Social and Interactional Practices for Disseminating Current Awareness Information in an Organisational Setting

Simon Attfield[a*]
s.attfield@cs.ucl.ac.uk
Tel: +44 (0)20 7679 0694
Fax: +44 (0)20 7679 0699

Ann Blandford[a]
a.blandford@ucl.ac.uk

Stephann Makri[b]
s.makri@ucl.ac.uk

[a] UCL Interaction Centre (UCLIC), MPEB 8th floor, University College London, Gower Street, London, WC1E 6BT, UK
[b] Department of Information Studies, University College London, Gower Street, London, WC1E 6BT, UK
**Corresponding Author

Abstract: Current awareness services are designed to keep users informed about recent developments based around user need profiles. In organisational settings, they may operate through both electronic and social interactions aimed at delivering information that is relevant, pertinent and current. Understanding these interactions can reveal the tensions in current awareness dissemination and help inform ways of making them more effective and efficient. We report an in-depth, observational study of electronic current awareness use within a large London law firm. The study found that selection, re-aggregation and forwarding of information by multiple actors gives rise to a complex sociotechnical distribution network. Knowledge management staff act as a layer of "intelligent filters" sensitive to complex, local information needs; their distribution decisions address multiple situational relevance factors in a situation fraught with information-overload and restrictive time-pressures. Their decisions aim to optimise conflicting constraints of recall, precision and information quantity. Critical to this is the use of dynamic profile updates which propagated back through the network through formal and informal social interactions. This supports changes to situational relevance judgements and so allows the network to 'self-tune'. These findings lead to design requirements, including that the system should support rapid assessment of information items against an individual's interests; that it should be possible to organise information for different subsequent uses; and that there should be back-propagation from information consumers to providers, to tune the understanding of their information needs.

1. Introduction

People acquire information in many ways. Understanding the variety of ways in which information is acquired and how these play out *in situ* is of key interest to researchers and system designers interested in understanding how to design for more effective information interactions. Whilst a good deal of information behaviour research has focussed on active information seeking, there is a dearth of literature on information behaviours surrounding passive information acquisition. Included within this is the way in which people access and use current awareness services. And yet systems that support current awareness monitoring provide an increasingly significant way in which users obtain information. Professionals, in particular, represent a

significant user-group, given their need to stay abreast of developments in their field. Information providers are reacting to this need through the provision of increasing numbers of electronic current awareness services (otherwise known as 'alerting services', 'auto alerts', 'selective dissemination of information' (SDI), or 'notification systems'). To take an example from the domain on which we focus in this paper, Lexis®PSL brings together a range of different information and resources specifically for lawyers working in different areas of law. As an addition to this service subscribers can receive a regular update which alerts them to legal developments in their area and new resources provided by the service.

Since current awareness services respond to information needs that remain relatively stable over time, they naturally lend themselves to the construction of relatively stable infrastructures for information delivery. In organisational settings these are particularly evident, their role being to control and incrementally adjust the flow of information, rather like irrigation systems controlling and diverting the flow of water. In modern organisations they are typically sociotechnical in nature, involving both people and technology in coordinated action to make the organisation as a whole more knowledgeable and effective. Understanding how these systems operate in their natural contexts represents a key resource for informing the development of user-technologies to make such systems more effective.

This provides the context for the current paper in which we report a naturalistic study of electronic current awareness distribution in a large, London law firm. Our aim was to understand the ways in which legal workers interact with and collaborate around the propagation of current awareness information in different parts of the company. Data was gathered through contextual inquiry observations with twenty one lawyers and knowledge management workers. We observed their interactions with current awareness information and used these situations to prompt discussion of broader contextual issues. From this we developed a systemic perspective on the distribution of current awareness alerts which provides insights into what people do and how they collaboratively organise action and interaction.

The remainder of the paper is structured as follows. In section 2, we discuss the research context relating to collaborative information seeking and electronic current awareness information. In section 3 we describe our data gathering and analysis method. In section 4 we report our findings and in section 5 we discuss these with particular reference to how automated systems might augment prevalent sociotechnical distribution systems.

2. Background

In order to explicate where gaps lie in current information seeking research, Bates (2002) mapped out some dimensions on which information seeking episodes can vary. One of these dimensions was the distinction between *active* and *passive* information seeking. Information seeking is active when an information end-user does something active to acquire information. Information seeking is passive when the information user is simply passively available to absorb it (Bates, 2002). To date, the literature on information seeking has focussed predominantly on *active* information seeking (such as searching or browsing an information system) (Bates, 2002). However, this is only one part of the picture when it comes to the variety of ways in which people acquire information. For example, Erdelez (2005) describes *Information Encountering*—instances of "…accidental discovery of information during active search for some other information" (Erdelez, 2005, p.180).

An information seeking activity which frequently involves passive information acquisition is *monitoring*. Monitoring is defined by Ellis and Haugan (1997) as "maintaining awareness of

developments and technologies in a field" (p. 396). Whilst this definition is useful, it is perhaps over-constraining, since it seems limited to understanding the state-of-the-art in fields of human endeavour; one might equally think of monitoring as applicable to understanding the current state of many kinds of situation or system (e.g. industrial processes or the physical environment). Nevertheless, a characteristic of monitoring is the way in which information acquisition is triggered. Ideally, systems are configured such that passive information acquisition is triggered by some change in the information environment (with information needs remaining relatively static). In active information seeking it is a change in information needs which triggers seeking (with the information environment remaining relatively static).

Current awareness monitoring—the kind that concerned Ellis and Haugan (1997) and forms the focus for this paper—can combine both passive and active information seeking (Bates 2002; Makri, et al. 2008). But given the role of environmental triggers it frequently lends itself to passive information seeking. In doing so, however, it is dependent upon some environmental infrastructure which can promote the right kind of information exposure (Bates, 2002). Such infrastructures can be physical (e.g. a wall barometer), electronic (e.g. a train controller's line display), or social (e.g. an old fashioned town crier) or a combination of these. In setting up and configuring these infrastructures, of course, users can commit different levels of effort (e.g. ordering a regular newspaper delivery or simply sitting in view of an ambient display). And so it is perhaps useful to consider the active/passive dimension as continuous rather than discrete, with current awareness monitoring typically occurring at the more passive end.

The significance of current awareness monitoring for professional groups has been highlighted in a number of naturalistic studies of information seeking by Ellis and colleagues (Ellis, 1989; Ellis, Cox & Hall, 1993; Ellis & Haugan 1997) and others (Meho & Tibbo, 2003; Makri et al., 2008). Ellis (1989) reported a series of information seeking behaviours of a group of social scientists. This included behaviours such as *starting* in a new area (e.g. using review articles), *chaining* (following citation connections) and *browsing* (semi-directed or semi-structured search). As part of this, Ellis noted the importance that *monitoring* had for social scientists in helping them stay up-to-date with developments in their field of study.

Since that time, a number of studies have explored the applicability Ellis' original model with other professional groups, each with similar findings. These include a study of research physicists and research chemists (Ellis et al., 1993), a study of engineers and research scientists (Ellis & Haugen, 1997); and a study of academic Lawyers (Makri et al., 2008). The question of the information behaviours of social scientists was re-addressed and updated by Meho and Tibbo (2003). What each of these studies demonstrated was that *monitoring* is a ubiquitous activity which can take many forms, frequently combining both formal routes (such as keeping up with specific journals or reviewing publisher's catalogues) with less formal routes such as the use of social networks (such as visits to conferences and dissemination by *information gatekeepers*).

Makri et al's. (2008) study is of particular relevance given its focus on the legal domain. These authors performed semi-structured interviews and naturalistic observations with 27 academic lawyers including law students, teaching staff and research staff. Using the Ellis (1989) model as a basic framework, they extended it by organising behaviours into three major categories: *identifying and locating*, *accessing*, and *selecting and processing*. They also qualified behaviours in terms of the *object* to which behaviour was directed; for example, one might *identify and locate* a resource, a source, a document or some content. They also identified the use of both active and passive monitoring strategies. Active strategies included manually conducting regular searches or browsing particular sources in digital law libraries or websites. Passive strategies

included subscribing to email alert services from sources such as government departments, publishers, think tanks or organisations responsible for different areas of law.

Despite the clear significance of current awareness monitoring demonstrated by Ellis (1989), Ellis, et al. (1993), Ellis and Haugan (1997), Meho and Tibbo (2003) and Makri et al (2008), it has rarely been the subject of studies in its own right. Predominantly, the focus has been on active information seeking, and this has been as true in the legal information seeking literature as it has been elsewhere. For example, Yuan (1997) looked at the effects of the end-user searching behaviour of law students over a period of a year. Kuhlthau and Tama (2001) looked at information seeking and use in the context of the work of a group of lawyers. And Komlodi and Soergel (2002) focused on lawyers' use of their memory and externally recorded search histories to inform their later searches. Focussing more closely on design, Dempsey, Vreeland, Sumner Jr. and Yang (2000) described the design and evaluation of two information retrieval systems specifically for supporting legal researchers in browsing and searching across legal websites. And Marshall, Price, Golovchinsky and Schilit (2001) used findings from a field study of legal research in law school Moot Court (simulated court) to design an e-book incorporating wireless access to information resources.

Some studies of legal information seeking have drawn attention to issues of collaboration. For example, Blomberg, Suchman and Trigg (1996) reported on their experiences of designing a system to support document search and retrieval following observations of the collaborative use of a lawyer's filing cabinet. They noted that information seeking in a law firm often involved 'walking the halls,' asking colleagues if they had ever drafted a particular type of document or one including specific provisions. Jones (2006) conducted a field study of Legal Aid workers resulting in a number of recommendations for design. She found the workers collaborated heavily (often using listservs and contacting practicing lawyers for assistance and advice) and concluded that future systems should support the social nature of legal research by acting as online repositories that facilitate the sharing, annotation and tagging of documents. And examining how aerospace engineers and lawyers shared their search histories with colleagues, Komlodi and Lutters (2008) found that individual search histories provided useful artefacts for discussion between lawyers and that sharing enhanced other aspects of their knowledge work, such as information seeking and use.

A few user-studies have taken current awareness monitoring as a focus although these have fallen outside the legal domain. Fernandez (2002), for example, reported on a small survey of cross-faculty researchers at a Canadian university. Among other things, the results demonstrated the use of a combination of *active* and *passive* methods for staying up-to-date, with the most common methods being accessing PubMed (the survey response rate from biology researchers was very high), scanning journal tables of contents, and receiving email alerts. Adams, Blandford, Budd and Bailey (2005) presented an account of the design and use of an organisational awareness tool that enabled clinicians to encounter current awareness information in 'bite-sized' chunks during less busy moments of the day. Hinze, Buchanan, Jung and Adams (2006) used findings from a series of studies with UK clinicians and patients to propose a novel alerting architecture. The architecture responded to the fact that users reported the desire to combine current awareness alerts with supporting background information. For example, to accompany public press releases about significant health issues, clinicians wanted to receive relevant research results that could support them in responding to patient enquiries. Also, patients wanted to receive related educational materials to help them interpret alerts from electronic health records reporting changes in their condition.

The role of current awareness alerts as tools to support collaboration within modern heterogeneous research communities was explored in a study by Farooq, Ganoe, Carroll, Councill and Giles (2007) who reported three cumulative user-studies of awareness mechanisms in CiteSeer. They found that in order to promote collaboration, their target users (predominantly researchers with a computer science background) wanted RSS feeds to notify them of new CiteSeer publication events with the alert presentation depending on the nature of the event. For example, for a feed alerting a user to papers citing one of their own papers, users wanted to see a title accompanied by the sentence(s) containing the citation(s).

In summary, there is strong evidence for the importance of current awareness monitoring for a number of professional and special interest groups, and yet this is an under-explored area in information seeking research. We present an account of current awareness use within a large law firm with a focus on how participants collaboratively distributed information, with this distribution forming a complex sociotechnical distribution network. Included in this is a description of different notions of relevance and how these were acquired. We conclude our review by discussing relevance briefly as background to subsequent sections.

Relevance is a central issue for information science and has been discussed and debated extensively. Emerging from these debates is the idea that there are two principal classes of relevance: objective or system-based relevance, and subjective or human-based relevance, corresponding respectively to approaches adopted within the tradition of IR evaluation and user-oriented studies (Borlund, 2003). The system based notion treats relevance as objective and static, whereas the human-based notions regard relevance as subjective and subject to restructuring (Borlund, 2003). Within the human-based notions there are further subtypes. In both system and human based relevances, the key distinction concerns what it is that an information object is being related to, whether to the query, the request, the information need, or the underlying situation (Saracevic, 1996). The question is: relevance in relation to what?

The idea of situational relevance was introduced by Eisenberg and Schamber (1988) and Schamber, Eisenberg and Nilan (1990) and corresponds to the utility of information as perceived by a user in a particular situation (Huukonen & Vakkari, 2006). It is highly context-dependent as well as potentially dynamic (Borlund, 2003). It relates the "utility or *usefulness* of the viewed and assessed information object(s) by pointing to the relationship between such retrieved objects(s) and the work task at hand underlying the information need as perceived by the user" (Borlund, 2003, p.915). As such, situational relevance contrasts, for example, with the objective or system based notion of relevance used in traditional IR evaluations which relates the topicality of information objects to that of an associated IR query.

3. Method

We conducted a contextual inquiry (Beyer & Holtzblatt, 1998) with lawyers and knowledge management workers within the London office of an international law firm. Contextual inquiry is a user-centred design method in which a researcher performs one-on-one observations with users whilst discussing their activity. Rather than following a fixed interview protocol, a contextual inquiry is structured by the activity itself. The role of the researcher is to observe and ask pertinent questions concerning what is being done and why (Beyer & Holtzblatt, 1998).

Participants were recruited through a combination of general email requests, snowball sampling (Johnson, 1990), and personal recommendation within the Dispute Resolution and Real Estate Departments of the firm. This provided a sample of nine fee-earning lawyers, eleven knowledge management workers and one trainee lawyer, all of whom received regular electronic current

awareness information. In the interests of observing 'actual' work, in advance of the sessions participants were asked to allow current awareness email alerts to accrue unopened in their inboxes during a period leading up to the session. They were then asked at the beginning of the session to work through these emails in the way that they normally would do.

In all but one session, audio recordings were made of the conversation. In one session it was possible to record the user's computer screen while she worked. All participants were observed once except for one knowledge worker who took part in three sessions. Sessions lasted forty five minutes on average and in total over sixteen hours of data were recorded.

The recordings were transcribed and analysed using Grounded Theory (Strauss & Corbin, 1998). Grounded Theory is a set of analysis techniques which provide a focused, structured approach to qualitative research. The methodology's flexibility can cope with complex data, and its continual cross-referencing between analysis of emergent themes and data allows for the grounding of theory in the data and the uncovering of previously unknown issues. Data is initially fragmented and then reconstituted in terms of underlying concepts using coding. Abstraction is achieved by grouping similar phenomena into higher-order 'categories' and associations are identified between them. Selective coding is used to prioritise major phenomena. There were few practical limitations on data-gathering and it was possible to achieve a high degree of theoretical saturation around major categories.

4. Findings

4.1 Participant Roles and their context of work

Our participants included fee-earning lawyers (or *fee-earners*) and knowledge management staff. Fee-earners are qualified solicitors who work directly on client cases (or 'matters'). They are so named because their time is used as a basis for charging clients. They may spend this time providing advice on legal rights and duties, preparing legal contracts for commercial transactions such as mergers and acquisitions, performing regulatory investigations or executing litigations.

Knowledge management staff provide an infrastructure of knowledge tools and resources within which this work takes place. A significant category of knowledge management staff is the Professional Support Lawyer (PSL) (similar roles are Knowledge Management Lawyer or Practice Development Lawyer, depending on role emphasis). PSLs are qualified lawyers who, rather than conducting fee-earning work, have responsibility for a firm's knowledge assets and systems (excluding the library). PSLs draft precedents (standard contracts) and practice notes (practical guides), manage the in-house knowledge management system, respond to legal queries from fee-earners, provide training, and deal with the provision of current awareness information. In many of these tasks they are supported by Knowledge Management Executives (KMEs) (also called Practice Development Executive, Practice Development Assistant), Paralegals, Researchers and possibly legal secretaries.

In common with many complex knowledge work domains (such as medicine, teaching and research), staff tend to specialise in particular areas of practice (e.g. corporate finance, employment law, intellectual property, family law etc). This specialisation applies to knowledge management workers as well as fee-earners and is reflected in practice area groups at the departmental level or at the level of departmental sub-groups. The firm dealt with a range of practice areas including dispute resolution, real estate, corporate, employment and tax.

At the site where we conducted the study, the firm employed around 1800 staff in total, of which 900 were lawyers. The offices were physically divided up to house a few large departments, each dealing with a particular area of law. Each department was arranged on one or more floors of the company's buildings. Each floor was laid out as a number of adjoining open-plan areas given over to administrational and secretarial staff, photocopiers, printers, water-coolers etc. These open-plan areas were surrounded by offices and meeting rooms, with each office occupied by between one and four lawyers and/or trainees. Generally, an associate lawyer shared with a trainee. Many lawyers kept their doors open whilst in the office allowing people to come in and talk. The open-plan areas were also locations for informal communication. There were also departmental libraries where knowledge management staff worked and lawyers visited. Corridors between open-plan areas also had connecting kitchens for making coffee or grabbing some fruit.

4.2   A Current Awareness Network

In the law firm, fee-earners and knowledge management staff subscribed to a wide range of current awareness services. These provided news and business information or more technical legal information (e.g. legislation updates, legal judgements) and materials (e.g. standard forms and practice notes). Services focussed on either information relating to specific legal practice areas or to specific industry sectors. Content was defined in terms of automated filtering expressions (i.e. a query) or was hand selected as corresponding to a pre-defined area of interest (e.g. food safety, nuclear energy, insurance law). Some services were subscription-based, others were free. For some alerts, staff arranged their own subscriptions including defining any filter expressions; in other cases these were set up by a single member of staff on behalf of a group of colleagues who had shared interests.

Email was the primary current awareness delivery mechanism (only one participant used RSS feeds). Current awareness emails typically presented users with a collection of items. These resembled automated search results (which many alerts were), with each item acting as a summary, or surrogate, of a longer document.

Participants received information with a view to informing their own work or to inform others they supported. Disseminating information around the firm was primarily the responsibility of knowledge management staff, although it was something that all staff might do (i.e. forwarding material to colleagues or clients). Dissemination happened in a number of ways. Bulletins and newsletters tended to be created collaboratively by small teams of knowledge workers led by a PSL and potentially including KME's, researchers and/or legal secretaries under their supervision. Within these teams, responsibilities were distributed for different aspect of distribution, such as selecting and authoring content, formatting, reviewing and circulating. This stratification of knowledge management staff, each with different qualifications, experience and responsibility, made this setting particularly prone to collaboration.

Bulletins were essentially re-aggregations of selected content from incoming alerts with minimal or no editing, and were sent out frequently (often daily). Newsletters might include more bespoke content such as articles written by a PSL, and were less frequent. Both were circulated around the firm via email mailing lists. In addition, ad-hoc emails about specific items might be sent to selected individuals depending on perceived significance to their work and to the work of others in a group.

The selection, re-aggregation and forwarding of current awareness information gave rise to a complex distribution network. Figure 1 shows this network as it involved participants in our study. Figure 1 is based on the regular services that people sent or received as noted from the

interviews and observations (ad hoc distribution is not included). In the figure, numbered circles represent regular updates. Circles outside the main square represent updates originating outside the company. Those inside the square were compiled in-house by knowledge management staff.

Each lettered square represents one of the study participants. In the figure, current awareness information moves from left to right. Lines coming into the left of each square (participant) show the updates that participant received. Lines coming out to the right show what that participant sent. Participants shown on the left (A to K) are knowledge management workers (PSL's, KME's and Researchers), and hence are active re-distributors. Participants to the right (L to U) are fee-earning lawyers. They also sent information on, but not as regular updates.

For example, participant D (top left) received (at least) four regular current awareness updates from outside the company (2, 29, 16 and 12). From these she selected information to compile a regular newsletter (54) which she then sent to (among others) participant N (the figure is limited to participants in the study). Collaboration in the production of bulletins and newsletters is shown by linked squares. For example, participant C (a PSL) was assisted by participant B (a PDE), to produce (43), a monthly newsletter circulated to the Financial Institutions Disputes Group.
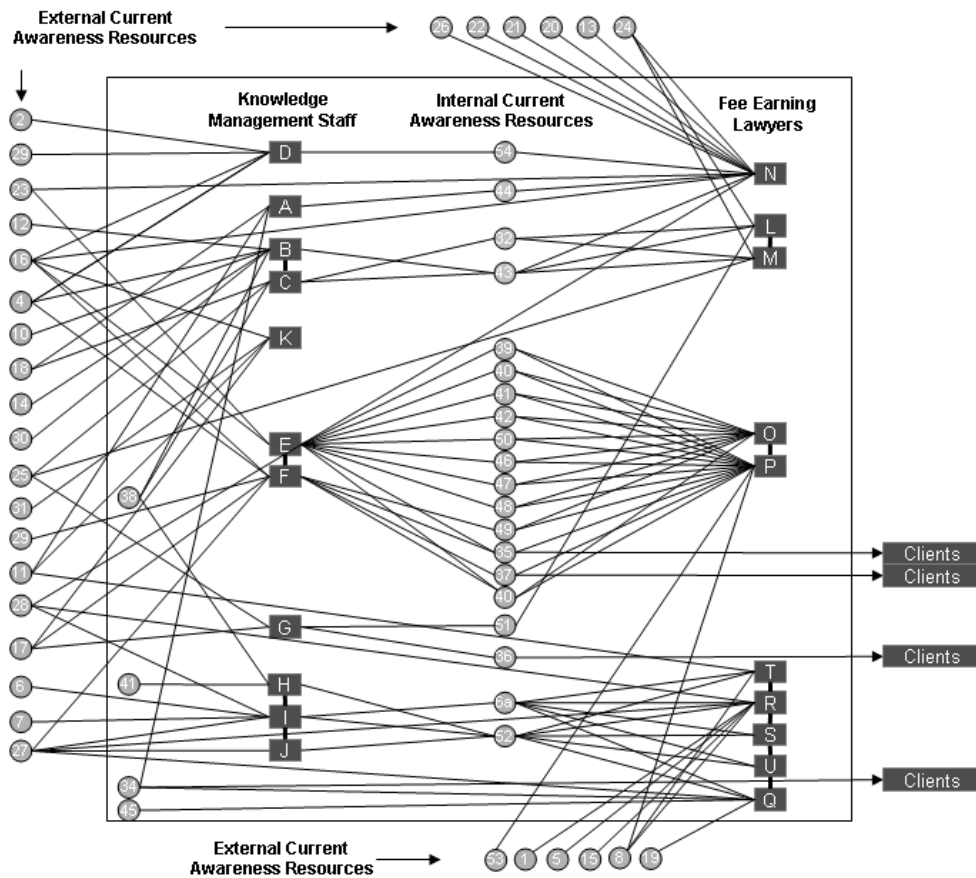

Figure 1. The distribution network of current awareness updates received and sent by
participants in the study as reported during the interviews.

Some in-house current awareness updates were sent to company clients (figure right). For example, E and F collaborated to produce (35), an Environment Planning and Regulatory newsletter, which they sent to selected clients and fee-earners. Providing current awareness information to clients was seen as important part of maintaining a relationship and procuring new business. This information provided clients with news and analysis of developments important to their area of interest. Alerting clients to legal issues that might affect them was considered a way of maintaining the client's confidence in the company.

Figure 1 shows some identifiable groupings (emphasised by linked squares). These correspond to practice area groups within our sample. Six practice areas are represented (see table 1).

Participants	Practice Area
D, N	Engineering, Procurement and Construction
A	International Arbitration
B, C, L, M	Financial Institution Disputes
E, F, O, P	Environment, Planning and Regulatory
G	Intellectual Property in IT
H, I J, T, R, S, U, Q	Real Estate

Table 1. The six practice area specialisations represented by our study sample.

Figure 1 also shows some variation in the number of regular updates circulated by knowledge management workers. To some extent these differences arose as a result of different levels of activity in an area of law. E and F, for example, worked in Environmental, Planning and Regulatory law which featured significant activity in a number of areas, such as food safety, emissions trading, renewable energy, the nuclear industry etc. There was also some fluctuation in how many updates a group produced for itself in an effort to find the right balance between too much and too little information (discussed in more detail below). And not all knowledge management staff sent out regular current awareness packages. K was a researcher who worked in a small department and distributed a lot of information on an ad-hoc basis. He felt strongly that he didn't want to send out group updates that might inundate fee-earners and so he only sent out selected items to individuals. Other knowledge management workers sent out occasional ad hoc emails to individuals depending on whether the information was considered urgent (and they wanted to ensure the recipient gave it their attention) and whether the information would only be of relevance to a small number of people within a group.

4.2.1 Back-propagating Dynamic Relevance Information
In Figure 1, current awareness information moves from left-to-right. At each point, distribution decisions were based on situational relevance factors concerning those downstream, informed by a rich and socially mediated knowledge of what those situations were. The acquisition of this knowledge had the effect of 'tuning' the network. In this section we discuss the variety of mechanisms through which profiles were updated.

The forward flow of information combined with back-propagating profile information between different roles in the network is summarised in figure 2. This shows that whilst information flowed forward from external information suppliers to knowledge management staff and fee-earners and from knowledge management staff to fee-earners and clients, profile information flowed in the reverse direction from fee-earners to knowledge management staff and from knowledge management staff to external suppliers.
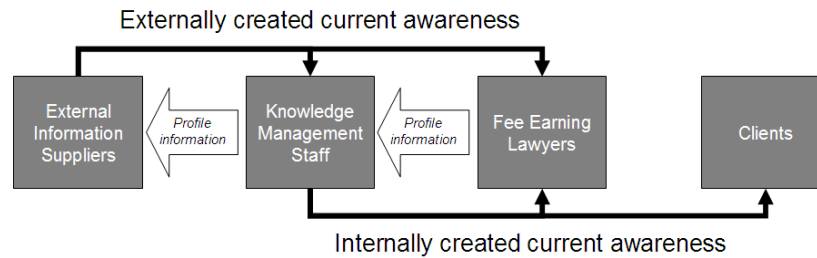
Figure 2. The network in overview showing current awareness information flowing from left-to-right, and socially mediated profile information flowing from right-to-left.

Profile information was communicated through interactions that took a number of forms ranging from formal, scheduled interactions to more spontaneous and informal discussion. On the more formal side, some knowledge management staff scheduled periodic interviews with fee-earners in their group in order to find out about their current work. Participant G, a Practice Development Executive in Real Estate, described how they did this:

G   We go through the cases that they are working on at the moment, so parties, what's happening in Counsel, officers, and also the stage that they're currently at, any major applications that they've made before. So,… because we see them every quarter and we spend about two hours going through all the information with them, I'm very au fait with where they're currently at with their projects.

These interviews were recorded and summarised in a written note. The extract indicates that a primary issue was the matters that members of the group were working on. This reflects the fact that this is the most important contextual relevance factor, and also the one most liable to change. G's comments also indicate that it is important to know what stage someone is at in a particular project.

PSLs also led monthly training sessions or 'Practice Development Meetings' and these provided an opportunity for lawyers to go round the room explaining what they were currently working on. Similar round-the-room discussions were conducted at monthly group lunch meetings. These provided a less formal forum for discussing current work and sharing know-how.

In addition to activities specifically designed for eliciting information about current work, PSLs also described some less formal and more serendipitous ways in which they gathered insights about group members' activities. Some geographically distributed groups used online discussion fora and email circulation lists to discuss their work. PSLs might take an active part in these, but might also simply watch the conversation. Enquiries from fee-earners were also an opportunity for finding out what they were interested in. PSLs often talked about sending new items to fee-earners which related to previous enquiries. Finally, informal conversations with fee-earners, group partners and other PSLs were all cited as embellishing an understanding of what people were doing. As participant A (PSL) explained,

A   Our group is fairly small and it's actually up there [indicates head], I know everybody and I go into their rooms and speak to them and they come and see me, so I know what's going on, roughly. […] I also keep in very close touch with all the other KMLs (PSLs) across the network so that we can swap information about what

10

our various fee-earners are doing. So I know what people in Paris are doing or in New York are doing, and they know what people in London are doing.

The richness of this communication depended in many ways upon the proximity that comes from working within a related group within a single organisation, and in many cases within close spatial proximity. Communication across organisational boundaries is typically less rich. Whilst it can be said that profile information flowed from participants to external information providers, in most cases this was impersonal and impoverished, such as setting up filter terms or selecting a new topic-based package. These kinds of service modifications were also rare and tended to lag significantly behind changes in interests.

Nevertheless, there was one notable example of rich and frequent profile information crossing organisational boundaries. Resource 6 was a highly regarded subscription-based daily update received by the Real Estate department. Each morning participant I, a researcher in Real Estate, received that day's version of the update as generated automatically by filters defined by the service provider. On receiving the alert, participant I triaged it by removing items that she viewed as non-relevant and emailed her selections back to the external provider. The provider then produced a formatted version of the edited alert and circulated this more broadly to staff in Real Estate (represented as if originating from I as 6a in figure 1).

One purpose of this process was to improve the precision of each update. However, it also provided the service provider with relevance feedback through which he could adjust the automated filters. Generalising from a set of exclusion instances to rules of exclusion, however, would require some inductive guesswork. Specifically, the service provider would need to know the grounds for each exclusion (e.g. inappropriate deal size, wrong jurisdiction etc.). And so the representative followed up the exchange with a phone call to participant I to find out why she had removed the items that she did. Hence, relevance feedback led to the adjustment of profile information.

To conclude this section, the forward flow of information within the network was controlled using knowledge of situational relevance which itself is acquired through rich, socially mediated profile updates propagating backwards through the network. This enabled the network to be sensitive to dynamically changing local interests. Hence, it was capable of tuning to the interests of its members. Proximity plays an important role in this communication, in particular, in updating knowledge of *project relevance*, the most dynamic form of situational relevance in this setting.

4.2.2 Information Overload and Time Constraints

It was evident from our observations that as an organisation the law firm placed a great deal of significance on the value of its employees staying up-to-date, and electronic current awareness services had a key role in supporting this. All participants received some form of alerting service, whether this was from external sources or from sources within the company. However, a frequently occurring theme for both fee-earners and knowledge management staff was the problem of information overload. Participants frequently found attending to current awareness information overwhelming. Participant R's perspective was common amongst fee-earners:

> R   It's constant in the sense that I received, for example, all of these I have left in my inbox, my e-mails in relation to things I ought to know; news in relation to Real Estate and Property; I never have time to read those so leave them there in my inbox ((laughs)), but nevertheless I have to eventually look at them.

11

Current awareness information was also a source of information overload for knowledge management staff, although their perspective was somewhat different, tending to see this as an inevitable part of their role. For example, participant K, a researcher in Intellectual Property said,

> K   Well in a sense it's kind of my job to suffer information overload to some degree I mean part of my role here is to act as an intelligent filter for all this information that's being produced on a daily basis so that the fee-earners can digest it, can take it in and utilise it so I probably do have information overload but I do see it as part of my role is to manage that for the Group so that they don't.

Knowledge management workers frequently saw their role as one of protecting fee-earners from information overload. Related to information overload, time was a key limiting factor on interactions with electronic current awareness, both in terms of the time available and when those times might occur. For fee-earners, interacting with current awareness information fell outside their main responsibilities and engagement was typically fleeting, intermittent and opportunistic. For knowledge management staff it formed a more significant part of their work schedule.

For fee-earners, decisions about engagement depended upon a balance between expected benefit and the temporal rhythms and demands of ongoing case work. As one associate said,

> L   So let's say if I was working on a case and I was at the stage where we instruct an expert, and I received this weekly e-mail and hypothetically speaking I had a spare ten minutes so I thought I'll have a look […] very often you're not going to have time to look into these things because you're so busy…

Knowledge Management staff were similarly affected by time pressures. As a PSL explained,

> D   Well, yeah, I mean, there's a lot of stuff that comes in by e-mail, but I've cut down quite a lot on what I do receive because I physically just don't have the time to look at everything that comes in.

4.2.3 Situational Relevance
Information overload and time constraints represented the major obstacles to engagement with current awareness information. Together with the need for up-to-date information, these pressures explain the critical role that knowledge management staff played in sifting, selecting, forwarding and providing continued access to relevant information. They offered an essential "personal touch" by understanding what information was important and reducing the need for fee-earners to spend time reviewing irrelevant information. However, relevance was a multifaceted and contextually bound concept which depended upon the appreciation of a number of factors.

For any one participant there were a number of ways in which interests could be characterised. Further, some areas of interest were more critical than others depending on how they related to work. Given these different levels of priority, different levels of effort were invested in passing on, or, from the end-users' perspective, engaging with different kinds of current awareness information.

Two ways in which information could be relevant related to the extent that it informed or related to a current case (lawyers call them 'matters'), or the extent to which information related to an individual's broader practice area interests. These two kinds of relevance and ways in which they differentially affected engagement are illustrated in the following extract. Here participant P was reviewing an item about genetically modified food in an internal, daily update (41 in figure 1),

> P   I'm actually on a case at the moment that involves food law and it's a growing area of my practice so I would read this with a little bit more attention than perhaps I read the others… although just having quickly glanced over it it's about GM and a challenge by the Friends of the Earth to the import of GM foods, [...] it's not something that directly impacts on what I'm doing at the moment. [...]. So having read that [indicates summary] I would then file it.

P's initial assessment of the item was that it related to a current *matter*. We refer to this as *project relevance*. He also saw it as relating to food law which is a "growing area of his practice". We refer to this as *practice area relevance*. This extends beyond a specific project by relating information to an individual's focus within a given area of professional practice. Significantly, an initial assessment of *project relevance* provoked attention from P, but a moment's reflection changed his assessment. On realising that the item was not relevant to his current work, but was nevertheless relevant in terms of his broader area of practice, P opted for reading the summary and moving on.

Another form of relevance arose in virtue of a lawyer's relationship with a particular client. Law firms develop long-term relationships with the clients. Partners take responsibility for fostering and maintaining these relationships, although Associate Lawyers can become involved on an ongoing basis. Consequently, information that might inform such a relationship was significant. We refer to the relevance of information in virtue of a client relationship as *client relevance*.

> S   …so like here [reading] '[client name] confirmed it will divide' […]
>
> Q   Is that interesting at all?
>
> S   Yeah that is pretty interesting; it's one of our clients actually [client name] and I do work for them so it's very relevant to me so I would try and read that article.

Given the aims of the firm to develop their business, *client relevance* carried a particular significance. For example,

> O   I was sort of trying to develop a relationship or get a client in the pharmaceutical sector, so I'm going to be on the lookout for issues that might affect them in the same way, to kind of saying, "We're here if you need us and by the way you might be interested to know about this, I thought this might interest you." Just as a way of keeping their lines of communication open and hopefully bring some work in.

Different kinds of relevance motivated different kinds and levels of activity, in terms of both personal engagement with information and forwarding it through the network. We reviewed the interview data in order to explore the range of different kinds of situational relevance expressed by all participants. This provided a basis for five different types of relevance based on different ways in which information could relate to participants' work. These are listed below in approximate order of significance with project relevance (top) being the most significant. In each case we give a quotation to illustrate how they are realised in participant's practice.

Project relevance – Relates to someone's work by virtue of an ongoing project.

> N   [reading] "Revisions to UNCITRAL rules" The arbitration that I'm currently working that's in a hearing at the moment is under the UNCITRAL rules, so I had a look at that.

Client Relevance – Relates to someone's work due to a client relationship (or potential relationship)

13

> S   I personally do work for them, so if there was something about them in the press I would try and read the article. There are lots of articles about them they are a very big company and they're often in the press, but if I see the article I usually read it.

Jurisdictional Relevance – Different legal jurisdictions operate under different rules and lawyers and legal workers specialise. Information is *jurisdictionally relevant* if it relates to the jurisdiction within which someone works.

> I   Well, yeah we're definitely interested in this; we've done work in Russia before.

Practice area relevance – Relates to someone's work by virtue of the general area of law in which they work.

> T   [Case name] which is all about insolvency and rights against guarantors under Leases for example, that was […] particularly relevant to our practice, so whenever I got an alert about it, I tried to use the cross-links and the hyperlinks and everything in the alerts to make sure that I'd had a quick scan

Industry Sector Relevance – Commercial lawyers also tend to specialise in a particular industry sector (e.g. transport, construction, the public sector). Current awareness information is *industry sector relevant* if it relates to a sector within which someone works.

> K   So I've gone in and selected specifically Intellectual Property but also things like telecoms and e-commerce and other things that we deal with within the Department.

Overall, users' interest profiles were multidimensional. It is notable that within these profiles the more significant relevance factors were ones that tended to change more frequently. Project relevance, for example, would change depending on the matter that someone was currently working on. A client list was more stable but also subject to change. Practice area relevance, however, might evolve little throughout a career.

The list serves to elaborate the kind of knowledge that was important when matching information against interests within the network. For example, when assessing a news item it would be important to be able to quickly assess which companies, if any, were being discussed, or to understand the broader industry sector or geographical area (i.e. jurisdiction). Moreover, the list indicates the contextually bound and dynamic knowledge that knowledge management workers needed in order to make matching decisions.

In the next section we discuss the way participants operationalised situational relevance through the constraints of recall, precision and quantity.

4.2.4 Balancing Recall, Precision and Quantity
Given notions of situational relevance, decisions about the distribution of current awareness information were affected by a number of potentially conflicting constraints that knowledge management staff used as guidance for the design of information packages (such as bulletins, newsletters and individual emails). These were *recall*, *precision* and *information quantity*. Recall and precision are of particular interest in this context because they are widely used as evaluation metrics within information retrieval benchmark tests as part of what has come to be known as the 'Cranfield Paradigm'. However, some question has been raised about the applicability of these metrics within a user-oriented perspective (Blomgren et al., 2004).

Within the IR context, recall is the proportion of retrieved relevant documents to all relevant documents in the database; precision is the proportion of retrieved relevant documents to all documents retrieved. However, since recall and precision typically trade-off against each other (i.e. improving recall usually negatively affects precision and vice versa), as do recall and minimising information quantity, together these constraints present a continual tension in the work of knowledge management staff.

Recall – Maximising recall is concerned with ensuring that, in relation to a set of interests (perhaps expressed as a request or query), all possible relevant information is made available. We operationalised this notion through evidence of knowledge management staff emphasising the importance of passing *all* relevant information that might be available to the people within the network who they supported. Achieving high recall tends to be a significant issue in many areas of legal information seeking given the law firms' responsibilities for due-diligence. Evidence for the optimisation of recall was found, for example, in comments that a PSL made when discussing the inevitability of information overload given her role,

> D   […] to a certain extent I think in <u>my</u> role that's [information overload] inevitable, and a certain amount of it is probably not a bad thing because it then makes sure that you pick up everything.

Although information overload was frustrating, participant D could see that by being exposed to more information she was less likely to miss anything important. The possibility of missing important information was a concern shared by many knowledge management staff. Some regarded the identification of *all* the important information as a matter of professional pride. In the following, participant E expressed her frustration at one of her automated alert services which failed to bring a important news article to her attention which she discovered later through other means:

> A   It should've included that article either way. It is very important. And I only found out that this article was missing because I happened to be in LexisNexis […] I saw that and thought 'that wasn't in there'. […] that's our job so if we're not doing that it's quite embarrassing.

These extracts draw attention to the significance for knowledge management staff of achieving high recall in relation to the interests of the fee-earners they support further along the network. In some ways this offers an explanation of the tolerance that knowledge management staff had for information overload. They didn't want to miss anything important.

Precision – Maximising precision is concerned with ensuring that all information that is made available is in fact relevant; in other words, with minimising irrelevant information. In terms of the context of our study, maximising precision was particularly significant given the common experience of information overload. We operationalised this concept through evidence of participants emphasising the importance of reducing the amount of irrelevant information that was received or sent.

Some evidence for this was provided by participant K in the comments (above) he made about his role of managing information overload on behalf of his group. We illustrate this further with an extract from the observation with participant C, a PSL in the Financial Institution Disputes Group with responsibility for insurance issues.

> C   Obviously on the time perspective, it would take people quite a long time to flick through all of the publications that there are within the insurance industry. Whereas I'm basically I suppose, the centre point for that exercise being done and that then

gets the sort of, hot topics, if you like, that I see being discussed within the journals, get flagged up for fee-earners within the department.

Like K, C saw her role as reducing the amount of irrelevant information with which fee-earners need to deal with the effect of reducing the time it takes for them to access important information,

For the most part, precision was managed by controlling the information that was forwarded within the network within regular bulletins and newsletters. An alternative strategy for ensuring precision (preferred by K) was to send information that was relevant to one or two people directly to them rather than to include it in a more general circular, thus reducing the amount of irrelevant information that everybody else needed to deal with. As A explained,

> A But we have to be careful that we don't overload fee-earners. I tend only, if I think something is relevant to one or two fee-earners, like for example these people I know are working on a Korean matter, I will send it just to them, I won't send it to the whole group because I think otherwise, (a) it's not relevant to 85% or 95% of the group and, (b) if I'm constantly sending e-mails people will stop reading them when they see my name flash up, is my feeling, so that's how I tend to deal with them.

Precision and recall were two of the constraints we saw in operation in dissemination decisions. A final constraint was information quantity.

Information quantity - In the last extract, A indicated an interest in limiting the amount of information she sends out for fear that people might stop paying attention to her. A number of knowledge management staff indicated this as a potential danger of poor precision. In addition to this motivating high precision, it was also cited as a reason for simply limiting the quantity of information that was circulated. Limiting quantity and ensuring precision are clearly closely related, since by minimising irrelevant information quantity is reduced. However, controlling information quantity was also found to operate independently of precision, and so is included here as an independent constraint. Two PSLs said,

> D When I came here first and started doing newsletters I used to put everything in […] I realised that people weren't reading the newsletters so an awful lot of effort was going in […]. So I just thought, well, this is silly, you know. […] So we decided not to put articles in the newsletter anymore either.

And,

> J Unless there's something exceptional going on, you shouldn't have more than about eight stories because people just can't read it and if you end up with too much then people don't read any it so it becomes counter-productive.

Where multiple constraints operate over a task, trade-offs may need to be found. Under conditions of uncertainty, precision and recall frequently conflict. Similarly maximising recall and limiting information quantity might independently lead to very different behaviours. Balancing these was part of the complexity of the knowledge management task. However, one way in which an optimal balance could be achieved was through a clear understanding of local situational relevance as discussed above, and the different levels of priority that these gave to information. In other words, given knowledge of the work and interests of people at any point downstream in the network, knowledge management staff were in the best position to manage these trade-offs effectively.

4.2.5 Information Collections

As current awareness information passed through the network, participants frequently selected from what they received to create new bespoke collections. These acted as local *caches* which

supported different kinds of use. In general, the need for local collections reflected separations between information acquisition and subsequent situations of use. These arose either because of the distribution of an individual's work tasks over time or the distribution of labour between collaborating team members, or because a situation of use had not occurred but might be anticipated at some point. These factors gave rise to two kinds of collection we refer to as: *known purpose collections* and *future reference collections*.

Known purpose collections - Collections (or lists) for a known purpose were used to set information aside with some specific goal in mind. Hence their shelf-life was limited to the period of the task. For example, a fee-earner might collect information to read later, a PDE might create a list of items to include in the next newsletter or bulletin, or a PSL might send one or two selected items to an individual. These collections took the form of dedicated folders within email clients, lists written on pieces of paper, or lists and items emailed between colleagues. Characteristically, it was the purpose that provided these collections with their identity (e.g. 'for Friday's bulletin').

Future reference collections - Future reference collections were more long-term. These were created in anticipation of some as yet unknown situation for which the information might be valuable. Here the separation between acquisition and situation of use occurred because the situation was dependent upon external factors and was not yet evident. In some cases future reference collections were physical files of printed documents, in others they were email folders. Participants were selective about what they kept; nevertheless, these collections could be very extensive, particularly around areas of key interest. Participant H, for example, kept around 30 folders of information relating to a single Act of Parliament,

> Res   So this is one single Act and you've got it must be thirty folders there or more on that single Act
>
> H   Yeah it was an Act that changed a lot of things about the way properties are dealt with so it had lots of implications. […]

Some fee-earners, however, chose not to create future reference collections given that nodes upstream in the network created such collections and that these could be called upon. In these cases, the reason for relying on upstream collections, such as those maintained by knowledge management staff or external providers, included the perception that the expertise and resources that these people had, combined with the lack of time that fee-earners had, meant that accessing upstream collections was a more effective and efficient way of recovering current awareness information. Hence, those staff that did create future reference collections frequently provided benefit to others in the network.

5.   Discussion

In this paper we have taken a systemic perspective on the distribution of electronic current awareness alerts within a corporate environment. From this perspective, current awareness information distribution operated as a collaborative activity in which actors across the organisation, who each operated locally and interacted with nearest neighbours, together created a network through which current awareness information was filtered, communicated and stored in local *caches*. Within the organisation, this network included knowledge management staff and fee-earners. However the network also extended beyond the organisation to include external current awareness service providers and clients.

We discussed a number of constraints under which people within the network operated. The significance of understanding constraints within any context of human activity is that they, at once, act as explanatory concepts for what people can be observed doing and at the same time offer parameters for judging the value of future design alternatives. Behaviour emerges from a confluence of constraints which together specify the dimensions of successful human-computer interaction (Vicente, 1999, p34). An obvious constraint was that people wanted to obtain information concerning developments in areas relating to their work. However, given the abundance of current awareness information that could be sourced and communicated, the effect of this could be significant information overload. Users often felt overwhelmed with information. This problem was particularly salient given constraints on the time available to interact with current awareness. For fee-earners, the availability of time fluctuated depending on other pressures.

The tension between a need for current awareness information and limited time explains the 'intelligent filter' role of knowledge management staff. They absorbed information overload on behalf of fee-earners and reduced the time and effort required by them to find important information from the incoming flow. However, they could only do this by obtaining and regularly updating rich understandings of what was important to the people they supported based on the work that they did. These understandings were informed by socially mediated profile updates propagating back through the network by a number of means; for the most part, these capitalised on proximity to the work itself. In one case, this kind of interaction occurred between an external information provider and an internal researcher, as described above.

Through this information, combined with background expertise (PSLs are themselves trained lawyers), information providers within the network were able to form situated 'theories of contextual relevance' which differentiated and prioritised information according to the multi-dimensional contextual relevance factors, each with differentiating levels of priority for the end-user. The system clearly depended upon human mediation, but this in itself was expensive and the system was under stress to find the right balance of cost against benefit.

The findings draw attention to some high-level requirements for future design interventions. Despite the variation of roles within the network, there was a good deal of overlap in terms of information behaviours. Forwarding was done mainly by knowledge management staff, but fee-earners did this as well. Knowledge management staff created local information collections but many fee-earners did this too. Fee-earners read and absorbed information to inform their work, but this was also true of knowledge management staff. Hence, there was some homogeneity in the network with each node potentially undertaking related activities. With this in mind, we can identify some high-level requirements for systems which are broadly applicable across the network.

Support the rapid assessment of information items against an individual's interests
Information distributors and information end-users need to be able to assess incoming information quickly as having relevance relating to the projects that people are working on, the clients they work with, and the jurisdictions, practice areas and industry sectors they work in. These factors are complex and multi-dimensional. Where a number of down-stream end-users are being considered with a range of interests, the assessment task has considerably greater complexity than, say, the case of *active* information seeking in response to one or maybe two information needs. And yet systems need to support fast decisions about whether to take action or not in relation to an item of information, and if so what that action should be.

This question is largely one of information presentation (i.e. how items are presented to users), and automated recommendation systems may well provide something of the answer here on the premise that situational relevance is reflected in day-to-day decisions that people make within the network. But for best effect a range of actions should be differentiated. Beyond simply reading items, actions which might be captured and leveraged in recommendations include adding information to particular bulletins or to particular newsletters, sending single items to individuals, or storing items in different collections designed for different purposes. All of these could result in differentiated recommendations with respect to future content in the interests of reducing the impact of information overload.

Support different information uses
Once a decision has been made about what to do with some information, systems need to support subsequent actions with minimum cost to the user. Depending on time factors, users sometimes read information fleetingly and sometimes they read in greater depth. Consequently the ideal presentation prioritises the most significant factors, so that users can extract gist information quickly, whilst also allowing them the option of engaging more deeply. Systems also need to help users re-aggregate and forward information easily. This includes supporting fast, ad-hoc communication of items to one or two individuals at a time or selecting and formatting new packages over extended periods of time. And where this work is distributed across a number of people, systems need to support collaborative decision making, compiling, editing and review.

Users need to be able to create, populate and manage different kinds of collection easily. The use of *known purpose* and *future reference* collections suggests a high-level organisation into two areas defined in terms of tasks and topic respectively. Known purpose collections were time limited according to task deadline, and so might incorporate such meta-data into their presentation or organisation. Being defined in terms of topic, future reference collections lend themselves to a hierarchical organisation. These collections can be extensive and the need to re-find content easily suggests the need for flexible browsing and search tools. Finally, from a collaborative perspective, tools might be provided that give users an interface for searching and browsing the collections of colleagues within a group, department or organisation, and in this way reduce the time needed to locate information distributed in collections throughout the network, particularly when others are not available to act as intermediaries to their own collections.

Support back-propagation of information consumers to providers to support the tuning of information provision against the general requirements of recall, precision and quantity.
The network we have described operated through a combination of human intelligence over an unintelligent electronic infrastructure (i.e. email) appropriated for the purposes of current awareness distribution. Whilst it seems unlikely that such a system could operate effectively free of human input and social interaction, we can theorise about how additional support from the technology itself could improve the performance to work ratio.

The forward flow of information was electronically mediated, and yet back-propagation relied upon social mechanisms which themselves depended upon proximity. Consequently, a network could be enhanced by providing users with downstream usage information and with associated summaries and automated recommendations. This would allow distributors to make informed decision and to assess the effects of past decisions in order to continually update their approach to subsequent distribution. Relevance feedback, however, is only useful if it is combined with a method for deriving a profile update. Relevance feedback makes a system or person aware of one or more specific judgements. But this is not sufficient for predicting future judgements. For this it is necessary to have a theory about *why* some information was relevant, and consequently why some different information might be relevant in the future. We have shown that items can be

relevant for a number of reasons (e.g. a particular project, client, jurisdiction, practice area, industry sector). Eliciting these details along with relevance feedback would support the generation of more accurate profiles.

The focus of this study has been on lawyers within one large law firm. We expect, however, that the sociotechnical network we have described might be found in other medium to large organisational settings. While some of the findings (e.g. the identification of priorities in each individual's interests) are particular to the legal domain, most are likely to be of wide applicability. This work contributes to the broader understanding of needs and practices for maintaining current awareness within knowledge communities.

Acknowledgements

References

Adams, A., Blandford, A., Budd, D. & Bailey, N. (2005). Organisational Communication And Awareness: A Novel Solution. In Health Informatics Journal. 11: 163 – 178

Bates, M. (2002). Towards an integrated model of information seeking and searching In Wilson, T.D. & Barrulas M.J. (Eds.) The New Review of Information Behaviour Research, 3 (pp. 1-15). London: Taylor Graham

Beyer, H. & Holtzblatt. (1998). Contextual design: Defining customer-centered systems. San Fransisco: Morgan Kaufman.

Blomberg, J., Suchman, L. & Trigg, R.H. (1996). Reflections on a work-oriented design project. Human computer interaction, 11, 237-265.

Blomgren, L, Vallo, H. and Bystrom, K. (2004) Evaluation of an Information System in an Information Seeking Process, Lecture Notes in Computer Science 3232, 57-68.

Borlund, P. (2003). The concept of relevance in IR. Journal of the American Society for Information Science and Technology, 54(10), 913-925.

Dempsey, B.J., Vreeland, R.C., Sumner Jr., R.G. & Yang, K. (2000). Design and Empirical Evaluation of Search Software for Legal Professionals on the WWW. Information Processing and Management 36, 253-273.

Eisenberg, M., & Schamber, L. (1988). Relevance: the search for a definition. In proceedings of the 51st annual meeting of the American Society for Information Science, 25, (pp. 164-168). Medford, NJ: Learned Information, Inc

Ellis, D., Cox, D. & Hall, K. (1993). A comprison of the infomration seeking patterns of researchers in the physical and social sciences. Journal of documentation, 49(4), 356-369.

Ellis, D. & Haugan, M. (1997). Modelling the information seeking patterns of engineers and research scientists in an industrial environment. Journal of documentation, 53(4), 384-403.

Erdelez, S. (2005). Information Encountering. In K. E. Fisher, S. Erdelez, and E. F. McKechnie (Eds.) Theories of Information Behavior. Medford, NJ: Information Today.

Farooq, U., Ganoe, C., Carroll, J., Councill, I. & Giles, C (2008). Design and evaluation of awareness mechanisms in CiteSeer. Information processing and management, 44, 569-612.

Fernandez, L. (2002). User perceptions of current awareness services: a faculty survey. Issues in science and technology librarianship, 33.
Available online at: http://www.library.ucsb.edu/istl/previous.html (accessed May 2009)

Hinze, A., Buchanan, G., Jung, D. & Adams, A. (2006). HDLalert – a healthcare DL alerting system: from user needs to implementation. Health informatics journal, 12(2), 121-135.

Huuskonen, S. & Vakkari, P. (2006). Situational relevance and task outcome. In: Ruthven, I. & al. (Eds.) Information interaction in context. (pp. 24-32). Copenhagen: ACM Press,

Jarvelin, K. & Ingwersen, P. (2004). Information seeking research needs extension towards tasks and technology. Information research, 10(1).
Available online at: http://informationr.net/ir/10-1/paper212.html

Johnson J. C. (1990). Selecting ethnographic informants. Newbury Park. CA: Sage.

Jones, Y. (2006). "Just the Facts Ma'am?" A Contextual Approach to the Legal Information Use Environment. In proceedings of the 6th ACM Conference on Designing Interactive Systems, 357-359. University Park, PA.

Komlodi, A. & Lutters, W. (2008). Collaborative Use of Individual Interaction Histories: Grounding a New Awareness Approach. Interacting with Computers, 20(1) 184-198.

Komlodi, A. & Soergel, D. (2002). Attorneys Interacting with Legal Information Systems: Tools for Mental Model Building and Task Integration. In proceedings of the 65th Annual Meeting of ASIS&T, 152-163. Philadelphia, PA.

Kuhlthau, C.C. & Tama, S. (2001). Information search process of lawyers: A call for 'just for me' information services. Journal of documentation, 57(1), 25-43.

Makri, S., Blandford, A. & Cox, A. L. (2008). Investigating the information-seeking behaviour of academic lawyers: From Ellis's model to design. Information processing and management, 44, 613-634.

Marshall, C., Price, M., Golovchinsky, G., & Schilt, B. (2001). Designing e-books for legal research. In Proceedings of the 1st ACM/IEEE-CS joint conference on digital libraries, Roanoke, Virginia, USA. 2001 (pp. 41-48).

Meho, L. I. & Tibbo, H. R. (2003). Modeling the information-seeking behaviour of social scientists: Ellis' study revisited. Journal of the American Society for Information Science and Technology, 54(6), 570-587.

Saracevic, T. (1996). Relevance Reconsidered. In Proceedings of the Second Conference on Conceptions of Library and Information Science. Copenhagen (Denmark), 201-218

Schamber, L., Eisenberg, M.B., & Nilan, M.A. (1990). A re-examination of relevance: Toward a dynamic situational definition. Information processing and managenemnt, 26, 755-775.

Strauss, A. & Corbin, J. (1998). Basics of qualitative research: Techniques and procedures for developing grounded theory. 2nd ed. London: Sage.

Vicente, K.J. (1999). Cognitive work analysis: Towards safe, productive and healthy computer-based work. Mahwah, NJ: Lawrence Erlbaum

Yuan, W. (1997), End-user searching behaviour in information retrieval: A longitudinal study. Journal of the American Society for Information Science and Technology, 48(3), 218-234