a variety of analytic activities: it is an effective tool for structuring information about a topic and it is also an effective tool for conducting an analysis, evaluation, or assessment.

A Catalyst *model* consists of "nodes" containing text that are organized into hierarchical tree structures. Catalyst's "trees" support natural decomposition of problems into sub-problems, systems into subsystems, and issues into sub-issues. Catalyst also supports links: any URL can be easily "dragged-and-dropped" from a Web browser into a Catalyst node, automatically creating a hyperlink within the node and thus allowing relevant Web-based resources and information to be organized directly within the context of the issue, problem, challenge, or system being considered.

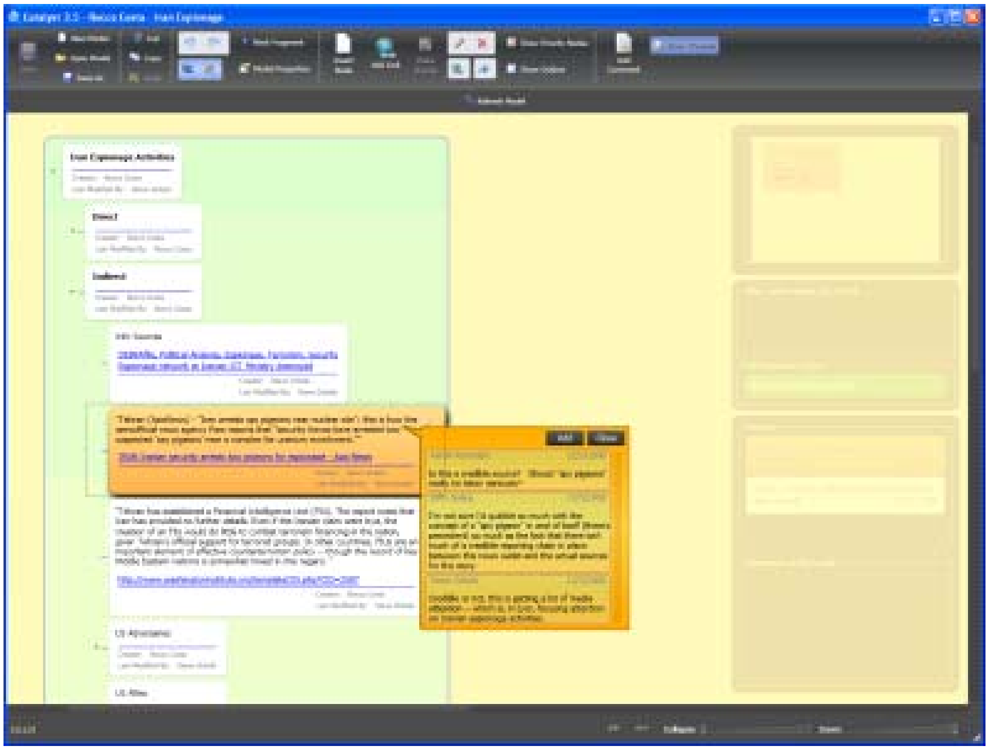

Figure 3. Catalyst

In addition to providing a flexible organizational and analytic framework for *individual* analysts at work, Catalyst also seamlessly extends analysis into the *collective* realm via a number of powerful collaborative features, beginning with access: as a framework for collaborative work, Catalyst is a client-server application that can store multiple Catalyst models created by multiple individuals on a single server. Any individual with the Catalyst client on his or her computer and with log-in access to the server is able to open and view anyone else's Catalyst models. Not only can analysts view the informational and analytic structures created by others in Catalyst, they can also "adopt" (copy and paste) each other's work—that is, node structures and associated links—directly into their own models, with full attribution. Catalyst thereby enables analysts to *augment* and *extend* their own work with relevant portions of the work of others.

Finally, any individual with the Catalyst client on his or her computer and with log-in access to the server can edit or update any model,[2] thus allowing not only "individual" analytics, but also shared and collaborative analytics as well. Each Catalyst node includes attribution, showing who created the node, and who edited it last.

In summary:

- Catalyst provides a framework that supports the analytic decision-making process.
- Catalyst enables analysts to develop problem- or task-specific organizational structures—"outlines"—that can be used for organizing existing, incoming, or discovered data and information. These structures can provide value not only to their creators, but also to others, as well (including via Catalyst's support for "adopting" content from one model to another).
- Catalyst supports a variety of workflows: analysts can develop and refine their structures in response to data and information that becomes available through an ongoing analysis, or they can develop their structures in advance to document a set of plausible alternatives and to provide a structure for organizing and assimilating incoming or discovered information relevant to those alternatives.
- Catalyst is both an "individual" and a "collaborative" tool.

2.2.2 Features and Functions

Analysts can use Catalyst to organize information into "nodes" of text that are joined via parent-child relationships into tree structures. A "root node" at the top of the tree generally represents the overall question, hypothesis, or subject being addressed.

The text in a Catalyst node can be short or long, ranging from catch-phrases and headings to multiple paragraphs of text quoting from source materials or laying out the details of an analytic argument.

[2] Previous versions of Catalyst had both "individual" models (which could be updated only by the model's creator) and "collective" models (which could be updated by anyone). In the current version of Catalyst, all models are "collective."

A single Catalyst "model" (that is, the units of data and information that are accessible via Catalyst's "New Model" and "Open Model" commands) can contain any number of trees, which can be arranged by users on the Catalyst "canvas." New nodes can easily be added as children of existing nodes (by a number of mechanisms, the simplest of which is probably to right click on a node and to select the "Add child" option). Also, new freestanding nodes can easily be added by double-clicking on the canvas where the new node is to be positioned. Any new (or old) node can then be dragged into an existing tree or can be left to stand alone—perhaps serving as the root node of a new tree.

Catalyst makes extensive use of a "drag-and-drop" mode of interaction for nodes and trees:

- A tree or single unattached ("free") node can be easily dragged and repositioned on the canvas.
- A sub tree can be detached from its parent node and dragged to a new location, including being dragged and "attached" to any other node, which then becomes its new parent.
- The nodes *within* a tree can be rearranged and reordered.
- "Links" (hyperlinks to Web resources) can be added to a node by drag-and-drop from any Web browser's address bar (by grabbing and dragging the small icon to the left of the URL in the browser's address bar). A node's links show up below the node's text as clickable hyperlinks.

(Any operation that can be accomplished via drag-and-drop can also be accomplished by copy-and-paste.)

Catalyst has another particularly powerful drag-and-drop feature: text "snippets" from a Web page can be highlighted and dragged onto the Catalyst palette, which then automatically creates a new node with that text in quotes. And, Catalyst captures the source URL and its title, which is automatically added as a link within the new node. This is shown in *Figure 4*: a portion of text from a Web page is shown in the upper left, and the Catalyst node that results when the text is highlighted and dragged into the Catalyst canvas is shown in the lower right. Note the blue hyperlink below the quoted text in the Catalyst node; this automatically-created hyperlink points back the source Web page from which the text was dragged.

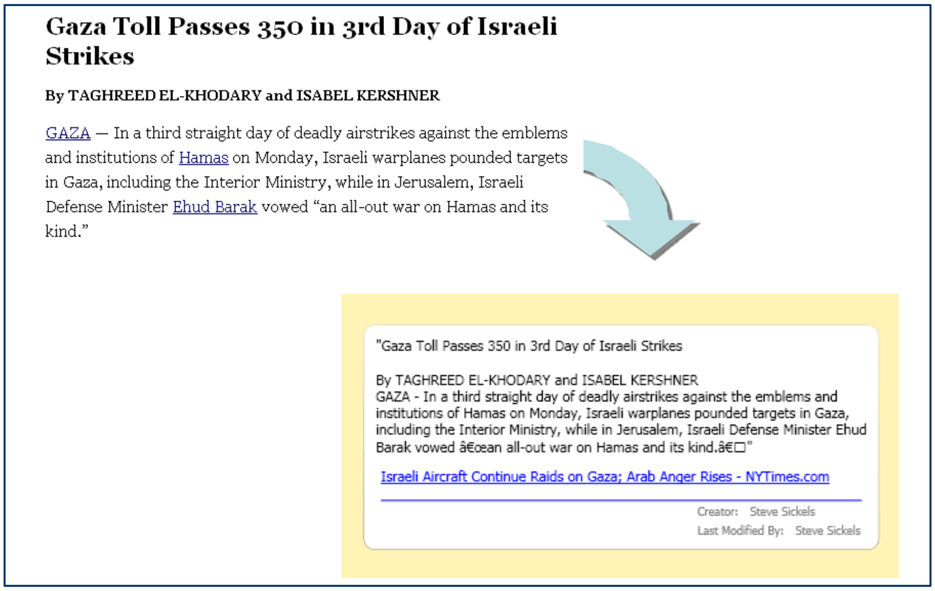

Gaza Toll Passes 350 in 3rd Day of Israeli Strikes

By TAGHREED EL-KHODARY and ISABEL KERSHNER

GAZA — In a third straight day of deadly airstrikes against the emblems and institutions of Hamas on Monday, Israeli warplanes pounded targets in Gaza, including the Interior Ministry, while in Jerusalem, Israeli Defense Minister Ehud Barak vowed “an all-out war on Hamas and its kind.”

"Gaza Toll Passes 350 in 3rd Day of Israeli Strikes

By TAGHREED EL-KHODARY and ISABEL KERSHNER
GAZA - In a third straight day of deadly airstrikes against the emblems and institutions of Hamas on Monday, Israeli warplanes pounded targets in Gaza, including the Interior Ministry, while in Jerusalem, Israeli Defense Minister Ehud Barak vowed â€œan all-out war on Hamas and its kind.â€□"

Israeli Aircraft Continue Raids on Gaza; Arab Anger Rises - NYTimes.com

Creator: Steve Sickels
Last Modified By: Steve Sickels

Figure 4. Catalyst's drag and drop of Web text

This feature is designed to make the mechanics of Web-based research (including on intranets such as Intelink) particularly fast and easy. Relevant information can quickly be added to a Catalyst model, with the source of that information automatically captured and available for easy access at any time in the future.

And finally, a key Catalyst feature is that individuals can “adopt” (via copy and paste) nodes from one model into any other model. If the copied node is pasted “onto” a node in the target model, it becomes a child of that node. If, on the other hand, it is pasted elsewhere on the canvas, it is copied as a “free node” to that new location.

As with all drag-and-drop and cut-and-paste operations in Catalyst, if the copied node has child nodes, the child nodes are automatically copied as well.

Catalyst provides a number of other features:

- Collapse/Expand: Catalyst includes a number of different mechanisms for controlling how much of and what portions of a model are in view at any time. There are collapse/expand points on each node that has children, identical to the directory trees shown when “exploring” directories in Microsoft Windows. There is also a slider control that allows all trees in a model to be expanded or collapsed incrementally.
- Attribution: Catalyst captures attribution for all nodes: all Catalyst nodes indicate the name of their creator. If someone adopts material from someone else’s model, then that material is forever attributed to its original creator. All Catalyst nodes also indicate by whom they were last modified. (The Catalyst database actually stores the full and complete history of every node, including all edits and who made them. However, currently in the Catalyst user interface we are presenting only each node’s present state, its creator, and who last modified it.)

- Outline Heading Nodes: Nodes in a Catalyst tree can be designated as “title nodes.” When a node is designated as a title node, two things happen: 1) its text is bolded; and 2) an entry for the node appears in the Model Outline, described below.
- Model Outline: The Model Outline is a separate summary view that is available for Catalyst models. It appears as a small window in the upper right corner of the Catalyst application, and is intended to provide an overview of the key topic nodes in a Catalyst model—that is, the nodes that a user has designated as “title nodes” (as described above). The model outline also facilitates navigation of the model. For example, selecting a node in the model outline causes the model on the main canvas to expand to show that node, if it had been hidden.

Several of Catalyst 3.0’s features are not available in the current version of Catalyst, Catalyst 3.5. (As we describe in Section 2.2.5, “*Moving Catalyst to the Web*,” we began developing Catalyst 3.5 only recently, and our focus to date has been establishing a lightweight Catalyst client with the core Catalyst functionality.) We hope to add these features in to Catalyst 3.5 when funding is available.

- “Copy as Wiki Markup” Function: Catalyst 3.0 provides a convenient mechanism for getting information from Catalyst into a Media Wiki-based wiki (such as Intellipedia). An analyst can right click on any node and can then select the “Copy as Wiki Markup” option, which copies that node and all its children to the clipboard. However, it copies the text in a special wiki “markup” format. In particular, when the copied Catalyst nodes are pasted into a wiki’s edit page and then viewed:
  - Title nodes in Catalyst become section headings in the wiki. And, furthermore, the “level” of the section headings corresponds to the indented structure of the title nodes in the Catalyst model.
  - Attachments in Catalyst automatically become proper references in the wiki (as in Intellipedia and Wikipedia). In particular, the titles of all attachments are listed below a “References” section header that appears at the end of the wiki entry, with reference numbers appearing in the body of the text (immediately following the text that came from each Catalyst node that contained an attachment).
- Full History: Catalyst maintains a full history of all model states as of each time any model was saved.[3] In Catalyst 3.0 there are two tools for revealing these histories:
  - A control icon allows users to select earlier versions of the model as indexed by revision number (starting with 1 and incremented by one for each save). If the control icon is selected and active, users can “scroll” through the revisions using the mouse wheel to see how the model evolved over time.

[3] As is discussed in the following section, the Catalyst database is updated only when changes to a model are “saved” and written from the Catalyst client to the Catalyst database on the Catalyst server.

- A “Show History” window is available from the “View” drop-down menu. When selected, it allows side-by-side comparison of any two versions of the selected node, much in the same manner is comparing versions can be done in a wiki. The General Dynamics Team would like to extend this capability to show differences across *models*, rather than just nodes.

2.2.3 Catalyst as a Shared Resource

As noted above, Catalyst is a client-server application: all Catalyst models are stored on a central server and are accessible by anyone with the appropriate access rights. And furthermore, Catalyst is designed such that multiple individuals can access—and potentially modify—a given model at the same time. This approach, however, raises two fundamental questions: 1) how are updates to a given model handled across the multiple users (each with their own client) who may have a particular model open, and 2) how are potential or actual conflicting edits to a model by multiple users handled?

<u>Regarding the first issue, updates</u>: We have always taken the position that we will not automatically “push” updates (which are on the server) out to all client computers. Instead, we let users make a choice to update their view of a model to show the work done by others on that model. This avoids the “moving target” situation that can occur with the “push” approach, in which an analyst is trying to read a model that is evolving (as a result of others’ updates) before his or her eyes.

<u>Regarding the second issue, conflicting edits</u>: Dealing with (or avoiding) conflicting edits is an issue for any system that allows multiple users to access and update the same data.

There has been much research into this issue and, in general, there unfortunately is no single “best” approach. There are, however, a number of alternatives, each of which has pros and cons.

One approach is to allow write access by only *one individual at a time*. A particular variant of this approach is sometimes referred to as “baton passing,” in which only one individual can possess a single symbolic “baton” at a time, and in which that individual is the only one who can make changes to the shared space. This is a popular approach for collaborative “shared whiteboards,” where only one person can have the “pen” at a time. We did not feel that this approach was at all appropriate for Catalyst, where it seems critical to let analysts get work done at their own pace, independent of what others might be doing.

We initially took a different approach with Catalyst (in Catalyst 3.0, the previous version): rather than avoiding conflicts by preventing simultaneous updates by multiple individuals, we decided to allow individuals to proceed at their own pace, and to have Catalyst provide assistance if and when conflicts occurred.

However, our mechanism for providing this assistance turned out to be more complex (and harder for users to understand) than we expected. We have therefore recently taken the decision instead to provide "awareness," and a fairly fine-grained locking mechanism—at the tree level, vs. at the coarse-grained level of an entire model, which may contain multiple trees. In particular, our approach is as follows:

- When Ann and Bill have the same model open, and Ann begins editing a node, that node gets a light red background out in Bill's display—thus providing Bill with an unobtrusive "awareness" of Ann's work.
- Not only does Catalyst show (with the red background) that someone else is editing or modifying a node, but it also displays the *name* of the person who is editing the node in the red background area. Furthermore, anyone who has the model open can click on the displayed name—Ann, in our example—to launch an Instant Messaging session with the person doing the editing. At a minimum, this feature enables people to ping the editor, to find out when they're going to be done.
- When Ann saves her work (thereby writing it to the Catalyst server), that same node on Bill's display changes from light red to light gray, indicating that Ann is done with that particular edit, but that Bill's view of the model is now out of date as a result of Ann's edits. Bill can then decide to update his view to reflect Ann's changes. So we are "pushing" *information* about updates, but not the *updates themselves*. The updates themselves are "pulled" only when Bill requests them, thus providing him with control over the canvas he is observing and working with.
- Furthermore, we adopt a "locking" strategy such that when Ann is editing a node, that node is "locked" to others, thus avoiding the "s/he who writes last, wins" situation, in which if Ann and Bill are both editing the same node in a model, and if Bill saves first and Ann saves second, she overwrites Bill's work.

2.2.4 Trees as a "Balance Point" Among Representations

We selected a tree structure as the fundamental organizing structure for Catalyst because we believe it provides a good balance between expressivity and "readability." For collaborative systems, in which people must at least to some degree be able to access and understand each others' work, getting this balance "right" is particularly critical.

<u>Regarding "readability":</u> Tree structures are familiar: we are all used to creating and reading outlines and other hierarchical decompositions. And in Catalyst, in particular, analysts can capture in upper-level nodes their "high level" approach to a particular issue, challenge, or problem.

In Catalyst we provide a number of UI mechanisms to allow the Catalyst tree structures to be collapsed and expanded, precisely so that individuals *can* quickly and easily get a high-level sense of what a particular Catalyst model is about, and can also drill in (by expanding selected sub-trees) to get more detail, as needed. Ideally, the high-level structure provides a "guide" on where to drill in, just as does the table of contents in a reference book. (This collapse/expand capability is valuable not only to the tree structure's creator, but also to individuals who want to see and leverage each others' work.)

<u>Regarding expressivity:</u> A tree structure is only one of many that can be used to organize information and thinking. Alternative structures are "general graphs" and "directed graphs" (which are general graphs that additionally specify directions—parent-child relationships—among nodes).[4]

"Concept Maps" are a particular form of general graph that are both expressive and flexible. However, they (arguably) lack readability. Concept maps aren't required to have "starting points," as do trees with their root nodes, and thus it can be difficult to get "oriented" in one. Also, as a related shortcoming, they lack the collapse/expand capability of Catalyst's trees.

In summary, there is no "right or wrong" with representations; there are only tradeoffs. For a collaborative system, we decided to place a strong emphasis on readability, arguably at the expense of expressivity, to ensure not only that the system is easy to learn and use, but also to ensure that individuals' work *is* "readable" by others. We believe that Catalyst's tree structures *do* provide a sufficient and powerful expressivity for a wide variety of analytic challenges.

A sobering account of the challenges in getting this balance "right" is provided by Conklin, et al,[5] who developed the Compendium collaborative software application. Compendium, which has roots going back the Issue-Based Information Systems (IBIS) developed in the 1970s, allows users to employ a small set of node types that can be connected into directed graphs via a small set of link types.

Compendium's developers attempted to develop an absolutely minimal set of formalisms to support their design goal of providing transparent and intuitive support to the thinking and design processes. Nonetheless, they concluded:

> "A primary lesson from these early experiments is that the effort required to think and represent hyper textually is comparable to the development of fluency in a new language—it is a whole new literacy."

[4] Trees are "directed graphs," with the restriction that each node has no more than a single parent. In a general directed graph, no such restriction holds, and, as a result, any node can have any number of parents, potentially creating cycles (loops).

[5] Conklin, et al., *Facilitated hypertext for collective sense making: 15 years on from gIBIS*, Proceedings of the twelfth ACM conference on Hypertext and Hypermedia, ACM Press, 2001.

Ultimately, Compendium became a tool targeted for facilitators' use in face-to-face meetings. This put the load of developing this "fluency in a new language," and of employing this fluency effectively, on the facilitators, who use the tool—projected on a screen—to capture and guide the group's collaborative thinking process.

Given our objective that Catalyst *not* require a facilitator, we believe that it is absolutely critical to keep the representation as simple as possible, hence our decision to use trees rather than more general relational structures.

2.2.5 Moving Catalyst to the Web

At the start of our CASE Connect effort, we had already (under a different contract) developed Catalyst up to version 3.0 with Microsoft's "Windows Forms" technology, which makes it a "thick client" client/server application that needed to be installed on client computers (via an installer). Catalyst 3.0 also included a rather rudimentary "Web View"—a read-only view of a Catalyst model from within a Web browser.

Even before the start of our CASE Connect effort, two things had become clear:

1. There is a strong demand within the Intelligence Community for browser-based access to applications.
2. A number of new developer tools are taking hold that allow far richer and more interactive content and applications to be hosted in a Web browser than was previously possible—even quite recently. These rich applications are sometimes referred to as RIAs, for "Rich Internet Applications.[6]"

We decided to aim for a Microsoft Silverlight-based reimplementation of Catalyst, Catalyst 3.5, which would run in a Web browser. As of the fall of 2007, only a very preliminary Beta release of Silverlight was available, so we switched to a related technology from Microsoft called WPF, for *Windows Presentation Foundation*. WPF applications are "lighter" than thick clients, but, as with Java Applets, they require a small download to the client computer. Both Silverlight and WPF use the same language for defining the UI, called XAML (pronounced "*zamel*"), but WPF is an older and therefore more fully-featured (and better supported and less buggy) technology. Our plan was to develop a new WPF-based version of Catalyst, which could later be (largely) transitioned to Silverlight when a more fully-featured version of Silverlight was released. (Silverlight supports a *subset* of the XAML that WPF supports, so as long as we used only "basic" XAML in our WPF development, the likelihood was high that we could easily use that same XAML for a later Silverlight release.)

[6] Two leading RIA development technologies are Microsoft's Silverlight and Adobe's AIR.

When Silverlight Beta 2 was released on 5 March 2008, we shifted from developing in WPF to developing directly for Silverlight. However, we soon encountered a substantial roadblock: because of security concerns, it is not—and likely will not be—possible for a Web application (Silverlight-based or otherwise) to support the rich automatically sourced drag-and-drop and cut-and-paste that we believe is a key Catalyst value proposition.

Given this situation, and given the short time left on our CASE Connect effort, we decided to shift Catalyst development back to WPF. We completed an initial version of Catalyst 3.5 in time for the start of our fall 2008 experiment, which began in the first week of November (as described in Section 3.4). We also developed a "read only" Silverlight Catalyst, which we hope can be extended to support the "Conversations" we describe in the following section.

2.2.6 Conversations In and About Catalyst Models

A key feature that we have added to Catalyst 3.5 is support for discussions and comments that are rooted in—and that refer to—Catalyst content. The need for this capability became clear in an experiment we conducted with an earlier version of Catalyst in 2006, in which the experimental participants used Catalyst nodes to "comment" on other Catalyst nodes—which had the undesirable side-effect of substantially cluttering up the Catalyst model by mixing content and commenting / discussion. It was clear that some form of discussion and commenting capability was required.

For Catalyst 3.5, we developed a capability where comments can be "attached" to a Catalyst node or a Catalyst model, much as Wikipedia and Intellipedia pages have discussion tabs where participants and interested parties can hash out issues regarding the "main content" of the sites. However, unlike with a wiki, these comments appear right next to the nodes to which they refer—much like sticky nodes. The presence of comments is indicated in two ways:

1. A small indicator icon on the node.
2. A separate consolidated listing of all comments in a model.

Clicking on either indicator causes the comment to appear, clicking on the comment's "close" button (or anywhere else) causes it to disappear.

The importance of allowing discussions or conversations about analytic content was a driving motivation in our development of "Context-Grounded Conversations," which we describe in Section 2.3, *Context-Grounded Conversations*, below, and in which discussions are available from outside of the context of Catalyst, specifically—at least for our first implementation, associated with Catalyst 3.0—in a separate Web-based discussion forum. We tested this initial implementation of Context-Grounded Conversations in our 2007 experiment. Feedback was generally favorable, but the presentation of the discussion content within the Catalyst UI (which was an initial implementation only) was not visually well-integrated with the Catalyst content to which it referred, which we believe was a serious shortcoming in our implementation.

The General Dynamics Team would like to continue to pursue this thread of research, but given our funding, we have chosen to focus for the remainder of our effort to support "Grounded Conversations" on an *integrated* discussion mechanism, vs. the decoupled Context-Grounded Conversations we describe in the following section.

Initial indications are that the "sticky note" presentation that we have developed for Catalyst 3.5 is a success, so a next logical step in our research would be to apply this presentation to the fully decoupled Context-Grounded Conversation approach, which was not visually well integrated with the Catalyst model content.

Finally, another avenue we would like to pursue is adding the "sticky note" commenting feature to our read-only Silverlight Web Catalyst. (This should be relatively straightforward to implement, but we do not have the funds remaining to do so.) This approach would allow anyone on the same network as the Catalyst server to view Catalyst models and to weigh in with comments and suggestions. Given the limitations of a Web-based Catalyst to support drag-and-drop (as described in Section 2.2.5, above), this seems like it would be the perfect role for a Web-based Catalyst:

- Expose Catalyst content to anyone (on the network) with a browser; and
- Allow anyone (on the network) to make comments and suggestions.

We anticipate that this approach will expose Catalyst content to users who do not have the full Catalyst 3.5 available on their computers and who may even be completely unfamiliar with Catalyst. Although we consider the Catalyst 3.5 interface to be relatively intuitive, the far simpler Web-based Catalyst client (being read-only) will, we anticipate, be quite straightforward for these new users to pick up.

2.3 Context-Grounded Conversations

2.3.1 Overview

As is mentioned in the preceding section, we made substantial progress in this area of research, but have taken the decision to pursue a somewhat less general approach in Catalyst 3.5.

Context-Grounded Conversations are computer-based discussions—or "conversations"—that are explicitly linked to *content* represented within an application. The conversation and the content are shown together, even though they are managed by two separate applications (that is, the conversation application, such as a Web-based discussion forum, and the content application, such as Catalyst).

A discussion that is linked to a particular application's content can be viewed from within *either* application:

- *Viewed from within the content application* (such as Catalyst): Any discussions linked to the application's content are viewable next to the content to which they refer.
- *Viewed from within the discussion application* (such as a Web-based discussion forum): A small "snippet" of application content is shown next to the discussion, thereby allowing readers to get a quick sense of what is being discussed without having to switch to the content application to see what is being discussed.

This explicit linkage, rendered from within both applications, provides a clear "grounding" of—as well as an explicit context for—discussions, hence the name "Context-Grounded Conversation," or "CGCs." These linkages allow the important "social infrastructure" of queries, requests, suggestions, challenges, and negotiations to be brought to bear in the development of application content. And they allow that social infrastructure to be enhanced and informed by the structure that can be provided by shared application content.

Furthermore, because the discussion applications are "loosely coupled" to the content applications, a given discussion application—such as Web-based discussion forums, Instant Messaging (IM), or email—can be linked to *multiple content applications*. And, the content within a given content application—such as Catalyst or Intellipedia—can be discussed via *multiple discussion applications*.

Our first key goal with CGCs is to *extend* discussion mechanisms that people *already use* to support explicit linkages between application content and discussions about that content. By integrating CGCs into a larger, pre-existing fabric of computer-based communications, discussions can be conducted naturally, organized by topic or theme, and can reference application content if and when appropriate. This stands in contrast to applications in which a discussion capability is *directly embedded* within the application framework and is designed only for discussing content within that application's framework.[7] Any direct embedding is *awkward* because users have to switch from their "standard" communication mechanisms to the embedded communication mechanisms whenever they want to easily discuss specific application content.[8] And a direct embedding is *limiting* because discussions embedded directly in content application frameworks are not accessible from general discussion contexts, and vice versa.

[7] This is the approach taken within wikis, such as Wikipedia and Intellipedia, where discussions are conducted on a special "discussion" tab associated with each wiki page.

[8] Users can make written references in their general discussions to "such-and-such a file," or "such-and-such a wiki page, but these are time-consuming to create and awkward to follow.

Our second key goal is that CGCs be generally available across a range of discussion mechanism and content applications:

- *Discussion mechanisms* can include not only discussion forums, but also IMs, email, and perhaps other communication modalities.
- *Content* that can be discussed will include not only Catalyst models and nodes, but also items tagged in tag|Connect. We anticipate developing a mechanism for discussing Intellipedia content, as well as. Our intention that both we and others can develop CGC capabilities for multiple content applications.

In the subsections below, we address our discussion forum implementation for CGCs, followed by our initial work on an Instant Messaging CGC application.

2.3.2 Implementation – Discussion Forum

General Dynamics' first implementation of Context-Grounded Conversations is a Web-based discussion forum (the discussion mechanism) and Catalyst (the content application).

- *Launching and Viewing Conversations from within Catalyst*: Catalyst users can launch discussions directly from within Catalyst by right-clicking on any node and selecting the "Discuss this…" option, which causes a special Catalyst window to pop up into which an analyst can type her comments. As soon as she hits "enter," her post is written—via the discussion forum's API—to the discussion forum database and is then immediately available to anyone viewing the forum. It is possible in Catalyst 3.0 to see all "Related Posts"—that is, a listing of all discussion forum posts that refer to the currently-open Catalyst model. Clicking on any of the listed post launches the full post in the Web-based discussion forum.

- *Viewing Conversations from within the Discussion Forum*: Any thread or post that is linked to a Catalyst model will automatically include a small snippet of the Catalyst model, rendered adjacent to the post, as shown in Figure 5. This snippet provides readers of the forum with a quick sense of what is being discussed, and also provides a hyperlink that launches the full Catalyst model (via the read-only Web view).

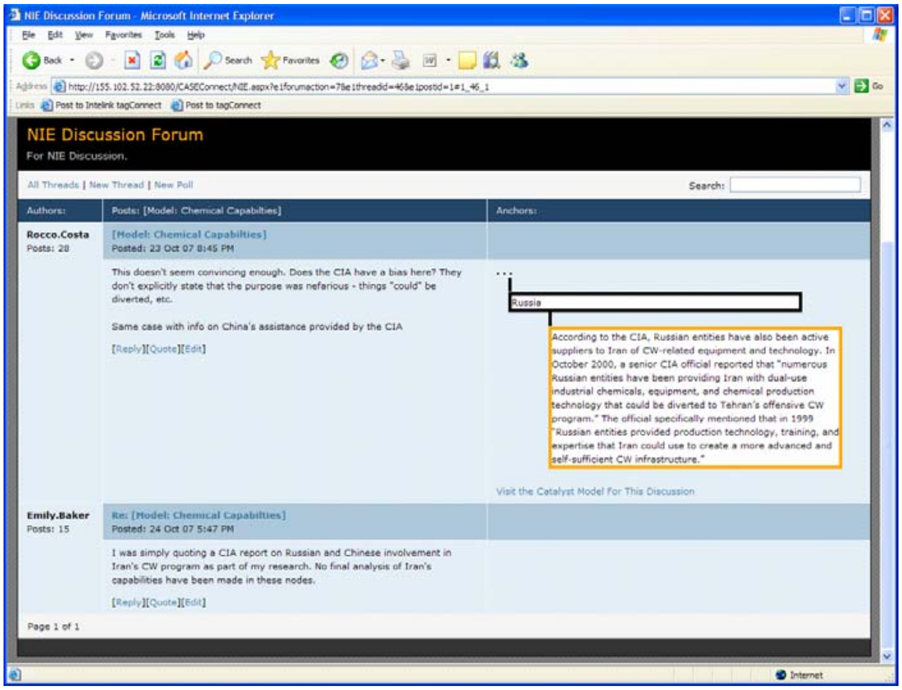
NIE Discussion Forum - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://155.102.52.22:8080/CASEConnect/NIE.aspx?e1forumaction=7&e1threadid=46&e1postid=1#1_46_1
Go
Links Post to Intelink tagConnect Post to tagConnect
Internet

NIE Discussion Forum
For NIE Discussion.

All Threads | New Thread | New Poll   Search:

Authors:	Posts: [Model: Chemical Capabilties]	Anchors:
Rocco.Costa Posts: 28	[Model: Chemical Capabilties] Posted: 23 Oct 07 8:45 PM	
	This doesn't seem convincing enough. Does the CIA have a bias here? They don't explicitly state that the purpose was nefarious - things "could" be diverted, etc. Same case with info on China's assistance provided by the CIA [Reply][Quote][Edit]	... Russia According to the CIA, Russian entities have also been active suppliers to Iran of CW-related equipment and technology. In October 2000, a senior CIA official reported that "numerous Russian entities have been providing Iran with dual-use industrial chemicals, equipment, and chemical production technology that could be diverted to Tehran's offensive CW program." The official specifically mentioned that in 1999 "Russian entities provided production technology, training, and expertise that Iran could use to create a more advanced and self-sufficient CW infrastructure." Visit the Catalyst Model For This Discussion
Emily.Baker Posts: 15	Re: [Model: Chemical Capabilties] Posted: 24 Oct 07 5:47 PM	
	I was simply quoting a CIA report on Russian and Chinese involvement in Iran's CW program as part of my research. No final analysis of Iran's capabilities have been made in these nodes. [Reply][Quote][Edit]	

Page 1 of 1

Figure 5. Context-Grounded Conversation - discussion forum implementation

- *Contributing to discussions from within the Discussion Forum*: If someone accessing the discussion forum wants to add a post that includes a link to a Catalyst model (perhaps responding to someone else's comments with a note about a relevant Catalyst model), he can get a special link from within Catalyst—again, by right-clicking on a Catalyst node, but this time selecting a "Copy Context Snippet" option, which adds the link to the user's clipboard. The user then pastes the content of his clipboard (containing the link) into a special field in the discussion forum's "create a post" input form. Anytime that post is rendered, the snippet is shown as well.

2.3.3 Implementation – Instant Messaging

The General Dynamics Team created the initial building blocks for a CGC implementation for Instant Messaging. We chose the "Jabber" instant messaging protocol—both because it is an open standard, as well as because it is currently in use within the Intelligence Community. (Technically, Jabber is a standard based on XMPP, or eXtensible Messaging and Presence Protocol. In practice the two terms are often used interchangeably.)

We developed an “extension” to the Jabber standard (via an extended namespace) that allows the same HTML snippets that we described above for our discussion forum implementation to be encapsulated within the XML of Jabber messages.

We also developed an initial prototype of a modified browser-based, open source Jabber client that understands the extension and that has two added windows: one for entering the snippets, and the other for displaying them. All users with this CGC-enabled Jabber client are able to send and receive Catalyst snippets; *and* users with “regular” Jabber clients (that is, Jabber clients that are *not* CGC-enabled) will still be able to send and receive messages, just without seeing the Catalyst snippets.

This client is shown in Figure 6.

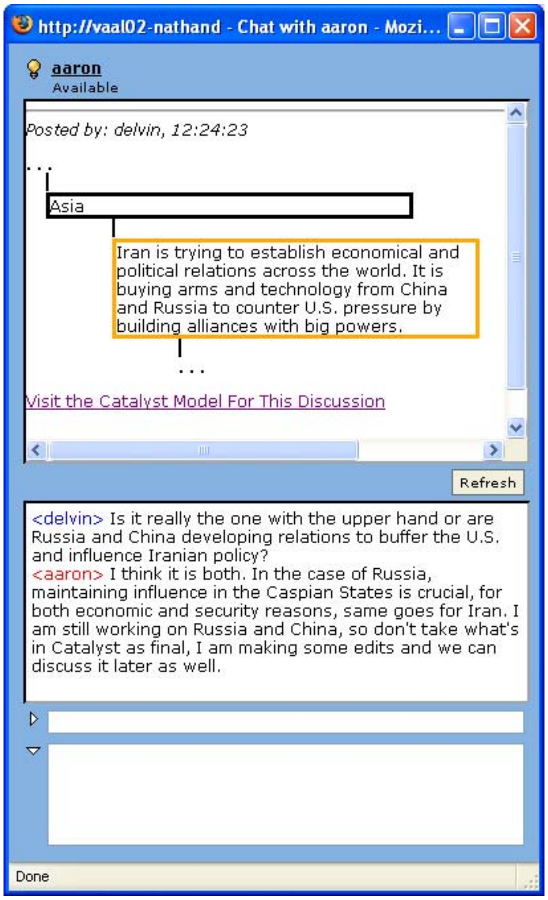

Posted by: delvin, 12:24:23

. . .

Asia

Iran is trying to establish economical and political relations across the world. It is buying arms and technology from China and Russia to counter U.S. pressure by building alliances with big powers.

. . .

Visit the Catalyst Model For This Discussion

<delvin> Is it really the one with the upper hand or are Russia and China developing relations to buffer the U.S. and influence Iranian policy?
<aaron> I think it is both. In the case of Russia, maintaining influence in the Caspian States is crucial, for both economic and security reasons, same goes for Iran. I am still working on Russia and China, so don't take what's in Catalyst as final, I am making some edits and we can discuss it later as well.

Figure 6. CGC - modified Jabber client

2.4 Relational Topic Models (RTMs)

We described in Section 1 that a conceptual foundation of our work is *inferring, rather than enforcing, structure*. With CASE Connect Team member Dr. David Blei of Princeton University, we have conducted research into using the techniques of Latent Dirichlet Allocation to infer “hidden” structures in the “tag space”—that is, subtle and hidden structures and relationships among tags, the information objects to which tags are applied, the people who applied these tags, and even the text of the tagged information objects themselves.

Dr. Blei has been instrumental in developing and applying the techniques of Latent Dirichlet Allocation[9] (LDA) to a range of data types with a focus on developing efficient inferencing and learning algorithms. The most common application of LDA is “topic modeling,” in which the words in a corpus of documents are mapped into an underlying probability space that characterizes the “topics” in the space. Each document in the corpus can then be characterized as a *mixture* of the topics in the corpus, which corresponds to our intuitive notions that a given document is likely to be “about” more than one thing. This approach provides a powerful and expressive means for characterizing documents, and for uncovering subtle interconnections amongst documents (via their mixed and sometimes shared topics).

[9] Blei, D., Ng, A., Jordan, M. “Latent Dirichlet allocation.” Journal of Machine Learning Research, 2003.

Because the mathematics of LDA lead to characterizations that correspond closely to our intuitive notions of "topics," LDA-based topic modeling is a particularly powerful technique for inferring relationships among information objects that can be leveraged in a variety of ways—and, particularly, to address the CASE goals of connecting analysts with high-value information objects relative to their analytic needs.

On our CASE Connect effort, Dr. Blei has been developing the foundations for a new application of LDA, "Relational Topic Models" (RTMs), which extend LDA-based topic modeling to also consider not only topics in documents, but also explicit links among documents—which is precisely what tags provide. (Two documents can be thought of as linked when a common tag—say, "Iraq"—is applied to both documents.) Just as "traditional" LDA-based topic models treat topics as being generated by an underlying "generative" probability distribution, the parameters of which are then estimated or inferred given a new document in the corpus, RTMs, additionally, treat the *links among documents* as also being generated by an underlying "generative" probability distribution, the parameters of which can also then be inferred or estimated. This leads to a richer characterization of the documents and other information objects, precisely because it leverages not only the words in the information objects, but also the concise summary descriptors (tags) that are applied by humans, and which likely represent key and important aspects of those documents.

It bears mentioning that RTMs can have broad applicability to many types of linked information spaces and corpora, including citation networks, linked Web pages, or social networks with user profiles.

Dr. Blei has applied initial RTM models to tag|Connect data sets developed in the 2007 experiment, described in Section 3.3, below. The results are encouraging. However, much remains to be done. The RTM software we have developed to date looks only at the links (shared tags) among documents and the words in the document's titles, whereas it will be important to consider also the text of the documents themselves.

And more to the point of the larger CASE objectives, much remains to be done in applying RTMs to the task of estimating the analytic value of information objects and of potential analyst-to-analyst connections.

3 “Understanding” Through Experimentation

At the heart of our “Develop / Understand / Improve” cycle, introduced in Section 1, are the processes by which we “understand” the collaborative, analytic, and technical foundations of the tools we develop. The knowledge and understanding we gain through these processes drives our development process, and is, we believe, critical and necessary to our success.

Experimentation has played a central role in our approach and in our success. With our subcontractors (and domain experts) the Center for Terrorism and Intelligence Studies (CETIS) and the Monterey Institute for International Studies (MIIS), we have performed a number of experiments with our tools that have provided us with invaluable insights and understanding, which have, in turn, informed and driven our development process.

3.1 Experiments - Overview

We have conducted a number of successful “experiments” with our subcontractors CETIS and MIIS, not only on our CASE Connect effort, but also on a previous effort under which we began our development of tag|Connect and Catalyst. On CASE Connect, we conducted experiments in 2007 and in late 2008:

- 2007: Beginning in October, 2007, we conducted a nine-week experiment with ten MIIS graduate students. The tools employed were Catalyst, tag|Connect, a discussion forum that we augmented to support CGCs, and a Media Wiki.
- 2008: We conducted a five-week experiment with nine MIIS graduate students in November and early December of 2008. A key goal for this experiment was to explore how social norms and policies (such as with Wikipedia) can guide a collective and collaborative effort.

Although in each experiment there has been a clear focus on the General Dynamics Teams’ tools’ usability, we have also been particularly interested in developing an understanding of *processes* by which the tools are employed and how the tools’ use does—or doesn’t—support the end-goal of effective collaborative analysis.

In particular, because a central goal of our work is providing effective support to *collaborative* analysis, we have been particularly focused on how the participants deal with the host of tradeoffs that any collaborative group must face, including individual vs. collective work, and diversity vs. consensus. Observing how the participants have—and haven’t—balanced these sometimes-competing needs, as well as the individual and collective strategies they employed in accomplishing their tasks, has been enormously informative for our development team.

So the sorts of questions we have addressed through our experiments have included not only rather concrete concerns regarding our tools, such as how to improve the intuitiveness of Catalyst's user interface, but also more abstract issues that have more to do with collaborative *processes* than with any specific details of our tools' implementations. For example, in our 2007 experiment we were particularly focused on the extent to which individuals could make use of each others' individual Catalyst models in the performance of a collaborative task.

Following each experiment, we have analyzed results and outcomes (which include not only the participants' work, but also questionnaire responses and follow-up interviews). We have then identified strengths and shortcomings along collaborative, analytic, and technical dimensions. And from those strengths and shortcomings, we developed prioritized lists of "next steps" for further development, which then drove our forthcoming work (and forthcoming experimentation).

3.2 Experimental Design

The goal of all of our experiments has been to learn as much as possible along three rather broad dimensions (collaborative, analytic, and technical) to inform and drive our software development process. Our experiments have therefore all been somewhat exploratory in nature, which stands in some contrast to a more rigorous experimental framework that would be predicated on clearly-defined hypotheses to be tested and a carefully-bounded experimental agenda.

And, furthermore, because of the central importance of the *collaborative* dimension to our research, we have always opted to maximize the number of participants who were using our tools, rather than dividing the available pool of participants into an experimental group and a control group.

3.3 2007 Experiment

Beginning in October, 2007, we conducted a nine-week experiment with ten MIIS graduate students, supervised and administered by the Center for Terrorism and Intelligence Studies (CETIS). The tools employed were Catalyst 3.0, tag|Connect, a discussion forum that we augmented to support CGCs, and a MediaWiki wiki.

Although it was our intention to learn as much as possible about all the tools, our primary interest was in Catalyst 3.0.

The tasking given to the participants was: "What are the most important and most likely threats to the security of the United States and its interests abroad that could arise from the Islamic Republic of Iran within the next 10 years?"

The participants were to use all the tools, and were to produce “a cogent, fully-sourced and cited, summary document reflecting their analysis of this issue” in the form of a National Intelligence Estimate (NIE), which was to be developed in the wiki, and which was to be no more that 10,000 words in length. A full set of Terms and References was provided to define terms and establish various parameters for the experiment.

The participants were allowed to develop their own schedule (to meet the specified due date) and to organize their collaborative work as they saw fit. Gary Ackerman and other advisors from CETIS were “on call” for any help the participants might request, as well as to monitor the proceedings and to intervene if necessary.

However, during training, a high-level workflow was suggested: begin by researching sources and tagging them in tag|Connect. Then shift to developing “descriptive” models in Catalyst. Then shift to developing “analytic” models in Catalyst. Finally, use Catalyst’s “Copy to Wiki” function to export the Catalyst models to the wiki, and continue work to produce the required final product. Given that high-level workflow, however, the participants had tremendous freedom to decide, for example, how and when to produce individual vs. collective models, and when to transition from Catalyst to the wiki.

3.3.1 Lessons Learned

Through the participants’ work (all of which was captured on the General Dynamics server that hosted the four applications), and through their survey results and follow-up interviews, we gained tremendous insights into the application of our tools to a collaborative analytic process in general, and many specific technical ideas as well.

An extensive analytical report on the experiment has been developed by Gary Ackerman, James Fobes, and Charles Blair, all of CETIS.

An abbreviated listing of a few key observations and conclusions is as follows:

Catalyst 3.0: The participants appreciated Catalyst’s power as a collaborative tool. However, they felt that Catalyst models became hard to read when the models grew large or when the nodes themselves contained extensive text. As a result, the participants for the most part moved directly from developing individual models to doing their collective and collaborative work in the wiki, thereby “skipping” the step of developing collaborative analytic models in Catalyst. In general, however, the participants felt that their work would have been better had they been able to and/or chosen to do more work in Catalyst, vs. in the wiki.

Context-Grounded Conversations (CGCs): Not all users utilized this feature, but several did. There was a consensus, though, that the CGCs that did occur improved the quality of the analysis, and that CGCs are a valuable analytic tool. There was strong agreement that CGCs need to be integrated more fully (from the UI perspective) into Catalyst.

Organization: Because of the rather broad nature of the participants' tasking, the participants chose individual focus areas that did not have much overlap. As a result, there was less chance for the students to collaborate on specific issues. The participants were consistent in their view that this hurt the quality of their analysis.

Production: As noted above, the participants switched to the wiki rather early in the overall process. And they had tremendous problems with editing conflicts in the wiki (an issue that Catalyst is specifically designed to minimize). They therefore selected three "writers" who were responsible for editing the wiki. This caused problems, though, because the wiki had grown to well above the 10,000-word limit, and quite a bit of cutting was needed. The cuts were not coordinated, though, which left some participants feeling that their contributions were misrepresented or left out.

Estimative Language: Related to the above issues were some difficulties with the NIE's estimative language. Some of the participants had proposed language that distinguished between uncertainties due to a lack of evidence vs. uncertainly due to a lack of consensus. However, the language that they ended up adopting—which was modeled after that in a declassified NIE—did not make this distinction clear. This conflation tended to wash out the presentation of disagreements in the group, which arguably lessens the value of the NIE to a policy maker or decision maker.

Conclusions: We believe that many of the issues encountered will be addressed by Catalyst when its ability to support viewing and understanding large models and large nodes is improved. In particular, we believe that Catalyst will be particularly effective at supporting the preservation of diversity within a collaborative task environment. And that improved Context-Grounded Conversations can provide excellent support in letting users work through the challenges of paring work down through collective participation.

3.4 2008 Experiment

Beginning in November, 2008, we conducted a five-week experiment with nine MIIS graduate students, supervised and administered by the Center for Terrorism and Intelligence Studies (CETIS). The tools employed were Catalyst and an open-source discussion forum.

The students were tasked with developing a simulated Nuclear Posture Review (NPR) product determining the relationship between United States nuclear deterrence and nonproliferation objectives. The specific tasking was: "*What is the relationship between United States nuclear deterrence requirements (preventing use), and nonproliferation objectives (preventing acquisition).*" The Terms of Reference (TOR) provided to the students laid out a number of Secondary Questions as well.

The focus of this experiment was on understanding how Catalyst can best support realistic collaborative development. The participants were divided into three groups: two "teams"—one focused on deterrence requirements, and the other on nonproliferation objectives—and one set of "free agents." The two teams were encouraged to collaborate with each other; however, their principal responsibility was to accomplish their own assigned work. The "free agents" were to pitch in to "help" as they saw fit. We believed that this set-up would encourage the participants to confront a variety of issues that often interfere with collaboration and sharing—even when people understand its value.

At approximately midway through the five-week experiment, we conducted a facilitated workshop with the participants in whom they were encouraged to explore potential policies and "norms" that might facilitate a balance between sometimes-conflicting individual and collective objectives. (Our view is that the "solution" to this dilemma is primarily social and organizational, rather than technical—as in the policies and norms that have evolved in Wikipedia, and which are arguably the real "fuel" that drives its success.)

3.4.1 Lessons Learned

As with our 2007 experiment, through the participants' survey results and follow-up interviews, we gained tremendous insights into Catalyst's strengths and weaknesses as a collaborative analytic tool, and many specific technical insights and feature ideas as well.

Our mid-experiment workshop yielded a number of specific feature suggestions, some of which we were able to implement quickly—and which we then received feedback on through our after-the-experiment questionnaires and interviews. The most significant change we made was moving the display of comments that can be "attached" to nodes from a sidebar display to "bubbles" that are attached directly to the commented-upon nodes themselves. These comments are displayed only when users click on the small comment icon that appears in any commented-upon node, or when users click on a comment title in the sidebar (which we redesigned to display only the most-recently-made comments for each node, sorted so the newest comments are at the top). Clicking on any comment title in the sidebar causes the model to expand and pan (if necessary) such that commented-upon node is visible, and to display node's comment bubble also.

Figure 7 (which is a zoomed-in view of *Figure 3*) shows a node that has been commented upon, with the comments in view.

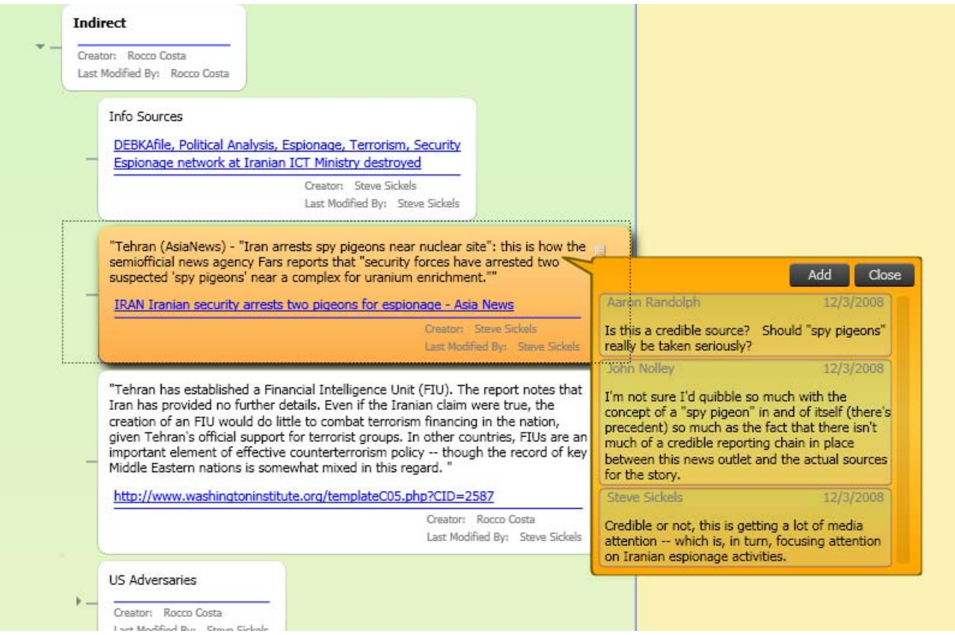

Indirect

Creator: Rocco Costa
Last Modified By: Rocco Costa

Info Sources

DEBKAfile, Political Analysis, Espionage, Terrorism, Security
Espionage network at Iranian ICT Ministry destroyed

Creator: Steve Sickels
Last Modified By: Steve Sickels

"Tehran (AsiaNews) - "Iran arrests spy pigeons near nuclear site": this is how the semiofficial news agency Fars reports that "security forces have arrested two suspected 'spy pigeons' near a complex for uranium enrichment.""

IRAN Iranian security arrests two pigeons for espionage - Asia News

Creator: Steve Sickels
Last Modified By: Steve Sickels

Aaron Randolph 12/3/2008

Is this a credible source? Should "spy pigeons" really be taken seriously?

John Nolley 12/3/2008

I'm not sure I'd quibble so much with the concept of a "spy pigeon" in and of itself (there's precedent) so much as the fact that there isn't much of a credible reporting chain in place between this news outlet and the actual sources for the story.

Steve Sickels 12/3/2008

Credible or not, this is getting a lot of media attention -- which is, in turn, focusing attention on Iranian espionage activities.

"Tehran has established a Financial Intelligence Unit (FIU). The report notes that Iran has provided no further details. Even if the Iranian claim were true, the creation of an FIU would do little to combat terrorism financing in the nation, given Tehran's official support for terrorist groups. In other countries, FIUs are an important element of effective counterterrorism policy -- though the record of key Middle Eastern nations is somewhat mixed in this regard. "

http://www.washingtoninstitute.org/templateC05.php?CID=2587

Creator: Rocco Costa
Last Modified By: Steve Sickels

US Adversaries

Creator: Rocco Costa
Last Modified By: Steve Sickels

Figure 7. Comments on a Catalyst model

An extensive analytical report on the experiment has been developed by CETIS. An abbreviated listing of a few key observations and conclusions is as follows:

Shared work artifacts: It became particularly clear during the mid-experiment workshop discussions that the participants were not entirely comfortable with the concept of a tool that allowed shared access to work artifacts. (Which is something that, broadly speaking, Catalyst, Wikipedia, and Intellipedia have in common.) This manifested itself in two ways: 1) the participants expressed a general reluctance to update others' work, and 2) the participants wanted some mechanism for preserving "their" work. In contrast to the Wikipedia/Intellipedia approach of working out conflicts over time (through sequential edits by multiple individuals), the participants wanted a mechanism by which they could stake out respective positions, potentially indicating explicitly where there were disagreements of conflicts. (We suspect that this approach, as proposed, would lead to fragmentation, a lack of coherence, and "agree to disagree" conflicts, whereas Catalyst, Wikipedia, and Intellipedia arguably promote a form of collective convergence. Fragmentation and "agree to disagree" conflicts are not inherently bad—indeed, they may indicate a healthy *questioning* process that is arguably necessary in intelligence analysis. Nonetheless, we suspect that an unavoidable consequence of this likely fragmentation would be a much harder-to-read product, which in turn would discourage collaboration, which would have the unintended and undesirable consequence of reducing the robustness and scope of the analytic process. Furthermore, we believe that the degree of fragmentation would be correlated positively with the number of users participating in a model's development, which in turn suggests that the proposed approach would not scale well for large numbers of users.)

Comments: The updated commenting feature was viewed as effective and useful. The only issue that some of the participants had with it was that they wanted a capability to remove comments (generally, after whatever issue they addressed had been resolved).

Overall effectiveness: In general, the participants felt that Catalyst was an effective tool for conducting and presenting an analysis. They felt that it was particularly useful to be able to see, via Catalyst's tree structure (when collapsed), a sort of summary of what their colleagues were thinking and doing. They felt that having a shared space with a built-in summary mechanism (the tree structures) contributed positively to the quality of their own work and to the work of the group. Some of the participants asked for a more free-form structure, where nodes could be linked arbitrarily to other nodes. It seemed that from some participants' perspectives, the tree structure was extremely useful, whereas for others it was a bit limiting. We address this tradeoff in Section 2.2.4, "*Trees as a "Balance Point" Among Representations*," above. All in all, Catalyst was viewed as an effective organizational and analytic tool. And, in particular, it seemed to provide the collective benefits—from "individual" work—that were a driving objective in its design.

4 Results and Accomplishments – Summary

We conclude this report with a summary of key results and accomplishments:

4.1 Blending the Personal and the Collective

We developed the concept of "blending the personal and the collective," in which tools are designed such that individuals' "own" work automatically has collective benefits. We developed two software tools that embody this concept: tag|Connect and Catalyst. This approach addresses the "zero sum" impediment to collaboration: people are unlikely to collaborate if doing so pulls them away from their own responsibilities or does not otherwise have direct "individual" benefits. tag|Connect addresses the ubiquitous (and "individual") need to organize Web-based resources. But, at the same time, it also provides tremendous collective benefits by connecting resources and individual together as an automatic outcome of this individual behavior. Similarly, Catalyst allows analysts to organize their "own" analytic frameworks and to conduct their "own" analyses. But because these frameworks are available for others to use and leverage (with automatic attribution), collective benefits can accrue from individually-centered analytic activities.

4.2 tag|Connect

We developed the tag|Connect social bookmarking application, which has been described by Dr. Thomas Fingar, DDNI/A, as "the intelligence community's social bookmarking service."[10] tag|Connect is deployed worldwide as a Core Service on all three Intelink networks (JWICS, SIPRNet, and Intelink-U), and is in use by tens of thousands of analysts and others in the Intelligence Community. tag|Connect is similar in many ways to del.icio.us (and other social bookmarking services on the open Internet), but—to meet the needs of the Intelligence Community—it places a stronger emphasis on connecting its users with other users. tag|Connect makes it easy to see who in the Intelligence Community is using—and who is making the most use of—particular tags or combinations of tags, which can easily open up new avenues of research and collaboration that would not have been available otherwise.

[10] See http://www.dni.gov/speeches/20070905_speech.pdf, pp. 74

4.3 Catalyst

We developed the Catalyst collaborative analytic application, which provides a flexible and user-friendly mechanism for analysts to both individually and collectively structure their analyses and to capture and organize resources and information relevant to those analyses. Catalyst employs a modular, "drag-and-drop" approach to organizing information: analysts create "trees" of nested "nodes" that contain text composed by analysts, links to web-based resources, and automatically sourced "snippets" of text from web-based text documents and resources. Nodes and "sub-trees" can be detached, reattached, and copied from one Catalyst structure to another. Trees can easily be collapsed and expanded, thereby hiding or exposing analytic detail—and making the trees more "browseable" when partially collapsed.

Not only is Catalyst a powerful "individual" analytic tool, it also designed specifically to support effective collaborative analytics. As on blogs and wikis, all Catalyst content is shared. And as on a wiki, anyone can modify, add, or delete Catalyst content—with full history, including attribution, maintained. Furthermore, Catalyst content can easily be copied, with full and automatic attribution, from any Catalyst "model" to any other Catalyst model, thus encouraging modular reuse, and leveraging and building upon the work of others. Finally, Catalyst provides an integrated "discussion" mechanism, where comments and discussion threads can be "attached" to specific Catalyst nodes or to entire Catalyst "models." The visual presentation of Catalyst discussions makes it readily apparent which Catalyst content is being discussed, and which comments and discussions refer to specific sections of the Catalyst "models." (This contrasts with wikis, in which discussions are relegated to a given topic's "discussion tab," and thus cannot be linked to a specific portion of the topic's content.)

4.4 Relational Topic Models (RTMs)

A central tenet of our approach is to provide flexible tools that are easy to use and that can be readily applied to new (and potentially unforeseen) analytic demands and challenges. Toward that end, both tag|Connect and Catalyst place minimal constraints on the information artifacts that analysts can create with them. (Any tag or tags can be applied with tag|Connect, and any sort of tree structure can be created with Catalyst.) This lack of constraints has tremendous benefits in terms of flexibility and adaptability, but the downside is that there are likely to be fewer explicit connections and interrelationships in the data. (In tag|Connect, for example, there is no guarantee that multiple taggers will adopt a unified approach—or even consistent spellings—to their tagging.) And connections and interrelationships—among information objects, among analysts or users, and between information objects and analysts—are the foundation upon which user models and assessments of analytic value can be built.

Our approach, then, is to infer interrelationships and structure in data that—by design—lacks formal and explicit structure. In particular, with Dr. David Blei we are developing Latent Dirichlet Allocation (LDA)-based generative topic and relation models, which provide compact and powerful descriptions of underlying structure in terms of "mixtures" of underlying components, factors, and relations. And we are using variational inferencing techniques[9] to estimate these underlying structures, given input data. These estimation techniques are extremely efficient and can be practically applied at massive scales.

For our CASE effort, we have begun developing a new statistical (and LDA-based) model of networks, the Relational Topic Model (RTM), which jointly models topics in information objects (as in "traditional" LDA-based topic models) and the relations or links among the information objects. Although we expect RTMs to have utility across a wide variety of networks and network data, we are developing the RTM approach specifically to uncover and leverage relationships among information objects induced, jointly, by shared "topics" in the information objects' text and by the shared and related tags that analysts apply to them (via tag|Connect). We anticipate that this approach will be particularly powerful because it explicitly leverages analysts' assessments of what is important about specific information objects, as conveyed by the tags they apply to those information objects.

Because RTMs capture "hidden" patterns and relations from text and related tags, they can be used as a basis for inferences regarding the expected utility of new information objects, and of potential analyst-to-analyst connections. Our research plans for 2009 and 2010 were to be focused not only on developing and refining our RTM approach, but also on developing, refining, and leveraging these further inferences as well. In particular, we anticipated exposing RTMs and the inferences that build upon them via Web Services for consumption by the CASE user modeling services. Initial tests with tag|Connect data are promising, but much research and testing remains to be done to fully develop, apply, and leverage this new form of information space modeling.

4.5 Conclusion and Next Steps

The next steps that we hope to see ensue from the research we have undertaken on our CASE effort follow directly from the *conceptual foundations* that we describe in Section 1, at the outset of this paper:

> <u>Blending the Personal and the Collective</u>: A design goal for all of our work on this effort has been to ensure that the tools we develop provide both *individual* and *collective* benefits. The success and popularity of tag|Connect—and, more broadly, of a number of popular "social bookmarking" applications on the open Internet—provides a validation that not only is this goal achievable, but also that tools implemented in this manner can be quite successful. Although the Catalyst application that we have developed on this effort has not (yet) been deployed for use by the Intelligence Community, all indications from our extensive testing are that this same sort of "blending" can be successfully accomplished in an *analytic* tool like Catalyst as well.

In general, we believe that not only can this conceptual foundation serve our own continued work to support the Intelligence Community, but also that it can serve as guidance for others in helping the Community in continuing to move toward being a fully collaborative enterprise.

Infer, Rather than Enforce, Structure: In contrast to our first conceptual foundation, we did not have the time necessary to develop this foundation to the degree we would have liked. Nonetheless, our initial work in Relational Topic Models (RTMs), as described in Section 2.4, above, is *very* encouraging. If this technology were funded for further development, we consider it very likely that it could substantially enhance analysts' ability to connect with others and with information they need in order to be successful in their challenging jobs. More concretely, we see tremendous opportunities to apply these techniques to the data and information that analysts create and interact with in and through tag|Connect and Catalyst—and then to expose valuable connections (to information, and to other people) to analysts from within these tools. As well as (via Web Services) through other analytic tools that can make good use of these connections.

List of Acronyms

ACM	Association for Computing Machinery
AFRL	Air Force Research Laboratory
API	Application Programming Interface
CASE	Collaboration & Analyst System Effectiveness
CETIS	Center for Terrorism and Intelligence Studies
CGC	Context-Grounded Conversations
CIA	Central Intelligence Agency
DDNI/A	Deputy Director of National Intelligence for Analysis
HTML	Hypertext Markup Language
IBIS	Issue-Based Information Systems
ICES	Intelligence Community Enterprise Services
IM	Instant Messaging
IMO	Intelink Management Office
Intelink-U	Intelink-Unclassified Net
JWICS	Joint Worldwide Intelligence Communications System
LDA	Latent Dirichlet Allocation
MIIS	Monterey Institute for International Studies
NIE	National Intelligence Estimate
NIPF	National Intelligence Priorities Framework
NPR	Nuclear Posture Review
ODNI	Office of the Director of National Intelligence
OSW	Open Source Works
REST	Representational State Transfer
RIA	Rich Internet Applications
RSS	Real Simple Syndication
RTM	Relational Topic Model
SIPRNet	Secret Internet Protocol Router Network
SOA	Service Oriented Architecture
SOAP	Simple Object Access Protocol
SPAWAR	Space and Naval Warfare Systems Center
TOR	Terms of Reference
UI	User Interface
URL	Uniform Resource Locator
WPF	Windows Presentation Foundation
XAML	Extensible Application Markup Language
XML	Extensible Markup Language
XMPP	Extensible Messaging and Presence Protocol