A Synchronous Multimedia Annotation System for Secure Collaboratories

Ronald Schroeter
The University of Qld
St Lucia, Qld
Australia 4072
617 33652097
ronalds@itee.uq.edu.au

Jane Hunter
The University of Qld
St Lucia, Qld
Australia 4072
617 33652097
jane@itee.uq.edu.au

Michael Henderson
The University of Qld
St Lucia, Qld
Australia 4072
617 33652097
mjh@itee.uq.edu.au

Imran Khan
The University of Qld
St Lucia, Qld
Australia 4072
617 33652097
imrank@itee.uq.edu.au

ABSTRACT

In this paper, we describe the Vannotea system - an application designed to enable collaborating groups to discuss and annotate collections of high quality images, video, audio or 3D objects. The system has been designed specifically to capture and share scholarly discourse and annotations about multimedia research data by teams of trusted colleagues within a research or academic environment. As such, it provides: authenticated access to a web browser search interface for discovering and retrieving media objects; a media replay window that can incorporate a variety of embedded plug-ins to render different scientific media formats; an annotation authoring, editing, searching and browsing tool; and session logging and replay capabilities. Annotations are personal remarks, interpretations, questions or references that can be attached to whole files, segments or regions. Vannotea enables annotations to be attached either synchronously (using jabber message passing and audio/video conferencing) or asynchronously and stand-alone. The annotations are stored on an Annotea server (extended for multimedia content). Their access, retrieval and re-use is controlled via Shibboleth identity management and XACML access policies.

Categories and Subject Descriptors

H3.1 [Content Analysis and Indexing]; H3.5 Online Information services.

General Terms

Collaboration, Annotation, Multimedia, Security.

Keywords

Collaborative Multimedia Analysis Annotation

1. INTRODUCTION

A *Collaboratory* is the term used to describe the networked environment that allows geographically dispersed teams of researchers to work together. Using information and communication technologies, researchers can access each other, as well as the remote data and information, software, tools and instruments that they require to carry out their research. Key components of collaboratories are: robust security mechanisms; multi-user video-conferencing services; shared data repositories; instant messaging services; application sharing and shared analytical tools and services. Many research communities are generating research data in multimedia formats - images, video, audio, 3D objects and animations (dynamic data visualizations, simulations and models). As a result, research communities are increasingly seeking tools that enable them to collaboratively analyse and annotate such multimedia content, either synchronously or asynchronously. Examples of communities generating multimedia content that requires collective analysis and knowledge capture through shared annotation tools include:

- Oceanography projects such as Visions 2005 [1], which is generating hundreds of hours of underwater video and data streams that require real-time analysis by teams of scientists;
- CIMA X-ray Crystallography portal [2] which is generating video streams, data streams and xray diffraction patterns that when processed, produce 3D crystal structures;
- The Paradisec project [3] - a web-enabled facility for collaborative digitisation, management and access to Australian researchers' ethnographic audiovisual recordings of endangered languages and music from the Asia Pacific region.
- The UK eScience Integrative Biology project [4] which is generating MPEG movies showing computer simulations of electrical propagation of the heart and the growth of cancerous tumours.

Currently available annotation systems vary widely with respect to the types of content they annotate, the extent of collaboration and sharing they allow and the communities which they serve [5]. Although they have been successfully applied to domains including education [6, 7], research, medicine [8] and neuroscience [9] in order to capture and exchange metadata, ideas, opinions and interpretations, evaluation of these applications indicates limitations in existing commercial and prototype systems. Current systems are limited by: lack of responsiveness, use of non-standard proprietary technologies; lack of authentication of the source; limited search capabilities; limited access control; inability to reply to/stagger annotations; asynchronous sharing only; support for limited media types; coarse granularity and unstructured annotations (single field, free text only).

The main focus of the work described in this paper is the development of a secure collaborative multimedia annotation system for collaborators within eResearch environments – that can be used either asynchronously or synchronously. An essential requirement for such a sector is the need to be able to authenticate the source of the annotation and to restrict access to a particular group of trusted colleagues - for reasons of privacy, confidentiality or protection of intellectual property. This is particularly important within eScience, where the annotation or interpretation of the raw document or data, is often more valuable than the target of the annotation.

Our implementation involves combining and extending a number of existing open source technologies based on open standards:

- Annotea – a Web-based annotation server developed by the W3C as part of the Semantic Web initiative [10] which we have extended to support annotation of fine-grained contexts within multimedia objects;
- Jabber – this provides the instant messaging required for the real-time application sharing and event logging [11];
- Vic/rat – these videoconferencing tools [12] are extended to enable the recording of separate participant's H.261 streams and their conversion to tiles within a single MPEG movie
- Shibboleth – an Internet2 middleware initiative that enables identity management and secure access to Web resources shared amongst a federation of organizations [13];
- XACML (eXtensible Access Control Markup Language) – XML-based language for defining and enforcing access control policies [14].

The remainder of this paper describes in more detail the secure annotation system that we have built. It is structured as follows:

- Section 2 describes previous related activities in the development of secure multimedia annotation systems;
- Section 3 describes the overall architecture and implementation details of our system;
- Section 4 illustrates the user interface and system's functionality;
- Section 5 provides an evaluation of the system to date and describes future plans for this work
- Section 6 provides a brief conclusion.

2. BACKGROUND AND PREVIOUS WORK

Existing annotation tools (which enable users to attach personal notes, questions, explanations, etc. to documents) can be categorized according to the media types which can be annotated (text, web pages, images, audio or video, 3D) and the extent of collaboration supported [5]. Table 1 provides an overview of the different products, tools, systems and projects according to these categories.

Table 1: Survey of Existing Annotation System Capabilities

	Non-Collaborative	Collaborative	
	private local annotations	**shared (asynchronous)**	**Logged/realtime (synchronous)**
Text/ HTML	- text processors like MS Word, Adobe Acrobat, etc. - Firefox Scrapbook extension - Multivalent browser (PDF) [15]	Annotea [10], Cadiz et al [11] Connotea [16] OntoMat Annotizer Annozilla [17] WebNote [18] Del.icio.us [19]	Churchill et al. [20], Collaborative Information Browser [21]
Images	Adobe PhotoShop stickynotes	PAIS [22], Photo Annotator (Annotea + SVG) [23] Flickr [24]	mimio classroom [25] SMAT [26]
Audio/ Video	ANVIL [27]	MRAS [28] IBM EVA, IBM VideoAnnEx [29] DIVER [30] MeFeedia MIT Shakespeare XMAS [31] iVAS	DTVI [32] eSports [33] VidGrid [34]
3D	QuickTime Player [35]	Jung et al. [36] RedLiner SpacePen	

Table 1 shows that there are a large number of systems available – particularly for the annotation of web pages and images. Annotations vary from simple semantic tags to rich, structured annotations (free text, hyperlinks, ranking, language, audiovisual). The systems also vary in their ability to attach annotations to fine-grained segments or regions. Some systems only allow annotations to be attached to whole files or specific types of segments (e.g., keyframes). The functionality of interfaces for searching, browsing and presenting annotations also vary considerably. None of the systems surveyed provide advanced searching mechanisms over different types of structured annotations or annotation metadata (e.g., creator, data, language) – requirements of many researchers.

The majority of the systems in Table 1 are asynchronous. Of the systems that do support synchronous annotations, they generally do it through shared whiteboards, rather than application sharing. A limitation of this approach is that the annotations are saved as an image which makes search and retrieval over the annotation's content highly problematic.

The systems that most closely resemble Vannotea are the collaborative video annotation systems: eSports [29], DIVER [26] and VidGrid [30]. Microsoft's Distributed Tutored Video Instruction (DTVI) [32] system enables students to replay and discuss videos of lectures collaboratively. But because it is based on NetMeeting, it only supports 5-10 fps video replay and does not support real-time synchronous annotations. eSports enables collaborative viewing of sports videos but only supports the annotation of single frame snapshots through a shared whiteboard. The UK eScience VidGrid project enables geographically dispersed users to synchronously view a video and allows users to draw on top of the projected video in real-time. However the annotations and discussions are not recorded for later search, retrieval and re-use.

None of the systems in Table 1 provides a flexible architecture that can support a wide-range of high quality media types (including text, html, images, video, audio, 3D) through easy incorporation of alternative embedded plugins. The Vannotea system covers all of the cells in columns 2 and 3 of Table 1.

The survey also shows that some systems support private local access only whilst others permit sharing amongst groups and/or public access through storage on a web-accessible server. However none of the surveyed systems provide the level of robustness or the fine-grained role-based access control mechanisms that is required by collaborative teams of scientists engaging in eResearch.

3. SYSTEM IMPLEMENTATION

Figure 1 shows an overview of the system architecture. Vannotea has been implemented in C# using the .NET 2.0 Framework. The color scheme in Figure 1 illustrates how the different server side services are being accessed by different UI components within Vannotea.

The content provider(s) on the left can either lie inside (secure) or outside the Shibboleth Federation (publicly accessible). Vannotea can easily incorporate search interfaces to both kinds of multimedia research repositories.

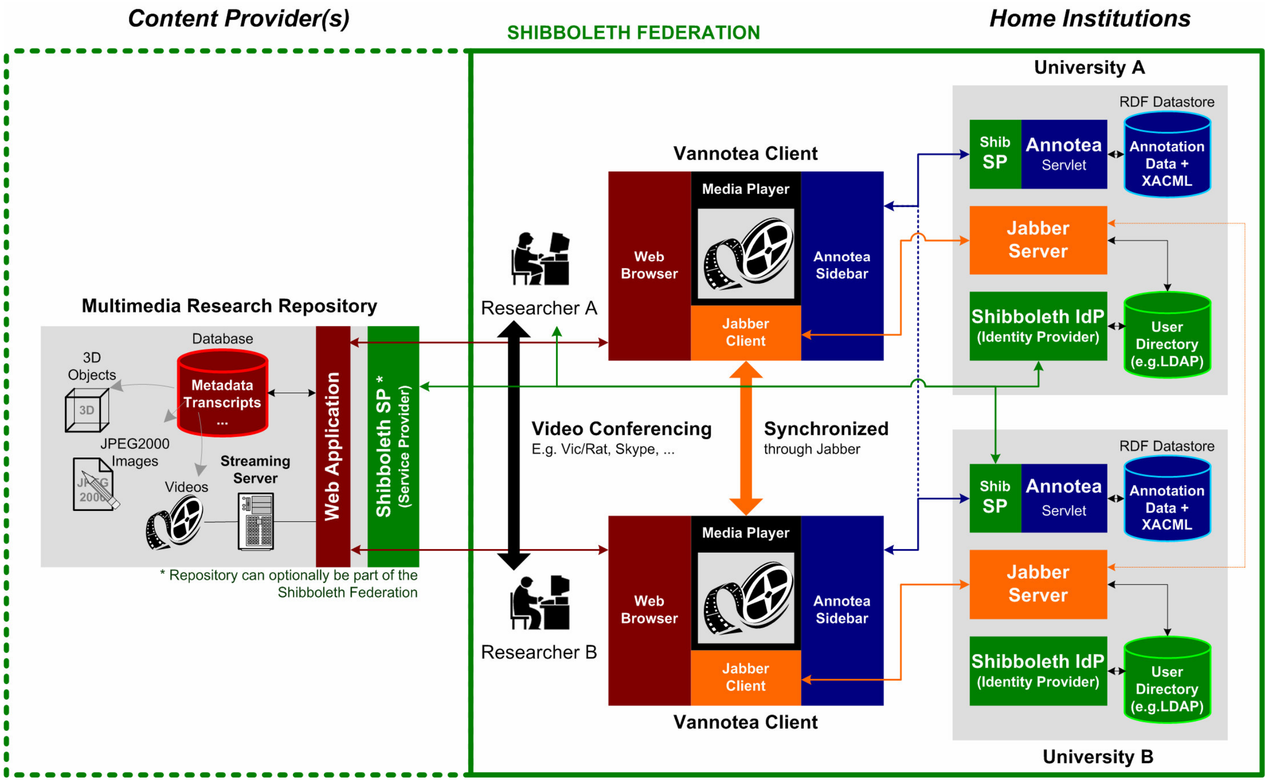

Figure 1: System Architecture

In the centre of Figure 1 are two researchers from different institutions that are part of a Shibboleth Federation. The researchers use video conferencing tools and the Vannotea Client to collaboratively discuss, analyze and annotate the scientific multimedia content they have retrieved from one of the content providers. Their Vannotea clients are synchronized through Jabber messages, which ensure that both researchers have exactly the same view of the multimedia content at the same time.

On the right we have the home institutions of the two researchers. Their universities provide access to services that include:

- the Shibbolized Annotea Server, where they can upload and retrieve annotations;
- a Jabber Server for instant messaging;
- and a Shibboleth Identity Provider (IdP), which authenticates users when they log on.

The IdP also releases attributes about the user to the Shibboleth Service Providers (SP) within the federation, e.g., to the content providers or the Annotea Servers, so they can make decisions regarding access rights.

3.1 The content provider

The content providers provide access to collections of domain-specific scientific multimedia data. They provide search, browse and retrieval interfaces to their collections through existing web portals or web applications using pre-existing metadata e.g. transcripts of linguistic data, technical specifications of microscopy images, provenance information etc.

The content may be publicly accessible or protected through site-specific user accounts. An emerging mechanism for restricting access to institutional research repositories is via Shibboleth. Institutions join a Shibboleth Federation and establish a Shibboleth Service Provider (SP). This controls sharing of data and services between institutions. Users can use SingleSignOn (SSO) from their home institution to log in and gain access to data or services within the Federation. We have chosen to use Shibboleth to implement the authentication and access control over our annotation servers.

Through the integration of a Wiki, the metadata of the content could be directly edited by the convening domain experts discussing the matter, e.g. to change the actual transcript of an interview.

The multimedia content can be either hosted on a Web Server, in which case the whole file is downloaded and cached on the client side by the appropriate player/viewer, or – for video content - can be streamed without downloading through any one of the available streaming servers, such as Darwin Streaming Server, Windows Media server, Video Lan Server, etc. Generally, the larger the file size of the video, the more essential it is that the video is streamed.

We deliberately leave the above decisions to the content providers. Vannotea's functionalities can be applied independent of the domain or discipline or their different security, quality, format and network requirements.

3.2 Shibbolized Annotea Server

Annotations are subjective, external comments, notes, reviews or references that can be attached to whole files, segments or regions. They can be in the form of free text, URLs or local files such as images or pdf files that are then uploaded to the server.

These annotations are stored on a HTTP enabled server. Our Annotea Server is implemented using Java Servlets hosted on a Tomcat Server. The RDF descriptions of the annotations are stored through the Jena API [37] in a MySQL database.

It is important to note that the annotations are stored **separately** from the content. This enables any community of users to discuss any collection but maintain control over access to these discussions. Annotation servers also vary in the extent of their distribution. They may be:

- public and centralized for general public access;
- private and centralized - installed on a single server and password protected by a specific group;
- distributed across a set of universities/institutions as shown in Figure 1. In this case new annotations are stored at the creator's home server i.e., the server of their home institution. Searches are performed across a list of servers and the results are aggregated on the client side.

3.2.1 Extended Annotea

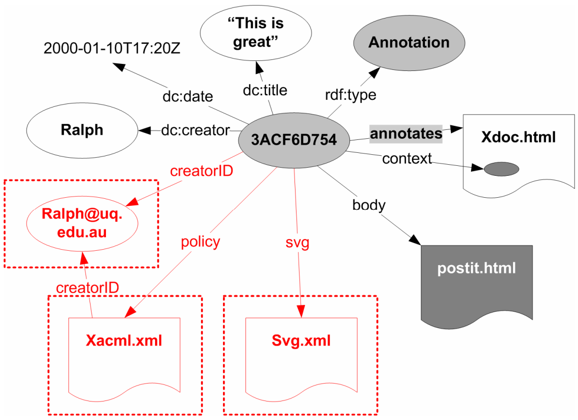

Figure 2: Extended Annotea Model (extensions in red)

Annotea is a simple but flexible Web-based annotation system that uses Resource Description Framework (RDF) to model annotations as a set of statements or assertions made by the author. Figure 2 illustrates the basic RDF annotation schema and our extensions: Every annotation stores basic metadata information including: the author, date of creation, title, type, which resource it annotates, the context within the resource (e.g. an XPointer to paragraph within a html document) and the body, the actual content of the annotation. Additionally we've added a unique creatorID field (provided by the Shibboleth IdPs), a link to an XACML policy that governs access to the annotation, plus an SVG field to store the coordinates of region outlines drawn on top of images, video frames, etc.

3.2.2 Shibboleth and XACML

In addition to the operations defined by Annotea (posting, querying, downloading, updating, replying and deletion of annotations) the Annotea server and protocol have been extended to support access control policies.

Users can define XACML policies (see 5.3.3 XACML Policy Creation) and attach them to their annotations. Policies are stored within the RDF repository, along with annotation bodies. They can be created either during the posting of a new annotation or independently of an annotation. Annotations are linked to particular policies through their policy property – which specified by a URL. This approach has the benefit of enabling multiple annotations to use the same policy. If a policy is modified, the changes will effect all those annotations associated with the policy

Policies are applied when a user tries to retrieve an annotation. The user's attributes are securely retrieved through the Shibboleth SP installation which queries the Shibboleth IdP for the user's attributes. The attributes are delivered through SAML assertions, which are translated into a XACML Request within the Annotea Servlet. Using Sun's XACML API [], an access decision is made based on the XACML Request and the XACML Policy, eventually either denying or granting the user access to an annotation.

Policies can be defined that allow users to be aware of the existence of annotations (e.g. list all annotations of a specific user) but do not reveal the actual content. This is important, because it allows the server to return results for statistical purposes, e.g., *how many annotations does this resource have*?

3.2.3 Integrating Annotation information on the content provider side

Although the Annotation servers are independent of the content provider, content providers may choose to embed annotation information retrieved from of a nominated list of Annotea servers into their web applications in some way. For example, to provide users with information about the distribution of annotations across a large collection, or to rank search results based on the number of annotations, content providers can post simple HTTP queries to the Annotea server.

3.2.4 Notification

A problem often identified in relation to the use of annotations in asynchronous collaboration [11, 38] is maintaining awareness of changes such as new annotations or replies to existing annotations. Cumbersome methods of dealing with this issue include manual checking or sending emails to notify creators of replies. To address this issue we have incorporated a simple feature which outputs changes to the annotation server as RSS by transforming the usual RDF output using XSLT. By querying the modified annotation server in the normal fashion (e.g. Give me all annotations by a creator for a particular resource) within a RSS Feed reader, users can keep up to date with any changes that may have occurred on an annotation server. This approach also ensures that the results of the RSS feed are filtered in exactly the same way as the secure annotation server, thus ensuring protection is maintained.

3.3 Jabber Server(s)

Within our setup we use the Open Source Wildfire Server [39], an instant messaging server, which uses the widely adopted open protocol for instant messaging, XMPP (also known as Jabber). Jabber enables any two entities on the Internet to exchange messages, presence, and other structured information in close to real time and in a secure manner. The messages pass through firewalls and proxies, so no additional setup is required by users or system administrators.

Jabber's distributed architecture is modeled after that of e-mail. Each user connects to a "home" server, which receives information for them, and the servers transfer data among themselves on behalf of users. Thus any domain can run a Jabber server. Each server functions independently of the others, and maintains its own user list. In addition, any Jabber server can talk to any other Jabber server that is accessible via the Internet (if server-to-server communications are enabled). The result is a flexible, controllable network of servers, which can scale much higher than the monolithic, centralized services run by legacy IM vendors such as AOL, Microsoft, and Yahoo. [40]

This fits nicely into the general philosophy of our distributed architecture. We also believe that Jabber represents the future for instant messaging. Wildfire (amongst other Jabber Servers) integrates nicely into existing authentication systems such as LDAP because the same LDAP directory that is used to authenticate users within the Shibboleth architecture can be used to authenticate users on the Jabber Server. Essentially the user only has to remember one username/password, the one at his home institution (Single Sign-on).

3.3.1 Synchronization over Jabber

Within our architecture, we use Jabber messages to keep current Vannotea Clients synchronized. The feature of Jabber that enables this is a *conference room* or *group chat*. In a group chat all messages that are posted to the conference room are being received by all the participating clients. Figure 3 illustrates how this is used within Vannotea.

For example, *Client A* pauses the video at 01:25.45. This information is passed to the embedded Jabber Client, which sends it as a message to the conference room on the Server. Jabber is based on a publish/subscribe architecture. By entering a conference room, the client subscribes to all the messages that are being sent to the room. Since all clients in Figure 3 have entered the room, they all receive the message that they should pause the video at 01:25.45. This information is passed to the MediaPlayer of each Vannotea client, which handles it accordingly.

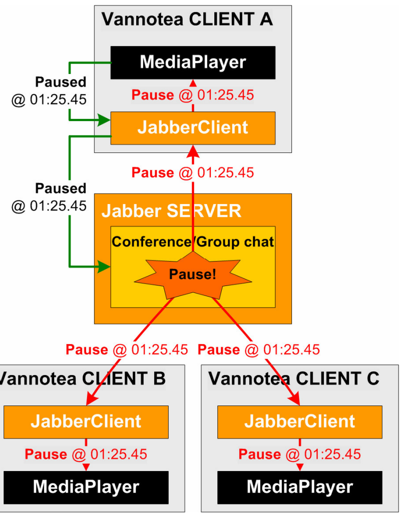

Figure 3: Synchronization over Jabber

It is important to note that even Client A, who sent the event, is handling the event through the Jabber Server. As a result the Jabber Server is the entity that keeps everything in sync, e.g. if two clients happen to send different events to the server almost simultaneously, they are still fired by the Server in the order they came in and sent to the client in that same order, with the last event describing the current state of the application.

4. USAGE SCENARIO

Below is an example of a usage scenario which illustrates how all of the key features of Vannotea are typically utilized within a collaborative eResearch/eScience Marine Science project.

A Marine Biologist from the University of Queensland (UQ) browses through the latest collection of deep water video footage, which was made available by the Visions'05 project at the University of Washington (UW). UQ and UW are both part of a Shibboleth Federation. Visions'05 has set up access policies that allow staff members from the Marine Biology department at UQ to gain full access to their online repository as part of their collaboration.

The Marine Biologist opens a particular recent video within Vannotea so he can bookmark and attach his personal notes to segments of the video, keyframes or regions within frames. To secure his notes, he reuses an existing policy that grants access to all of the participants of this collaboratory. Through this policy, his notes are securely stored on an Annotation Server at UQ. The Annotation Server is shared amongst other departments at UQ and other members of the Federation but his annotations are only visible to those individual participants of this collaboratory.

The Marine Biologist notices a close up of a unrecognizable tubeworm species in a hydrothermal vent being filmed in one of the videos. He highlights it with one of the drawing tools and posts a question "unknown species of tubeworm?" to the annotation server. It is early in the morning, and his Jabber contact list indicates that his colleague, an Oceanographer at UW, is still online. He fires up a videoconferencing tool to speak to him. The Oceanographer starts Vannotea, and is invited to join a Jabber conference room by the Marine Biologist. This triggers an event that opens up the same video at the same location inside the Oceanographers Vannotea Client. It also retrieves all the annotations of everyone within the collaboratory about this video, including the "unknown species of tubeworm" annotation the Marine Biologist posted earlier.

The Marine Biologist hits the record button. This records their conversation and the application events they fire whilst collaboratively watching and browsing the same video content. They start looking for other occurrences of the tubeworm while discussing the matter. Every time they see one, they pause the video, highlight the region and store it as an annotation reply of type bookmark to the "unknown species" annotation on the Annotation server. As a result of their discussion, they not only get a list of all locations of the unidentified tubeworm within the video, but they also narrow the classification down to three different possible species of tubeworm. They terminate the session and the Marine Biologist uploads the recorded session, including the video/audio and a time stamped log file of the application events, to a shared repository. He also creates an annotation that links to the audiovisual recording of the videoconference session.

A few hours later, another member of the collaboratory at the University of London starts his work day. He checks his RSS feeds and notices new annotations about an "unknown species of tubeworm" on the annotation servers that he has access to. A simple double-click opens the annotation and the video inside Vannotea and jumps to the frame that shows the worm. He goes through the list of bookmarks that point to the different occurrences. Still unsure, he retrieves the discussion and replays the complete session between his two colleagues earlier. Curious, he starts researching online publications relating to the three possible worm species and is able to reject them through his findings, which he links by posting further replies.

By using Vannotea, the globally distributed collaborators are able to analyse, annotate, and share their knowledge about the multimedia research data, either synchronously or asynchronously without compromising security. The possibility that they have discovered a new species of tube worm that only survives in high temperatures and deep sea conditions of hydrothermal vents is kept confidential until the collaborators are ready to publicize their findings.

5. USER INTERFACE

Actual screenshots of Vannotea's UI are attached in Appendix A. Figure A.1 shows Vannotea being used within the Paradisec Ethnographic Analysis Project [3]. Figure A.2 shows the Annotea Sidebar being used by protein crystallographers based at the Institute of Molecular Biology (IMB) at the University of Qld.

The UI consists of the following embedded components:

- A Web Browser (Internet Explorer),
- Various embedded media players such as Quicktime, Windows Media Player, Video Lan Client, etc.,
- an Annotea Sidebar, and
- a Jabber Client.

5.1 Web Browser

The Web Browser is used to access the web interface provided by the content provider. Alternatively users can choose to use their external browser and drag & drop the content into Vannotea. The reason we embedded a Web browser was to enrich user-friendliness and functionality by tailoring the web application for Vannotea. Simple JavaScript can be added to the web application that allows communication between the web site and Vannotea, e.g. to enable browsing through different segments of a video. By clicking a segment, the web application passes the information about the media file (URL), and the segment (id, start and end point) to Vannotea, which then jumps to the start and plays the segment. Communication also works in the other direction, e.g. Vannotea can call scripts to retrieve and highlight current segments.

5.2 Media Players

The media players display the multimedia content that the users want to annotate. The mediafiles are rendered by various embedded ActiveX controls, e.g. the Quicktime 7 control [35], the Windows Media Player 10 SDK [41] and the VideoLan Client ActiveX control [42]. This ensures that the majority of video and audio formats are supported. We are currently in the process of embedding players to display 3D formats e.g., 3D crystallographic structures through JMOL [43].Because Quicktime is embedded, 3D objects in QTVR format are also supported.

To view different media types, any third party ActiveX controls/plug-ins for that particular type can be easily embedded into Vannotea. The only prerequisite is that the control allows us to programmatically get and set the view of the current media. In terms of videos, we need to be able to get and set the current media file, the current time and the current state (play, pause, etc.). For 3D objects, we need to be able to do the same for the zoom factor, camera angle, pan or tilt angle and time if it is part of an animation. For images, the zoom factor and the coordinates of the currently enlarged section of the image.

Commonly required functionalities (that depend on the media type) are provided, e.g. for videos the user can Play and Pause, step a frame forward/backward and seek to a specific position using a slider. For 3D objects, users can pan around, zoom in and out and rotate the object.

Additionally a transparent drawing panel is layered on top of every player enabling the user to draw simple shapes on top of the content to point out or highlight important bits. The drawing information is stored as SVG. The drawing can be done collaboratively, in which case the SVG string is exchanged over Jabber and drawn respectively. Every user participating in the collaboration can see each others drawings similar to collaborative whiteboards, aiding the ongoing discussion about the content. The user stores the coordinates of these drawings as part of an annotation (see Figure 2). It is important to note that only the drawing information and the reference to the file/view are stored. This avoids copyright infringements by making copies through screen captures, the approach used by many whiteboard based annotation systems. When an annotation is retrieved, the player opens the file (*annotates*), seeks to the specified view (*context*), and draws the shapes (*svg*) on top.

For the 3D crystallographic structures, JMOL allows users to select particular uniquely identified atoms for annotation. We are currently investigating 3D volume specification methods to allow users to specify multiple strands or ligands.

5.3 Annotea Sidebar

The Annotea Sidebar has also been implemented in C# and .NET. It can be installed as a sidebar plug-in for Internet Explorer to annotate web pages asynchronously and is embedded as a control in Vannotea.

5.3.1 Annotation Retrieval

The Annotea Sidebar automatically displays a list (or tree for discussion threads) of all the annotations that refer to the current active media file. Using the open Annotea Protocol over HTTP, it retrieves the annotations from a list of servers that the user has specified during system configuration. The sidebar only shows the annotations that the user has access to.

For time-continuous media files, the annotations' *context* points to specific timestamps or temporal segments. Because the time information is not reflected in the list/tree view, we implemented a timeline (Figure 4) to display the annotations across the duration of the current media file, e.g. audio, video, or an animation, and added it as an optional view to Vannotea. The timeline is constantly synchronized with the annotations listed in the sidebar. This gives the user an instant overview of the annotations on the video or animation. This provides an indication of the important and most relevant segments of a video or animation.

Figure 4: Annotations on a timeline, grouped by authors

The different tracks on the timeline are by default grouped by the different authors, but they can also be grouped by annotation type or creation date.

The annotation server supports searching the body of annotations as well as complex queries across the annotation metadata e.g., *show me annotations by a particular author, between given dates and containing the term "Ethnography"*. This functionality on the server side is accessible through the search interface (Figure 5) which allows users to specify searches quickly and intuitively. Although we have chosen to implement the search user interface within the browser sidebar, we are also considering the development of a web portal style interface which users can access without having the sidebar installed. This web portal will provide a broader overview of the distribution of annotations across collections – thus indicating information which is of greater interest and relevance.

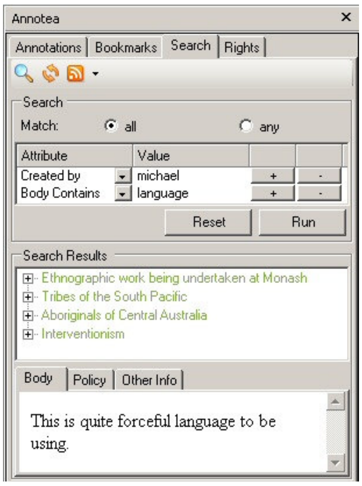

Figure 5: Search Interface for annotations

5.3.2 Annotation Creation

Figure 6 illustrates the user interface for creating and attaching an annotation to a digital multimedia object. We have extended Annotea to support structured annotations that contain a number of fields including hyperlinks, files, free text or controlled vocabularies through drop down boxes.

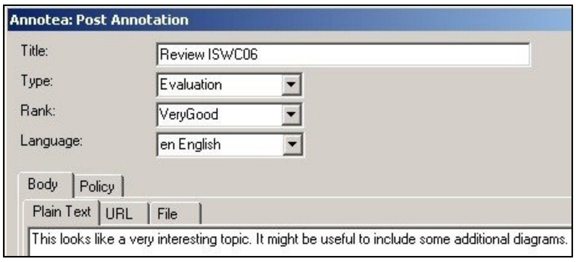

Figure 6: User Interface for creating/editing an annotation

5.3.3 XACML Policy Creation

Figure 7 shows the interface which was developed to define policies. It consists of two main parts; the definition of access control rights to a particular user group and the definition of user groups (based on particular eduPerson attributes [44]).

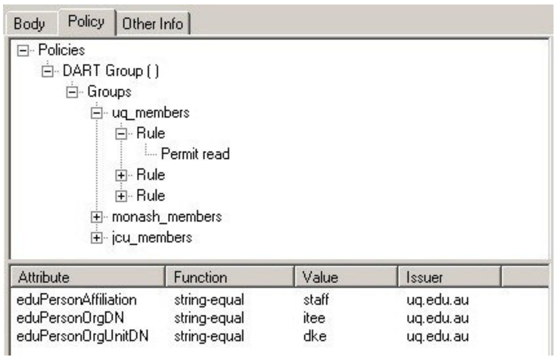

Figure 7: User Interface for defining access policies

For more complete details of the access policy implementation refer to [45].

5.4 Jabber Client

The embedded Jabber Client (see Appendix A.1), like any other chat client, manages a list of contacts and their presence information. It provides all of the common functionalities such as adding new contacts and deleting them, allowing/disallowing users to see presence information, and inviting other users into a conference room or group chat for collaboration and discussion over the same content. We used Jabber-NET [46], a library written in C# to access the Jabber Server and handle Jabber messages.

5.5 Event Logging and Replay

The event logging and replay is part of the Jabber Client. By clicking the record button during a collaborative session, all the incoming Jabber events are being time stamped and stored in a simple log file. The information of the log file is displayed along a user friendly timeline (see Figure 8). Clicking on any of the events displayed in the timeline, replays the session from that event.

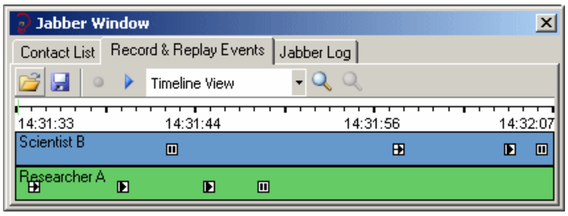

Figure 8: Record & Replay of events

5.6 Video conference recording

The Vic video conferencing tool is being extended to support session recording. Vic receives an RTP stream of H.261 encoded video from each site participating in the video conferencing session. The incoming RTP packets may be received in any order, so timestamps and site identification information are extracted. The packet is then decoded, synchronised and used to update the video stream for the associated session participant. Vic has been modified to record an output video consisting of the video streams from all session participants. The decoded frames from the individual sites are tiled into a single image and used to create an output video. ImageMagick [47] is used to tile the images and pack the frames into an animated sequence. Rat is used to record the audio streams. The recorded session is stored in MPEG format so that it can be replayed using a generic movie player.

6. EVALUATION

System testing has involved testing the creation, editing and deletion of (the complete set of possible) policies and annotations on images, video, audio, HTML and 3D objects. We tested policy enforcement by logging on as users with different attributes and modifying attributes directly in the LDAP directory. In all cases the annotation server behaved as expected – restricting user access to annotations in compliance with the policies. However the testing phase did reveal a number of problematic issues. These included:

- Deletion and update of annotations can lead to 'hanging references' where replies refer to annotations which have been updated or deleted.
- The use of URLs to identify policies enables them to be re-used and applied to multiple annotations. However this may cause problems when a policy referred to by multiple annotations is deleted.

A previous version of the system used .NET Remoting for handling and synchronizing remote events. Comparison of Jabber and .NET Remoting indicated the following:

- .NET Remoting handled events faster. The parsing of Jabber messages etc. can cause longer delays, but we expect that this will become more efficient in the future.
- Jabber made Vannotea significantly more robust. Dealing with remote events, multiple threads over network boundaries, proxy and firewall issues, caused many problems
- Jabber is much more flexibile. Users can reuse their contact lists, and the string based messaging allows easier integration into record & replay tools or the Compendium tool for meeting archival.

We are currently working with researchers from the Paradisec project [3] and the Crystallography portal [2], carrying out user evaluation studies of the Vannotea system within these two very different disciplines. We are also in the process of arranging deployment of the system to the oceanography [1] and integrative biology [4] projects for their use. User feedback has on the whole been very positive. The ethnographers requested the ability to display multiplexed annotations on the same segment. The crystallographers requested the ability to annotate strands or ligands (as opposed to just single atoms).

Although we were able to apply our annotation tools to 3d crystallography structures quickly and easily, the context describing the annotated region is currently an application-specific string. Using Jmol specific camera coordinates allowed us to demonstrate 3d structure annotation immediately. However, further work is required to standardize that information and make it accessible to a wider spectrum of 3d crystallography applications.

The Vic/Rat recording of the videoconferencing session corresponding to an annotation session is proceeding. Although the precise synchronization of the separate audio and video streams is proving difficult.

7. FUTURE WORK AND CONCLUSIONS

In the immediate future we are planning to investigate the following aspects of and extensions to Vannotea:

- More detailed and extensive user evaluation studies;
- Annotation of PDF files, databases and spreadsheets. The diversity and commercial nature of these file formats may prove problematic;
- Semantically enriched annotations through ontologies;
- Extension of the existing Open Source Annozilla plugin for Firefox to provide similar functionalities to the IE sidebar to make secure annotation services available to Firefox users;
- Searching the RDF annotations via a SPARQL [48] interface. This approach would be more efficient and flexible than using Algae [49] (the query language developed for Annotea);
- Allow arbitrary creation of classified links to other resources e.g., an RDF graph of connections that can be traversed;
- A repository of synchronized event logs and video conference recordings that can be searched and replayed.

To conclude, this paper describes a secure, collaborative synchronous annotation service for multimedia that we have developed by combining and extending a number of existing open source technologies. The proliferation of eResearch activity has led to a demand for such a service across many disciplines. By providing researchers with the necessary support for authenticating the source and protecting the confidentiality and intellectual property of their annotations, they will be more willing to share their views and engage in inter-organizational collaborations with trusted colleagues. Moreover, the modular design and interoperable technologies that we have adopted, makes it easy to quickly adapt the server to a variety of different media types, different domains and different communities.

8. ACKNOWLEDGMENTS

The work described in this paper has been funded by the GrangeNet program which is a component of the Australian Federal Government's BITS Advanced Network Initiative (Department of Communications, Information Technology and the Arts).

9. REFERENCES

[1] "Visions 05 Project." Available at: http://www.visions05.washington.edu/.
[2] "CIMA X-ray Crystallography portal." Available at: http://iri164.informatics.indiana.edu:8080/gridsphere/gridsphere.
[3] "The Paradisec Project." Available at: http://paradisec.org.au/.
[4] "The UK eScience Integrative Biology Project." Available at: http://www.integrativebiology.ac.uk/.
[5] R. M. Heck, S. M. Luebke, and C. H. Obermark, "A Survey of Web Annotation Systems," Department of Mathematics and Computer Science, Grinnell College, Grinnell, Iowa 1999.
[6] E. Desmontils, C. Jacquin, and L. Simon, "Dinosys: An Annotation Tool for Web-Based Learning," in Advances in Web-Based Learning – ICWL 2004: Third International Conference, Beijing, China, pp. 59 - 66, 2004.
[7] M. Kirsch-Pinheiro, M. R. S. Borges, and J. V. d. Lima, "A framework for awareness support in groupware systems," Comput. Ind., vol. 52, pp. 47-57.
[8] M. Lewkowicz, G. Lortal, A. Todirascu, M. Zacklad, and M.-F. Sriti, "A Web-based Annotation System for Improving

Cooperation in a Care Network," in ICWE Workshops, pp. 227-239, 2004.
[9] M. Gertz, K.-U. Sattler, F. Gorin, M. Hogarth, and J. Stone, "Annotating Scientific Images: A Concept-Based Approach," in 14th International Conference on Scientific and Statistical Database Management 2002.
[10] J. Kahan, M.-R. Koivunen, E. Prud'Hommeaux, and R. R. Swick, "Annotea: An Open RDF Infrastructure for Shared Web Annotations," in Proceedings of the WWW10 International Conference, Hong Kong 2001.
[11] J. J. Cadiz, A. Gupta, and J. Grudin, "Using Web annotations for asynchronous collaboration around documents," in Proceedings of the 2000 ACM conference on Computer supported cooperative work. Philadelphia, Pennsylvania, United States: ACM Press, 2000, pp. 309-318.
[12] "Mbone Conferencing Applications," UCL Network and Multimedia Research Group. Available at: http://www-mice.cs.ucl.ac.uk/multimedia/software/.
[13] "Shibboleth Project," Internet2. Available at: http://shibboleth.internet2.edu/index.html.
[14] M. Lorch, S. Proctor, R. Lepro, D. Kafura, and S. Shah, "First experiences using XACML for access control in distributed systems," in Proceedings of the 2003 ACM workshop on XML security Fairfax, Virginia 2003.
[15] T. A. Phelps and R. Wilensky, "The Multivalent Browser: A Platform for New Ideas," in Document Engineering 2001, Atlanta, Georgia 2001.
[16] B. Lund, T. Hammond, M. Flack, and T. Hannay, "Social Bookmarking Tools (II)," D-Lib Magazine, vol. 11, 2005.
[17] "Annozilla (Annotea on Mozilla)," Mozdev.org. Available at: http://annozilla.mozdev.org/.
[18] "Webnote - an online tool for taking notes." Available at: http://www.aypwip.org/webnote/.
[19] "del.icio.us." Available at: http://del.icio.us/.
[20] E. F. Churchill, J. Trevor, S. Bly, L. Nelson, and D. Cubranic, "Anchored conversations: chatting in the context of a document," in Proceedings of the SIGCHI conference on Human factors in computing systems: ACM Press, 2000, pp. 454-461.
[21] Q. Jin, N. Furugori, K. Hanasaki, N. Asahi, and Y. Ooi, "Collaborative information browser: collecting and organizing information through group collaboration," in 2002 IEEE International Conference on Systems, Man and Cybernetics, pp. 5 pp. vol.6, 2002.
[22] W. B. Lober and J. F. Brinkley, "A Portable Image Annotation Tool for Web-based Anatomy Atlases," in Proceedings AMIA Symposium 99 1999.
[23] J. Ley, "Photo Annotator - Annotating Images with SVG." Available at: http://jibbering.com/svg/AnnotateImage.html.
[24] "Welcome to Flickr - Photo Sharing." Available at: http://www.flickr.com/.
[25] "mimio classRoom," Virtual Ink Corporation. Available at: http://www.mimio.com/meet/classroom/.
[26] M. P. Steves, M. Ranganathan, and E. L. Morse, "SMAT: Synchronous Multimedia and Annotation Tool," in 34th Annual Hawaii International Conference on System Sciences, 2000.
[27] M. Kipp, "Anvil - A Generic Annotation Tool for Multimodal Dialogue," in 7th European Conference on Speech Communication and Technology (Eurospeech), Aalborg, pp. 1367-1370, 2001.
[28] D. Bargeron, A. Gupta, J. Grudin, and E. Sanocki, "Annotations for Streaming Video on the Web: System Design and Usage Studies," Computer Networks, vol. 31, pp. 1139-1153, 1999.
[29] "MPEG-7 Annotation Tool," IBM. Available at: http://www.alphaworks.ibm.com/tech/videoannex.
[30] R. Pea, M. Mills, J. Rosen, K. Dauber, W. Effelsberg, and E. Hoffert, "The Diver Project: Interactive Digital Video Repurposing," IEEE MultiMedia, vol. 11, pp. 54-61, 2004.
[31] "XMAS: Cross-Media Annotation System," MIT Shakespeare Project. Available at: http://icampus.mit.edu/projects/xmas.shtml.
[32] J. J. Cadiz, A. Balachandran, E. Sanocki, A. Gupta, J. Grudin, and G. Jancke, "Distance learning through distributed collaborative video viewing," in Proceedings of the 2000 ACM conference on Computer supported cooperative work. Philadelphia, Pennsylvania, United States: ACM Press, 2000, pp. 135-144.
[33] G. Zhai, G. C. Fox, M. Pierce, W. Wu, and H. Bulut, "eSports: Collaborative and Synchronous Video Annotation System in Grid Computing Environment," in Seventh IEEE International Symposium on Multimedia (ISM'05), pp. 95-103, 2005.
[34] M. Fraser, G. Biegel, K. Best, J. Hindmarsh, C. Heath, C. Greenhalgh, and S. Reeves, "Distributing Data Sessions: Supporting remote collaboration with video data," in Proceedings of the First International Conference on e-Social Science, Manchester, UK 2005
[35] "Scripting the COM/ActiveX Control in QuickTime 7 for Windows," Apple. Available at: http://developer.apple.com/quicktime/activexcontrol.html.
[36] T. Jung, M. D. Gross, and E. Y.-L. Do, "Annotating and sketching on 3D web models," in 7th international Conference on intelligent User interfaces (IUI '02), San Francisco, California, USA, pp. 95-102, 2002.
[37] B. McBride, "Jena: A Semantic Web Toolkit," IEEE Internet Computing, vol. 6, pp. 55-59, 2002.
[38] A. J. B. Brush, D. Bargeron, J. Grudin, and A. Gupta, "Notification for shared annotation of digital documents," in Proceedings of the SIGCHI conference on Human factors in computing systems. Minneapolis, Minnesota, USA: ACM Press, 2002, pp. 89-96.
[39] Wildfire Server, Jive Software. Available at: http://www.jivesoftware.org/wildfire.
[40] "Technical Overview," Jabber Software Foundation. Available at: http://www.jabber.org/about/techover.shtml#distributednetwork.
[41] "Windows Media Player 10 SDK," Microsoft. Available at: http://www.microsoft.com/windows/windowsmedia/forpros/platform/sdk.aspx.
[42] "VideoLAN." Available at: http://www.videolan.org/.
[43] "Jmol." Available at: http://jmol.sourceforge.net/.
[44] "EduPerson Object Class Specification (Draft)," Internet2 Middleware Architecture Committee for Education, Directory Working Group (MACE-Dir), 2006. Available at: http://www.nmi-edit.org/eduPerson/draft-internet2-mace-dir-eduperson-latest.html.
[45] I. Khan, R. Schroeter, and J. Hunter, "Implementing a Secure Annotation Service," in IPAW06, Chicago 2006.

[46] "Jabber-NET," JabberStudio. Available at: http://jabberstudio.org/projects/jabber-net/project/view.php.
[47] "ImageMagick." Available at: http://www.imagemagick.com/.
[48] E. Prud'hommeaux and A. Seaborne, "W3C Working Draft - SPARQL Query Language for RDF," 2004. Available at: http://www.w3.org/TR/2004/WD-rdf-sparql-query-20041012/.
[49] "Algae HOWTO," W3C. Available at: http://www.w3.org/1999/02/26-modules/User/Algae-HOWTO.html.

Appendix A

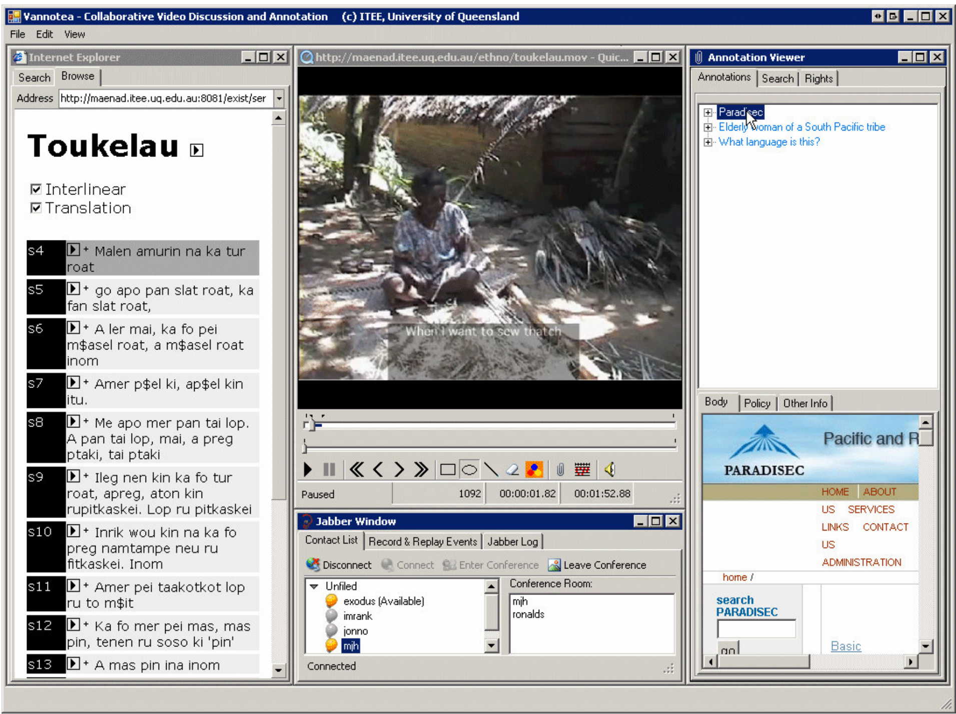

Figure A.1: Screenshot of Vannotea within Paradisec Ethnographic Analysis Project

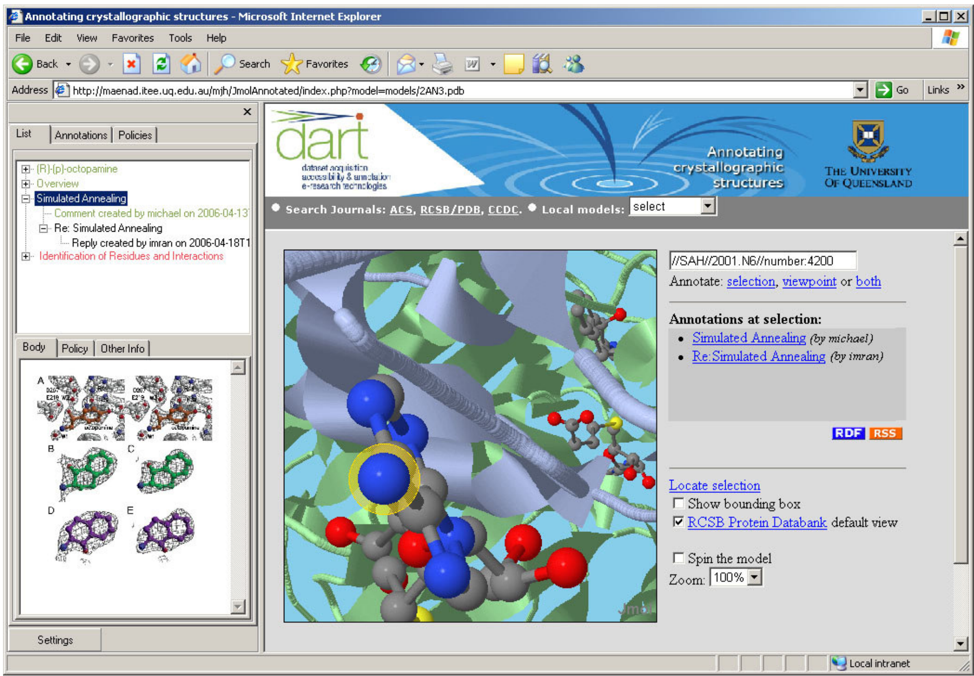

Figure A.2: Annotea Sidebar used within Institute of Molecular Biology (IMB)