Designing A Cross-channel Information Management Tool For Workers in Enterprise Task Forces

Gregorio Convertino, Sanjay Kairam, Lichan Hong, Bongwon Suh, and Ed H. Chi
Palo Alto Research Center (PARC)
3333 Coyote Hill Road, Palo Alto, CA 94304, USA
{convertino, kairam, hong, suh, echi}@parc.com

ABSTRACT
This paper presents a research project on the design of a cross-channel information management tool for knowledge workers: we focus on IT services professionals in a large enterprise who work in multiple ad hoc task forces. Through three rounds of investigation, we characterized their work practices and needs, specified their requirements for a cross-channel information management tool, and designed and evaluated a prototype to address these needs. We found that these workers shared the problem of managing information across multiple channels, requiring better support for aggregating, filtering, and organizing this information. We report the requirements elicited and the prototypes built during the design process.

Author Keywords
Cross-channel, cross-media, information management, design, CSCW, enterprise, task force

ACM Classification Keywords
H.5.3. [Information interfaces and presentation] Group and Organization Interfaces: Collaborative computing

INTRODUCTION
We present the *design* of a cross-channel information management tool. The *target users* are knowledge workers collaborating in multiple ad hoc task forces and managing multiple information streams. The prevalence of such task forces has increased as companies have recognized that group work increases productivity in organizations (e.g., [2]). Since the mid-1990s, as the IT services market has grown [3], task forces have become increasingly common in IT enterprises. Forming ad-hoc cross-functional task forces has been a key strategy for these companies to compete in this new market.

We have adopted Zimmerman, et al.'s definition of design research as "an intention to produce knowledge and not the work to more immediately inform the development of a commercial product" [13]. We focus on a special class of knowledge workers, learning about their content management needs and concurrently designing prototypes in order to iteratively identify and ecologically test 'tangible' hypotheses on how to support this user class.

The primary contribution of this paper is a characterization of an *emergent class of knowledge workers* and *their requirements* for a tool to support their need to manage information across multiple channels. In addition, it describes the process whereby such a tool was designed. Details on the implementation and software architecture of this tool can be found in our companion short paper [15].

We grounded the design on a field study with two groups of senior professionals in a large IT services enterprise: Alliance Professionals (APs), who craft strategic business alliances with partner companies, and Bid Managers (BMs), who write bid proposals for large service deals (10M-50M$ US). Members of both groups manage multiple task forces, each of which must generate shared products efficiently (e.g., a proposal or an alliance plan). In doing so, they coordinate multiple stakeholders from different locations and multiple streams of information.

We found that these busy professionals must handle large amounts of content and, more importantly, face *the challenge of managing multiple information channels*. This challenge is emphasized by two basic tasks that are recurrent and central in their job: (1) foraging and organizing new information relevant to each task force, and (2) monitoring status and progress updates about individual and group activity. Currently, incoming content and process updates are dispersed across multiple channels with little support for aggregating and managing this information.

In this paper, we present the procedure and findings of our three rounds of characterization, design, and evaluation, including the elicitation of requirements for advanced functions. Finally, we discuss future applications for our methodology and findings, including our ongoing work.

RELATED WORK
In addressing knowledge workers' needs for information management tools, some researchers have focused on the fragmented structure of the work and costs associated with interruptions [4]. Others have focused on the increasing quantity of information that knowledge workers must

AVI'10, May 25–29, 2010, Rome, Italy.

manage, overloading cognition and threatening productivity [5, 11]. Information overload has been recognized as an acute problem by several IT enterprises. Companies such as Microsoft, Google, IBM, and Intel have recently launched a research group to address this problem [5].

It is increasingly common for knowledge workers in organizations to manage multiple tasks, projects, or "working spheres" at any moment in time. Gonzalez et al. [4] studied analysts, software developers, and managers working in a company. They found that people simultaneously engaged in an average of 10 tasks per day and constantly moved from one task to another (every three minutes on average). They also observed that people used various tools and artifacts to maintain an awareness of the status of their tasks in order to assess whether they should switch to a new task or remain focused on the current task.

Recently, tools have been developed to address the problem of managing *multiple channels* of information. Moran, et al. [9] suggested organizing information and activities around tasks instead of tools and artifacts. Laqua, et al. [6] proposed an email plug-in that aggregates relevant information from diverse corporate sources. Aizenbud-Reshef, et al. [16] tackled the problem of information overload by proposing a feed aggregator with collaborative features allowing workers to share feeds and divide reading tasks. In the consumer domain, web tools such as FriendFeed (friendfeed.com) and Google Reader (reader.google.com) aggregate RSS feeds from various web-based sources in a single location. However, busy workers need more than reverse-chronological lists. While these tools support the aggregation of information, they do not support easy "slice and dice" functions to filter, monitor, and organize the various information streams.

STUDY BACKGROUND

A recent field study conducted with the community of BMs describes the challenges faced as a result of handling simultaneously multiple projects [10]:

> "[…] several different bids, possibly at different stages, may have to be managed simultaneously and shifted between in a more or less conducive fashion, in the course of any one working day. In addition, bid management does not stop at the management of specific bids. There are a number of other activities that occupy the time of bid managers and their associated staff and these activities also have an impact on how the work as a whole gets accomplished." [10]

We confirmed that these professionals manage multiple projects and task forces. BMs work on an estimated average of 3 proposals at the same time (from 2-5 depending on deal size) and APs often work on an even greater number. Each project involves discovery, aggregation, and organization of information from multiple sources. For example, an AP initiating an alliance might forage information from the Web (e.g. company web sites and online news reports), from internal databases and knowledge systems, and from other members of the company.

In addition, information and activity coordination within teams is critical as tasks progress. In responding to a bid opportunity, a BM must build a support team, assign tasks to team members, monitor progress, assemble contributions, and collaboratively write a bid proposal. For both APs and BMs, task management occurs mostly via email and phone, where email is also used to share drafts and company-level communications. This often leads to information overload, echoing Whittaker et al.'s work [11].

STUDY METHOD

This design research involved three rounds of investigation. During each round, the study data were collected using online questionnaires and/or semi-structured interviews. Figure 1 summarizes the work during the three rounds.

Each round of data collection had a distinctive purpose: Round 1 involved characterization of work and needs; Round 2 focused on analysis of basic requirements and paper prototyping; and Round 3 involved prototype evaluation and analysis of advanced requirements. Participants varied in number at each stage, depending on the availability of these busy professionals at each time; all rounds, however, included both APs and BMs and many of the same individuals. We describe the methods and the sample sizes for each round as follows:

1. Characterization Round. Starting in September 2008, a group of 23 APs and 9 BMs completed an online questionnaire about job profile, tools usage, and tool preferences, with an overall participation rate of about 60%. A more focused questionnaire was then given to a subset (13 APs, 6 BMs) of this initial group. The surveys were designed to allow comparison with national baseline statistics on IT tool usage in the workplace [7]. A first 1-hour interview was conducted in March 2009 with 4 APs and 4 BMs who had participated in the survey earlier. The interview collected information about the tools, tasks, stakeholders and a detailed description of a recent project, including the processes, tools, people, and timeframe involved. After the interview, a two-page summary and workflow diagram of a typical project were generated, and then validated and corrected with the participant in a follow-up call and/or email exchange.

2. Requirements Analysis and Paper Prototype. In June 2009, a second 1-hour interview was used to analyze requirements with 6 APs (including the original 4) and 1 BM who had also been interviewed earlier. The BM held the role of knowledge manager for that community and was attuend to the information sharing practices in the BM community. In the first half of each interview, we collected key requirements for information-management tool. In the second half, we showed them two paper prototypes (Figure 4) and refined some of the requirements collected earlier.

3. SW Prototype Evaluation, Requirements Analysis. A software prototype (described in the Interactive Prototype section) was developed over a period of 2-3 months, informed by knowledge generated during the prior rounds

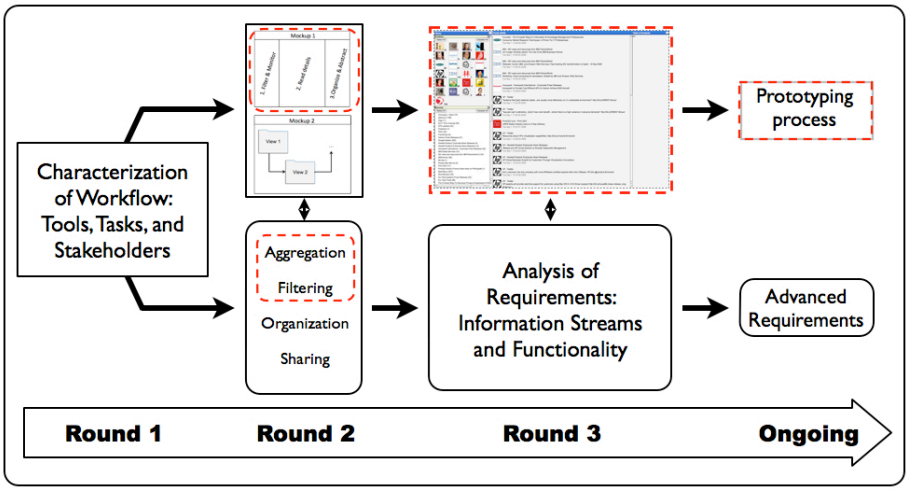

Figure 1. Our two-tiered approach for eliciting requirements (bottom) while designing a prototype for a system to meet these requirements (top), from characterization to future

of investigation. In August 2009, we ran a third 1-hour interview (4 APs, 4 BMs) to evaluate a software prototype designed according to these requirements. Interviewees first provided common information streams (10-15 min), then gave feedback about existing system components and functions (30-35 min), and finally gave suggestions for improvements and new requirements for features not yet implemented (10-15 min). Seven of the professionals (4 APs, 3 BMs) completed a questionnaire evaluating the software prototype after this final interview. Responses included ratings scales and free-form answers.

STUDY RESULTS

Round 1: Task Forces' Work: Characterization

The data from the surveys and first interview revealed that the two groups were very similar in seniority and expertise: senior workers with high self-ratings of expertise (Table 1).

	AP		BM	
	Avg.	*Sd.*	*Avg.*	*Sd.*
Company Experience (years)	9.8	8.2	9.2	9.0
Overall Corporate Age (years)	**16.1**	**7.7**	**14.2**	**6.0**
Expertise Self-Rating (1-5 scale; 1 = novice)	**4.3**	**0.9**	**4.3**	**0.8**
Computer usage (hours per week)	28.5	10.1	33.7	8.5
Online Tools usage (hours per week)	**15.5**	**12.5**	**13.8**	**10.0**

Table 1. Profiles of the APs and BMs: busy, expert professionals relying heavily on computers and online tools.

Tools. The professionals' primary tools were *email*, *phone*, and *local folders* on their computers. Email and phone were key communications means, used primarily to coordinate meetings and exchange information. All utilized email at least several times an hour. Almost all indicated daily phone use, though the APs (cell: 100%, office: 67%) were slightly more mobile than the BMs (cell: 92%, office: 85%). All used personal laptops several times an hour. Document, spreadsheet, and slide editors were main production tools and the local file system was the main storage tool. They also used less frequently, though significantly, other tools such as Instant Messaging (daily, AP: 62%, BM: 67%) and collaboration tools such as wikis and a content management system (i.e. DocuShare). Overall, their level of technology adoption (see Table 1) far surpassed averages for networked workers in the US [7].

Tasks. Tasks for APs and BMs involve managing ad hoc task forces of service professionals to build a set of documents, spreadsheets, and slide decks. An AP's tasks center on forging alliances with large IT industry partners (e.g., IBM, Dell), which includes identifying opportunities with an alliance partner, negotiating a value proposition with various stakeholders (such as the executives of a partner company), and generating an operational business plan (Figure 2). The BM's tasks focus on identifying and responding to bid opportunities with other corporations, which includes evaluating potential bid opportunities, building a support team to respond to an opportunity, and assembling a bid proposal document (Figure 3).

Both groups follow a similar, loosely-defined, collaborative workflow (Figures 2 and 3). One shared property was the long-term, asynchronous nature of the project. Bid engagements for BMs last up to 6-8 months (including monitoring, not represented in Figure 3), while forming a new alliance for an AP requires up to a year. Another shared property is the wide geographical distribution of the task forces, requiring coordination across multiple countries and continents. Therefore, they spend 30-40% of their time traveling, working at a client side, or working from home.

Stakeholders. Because teams are custom-built for each opportunity, each engagement involves collaborating with different sets of stakeholders, and within a project, different subgroups of members are engaged over time. For example, after first establishing the plan with a team of executives

through workshops (5-20 people), an AP manages the deployment of the plan by working with teams of regional managers from the involved companies (5-10 people).

Design Implications from Round 1 Interviews

Pertaining to tools for managing information, a common theme that emerged is that *members of both professional communities struggle with managing information across multiple channels*. Information and documents are fragmented across numerous tools (email, file system, phone, document editors, calendar, IM/SMS, databases, CMS, Wikis, etc.). A key source of strain is managing document versions across remote and local archives, email, and other repositories. Thus, this general situation points to the need for better cross-channel information management, and such a tool would need to be flexible in order to accommodate the constantly changing set of technologies used by the teams. As one of the BMs expressed, *"There are so many different tools that we have…I'd just like to have one point where I can go and everything's there."*

Email emerged as the central tool for managing and exchanging information, transferring documents, and coordinating; as such, it is critical that any new system either incorporates email or at least offloads some information normally transmitted via email. Reliance on idiosyncratic *foldering* patterns was common for managing information, suggesting that flexible foldering options for organizing documents should be supported. Finally, persistent needs for management of multiple projects and re-organizing content around each new ad hoc task force should be addressed in the design of an effective cross-channel information management tool.

Round 2: Requirements and Paper Prototyping

A second round of interviews with 6 APs and 1 BM was used to elicit requirements for the functionality and design of a tool for cross-channel information management. For the BMs, in this round, the data was collected from the group's knowledge manager and was integrated with the findings of a prior ethnographic study with the same group of professionals, in 2008 [10]. The requirements and paper prototyping data below represent both BMs and APs.

Information Management Needs

During the requirements-gathering interview, we asked each professional to envision the hypothetical condition of carrying out their usual tasks with normal tools such as email, phone, and word processor while also imagining access to an additional web-based tool called *Workstreams*. The proposed tool was presented as an *augmentation of their current toolset* to help collect and organize content, monitor various streams of content in their projects as well as content their collaborators might be processing. In this context, the participant indicated key general functions that the tool could support to improve productivity or reduce workload. Table 2 summarizes the results.

We anticipated strong, shared desires for content aggregation, filtering, and organization, given findings from our first round of interviews, but anchoring these requirements to a specific tool gave additional insight into these needs: (1) First, the professionals highlighted the need for a single location to search for content pulled from multiple sources (email, DocuShare/CMS, web, etc.); (2) They needed precise and powerful ways to filter through large amounts of data and ways to easily customize these filters for different types of search tasks; (3) They wanted the content organized in clear and consistent ways to support awareness and sensemaking. An additional finding was that while status updates about projects and people were crucial to the work of these professionals, they were severely lacking tools for managing status updates, stating that they were primarily relying on email for this purpose. This finding pointed to the need for tools that incorporate tokens of both content information and status updates.

Feedback on Paper Prototypes

In order to elicit requirements specifically on the user interface of the envisioned Workstreams tool, we used two paper mock-ups. The first mock-up (Figure 4, left) depicts a user interface similar to a browser window with three

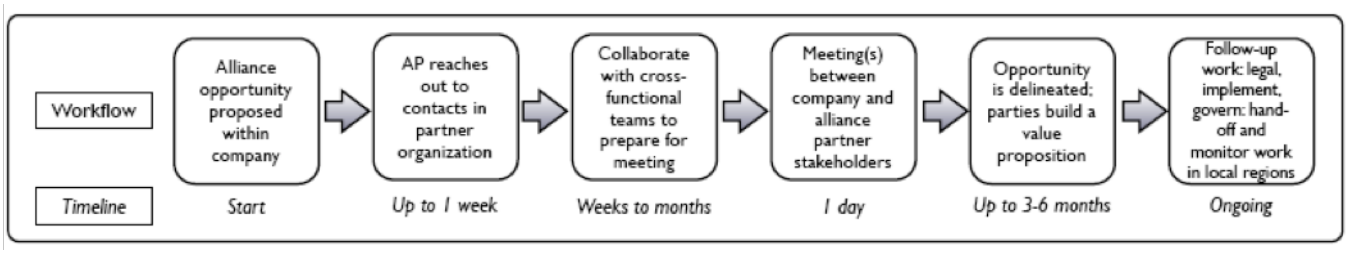

Figure 2: Alliance Manager Workflow

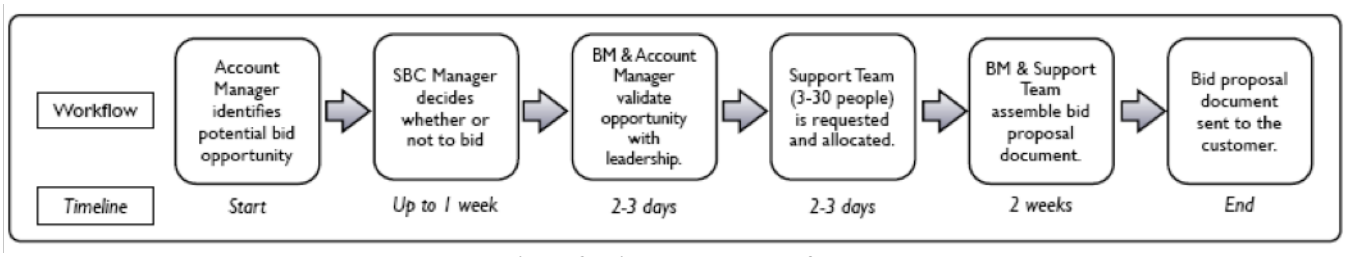

Figure 3: Bid Manager Workflow

Content Aggregation (Foraging)	• Single, shared repository for content • Consolidate relevant project information (e.g. company financials) • Push "information" that I don't know that I need to know but I ought to know"
Personalization and Customization (Filtering)	• Personalized information for users • Alternative views on the fly • Easy filtering and configuration • Search, ability to use complex queries
Content Organization (Sensemaking)	• Duplicate organization scheme across tools (email, local folders, etc.) • Link data to better see opportunities • Content displayed in a consistent format
Managing Stakeholders (Awareness)	• Increase visibility for others regarding status of project workflow • Provide views for different stakeholders • Easier contribution; information sharing
Managing Projects (Awareness)	• Support managing multiple projects • Streamlined status monitoring for deals • Manage and link document versions

Table 2. Requirements and Key Functions.

coordinated views or panes designed to support basic functions as listed in Table 2. The left pane supports filtering and monitoring, the center pane presents detailed information as it is filtered via the left pane, and the right pane supports organization and abstraction over the information aggregated by the tool. The second mock-up (Figure 4, right) represented a tool with multiple alternative views on the content. Inspired by the professionals' reliance on the folder paradigm for organizing information, this second mock-up resembled a file-foldering system, where information can be aggregated and filed under a hierarchy.

Paper Prototype 1: Regarding the first mock-up, the professionals indicated that the 'facets' in this design would be useful, especially in allowing multiple views on information (e.g., some situations might benefit from person or company-centric views, whereas other situations might benefit from views grouped by topic (as for emails), task, or sub-process). When reading details in the central pane, they would be interested in seeing critical elements extracted from documents, such as the date, title, or an executive summary. Finally, some professionals indicated that the ability to drag and drop content in folders would be useful for organizing and abstracting content, as it would match their filing practices.

Paper Prototype 2: Regarding the second mock-up, the professionals reported similar preferences about the *hierarchy* of attributes and views that they would use to monitor and organize the content. Most often they chose 'project' as the top-level organizer and 'topic' as the next level. Company or client name was also indicated as a possible organization criterion. These preferences mirrored their filing strategies in the file system or the email client.

Overall, the professionals agreed that the second prototype was more familiar, but ultimately less useful than the first. Beyond describing it as *"intuitively better"* and as *"the one that I would be more comfortable using"* (Mr. P, AP), they preferred the first for corralling information:

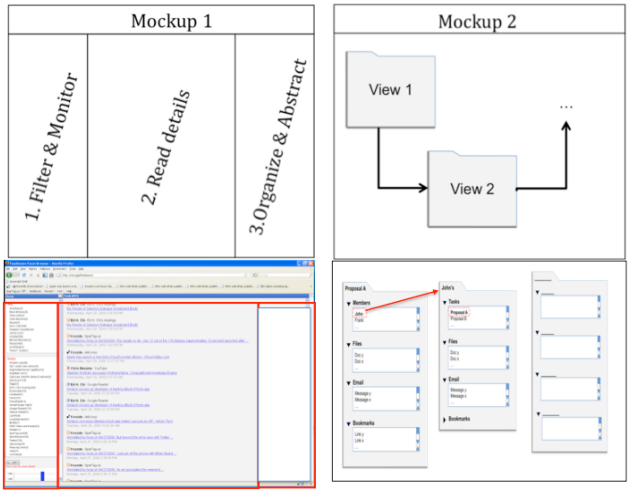

Figure 4. Paper Prototypes 1 and 2: schema (top) and paper mock-up of the interface (bottom)

> "The critical aspect is the aggregation of information in one place. The key problem is that if I cannot bring all the information that is relevant in one place accessible with 1 click then I am back to square 1. Mockup 2 is more for exploding information, which is less critical than what Mockup 1 does." (Mr. H, AP)

Round 3: Interactive Prototype

Design Rationale

The second round of investigation found a clear preference for the first paper prototype. Thus, we build the interactive software prototype extending this design (Figure 5). The goal here was to build a tool that implemented an initial set of the requirements (functions in Table 2). Table 3 summarizes the requirements implemented (see [15]).

Content Aggregation (Foraging)	• Easily accessible web-based portal. • Integrates multiple existing information channels. • Delivers tokens in a single, searchable location where users can forage.
Personalization and Customization (Filtering)	• Four facets (creators, sources, topics, and time) enable users to filter out non-essential information • Different stakeholders can 'slice & dice' information in many different ways.
Content Organization (Sensemaking)	• Aggregates information & status tokens. • Displays tokens in a consistent format. • *Table of Contents* area (unimplemented) to support content organization.
Managing Stakeholders (Awareness)	• Integration of information and status tokens allows for ambient, simultaneous awareness of stakeholders and projects.

Table 3. Requirements as implemented in software prototype.

Interactive Prototype Evaluation

In the final interview, participants were situated in their usual work environment and asked to describe their usual information streams. They were introduced to the system

components in a standard sequence, and then asked to give feedback on system functions while envisioning using the tool to manage typical streams of information.

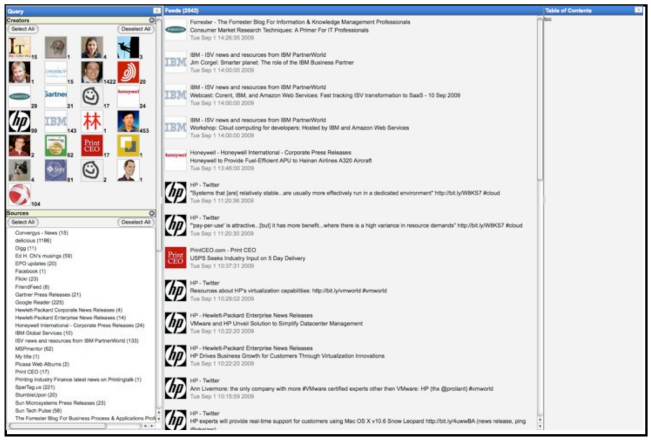

Figure 5. Software Prototype (Workstreams System)

Information Channels. As expected, the APs and BMs listed a diverse set of potential information streams, summarized in Table 4. The external business sites described in the table include analysis sites such as forrester.com and gartner.com, and the company websites were usually those of partners and competitors. Other input streams included RSS feeds, newsletters and mailing lists, and news aggregators such as Google Finance.

Internal CMS	**8/8**
Email	**7/8**
Wikis	**4/8**
External Business Sites	**4/8**
Specific Company Websites	**4/8**
Internal Web Resources	**3/8**
External News Aggregators	**2/8**
Status Updates	**1/8**
RSS Feeds	**1/8**
Document Updates	**1/8**
Newsletters / Mailing Lists	**1/8**

Table 3. Main Information Channels as Sources for Aggregation (N=8).

Facets. The tool had been pre-populated with a realistically rich set of streams, which was kept constant across the participants. The information could be filtered according to four facets: creators, sources, topics, and time. A few example channels were provided to help illustrate system functionality, but each participant was encouraged to envision that the specific creators, sources, and topics had been derived from his/her own channels that s/he had described earlier in the interview. They were encouraged to try filtering the content while giving feedback on the utility and usability of each facet and preferred combinations.

Across the professionals, topics emerged as the preferred primary facet for filtering. All of the BMs and three out of the four APs indicated in interviews and survey that they would use this facet to filter for content and status updates. Table 5 summarizes the 'usefulness' ratings given by the professionals for the facets in the survey (note that the 4th AP rated topics significantly lower, contributing to the low mean and high standard deviation).

Usefulness (N=7) (Questionnaire 1-7 scale)	**BM**		**AP**	
	Avg.	*Sd.*	*Avg.*	*Sd.*
Creators	4.3	0.6	**4.5**	2.1
Sources	**5.3**	1.2	**4.8**	2.1
Topics	**5.7**	1.2	4.0	2.6
Time	4.3	0.6	3.8	2.8

Table 4. Average ratings of 'usefulness' of facets from survey.

The professionals spontaneously formulated a variety of strategies for combining filters in order to find items. This demonstrates the flexibility that the tool might offer to users with different information needs. The most common strategy used was combining *topic and time* to gain access to the most recent and relevant items on a particular topical subject. Three of the five professionals who attempted combinations did so. Another common combination was using *creator* and *time* to monitor status notifications.

Main Window. As described previously [15], our prototype displayed feed items matching the filters in a scrollable, reverse-chronologically ordered list with information about the creator, source, and time, and the title or subject of the item. The professionals were generally positive about the presentation of the information and status tokens in the main window, as shown in the survey ratings of usefulness and usability presented in Table 6:

Main Window (N=7) (Questionnaire 1-7 scale)	**BM**		**AP**	
	Avg.	Sd.	Avg.	Sd.
Useful	5.0	1.0	5.0	2.2
Easy to Use	5.7	0.6	5.0	0.8

Table 5. Average ratings on presentation in the main window.

One common suggestion was to parse and display more information from feed items. One AP indicated that providing the subject was insufficient by referencing a common email pattern (*"Sometimes the subject won't be useful…it's just gonna say 'RE:' the same subject over and over and over again."*). Another subject suggested the inclusion of a mouse-triggered pop-up previewing more text from the body of the feed item.

Cross-Channel Aggregation Needs Satisfied

Though individual preferences varied somewhat, all indicated that a cross-channel aggregator such as the system demonstrated would be extremely useful in helping to improve productivity and reduce workload. Table 7 shows their responses to the question: *"How much would the system help your team to perform the following functions?"* These ratings hint that the benefits provided by a cross-channel aggregator may differ for groups with different workflows and responsibilities. Through surveys and interviews, 3/4 BM's saw the greatest utility in the initial information-gathering phase of their workflow, and 2/4 BMs also indicated it would be useful in consolidating information at the end of the workflow for status

monitoring while creating the final products. 2/4 APs saw the most utility in monitoring status, consolidating action items, and preparing 'status documents'.

(N=7) (Questionnaire 1-7 scale)	BM		AP	
	Avg.	Sd.	Avg.	Sd.
Gather Information	3.3	1.5	**4.0**	2.2
Discover Information	**3.7**	2.1	**4.0**	2.2
Organize Information	**4.0**	1.7	3.8	2.4
Discuss Issues	3.0	0.0	3.8	2.2
Produce Response Document	2.3	0.6	3.0	1.4

Table 6. Average ratings for support of high-level functions.

New Requirements on Advanced Functions

After evaluating the current features, we used the prototype to elicit new requirements on more advanced functions.

Organization. On the right-panel was an area labeled "*Table of Contents*" which was deliberately left open-ended (see Figure 5, rightmost pane), so that the professionals could envision how they might want to utilize this space to organize information items once found. The folder layout (or some variation) was a preferred organization method, suggested by 2/4 BMs and all APs; these might mirror those currently used for information sharing via the CMS or the personal folder structures from their desktops or email.

Sharing, Search, and Other Functions. The main additional advanced functions requested by multiple professionals (n=8) after evaluating the system include Sharing (6 out of 8), Search (5/8), Rating (3/8), and History (3/8). Other functions requested were: grouping of results, token prioritization, workflow integration, customization, built-in Intelligence, and output formatting).

Sharing, e.g., pushing items or entire feeds to other individuals, distribution lists, or shared repositories, was highly desired. Familiar sharing paradigms such as '*drag items into folders*' or '*right-click for a 'send-to' option*' were referenced. This suggests opportunities for integrating collaborative features, such as in [16]. Though the tag cloud representation for filtering on topics was generally liked, the tag list provided was viewed as incomplete, and 3/8 professionals suggested complementing it with a free-text *search* box. *History* was envisioned by some as a record of past interactions with the tool ranging from saved queries to templates. *Ratings* were an area of disagreement: 3 BMs imagined rating items or feeds either for personal consumption at a later time or for others, but 1 BM and 1 AP each explicitly said they would not rate items, expressing doubts about the enduring value of such ratings and interaction costs. We hope to integrate some of these features in future versions of the prototype.

DISCUSSION

As shown in Figure 1, our design research was structured in three rounds, following a two-tiered approach after the first round (characterization of work and needs). One tier involved eliciting requirements about information management functions, such as aggregation and filtering. The other tier aimed at building a tool to support cross-channel information management. The requirements were iteratively translated into prototypes, which were then, in turn, used to further specify the existing requirements, discover new ones, and define priorities among them.

Aggregate Content and Process Across Channels

Prior work on web-based aggregators has focused either on managing content (e.g., CMS, see [1, 8]) or monitoring status and process updates (e.g., task management tools, see [9]). In contrast, our prototype for cross-channel management assembled content bits and process updates in one place with flexibility in filtering and monitoring this information via the four facets. The professionals affirmed that such a system could serve their needs well, potentially by reducing interruptions (i.e., fewer unnecessary switches between channels) and easing information overload costs resulting from interruptions due to these switches. Based on prior work (e.g., [8], [14]) about layering intermediate representations over low-level information items, we also expect that introduction of summary representations such as tag clouds will aid in individual and collaborative sensemaking over aggregated content [16].

Confirming Whittaker et al.'s findings [11], these knowledge workers pointed to email as the most popular content management tool, used heavily for managing status updates in addition to transferring information. We noted a clear need for *offloading non-communication messages, such as 'FYI' messages, from email*. Thus, our design provides a possible solution to alleviate the email overload problem, which arises from making email the hub of every transaction. Note that this solution is different from those that have attempted extending functionalities of email [6].

The focus on aggregation is also consistent with a recent trend towards integration. Web 2.0 products such as Google Wave for consumers and Lotus Live in enterprises suggest interest in removing barriers between information-sharing tools currently used in isolation (e.g., email, RSS feed reader, social network). Our prototype goes beyond existing systems by (1) utilizing document features to automatically generate filters and (2) allowing for the combination of filters for more efficient browsing. As more tools are developed to facilitate collecting, reviewing, and sharing of information, the need for tools to aggregate, filter, and organize multiple information streams will inevitably become stronger.

Empirical Work: Assume No Silver Bullet

There is no silver bullet for addressing information management problems for knowledge workers, nor is there even a 'typical' knowledge worker against whom design solutions can be tested. We believe that it is sensible to assume multiple classes of knowledge workers, differentiable based on the kind of work they do. Each class may share needs and requirements for new information management tools and then--when these tools are shown to them--a specific assessment of the expected benefits.

We see a disconnect in HCI research between studies characterizing specific classes of users or settings and the design of innovative interaction techniques. Whittaker, et al. proposed that researchers should moderate emphasis on radical innovation and develop 'reference tasks' for having a shared focus [12]. Extending this claim, we propose that the characterization of a 'reference class' of knowledge workers is essential to the development of new tools aimed at such a user class. Thus, we characterize IT service workers, relating their properties to previous studies of knowledge work [4, 7, 10]. Our approach here, combining iterative design and evaluation, can similarly be generalized to help characterize and design for other groups.

Despite the growing presence of task forces in organizations, we are not aware of prior systematic studies focusing on the needs of this group. We found that they are experiencing, perhaps earlier than other knowledge workers, new challenges in managing information. They must handle large amounts of content, and more importantly, manage it across multiple information channels. This is particularly evident in two tasks which are recurring and central to their work: (1) foraging and organizing new information for each task force, and (2) monitoring status and progress updates about task force activity. The present lack of support for aggregation and management requires extra effort to cope with the fragmentation of information across channels, ultimately limiting performance. Finally, while evaluating our prototypes with these professionals we learned about informative job-related differences and similarities (APs vs. BMs) on work functions that the tool can support (Table 7).

CONCLUSION AND FUTURE WORK

Studying, as exemplars, two communities from a large IT enterprise, we characterized the work practices and information-management needs of a growing class of busy knowledge workers. We found that they need (1) information aggregated across multiple channels, including the combination of content and status updates, (2) filters that help to easily find important content, and (3) organization and sharing functions for individual and collaborative sensemaking. We elicited requirements through interviews and iteratively designed with these workers a tool for cross-channel information management. They confirmed the tool's utility and specified priorities for new advanced functions, which will inform future design.

As it is important to explore how well this design research applies to other knowledge workers with similar information management challenges, we have continued requirements specification and prototype evaluation with seven senior managers in a mid-sized research organization (see [15]). They are also responsible for managing multiple information streams and teams of people or projects. We consolidated the three rounds into a single interview (60-90 minutes), surveying typical information streams, eliciting design requirements, and having them evaluate the software prototype (following up with the usability questionnaire). Their feedback was consistent, overall, with that of the APs and BMs, suggesting that the current findings could generalize beyond the groups characterized in this paper.

REFERENCES

1. Berlin, L.M., Jeffries, R., O'Day, V.L., Paepcke, A., and Wharton, C. (1993). Where Did You Put It? Issues in the Design and Use of A Group Memory. *Proc. INTERCHI'93*, 24-29.
2. Cleland D.I. *Strategic management of teams*. Wiley and Sons: New York, NY. 1996.
3. Cusumano, M.A. (2008) The Changing Software Business: Moving from Products to Services. *Computer*, 41 (1), 20 – 27.
4. Gonzalez, V. and Mark, G. (2004). "Constant, Constant, Multi-tasking Craziness": Managing Multiple Working Spheres. *Proc. CHI'04*, 113-120.
5. Information Overload Research Group. iorgforum.org.
6. Laqua S., Sasse A., Gates C., and Greenspan S., (2009). Making Sense of the Unknown: Knowledge Dissemination in Organizations. In *2nd Sensemaking Workshop at CHI 2009*, Boston, MA. April 4-5, 2009.
7. Madden M. (2008). Networked Workers. Pew Internet & American Life Project. September 24, 2008.
8. McDonald, D., Ackerman, M. (1997). Collaborative Refinery: A Collaborative Information Workspace for the World Wide Web. Tech Report 97-03, Information and Computer Science, Univ. of California at Irvine.
9. Moran, T., Cozzi, A., and Stephen, P. F. (2005). Unified Activity Management: Supporting People in E-Business. *Comm. of ACM*, 48, (12), 67-70.
10. Tolmie P. (2009). Identification of Real World Issues in the Work of Bid Management. Xerox report, 4/2009.
11. Whittaker, S. and Sidner, C. (1996). Email Overload: Exploring Personal Information Management of Email. *Proc. CHI'96*, 276-283.
12. Whittaker, S., Terveen, L., & Nardi, B.A. (2000). Let's Stop Pushing the Envelope and Start Addressing It: A reference task agenda for HCI. *HCI* 15, 75–106.
13. Zimmerman J., Forlizzi J., Evenson S. (2007). Research Through Design as a Method for Interaction Design Research in HCI, In *Proc. of CHI 2007*, ACM Press.
14. Convertino, G.; Hong, L.; Nelson, L.; Pirolli, P. L.; Chi, E. H.(2009) Activity Awareness & Social Sensemaking 2.0: Design of a Task Force Workspace. Proc. HCII 09.
15. Hong, L., Convertino, G., Suh, B., Chi, E. H., Kairam, S. (2010). FeedWinnower: Layering Structures Over Collections of Information Streams. In *Proc. CHI 2010*.
16. Aizenbud-Reshef, N., Guy I. Jacovi M. (2009). Collaborative feed reading in a community. In *Proc. of GROUP* 2009. ACM Press. 277-280