SUPPORTING DIGITAL COLLABORATIVE WORK THROUGH SEMANTIC TECHNOLOGY

Simon Scerri, Gerhard Gossen and Siegfried Handschuh
Digital Enterprise Research Institute, National University of Ireland Galway, IDA Business Park, Galway, Ireland

Keywords: Computer Supported Collaborative Work, Semantic Email, Workflow Management, Email Visualisation, Social Semantic Desktop, Intelligent User Interface, Speech Act Theory.

Abstract: Taking advantage of the fact that knowledge exchanged within digital working environments can be made persistent, a lot of research has strived to make sense of the ongoing communications in order to support the participants with their shared management. Semantic technology has been applied for the purpose as it ensures a shared understanding of the underlying collaboration, between both humans and machines. In this paper we demonstrate how, coupled with appropriate information extraction techniques, robust knowledge models and intuitive user interfaces; semantic technology can provide support for digital collaborative work. As a virtual working environment, e-mail was a natural contender for testing our hypothesis. Taking a workflow management-based approach, we demonstrate how semantics can indeed support email-based collaboration via Semanta – a tool extending popular email clients enabling semantic email. In particular we present a novel workflow-based email visualisation, the tool’s summative evaluation, and discuss the odds of semantic applications like Semanta evolving beyond research prototypes.

1 INTRODUCTION

The vast amount of heterogeneous information reaching the users' desktops outstrips their abilities to correctly manage and exploit it. This results in widespread information management problems that especially affect those users that thoroughly depend on electronic collaboration to carry on with their daily work. With email persisting as the most popular digital communication medium, email users are not spared the effects of this problem. These are aggravated by the fact that the uses of email have evolved beyond its original intended design (Whittaker, 2007), and the majority of these uses are either not supported at all or only to a very limited degree. The problem of information overload within email, or simply *Email Overload*, is particularly notorious. Numerous research efforts in various computer science sub-domains have attempted to alleviate the problem, by attempting to enable machines to support the user with the management of their email data. Some have taken a direct approach, through the development of technologies for email classification, search and retrieval; whereas others have taken less direct approaches to solving the problem, e.g. by facilitating email visualisation. Central to our own approach is the idea that email overload can be reduced by providing automated support for the underlying workflows in email applications.

In the next section we provide the motivation for our implementation in Semanta, focusing particularly on the epistemological gap between the user’s mental visualisation of email collaboration and its visible representation on the desktop. In Section 3 we provide an overview of the underlying knowledge models enabling Semantic Email. Here we will review our earlier presented models (Scerri, 2008a) (Scerri, 2008b) enabling the representation of email workflow knowledge in a machine-processable language. These models enable machine support for the user's email information management and promote the sharing and integration of email data over a network of social semantic desktops. Details as to Semanta's final architecture will be presented in Section 4 whereas in Section 5 we will describe how the latest version of Semanta utilises information extraction techniques to elicit workflow knowledge. A number of workflow supportive features have previously been demonstrated in (Scerri, 2009). In Section 6 we provide a brief overview of these features before presenting the

Scerri S., Gossen G. and Handschuh S..
SUPPORTING DIGITAL COLLABORATIVE WORK THROUGH SEMANTIC TECHNOLOGY .
DOI: 10.5220/0003103300920101
In *Proceedings of the International Conference on Knowledge Management and Information Sharing* (KMIS-2010), pages 92-101
ISBN: 978-989-8425-30-0

novel workflow-based email visualisation technique which has now been incorporated in this tool. In Section 7 we describe the process and results of the summative evaluation of Semanta as a whole. After an overview of the most relevant related work in Section 8, we provide guidelines for future works and some concluding remarks on the prospects of semantic technology in supporting the management and sharing of collaboratively-generated knowledge.

2 EMAIL OVERLOAD

Email can be considered an extension to the collaborative workers' working environment, serving as a virtual workplace where they collaborate, carry out tasks, etc., generating and sharing new personal information in the process. From this perspective email overload can be considered as a workflow management problem where, faced with an increasing amount (and complexity) of co-executing workflows, users become overwhelmed and lose their control over them.

Our approach is to identify and place patterns of email communication into a structured form, without changing the email experience for the end-user. We start by considering *Action Items* embedded in email content (e.g. Task Assignment, Meeting Proposal). Sequences of related action items exchanged in email messages are then treated as implicit but well-defined *Ad-hoc Email Workflows* (e.g. Task Delegation, Meeting Scheduling). The nature of these workflows is such that they occur spontaneously and evolve dynamically and to an extent unpredictably with time. Besides their lack of support, the way these workflows (or rather their implicit components) are represented on the user's conventional desktop system is too different than the way the user would visualise them through their mind's eye. In fact, we say that there is a ***huge epistemological gap between the way users conceive email workflow knowledge and the way it is represented on their desktop***.

We will explain this situation via the example in Fig. 1, which illustrates how Martin conceives an email conversation (workflow) in his mind and how he can see the corresponding fragmented information physically on his desktop. At time t1, Martin writes an email (1) to Dirk and Claudia, which amongst other things contains a *Meeting Proposal* action item asking about their availability for a group meeting. This initiates an implicit *Meeting Scheduling* email workflow, which splits in two co-executing paths at time t2 – control of which is passed to Dirk and Claudia individually. Dirk reacts to the meeting proposal immediately by sending an email (2) with his feedback (*Deliver Feedback* action item) back to Martin. Claudia instead, is not sure about the purpose of the meeting and thus sends an email back to Martin (3) with her inquiry (*Information Request* action item). This is considered a sub-workflow of the currently executing workflow. Martin deals with this sub-workflow at time t3 by replying with an *Information Delivery* in Email 4. Martin's answer to Claudia's query terminates this sub-workflow, upon which Claudia can get back to the initial workflow. At time t4, she also sends her feedback back to Martin (Email 5). At this point Martin has all the required information for the meeting proposal he sent in Email 1. Thus at time t5 the two parallel workflow paths to merge back together and Martin is passed back its control. He decides on a specific date and time for the meeting right away and sends another email (6) with an *Event Notification* to both Dirk and Claudia. Upon sending the email, an event involving Dirk, Martin and Claudia has been generated for Martin. After both Dirk and Claudia have acknowledged the *Event Announcement* action item at time t6, the same shared event has been generated for all of them.

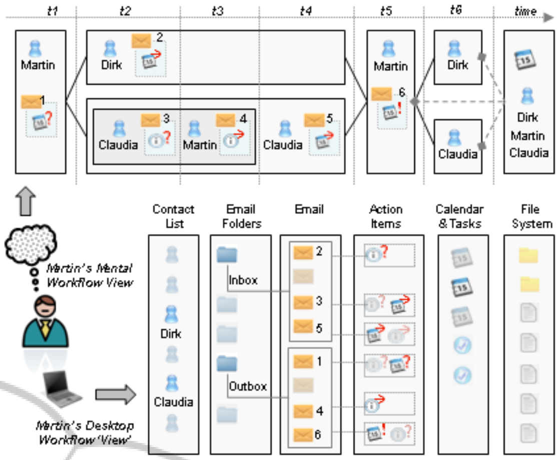

Figure 1: Martin's Workflow Views.

Unfortunately for Martin, the email workflow knowledge as presented above is in no way similar to what he can visually gather through a conventional email client on his desktop. Unless the email conversation is fresh in his mind, there is no straightforward way for him to quickly get an overview. Especially if there are many such workflows running at the same time, within tens or hundreds of email messages, and varying in priority and complexity. Fig. 1 shows the fragmented physical view of the same workflow with which

Martin would have to contend. The main workflow components are scattered within a number of separate, largely unconnected, data 'islands'. The action items making up the workflow are obscurely strewn across a number of usually (physically) unrelated email messages, belonging to different email folders. People in the contact list are only associated with these emails, and their roles in the contained workflows remain unspecified. The workflow artefact generated at the end of the example is stored in the Calendar, with little or no connection to the email or the email thread wherein it was generated. Workflow artefacts can also be dispersed in additional data islands, such as generated tasks which end up in a separate task list or having attached documents propagated onto the file system without keeping any connection to their source emails.

3 MODELLING WORKFLOW KNOWLEDGE

To provide for the elicitation, support and visualisation of email workflows, we required robust knowledge representation. We will next review the models enabling *Semantic Email*, that is, email enhanced with machine-processable metadata about the underlying workflow knowledge. Here we intend only to provide an overview of this research[1], as the details have been covered in existing publications (Scerri, 2008a) (Scerri, 2008b). However, this is essential to appreciate the non-trivial technology behind the functions provided by the latest Semanta prototype as presented and evaluated in this paper.

The first milestone of this research was reached with the design of a concise but expressive model with which various email action items could be represented. The model is based on aspects of the speech act theory (Searle, 1969), which is based on the idea that every explicit utterance, or email statements in this case, implies one or more explicit or implicit actions. The model is thus aptly referred to as the *Speech Act Model* (Scerri, 2008b), and has at its core an action item (or speech act) concept consisting of the triple (action, object, subject), where the action defines the nature of the action item, the object (of the action) defines what the action is in relation to, and the subject(s) (of the action) corresponds to the people implied by the action. The model provides for seven different actions (Request, Assign, Propose, Suggest, Deliver, Abort, Decline) and five objects (Task, Event, Information, Feedback, Resource). The subject depends on the email participants, and is a member of the power set for the email sender and the recipient(s). Thus a request from Claudia to Dirk for a joint task can be represented as (Request-Task-{Claudia, Dirk}); a permission request for an event from Claudia as (Request-Event-{Claudia}).

The second milestone was the design of the *Speech Act Workflow Model* (Scerri, 2008a). Here we took our approach one step further by considering each action item as the start (or the continuation) of a workflow as depicted in Fig. 1. Although ad-hoc email workflows are spontaneous by nature, there exist trends which enable the prediction of what occurs after certain action items are received or sent. Our workflow model is based on a statistical study of real email threads, whereby human annotators annotated email threads in the Enron corpus with sequential speech acts[2]. Thus the model is considered a formal representation of the ad-hoc workflows taking place over email communication, outlining the most likely reactions to incoming email action items while allowing for email's characteristic flexibility. The model is grounded on key research in the area of control flow workflow patterns, utilising a number of patterns from a standardised workflow language (Voorhoeve, 1997).

4 IMPLEMENTATION

Fig. 2 depicts how the conceptual framework for semantically-enabled email was put into practice to provide additional support to the user via Semanta[3]. Knowledge expressed in the semantic email models at the conceptual level is exploited within the knowledge representation (KR) level via a semantic email ontology[4]. This ontology re-uses knowledge from within additional ontologies on the Semantic Web, especially those designed for the *Social Semantic Desktop* project (SSD) (Groza, 2007). In fact, Semanta is one of the semantic applications conceived within this project. Although Semanta still functions on a normal desktop, the semantic desktop has the added benefit of desktop data integration – whereby machine-processable data generated by multiple semantic applications can be shared across multiple machines and desktops. The rich knowledge representation models provided by

1 http://smile.deri.ie/projects/smail/

2 http://www.cs.cmu.edu/~enron/

3 http://smile.deri.ie/projects/semanta/

4 http://ontologies.smile.deri.ie/smail#

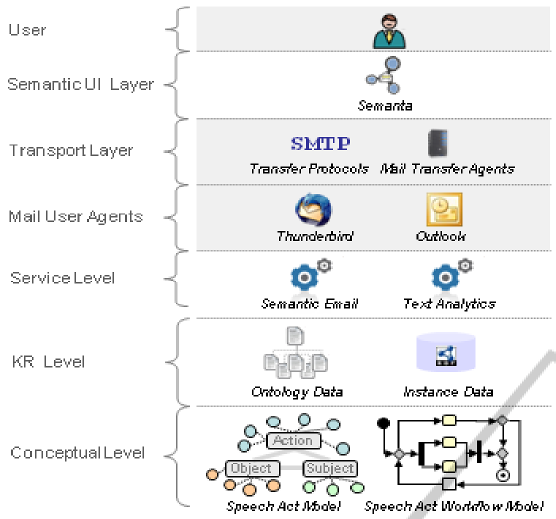

Figure 2 : Semantic Email across the different levels.

SSD mean that the representation and integration of workflow components can be extended to the whole user's personal information model. Thus, the objects in Fig. 1 could be linked and related to other physical and abstract personal concepts, e.g. to a concept representing a project to which the scheduled meeting is related. Additionally, the social aspect of the SSD makes it possible for all those involved in the meeting to actually share the same instance of the meeting across their desktops. The merits of Semanta as one of many interoperable SSD applications have been discussed in (Scerri, 2008a). In this paper we will instead focus on the personal information management support provided by Semanta as a stand-alone semantic application.

Semanta is empowered by the services in the service level, which provide for all the business logic of the system. The text analytics service performs email action item classification whereas the semantic email service is responsible for the running of most of Semanta's underlying technology, including the generation, retrieval and querying of all metadata. Data elicited and generated through this service is expressed in the machine-processable RDF[5] format, serving as the instance data stored in the system's RDF store (KR level). The semantic email service then acts as an intermediary between the knowledge in the KR level and the enhanced semantic user interface. Although this interface is only the tip of the iceberg, this semantic UI is what the user perceives as Semanta. The enhanced UI is built on top of two popular email clients – Microsoft Outlook and Mozilla Thunderbird, which utilise standard email transfer technology. The shaded levels in Fig. 2 are in fact to stress that Semanta relies on the existing mail user agent and email transportation layers. As a result, aside from the additional functionality provided by Semanta's UI extensions, the user's email experience also remains relatively unchanged. Thus most of Semanta's technology and the generated semantic data are conveniently hidden from the user beneath an enhanced UI that is itself integrated within the existing technological landscape.

5 ELICITING WORKFLOW KNOWLEDGE

In the next section we will present the many ways in which Semanta is able to support the user with managing their email workflows. However, for this support to be provided the system needs to be aware of the workflows in the first place. The bottleneck is thus the recognition of action items not executing in an existing workflow, e.g. a new meeting request, rather than an amendment to an existing one. The text analytics service in Fig. 2 is an important component as it provides for the semi-automatic classification of action items in email text. This service implements a rule-based classification model[6] that classifies email segments into action item instances from the speech act model. The results of a separate evaluation of this classification technique indicate an accuracy level of around 60%. Although this is considered a low score, this has to be seen in the light of an earlier evaluation (Scerri, 2008b) which highlighted the difficulty of the classification task, even when performed by humans. In fact this experiment calculated a human inter-annotator agreement rate (via the Kappa statistic) of 81%. As pointed out in the evaluation of MailCat (Segal, 1999), an error rate of over 20% is completely unacceptable in automated processes. Thus our semi-automatic classification is meant to facilitate, rather than completely automate email annotation. Consequently, when the user hits the send button for a new email, the results of automatic classification are highlighted in the content and presented to the user for review. Additionally, an annotation wizard facilitates this task, supporting the user with the easy creation or modification of action item annotations while hiding the complexity of the speech act model. The resulting action item

[5] http://www.w3.org/RDF/

[6] Details of this model and technique are sufficiently covered in (Scerri, 2010).

annotations together with other harvested metadata are then encoded in RDF within the email headers and transported alongside the email. Once Semanta is aware of the initial action items in a workflow the email workflow model is employed keep track of subsequent action items (updates to the workflow) and to support with their management on both the sender and the recipient(s) side.

6 SUPPORTING WORKFLOW MANAGEMENT

When an email reaches the inbox, Semanta displays the number of pending action items alongside the other email information in the inbox (last column, Fig. 3). This count adjusts dynamically when action items are taken care of. When an email is selected, action items are highlighted in the content (red italic). Users can interact with each item, where a number of relevant options are provided given the item's type and the knowledge provided by the workflow model. For the *Task Request* shown in Fig. 3, Claudia can approve or disapprove the task or alternatively amend its properties.
Some of these options result in a reply email item being generated, but this is not always necessary.

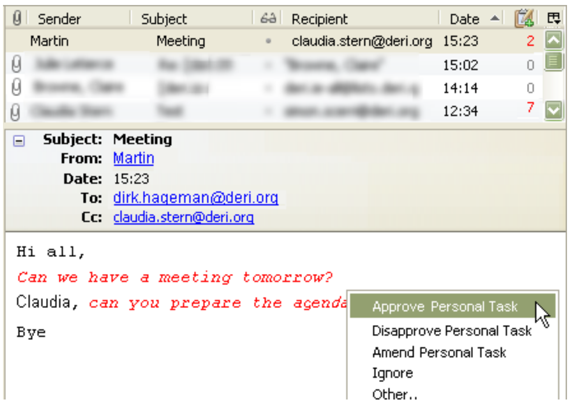

Figure 3: Semanta's email action item support.

For example upon receiving a *Task Assignment* instead of a task request, the user can simply acknowledge the assignment without the need for a reply, as per the semantics of this type of action item specified by the workflow model. Action items can also be ignored (and later unignored) indefinitely. Most importantly Semanta's support for ad-hoc workflows allows the user to react to an action item in additional ways. For example in Fig. 3, before submitting her availability for the proposed meeting, Claudia decides to question its purpose. After selecting the 'Other..' option, Claudia writes her question (e.g. "I thought the meeting was cancelled?") and is then assisted once again with the annotation wizard to select an appropriate action item for this text (*Information Request*). This results in an ensuing email reply and constitutes a subworkflow of the original, control of which is passed back to Martin when Claudia sends the automatically generated email. This information is stored and updated with each workflow update that ensues. Semanta detects events/tasks generated when writing or reacting to incoming email, in which the current user is implied. For example, when Claudia approves the requested task in Fig. 3, Semanta support her with storing the workflow artefacts directly to the associated Tasklist. The default Outlook task list and calendar are used for Outlook, whereas the Lightning add-on is required for Thunderbird[7]. Semanta auto completes some of the properties of so-generated tasks/events. The subject of the task generated from the email will carry the textual excerpts from the workflow (e.g. Martin wrote: "can you prepare the agenda?": You replied: "Yes"). The contacts implicated in the activity are also known; in this case Claudia has the sole responsibility.
Links between email messages and the tasks/events generated from within are stored and exploited for the user's benefit. Fig. 4 shows three items related the workflow which ensued following the task request in Fig. 3. This request (Fig. 4 - 1) was answered via an email reply (Fig 4- 2). These two emails therefore belong to the same thread, and are linked via the 'Previous Email' and 'Next Email' buttons. The last item (Fig 4- 3) is the task generated by this workflow, specifically from the second email (Fig 4 - 2). The user can jump to this event from this email via the 'Related Activity' button. Semanta extends the display of the task item by a 'Conversation' panel which shows the history of the workflow until the generation of this task. In the example, the workflow before the task generation consisted of two action items shown. The user can also directly jump from these items to the emails in which they were exchanged, i.e. (Fig. 4 - 1) and (Fig. 4 - 2) respectively.
We will now introduce Semanta's latest and most novel latest feature – the workflow-based visualisation of email. As discussed in Section 2, email users have so far only been able to view scattered fragments of email workflows when

[7] http://www.mozilla.org/projects/calendar/lightning/

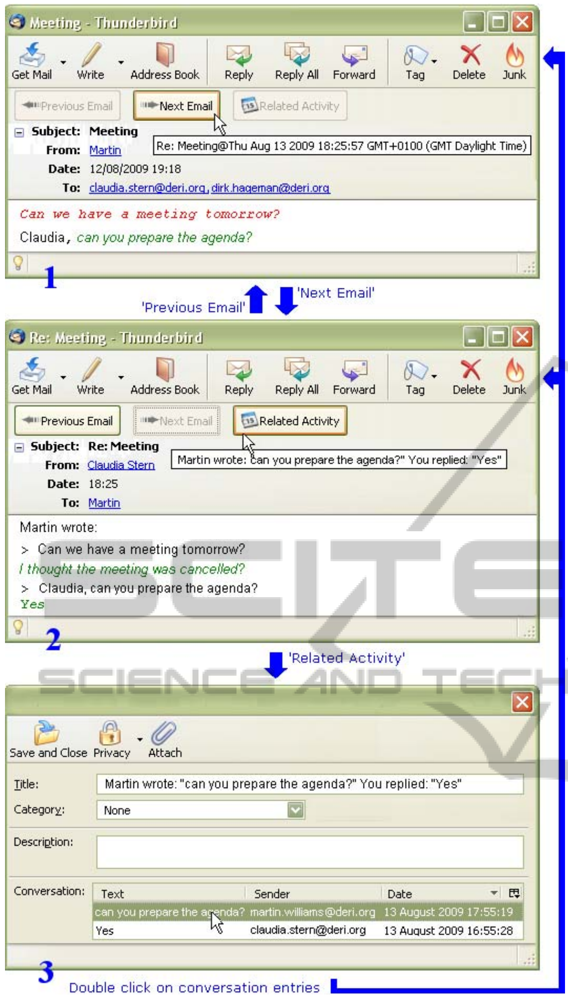

Figure 4: Linking Workflow Artefacts.

looking at messages in their email folders. The most groundbreaking feature of Semanta is to provide a novel workflow-based email view that is more akin to the user's mental conceptualisation (Fig. 1). The knowledge elicited and gathered by Semanta means that the system is aware of all exchanged action items, their position within a workflow as well as their status. Semanta's semantic UI exploits this knowledge to generate a view wherein users can visualise these workflows, and also navigate to the email message within which each individual action item in the workflows is contained. Semanta's *Workflow Treeview* (Fig. 5) is available alongside Thunderbird's default email treeview on the left-hand side. The treeview provides for three views, the selection of which enables the UI components on the right-hand side. These components offer a form of visualisation which functions like faceted-search – the user restricts the field of view to a particular email, starting from a workflow. The main view ('All') displays a list of all workflows that have taken place or are still running/pending (displayed in bold) in the *Workflow List*, ordered by start date. When a workflow is selected its details are shown in the *Workflow Details* below. This component shows the sequence of individual action items in the workflow. Finally when an action item is selected, Semanta retrieves the email within which it has been exchanged and displays it to the user in the *Email Message* component below. The example shown in Fig. 5 is more akin to Martin's view of the workflow in Fig. 1. In fact, the workflow selected in the workflow list originates from the *Meeting Proposal* sent to Dirk and Claudia. The workflow details below show that whereas Dirk provided his availability right away (4th action item), Claudia asked for further information before providing hers. This sub-workflow is represented by the two indented action items – the information request (2nd item), followed by an information delivery (3rd item). The email within which this action item was exchanged (in Martin's Outbox folder) is displayed in the email message view below. The workflow is still marked as pending in the workflow list because although Martin has received the feedback from both the other two meeting participants, he has yet to announce the meeting at that stage. Alongside the main view, the workflow treeview provides two other specific views. The *Incoming* view shows all incoming action items (e.g. requests, assignments, suggestions) which remain pending. In this case, rather than displaying a list of workflows, Semanta displays a list of pending action items, shown in the context of their workflow. The user can then directly resume the workflow by reacting to the pending items. Alternatively, the *Outgoing* view shows all outgoing action items (e.g. requests) for which the user is still awaiting a reply.

After viewing these items the user can decide whether to send a reminder urging the correspondents to reply (and resume the stalled workflow).

7 EVALUATION

Our evaluation methodology follows the guidelines outlined in (Gediga, 2001). All material used for the evaluation, including the full results, is available online[8]. The process consisted of a Formative stage – where the initial system prototype was improved

[8]http://www.smile.deri.ie/projects/semanta/evaluation/

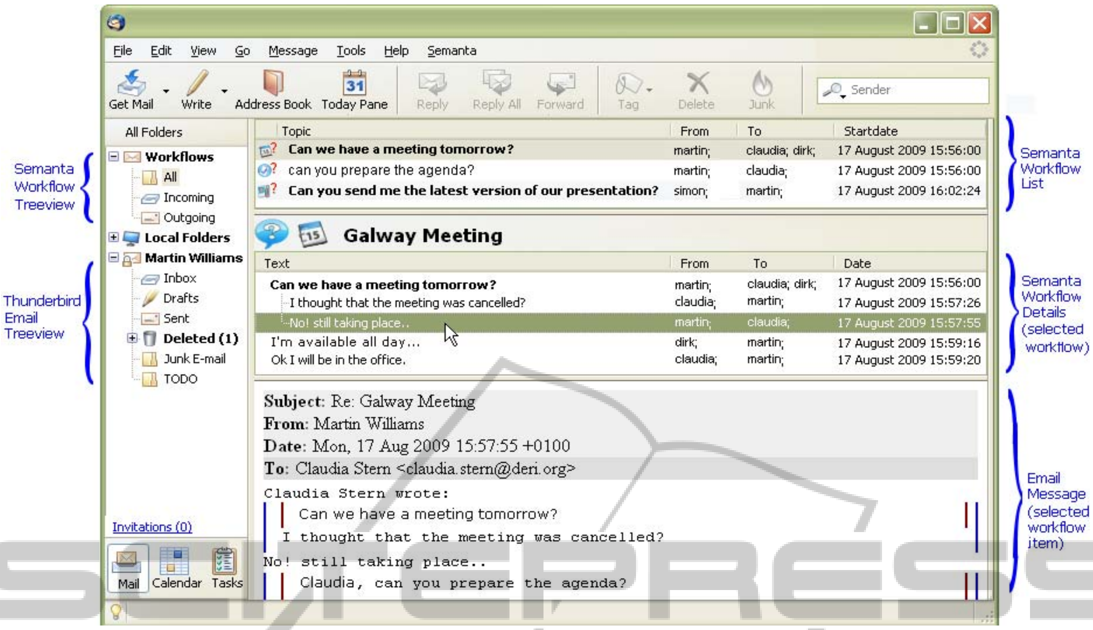

Figure 5: Screencast showing Semanta's workflow-based email visualisation. The user navigates from an initial action item (top right) to the ensuing workflow (middle right) and finally a specific email (bottom right).

following a controlled study; and a Summative stage – where users tried the improved prototype in their actual day-to-day email work. Given its platform independency, this stage was based on Semanta's Thunderbird add-on. The results of the formative stage were published in (Scerri, 2009). In this paper we will report the findings of the second evaluation stage. The purpose of the summative evaluation was to compare Semanta with an alternative – the standard Thunderbird with no add-ons. As most of Semanta's features can only be appreciated when exchanging email between Semanta users, our hypothesis – that the use of Semanta improves the email experience over the use of a standard Email Client; needed to be tested within such groups. Thus the evaluation involved a total of 18 users, collaborating in subgroups of between 2 and 6 people. The users consisted mostly of Computer Science researchers within three universities (including our own) where English is used as the first language; but also included a few industrial partners with whom they collaborate. The evaluators were introduced to the evaluation via a web page[9] and supported by a detailed user manual[10]. They were instructed to use Semanta for 10 days, at the end of which they sent their automatically-generated usage statistics. On a per-person average, 40.42 action items in 29.29 semantic emails were exchanged in 11.83 days. An average 6.57 incoming and 9.29 outgoing action items remained pending at the end of the evaluation. Semanta also assisted the users with handling an average of 3.29 email-generated tasks and 2.14 events. The evaluation included a questionnaire[11], starting with a reproduction of the standard USE questionnaire[12], measuring the usability of the system across four dimensions: usefulness, user satisfaction, ease of learning, ease of use. Results of this part of the questionnaire are averaged in Fig. 6a.

The next part of the questionnaire tried to quantify the performance of Semanta over the

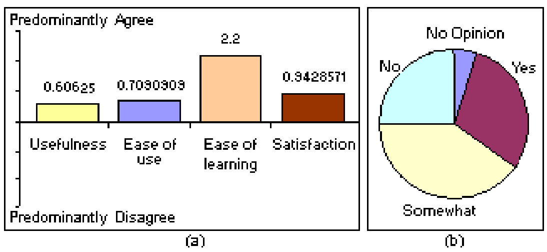

Figure 6: Main results of the evaluation.

[9]http://www.smile.deri.ie/projects/semanta/semantaevaluation2009
[10]http://www.smile.deri.ie/projects/semanta/usermanualthunderbird

[11]http://www.surveymonkey.com/s.aspx?sm=hwJLbdf_2bZdyUL6hXw4dhiQ_3d_3d
[12]http://www.stcsig.org/usability/newsletter/0110_measuring_with_use.html

Table 1: Main results of hypothesis testing.

Null Hypothesis H_0	*Mean*		*T-test*	*Outcome*
	T.B.	*Semanta*		
H1 : No change in time required to write email	0	-0.75	-3.000	Rejected
H2 : Annotation process doesn't effect email writing experience	0	-0.08	**-0.173**	Accepted
H3: Flexibility of email replies is not effected	0	-0.85	-3.091	Rejected
H4: Difficulty of keeping track of pending received action items is unchanged	0	2	11.015	Rejected
H5: Difficulty of keeping track of pending sent action items is unchanged	0	2	11.832	Rejected
H6: No effect on the mental visualisation of email workflows	0	2	6.36	Rejected

standard Thunderbird. Questions were based on a 7-point Likert scale, ranging from -3 (Predominantly worse) to 3 (Predominantly better), with 0 signifying no perceived changed (thus the ratings for Thunderbird are zero by default). A one-sample t-test (two-tailed, 99% confidence interval) was then performed to interpret the ratings. Here we only provide the highlights, the full results being available via the evaluation page. The first result (Table 1 - H1) rejects the hypothesis that the same amount of time is required to write email with Semanta. This is expected due to the annotation reviewing stage. However, the users feel that this stage does not harm the email writing process, and H2 is (H2). Optional comments in the questionnaire suggest that although some see it as an annoyance, others like the idea of annotating email if it helps with getting things done. H3 was rejected, shows that the flexibility of email replies was somewhat jeopardised by Semanta. In fact, in an additional best and worst feature fields in the questionnaire, the email reply interface got the highest number of negative votes. H4-H5 were rejected in Semanta's favour, concluding that keeping track of both incoming and outgoing action items is significantly easier with Semanta. Finally the hypothesis that Semanta does not help the user with visualising email workflows (H6) is also rejected, implying that workflow-based view of email was successful in this regard. Additional results confirmed that the users appreciate Semanta's ability to link tasks and events to the email threads wherein they were generated and the possibility to traverse independent email messages in a thread. This was expected, given that the standard Thunderbird lacks these features.

The final part of the questionnaire posed the following bottom-line question: "Are Semanta's functionalities worth the effort to review automatic annotations or manually create them?". The results, shown in Fig. 6b seem to suggest that whereas the majority of user felt that the time sacrificed reviewing email annotations was worth the subsequent email support to an extent or another, around 25% of the evaluators seemed to think otherwise. This can perhaps be summed up by the following comment provided by one user: "leaving aside the fact that Semanta is a research prototype, for a new email tool to be accepted by a broad set of users as beneficial, it will need to provide benefits that are at multiple orders of the additional cost that it imposes on users".

8 RELATED WORK

As there have been numerous attempts at supporting computer collaborative work, we will here also stick to the email use case, providing a number of approaches that are most relative to the work presented in this paper. One of the most well known initiatives in this area was IBM's ReMail (Rohall, 2004) – a reinvented email prototype focusing on email visualisation, calendar entry discoveries and user attention management. In contrast, we seamlessly integrated our technology into the existing technical landscape, using existing transport technology; while hiding complex workflow models and semantics beneath intuitive GUI extensions to existing email clients. Other initiatives have focused on improving the user's email experience by targeting specific email tasks and features e.g. reply prediction, attachment reminders, automatic foldering and recipient prediction (Dabbish, 2005) (Dredze, 2008a) (Dredze, 2008b) (Segal, 1999). We consider these solutions to be a patching-up exercise to the underlying problem, i.e. the lack of support for email workflows. In contrast, the comprehensiveness of our approach allows for the indirect provision of most of these features.

Speech Act Theory was applied to email communication a couple of times, in particular to ease the management of email-generated tasks (Corston-Oliver, 2004) (Khoussainov, 2005) and for email classification (Carvalho, 2005) (Goldstein, 2006) (Khosravi, 1999). The speech act model itself is based on an earlier one provided by Carvalho et. al. (Carvalho, 2005), which considered a speech act as the pair (verb, noun), e.g. (Request-Task). In particular, we extended this model to also refer to the speech act subject. The human inter-annotator agreement experiment mentioned earlier was also applied to this model. The results conclude that our model is more intuitive for the classification of email action items. The research conducted by Carvalho et. al. also computed transition diagrams for sequential speech acts, for the prediction of successive acts. In (Singh, 1998), the author investigated the condition of satisfaction for individual speech acts. In our research, we extend these conditions of satisfaction to workflows.

Apart from speech act theory, our work is also directly inspired by the research contributions of *Semantic Email Processes*. In fact Dowell et. al. (McDowell, 2003) first used the term ‘semantic email’ to refer to an email message consisting of a structured query (or an update to the query) coupled with a corresponding explanatory text. Their approach was based on the provision of a broad class of semantic email processes that represent commonly occurring workflows within email (e.g. collecting RSVPs, coordinating group meetings). Implemented within Mangrove the system provided templates which exposed structured knowledge about these scenarios to both humans and machines. The ultimate goal was to support the user with common email-related tasks such as collecting information from a group of people, handling event information, etc. Although we believe that the option of fixed templates taken in (McDowell, 2003) is in some cases useful, our approach is more oriented towards the handling of ad-hoc email workflows.

9 CONCLUSIONS

In this paper we demonstrated how semantic technology can enable automated support for digital collaborative work, focusing on the email use case. In this context, our approach has been to identify and place patterns of email communication into a structured form, such that machines can support the user with email workflow management. In turn, this knowledge is employed to reduce the epistemological gap between the way users conceive collaborative workflows and the fragmented way in which these are currently ‘displayed’ in the respective digital working environment.

The concept has been implemented and showcased via Semanta: a user-supportive email extension for popular email clients. If the average email user is sacrifices minimal extra time to review the automatic action item annotations when writing new email, Semanta in return:

- is aware of the existence and status of (otherwise implicit) email action items within email
- is able to support the user with reviewing incoming action items and the semi-automatic provision of replies
- detects tasks and events generated within email messages, and provides contextual information and links from both directions
- provides an alternative workflow-based email visualisation that is more akin to what the users conceive conceptually when carrying out their email tasks
- provides ancillary features such as linking email within the same thread and file attachment reminders, as well as social semantic desktop integration;

Following the results summative evaluation of Semanta, we are happy with the acceptance of our tool but acknowledge that in order for Semanta to jump over the research fence into the real world, the extra cost imposed on the user needs to be further reduced. The latest evaluation has outlined further room for improvements. We intend to extend the text classification grammars to enable the recognition of more information, e.g. matching person names in text to the user’s email contacts, recognition of dates and times related to upcoming events or task deadlines, etc. We are also investigating the use of ML techniques to improve both precision and recall of the automatic annotation. GUI-wise, we are considering the suggestions received to improve the least attractive features. Semanta will be extended to work also when the corresponding users are not using Semanta, so that non-semantic email can still be mined for action items. Finally, the workflow views will be extended to incorporate any resulting events/tasks. The status of tasks can then also be dynamically updated when the responsible participant(s) update it as such.

The lessons learnt from Semanta can, to a large extent, be projected onto general approaches that employ semantic technology to provide support for digital collaborative work. Our experience demonstrates that although semantic applications are indeed able to provide the envisioned additional support to the collaborative knowledge worker, this support comes at a cost. The extent of this cost is controversial. For the email use case, whereas some people were more than willing to spend a little more extra time reviewing and adjusting email action item annotations in view of the rewarding support provided, others considered it as yet another email chore.

ACKNOWLEDGEMENTS

The work presented in this paper was supported (in part) by the Lion project supported by Science Foundation Ireland under Grant No. SFI/02/CE1/I131.

REFERENCES

Carvalho, V., Cohen, W., 2005. *On the collective classification of email speech acts*. In Proc. SIGIR-2005. 345—352.

Corston-Oliver, S., Ringger, E., Gamon, M., Campbell, R., 2004. *Task-focused summarization of email*. In Proc. Text Summarization Branches Out workshop ACL 2004.

Dabbish, L., Kraut, R., Fussell, S., Kiesler, S., 2005. *Understanding email use: predicting action on a message*. In Proc. SIGCHI conference on Human factors in computing systems.

Dredze, M., Wallach, H., Puller, D., Brooks, T., Carroll, J., Magarick, J., Blitzer, J., Pereira, F., 2008a. *Intelligent Email: Aiding Users with AI*. In Proc. AAAI 2008.

Dredze, M., Wallach, H., Puller, D., Pereira, F., 2008b. *Generating summary keywords for emails using topics*. In Proc. 13th international conference on Intelligent user interfaces. New York, NY, USA.

Gediga, G., Hamborg, K., 2001. *Evaluation of Software Systems*. Encyclopedia of Computer Science and Technology. Vol. 45, 166—192.

Goldstein, J., Sabin, R.E., 2006. *Using Speech Acts to Categorize Email and Identify Email Genres*. In Proc. System Sciences, HICSS2006.

Groza, T., Handschuh, S., Moeller, K., Grimnes, G., Sauermann, S., Minack, E., Mesnage, C., Jazayeri, M., Reif, G., Gudjonsdottir, R., 2007. *The NEPOMUK Project - On the way to the Social Semantic Desktop*. In: 3rd International Conference on Semantic Technologies (ISEMANTICS 2007), Graz, Austria

McDowell, L., Etzioni, O., Gribble, S., Halevey, A., Levy, H., Pentney, W., Verma, D., Vlasseva, S., 2003. *Evolving the Semantic Web with Mangrove*. UW Tech Report.

Khosravi, H., Wilks, Y., 1999. *Routing email automatically by purpose not topic*. Natural Language Engineering, Vol. 5, 237–250.

Khoussainov, R., Kushmerick, N., 2005. *Email task management: An iterative relational learning approach*. In Proc. Conference on Email and Anti-Spam.

Rohall, S.L., Gruen, D., Moody, P., Wattenberg, M., Stern, M., Kerr, B., Stachel, B., Dave, K., Armes, R.,Wilcox, E., 2004. *ReMail: a reinvented email prototype*. In proc. CHI '04 Extended abstracts on Human factors in computing systems.

Scerri, S., Davis, B., Handschuh, S., Hauswirth, M., 2009. *Semanta – Semantic Email made easy*. In Proc. European Semantic Web Conference 2009. Crete, Greece.

Scerri, S., Gossen, G., Davis, B., Handschuh, S., 2010: *Classifying Action Items for Semantic Email*. In proc. 7th International conference of Language Resources and Evaluation.

Scerri, S., Handschuh, S., Decker, S., 2008a. *Semantic Email as a Communication Medium for the Social Semantic Desktop*. In Proc. European Semantic Web Conference 2008. Tenerife, Spain.

Scerri, S., Mencke, M., Davis, B., Handschuh, S., 2008b. *Evaluating the Ontology underlying sMail – the Conceptual Framework for Semantic Email Communication*. In proc. 7th International conference of Language Resources and Evaluation. Marrakech, Morocco.

Searle, J., 1969. *Speech Acts*. Cambridge University Press.

Segal, R.B., Kephart, J.O., 1999. *MailCat: an intelligent assistant for organizing e-mail*. In Proc.3rd annual conference on Autonomous Agents. New York, NY, USA.

Singh, M., 1998. *A Semantics for Speech Acts*. Annals of Mathematics and Artificial Intelligence. Vol. 8, 47—71.

Voorhoeve, M., Van der Aalst, W., 1997. *Ad-hoc workflow: problems and solutions*. In proc. 8th International Workshop on Database and Expert Systems Applications (DEXA07).

Whittaker, S., Bellotti, V., Gwizdka, J., 2007. *Email and PIM: Problems and Possibilities*. In Proc. Communications of the ACM CACM07.