

Collaborative environment of the PROMISE infrastructure: an "ELEGantt" approach

Marco Angelini
Sapienza University of Rome
Italy
angelini@dis.uniroma1.it

Guido Granato
Sapienza University of Rome
Italy
granato@dis.uniroma1.it

Claudio Bartolini
HP Labs
USA
claudio.bartolini@hp.com

Preben Hansen
SICS
Sweden
preben@sics.se

Gregorio Convertino
Xerox Research
Centre, Europe
convertino@xrce.xerox.com

Giuseppe Santucci
Sapienza University of Rome
Italy
santucci@dis.uniroma1.it

ABSTRACT

This paper focuses on developing lightweight tools for knowledge sharing and collaboration by communities of practice operating in the field of information retrieval. The paper contributes a motivating scenario, a characterization of these communities, a list of requirements for collaboration, and then a system design proposed as a proof-of-concept implementation that is being evaluated.

1. INTRODUCTION

This paper focuses on the problem of supporting knowledge sharing and collaboration in communities of practice that operate in the field of information retrieval (IR). These communities include developers, researchers, and stakeholders who periodically collect and use scientific data produced by the experimental evaluation of IR systems. Specifically, the communities considered include those involved in three specific IR domains: Patent, Cultural Heritage, and Radiology.

The research context of the work reported in this paper is the PROMISE NoE. This project aims at advancing the current tools for IR communities to perform experimental evaluation of complex multimedia and multilingual information systems. The ultimate goal of the project is to develop a unified infrastructure for the community to efficiently collect and reuse data, knowledge, tools, methodologies, and communities of end users. In this context, providing adequate support for collaboration is crucial. Herefrom the specific goal of the work reported in this paper: designing and evaluating lightweight support for knowledge sharing and collaboration.

Currently, the following problems result from lack of suitable collaboration tools:

- 1) Greater **effort** is required by individual members, who contribute as volunteers, for sharing knowledge and collaborating. In the long term, this discourages broader participation.
- 2) Poor **reuse of content** and process information across the multiple instantiations of similar experimental evaluation processes. Over time, this leads to inefficient processes: e.g., content is al-

ways recreated from scratch, successful processes (best practices) cannot be reused, novices cannot be easily trained based on shared experience.

- 3) The overall community cannot easily reflect on (and thus re-engineer) its own **workflow** around specific TRECs.

2. MOTIVATING SCENARIO

The starting point of our analysis is a typical IR evaluation campaign (lab). In a typical scenario, Adam (lab organizer) is preparing an IR experiment and evaluation task and spends time and resources for coordinating, communicating and assembling people and resources in order to proceed with the overall evaluation task, e.g. recruiting people that will be responsible for different evaluation task(s). Communication and sharing of information may be different within different across sub-tasks. Furthermore, they may be different between labs without any awareness among actors of the similarities/differences in the evaluation task processes. Thus, it is important to identify the stages in the evaluation task process as well as how collaborative and information sharing activities are manifested.

3. CHARACTERIZING IR COMMUNITIES

The CLEF experimental platform involves a series of CLEF Labs and one or more tracks within each Lab. Each Lab as well as each track involves a certain set of tasks that could be considered as a task process or workflow. In order to define and describe these tasks we have investigated the lab and track organizers of a CLEF experiment, how they performed their work and what steps they went through during their work. Furthermore, we have extracted requirements for collaborative information handling and information sharing activities specifically [3, 5]. An evaluation campaign is an activity intended to support IR researchers providing a large test collection and uniform scoring procedures. An evaluation campaign is organized within an evaluation framework like TREC or CLEF and can involve different domains (cultural heritage, patent, radiology and so on). Within an evaluation campaign there are many tracks, such as multimedia, multilingual, text, music, images, etc. A track can be organized differently according to a specific domain and include, in turn, several tasks. A task is used to define the structure of the experiment, specifying a set of documents, a set of topics, and a relevance assessment. For each task the set of documents can be structured defining, for example, a title, keywords, images and so on. A topic represents an information need. Documents can be assessed as being relevant or not (or more or less relevant) for a given information need (topic).

Some of the most common tasks that we observed as part of a




Figure 1: Task stages of organizing an experimental IR track in CLEF.

typical evaluation campaign include the followings: submission, preparation, track definition, topic creation, data set, relevance assessment, summarizing, and finalizing.

3.1 Observed roles in IR Communities

Within an evaluation campaign many people are involved in different tasks, such as organizing, creating topics, managing collections, handling participants and submission, choosing measures, and running the final evaluations.

The set of actors involved in PROMISE activities is not homogeneous and depends on the domain which is taken into account. Looking at three domains (patent, medical, and cultural heritage), we defined the following actors:

- **organizers:** people who are in charge of preparing a campaign; it is possible to distinguish domain organizers and track organizers;
- **participants:** people who run their algorithm(s) according to the actual tasks;
- **relevance assessors:** people who make the relevance assessment;
- **topic creators:** people who define topics for a given task;
- **site administrators:** (e.g. system administrators);
- **other researchers**
- **annotators:** people who annotate resources to highlight some hidden information.

Each of these actors is described along with a set of activities or tasks. Moreover, a user can have more than one role.

3.2 Observed tools of IR Communities

The collaborative work may be performed in a non-structured manner using basic tools for collaboration in an *ad hoc* fashion.

The following are the most commonly used tools for collaboration:

- **E-mail:** the most common way to organize the work and spread information;
- **Face-to-face meetings:** useful to discuss more effectively about problems and solve them; it is complex if people don't work in the same building;
- **Video conference tools** (e.g. Skype): used instead of face-to-face communication;
- **Shared workspaces** (e.g. shared document editor, desktop sharing): useful to share documents. However an issue may arise where many people work on the same document.

4. COLLABORATION REQUIREMENTS

As mentioned in the introduction, a more suitable collaborative tool is needed to help researchers to accomplish their tasks. To realize it we have to overcome some limitations.

The first one is the impossibility to define a common detailed workflow due to the presence of different domains, each of them with specific needs. This makes it difficult to realize a collaborative environment completely specified. Despite this limit, it is possible to individuate some common needs such as: communication with other actors, access to data of previous campaign, sharing of task flow of actual evaluation campaign, and sharing workspace with actors involved in the same tasks.

Another aspect that characterizes the work of people involved in a lab is that there is an alternation between individual and collaborative work, which is in contrast to a too rigid environment.

The basic idea of our system is to improve the actual tools used in IR community without defining the collaborative environment in a too rigid way. Our purposes in this paper are:

- 1) Collecting the tools actually used (e.g. Skype, GoogleDocs, etc.) in a structured environment.
- 2) Making available to users other collaborative tools (polls, news).
- 3) **T@GZ:** a social software system for organizational information sharing.

4.1 Requirements by role

In order to get information on the actual tasks users performed, we made requirements elicitation through a number of questionnaires to track organizers within the three predefined domains of Patent, Cultural Heritage, and Radiology in the CLEF platform. Working through these questions we identified: a) different **roles** of actors involved in the CLEF experiments, b) **requirements** for collaborative information handling activities and information sharing and c) **links** between roles and collaborative events. Furthermore we describe each task stage regarding subtask involved (fig.1):

- **Submission task:** preparing a lab proposition including sub tracks, acceptance of the lab or not.
- **Preparation task:** preparation of a CLEF lab flyer including details of each lab, obtaining databases, preparing a copyright agreement, preparation of the web page.
- **Track definition task:** definition of broad tasks, start of registration for the participants.
- **Topic creation task:** preparation of detailed topics for each task, release of topics, checks of the copyright forms.
- **Data set:** data access for all registered participants with signed copyright forms for their tracks.
- **Relevance judgment task:** preparation of the judgement system; finding qualified judges; submission of all runs by participants; pooling of the runs to create documents to judge; judgement of the documents for relevance; evaluation of all submitted runs.
- **Summarizing task:** release of ground truth; submission of participants' papers with results and technical descriptions; analysis of the results and submission of overview paper.
- **Finalizing task:** CLEF workshop and labs with discussion; feedbacks on CLEF; preparation of next year of CLEF; distribution of responsibilities.

4.2 Implicit requirements

Support the community in **managing the process:**

- **Tasks and roles.** Various roles are involved in communities work on experimental evaluation of IR systems. Multiple tracks and tasks are part of a process occurring in an evaluation experiment such as CLEF. Within a track, some of the tasks are interdepend. Specific roles perform specific tasks. One member can play multiple roles. The set of the role and tasks changes across different IR domains.
- Assume both individual and collaborative works. Many tasks are conducted individually and the individual work is interleaved with collaborative work. Support both individual and collaborative works and faster transitions.
- Only some of the steps in the work-flow are **fully specified** before or at the outset of the process, others are defined during the process. Needs related to **existing work tools:**
- As for other communities of knowledge workers [1], the members of these communities use email as their primary communication tool. It would be helpful for any new tool for knowledge sharing and collaboration to build on the central role of email.

- Clear added value. The user should be required to log onto a new system only if such new system supports additional functions that are useful and are not already supported in email or other general-purpose media already in use (e.g. VoIP).

- Difficult tracking and reuse. Perhaps a useful role can be played in facilitating a smoother integration across the multiple existing tools.

Needs related to **new collaborative functions** that might support collaboration:

- **Groups.** Group creation is tied to the task.

- **Polls.** The polls component should be integrated with the other components (this is an example of function not available in email).

- **Collaborative workspace.** It is desirable for the members of the community to have easy access to the shared resources and typical steps, which, ideally, should be made available all in one place.

- **Process visualization.** We observed visible differences in the different instantiations of the work-flows elicited from different sub-communities. It allows people to become aware of differences and similarities in the way sub-communities go about performing the same process.

5. CONTEXT AND RELATED WORK

Recent research has pointed to the importance of investigating and supporting collaboration in the field of Information Retrieval (IR). For example, [4] reviews the studies of collaboration relative to this field and concludes that the IR field needs to better understand and improve the systems that support both direct and indirect collaboration during information tasks. This is supported by studies in various IR settings. [6] investigated patent engineers, which is a specific community conducting IR processes, and found that they were involved in various collaborative activities. Overall, existing studies have observed that collaboration is indeed endemic to the broader activities of individuals who perform information seeking, searching, and retrieval tasks. However, with our work we aim to address two specific limitations of the existing literature. First, the prior studies of collaborative practices in IR have focused on describing the practices of teams or groups of users in different settings and user communities (e.g., academia, industry, medicine, patent offices; see [4] or have developed new tools to support these practices [8], but have not yet systematically investigated how to support collaboration at the level of a large communities of practice. Second, while there has been research on how to support communities of practices of professionals (e.g., [10]) or scientists (e.g., collaborators, see [6], we focus specifically on how to support lightweight knowledge sharing and collaboration in the community of IR researchers and professionals who develop and evaluate IR tools. This is a community with unique needs for collaboration and types of workflows: recurrently, specific sub-communities of volunteers need to agree on, build, and refine evaluation campaigns for testing IR systems.

A key distinctive property of the IR communities is that their workflows cannot be fully specified a priori. That is, if we consider a continuum from highly specified to highly unspecified processes, then we could classify the instances of workflows of the IR communities as intermediate cases along this continuum. [2], who named this continuum the Specificity Frontier, observed a gap between two existing approaches for supporting collaboration: most collaborative systems have focused on either automating fixed work processes (e.g., Enterprise Resource Planning tools) or simply supporting communication in ad-hoc processes (e.g., email). To adequately support the collaboration in IR communities we need to bridge the gap between these two approaches using lightweight tools that are compatible semi-structured workflows. While the




Figure 2: Specificity Frontier

specific instances of these workflows share several of the tasks and roles, the specific instances will inevitably vary across IR domains (and data types), evaluation campaigns, and over time (because the process is refined by the IR community in a collaborative manner as it is repeated over the years). Interestingly, recent research on communities of professionals pointed to the same need for collaborative tools that are able to support flexible realizations of the processes rather than forcing the community into hard-coded processes [9].

Articulating further the design requirements for supporting semi-structured workflows, [2] divides the specificity frontier into four sub-spectra: providing context for enactment, monitoring constraints about the task, providing/planning options to reach a goal, and guiding through a given script. Building on this classification, in this paper we focus on providing support to the IR communities in the first two sub-spectra.

6. DESIGN AND ARCHITECTURE

To fulfill the requirements described on Section 4, we devised the architecture shown on Figure 3. It refers to the classical CLEF experiments organization, that is arranged in terms of different domains (e.g., Medical, Patents, etc.). For each domain one or more tracks are available. As described on Section 3, the organization of a track is a complex collaborative process and encompasses several tasks that exhibit some precedence relationships. The task flow of a track is formalized using an Extended Light livE Gantt (EleGantt) and is shown within the CUI to the user, acting as the main entry point for collaborative activities (Figure 4). A suitable administrative interface allows for adapting the task flow of a track to procedural changes. According to [2] (Figure 2), EleGantt is in charge of providing a part of the process context, i.e., a structured to-do list. The second part of the context, i.e., a shared common space, is provided by T@GZ [7]. EleGantt is an extension of the traditional Gantt chart. It allows for:

- **attaching a rich set of meta-data** to tasks with the goal of supporting collaborative activities: involved people, involved roles, associated tags, kind of collaboration activities needed to accomplish the task, and the list of other processes that share the same activity;
- **expressing the non overlapping constraint** between tasks that must be executed in sequence;
- **specifying temporal uncertainties** (e.g., minimum and maximum duration of an activity), and degree of freedom for milestones and deliverable releases. Moreover, the EleGantt visualization is both a visualization of the task flow and an interactive interface that allows for exploring and accessing task flow associated information, like roles, people, similarities with other task flows, etc.

T@GZ is a social software system for organizational information sharing. In T@GZ the user can share by simply sending an




Figure 3: The architecture of the Promise collaborative system

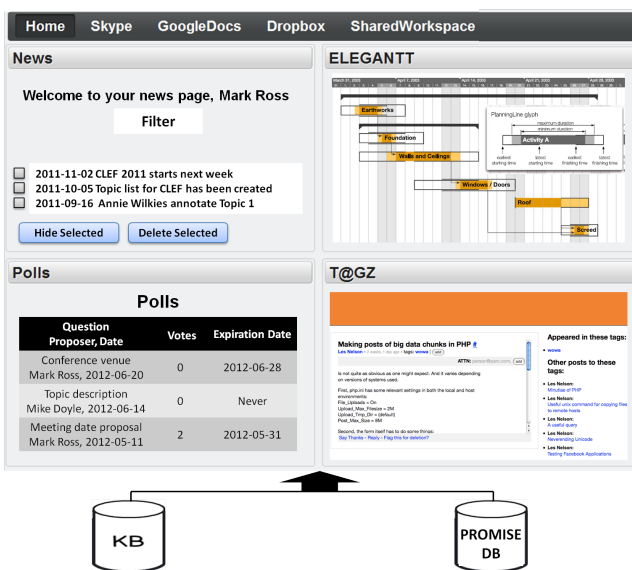


Figure 4: The Promise Collaborative User Interface (CUI)

email message with the content to be shared, and addressing the message to one or more topic specific keywords. For example, one might use the address, bizdev@share.X.com, for referring to information related to "business development" topic (see Figure 4, top left). Thus, the content of that email is 'tagged' by the keyword 'bizdev'. Any mail may have multiple tags attached in this manner, in the 'To' or 'CC' fields, using any client. While enabling easy publishing and re-finding of this information, the system does not induce people to send additional emails other than those that they are already sharing. Focusing now on the implementation of T@GZ in the whole system, using a set of predefined tags (i.e., the tags associated to the elegantt's tasks), T@GZ provides a means for indexing the emails that are exchanged among the organizers of the tracks, including links to smart attachments. The work-flow engine is aware of the elegantts and using time information and inspecting the KB sends through email different kinds of notifications (e.g., a deadline is approaching, it is time to move to the next step, etc.) to people involved in the tracks organization.

The Collaborative User Interface (CUI) is the Web based access point to all the collaborative activities (see figure 4). It is basically split in two subcomponents. The first one allows for managing personal user collaborative information (left part of the picture), e.g., messages, polls, etc. The second one refers to the whole process and allows for both exploring it using EleGantt, discovering people, roles, and tags and browsing the whole set of tagged emails. Moreover, the CUI contains a set of tabs that allows for accessing the collaborative tools that have been specified in the EleGantt. In order to provide the user with a unique access point to the process resources, the CUI sends tagged emails to T@GZ, containing a link to the collaborative resources (e.g., link to dropbox folder or to Google Docs document)

7. CONCLUSION AND FUTURE WORK

As result of our investigation we identified some general challenges or open issues for this domain: 1) find a good balance between the need for flexibility to fit various, partially-defined processes and the need for enough specification in order to allow automation 2) identify the set of predefined tags (some common and some domain specific) 3) semi-automate the tagging process (e.g., an intelligent assistant).

Acknowledgements

The work in this paper has been supported by the PROMISE NoE (contract n. 258191) project as a part of the 7th Framework Program of the European commission (FP7/2007-2013).

8. REFERENCES

- [1] V. Bellotti, N. Ducheneaut, M. Howard, I. Smith, and R. Grinter. Quality versus quantity: e-mail-centric task management and its relation with overload. *Human-Computer Interaction archive*, 20(1), June 2005.
- [2] A. Bernstein. How can cooperative work tools support dynamic group process? bridging the specificity frontier. In *ACM conference on Computer supported cooperative work (CSCW '00)*, 2000.
- [3] M. Croce, E. Di Reto, G. Granato, P. Hansen, A. Sabetta, G. Santucci, and F. Veltri. Collaborative User Interface Requirements. In *PROMISE deliverable D5.1*, 2011.
- [4] J. . Foster. Collaborative information seeking and retrieval. In *Annual Review of Information Science and Technology, Volume 40, 2006, 329-356*, 2006.
- [5] P. Hansen, G. L. Granato, and G. Santucci. Collecting and Assessing collaborative requirements. In *Collaborative Information Seeking (CIS 2011) workshop*, 2011.
- [6] P. Hansen and K. Jarvelin. *Collaborative information retrieval in an information-intensive domain. Information Processing and Management (IPM)*. 2005.
- [7] P. M. Joshi, C. Bartolini, and S. Graupner. T@gz: intuitive and effortless categorization and sharing of email conversations. In A. Mille and F. L. e. a. Gandon, editors, *WWW(Companion Volume)*, page 365. ACM, 2012.
- [8] M. Morris and E. Horvitz. SearchTogether: an interface for collaborative web search. In *20th annual ACM symposium on User interface software and technology (UIST '07)*, 2007.
- [9] H. R. Motahari-Nezhad, C. Bartolini, S. Graupner, S. Singhal, and S. Spence. IT Support Conversation Manager: A Conversation-Centered Approach and Tool for Managing Best Practice IT Processes. In *Hewlett Packard Laboratories Palo Alto, USA, 2010*, 2010.
- [10] R. S. W. . Wenger, E.; McDermott. *Cultivating Communities of Practice*. Harvard Business Press, 2002.