

Contact Management: Identifying Contacts to Support Long-Term Communication

Steve Whittaker

AT&T Labs Research,
Florham Park, NJ, 07932
+1 973 360 8339
steve@research.att.com

Quentin Jones

New Jersey Institute of
Technology,
Newark, NJ
qgjones@acm.org

Loren Terveen

Computer Science &
Engineering
University of Minnesota
Minneapolis, MN
terveen@cs.umn.edu

ABSTRACT

Much of our daily communication activity involves managing interpersonal communications and relationships. Despite its importance, this activity of *contact management* is poorly understood. We report on field and lab studies that begin to illuminate it.

A field study of business professionals confirmed the importance of contact management and revealed a major difficulty: selecting *important contacts* from the large set of people with whom one communicates. These interviews also showed that *communication history* is a key resource for this task. Informants identified several history-based criteria that they considered useful.

We conducted a lab study to test how well these criteria predict contact importance. Subjects identified important contacts from their email archives. We then analyzed their email to extract features for all contacts. Reciprocity, recency and longevity of email interaction proved to be strong predictors of contact importance. The experiment also identified another contact management problem: removing 'stale' contacts from long term archives. We discuss the design and theoretical implications of these results.

Keywords

Contact management, asynchronous communication, communication history, address books, PDAs.

INTRODUCTION

Theorizing about asynchronous communication has been dominated by comparisons with face-to-face communication [8]. Early asynchronous theories emphasized media differences arguing that asynchronous communication differs from face-to-face communication because of the absence of non-verbal information afforded by gaze and gesture. However, the emphasis on media differences leaves other crucial aspects of asynchronous communication unexamined, particularly those that stem from its persistent nature [2,14,15]. We explore those

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW'02, November 16–20, 2002, New Orleans, Louisiana, USA.

persistent aspects of asynchronous communication in this paper.

Research on email [15], voicemail [13], and Usenet [16] has revealed various critical features of asynchronous, technologically mediated interpersonal conversations. These conversations consist of multiple messages exchanged over a fairly extended period of time: days, weeks, or even months. This extension of conversations over time implies that people are typically engaged in multiple conversations at any given time. And each conversation often involves multiple people. These properties lead to significant problems of conversation management. People find it difficult to keep track of the content and status of their multiple conversations, as well as the identity, contact¹ information, and expertise of all their conversational partners. Maintaining knowledge of one's contacts is a significant problem in its own right [14]: we refer to this problem as *contact management*.

Contact management is clearly complex. A major problem is that people are exposed to an unmanageable number of potential contacts. This is exacerbated by widespread use of distribution lists [15]. It would be both onerous and unnecessary to store detailed information about all these potential contacts. As a result, individuals must decide: (a) which of these potential contacts are important enough to retain information about; and (b) what sorts of information to retain about these chosen contacts.

Several benefits follow from a better understanding of contact management, and the improved ability to identify important contacts: (1) Messaging systems will be improved if messages are filtered and prioritized based on the importance of the sender. [15]; (2) Tools such as electronic address books will be more effective if they are based on important contacts and associated information users truly need; (3) Contact information enables additional types of software support, such as social

¹ A contact is defined as someone who the user has been exposed to during communication. For email and voicemail, contacts not only include people that the user has communicated with directly, but also people included on distribution lists, ccs, and bccs.

recommender systems [10]; (4) Theories of asynchronous communication will be better informed.

To elucidate these contact management processes, we studied the use of current tools. People track their contacts using a variety of tools, including paper address books, PDAs, email and voicemail archives. These tools are ubiquitous, they play a central role for communication, and people invest large efforts in maintaining them. Nevertheless, there has been relatively little study of how they are used.

Our interviews indicated various criteria people use for selecting important contacts. One set of criteria relate to the *history of communication* between a person and a potential contact. These include *frequency* and *reciprocity* of communication exchanges, and temporal factors such as *longevity* and *recency* of communications. The interviews also suggested differences among users in the number and type of contacts they recorded information about. These differences seemed to relate to *interaction style*. We conducted a lab study to test which of these criteria predicted perceived contact importance.

The structure of the paper is as follows. We first present the details of a field study of contact management practice. Next, we describe a lab study that tested predictions about how communication history and individual differences predict perceived contact importance. Finally, we discuss the design and theoretical implications of these findings.

FIELD STUDY OF CONTACT MANAGEMENT

We present new findings about contact management derived from a field study of workplace communication practices. Other aspects of this study have been reported elsewhere [14]. The study had two related goals: (1) to identify the main problems informants experienced with current communication applications; and (2) to document the key strategies that users had evolved to address these problems. The study consisted of semi-structured interviews and observations of 20 business professionals. They included financial analysts, lawyers, brokers, estate agents, bankers, IT managers, academics, researchers, secretaries, administrators, marketing managers, conference organizers, and public relations specialists. They worked in a variety of settings from multinational corporations to personally owned small businesses. We asked them what communication tools they used, to explain how they used these tools, to describe the main problems with these tools, and to identify strategies they used to cope with the problems. The main tools used were email, voicemail, IM, fax, phone, and written documents. We observed the informants using their tools, also focusing on their use of communication support tools (such as address books, PDAs, and post-it notes) to manage contact information.

We first describe the nature and value of contact information and the ad hoc set of tools used to manage it. We then elaborate: (a) the problems people experience in deciding whom to maintain contact information about; and (b) the onerous nature of data entry for the large number of contacts that most people possess. Finally we document the criteria people use to decide which of this huge set of contacts to keep track of.

The Value of Contact Information

We observed a wide variety of tools being used to store and retrieve contact information. They included: dedicated tools such as personal address books (both digital and physical); corporate directories; organization charts; “tool-specific” address books in email and speed-dial lists for phones; business cards – either in rolodexes or kept loose; ‘hotlists’ – small sets of frequently called numbers placed in salient locations; pieces of paper on refrigerator doors, post-it notes, notes on calendars.

A first question we put to informants was why they thought it was so important to maintain their own *personal* contact information, when much of the information they stored was publicly available. This is particularly true for employees of large corporations, who have access to corporate directories and organization charts. Three features of current business practice led people to keep personal contact information: (1) Informants often worked with partners or clients from other organizations, and they did not have access to corporate directories for these people; (2) They often needed access to contact information while on the move. It is much easier to take one’s contact information along in a PDA or filofax than to access a corporate database from a hotel room or client’s office; (3) Corporate databases do not contain the esoteric, personal information needed to maintain a relationship with a contact (birthdays, universities, sports team allegiances, number of children, and so on).

Informants were unanimous about the value of their contact information. This was evident not only from their comments, but also from the time they invested in creating and maintaining contact archives. As one informant, Mary, a freelance researcher, pointed out, her personal contact list was a resource that pervaded all of her work:

“I cannot work today unless I have some source of contact information, some organized source so I can actually actively search for people. I use this list all the time just to browse it to find people when I need somebody to do a particular task.”

Problems: Contact Selection, Data Entry, Tool Diversity

However, contact management has a number of associated difficulties. At first glance, the main problem informants had was the *number* of contacts they needed to manage. We estimate that this number varied from a low of several hundred to well into the thousands, although reliable

estimation was hard given the large number of contact management tools people typically used, and the fact that there was often duplication between these. Upon further examination, though, deeper problems concerned: (a) the need to make an explicit decision that someone was a valuable contact; (b) the diversity of tools used; and (c) data entry.

Contact Selection. When someone calls you on the phone, leaves you voicemail, sends you email, or hands you a business card, what do you do? Do you record their contact information or not? The difficulty is that it is hard to anticipate whether, and to what extent, you will need to communicate with that person in the future. Whether someone is an “important contact” only becomes clear over time. The ease of electronic communication, especially the ability to broadcast messages to large numbers of people at little cost, exacerbates this problem: you may be cc’ed on messages, get email from various distribution lists, or receive mass mailings. To be safe, our informants often “over-saved” information, leading to huge rolodexes, overflowing booklets of business cards, and faded post-it notes scattered around their work areas. Despite this strategy, participants were still exposed to many more contacts than they recorded information about. One reason for this was the laborious nature of recording contact information.

Data Entry. Informants made it clear that contact information is costly to acquire and especially hard to maintain. They often wanted to record various types of addressing information for a particular contact: work, home, and mobile phone numbers, fax number, email address, postal address, instant messaging alias, as well as the IM system it was good for, and so on. And, as we mentioned earlier, some people found it important to include detailed personal and social information that was useful in maintaining an effective relationship with that contact.

What makes data entry especially onerous is that much information has to be copied verbatim – and must be absolutely correct. A phone number with one wrong digit or an email address containing an incorrect character is of little use. Our users were positive about email systems or mobile phone software that allow addresses to be added with one button click. However such applications only enter information that is specific to a particular communication device or software package. This was problematic because many of our participants used several communication modes with the same contact and thus needed access to multiple addresses for that person.

The cost of recording verbatim contact information makes contact selection even more important and frustrating – do I make the effort to capture this information, or do I risk losing track of a potentially important contact?

Diversity of Tools. All the informants used ad hoc combinations of tools, with some people evolving highly complex and idiosyncratic systems. For example, Mary, the freelance researcher, had over 1000 people in her email address book, a 60 page Word document containing over 1200 people, over 400 people in her PDA, as well as miscellaneous people in Christmas card lists. Ollie, a corporate research scientist, kept 7 different address books, using 2 PDAs, Microsoft Outlook, and 4 independent email address books. He also wrote key work numbers on his office blackboard. One reason why these complex systems evolved was that informants seldom ‘cleaned up’ their contact information. People were loath to delete any contact information. This seemed to be motivated both by the effort of data entry, along with the belief that even little used contact information may be relevant at some future time.

In summary, it is no surprise that many informants were not satisfied with their methods of contact management. For example, Mary was unsure which contacts she had recorded; her “system” did not make it easy to manage or access contact information. And everyone reported losing information for valued contacts from time to time. This was partly because they had information stored in a diverse, unintegrated array of records, media, and tools. And, of course the problem was made worse, because much of the information – business cards, post-it notes, email addresses – that was stored “just to be safe” turned out to be irrelevant. Thus, it just got in the way, making search for truly important contacts both difficult and error prone.

Criteria for Determining Contact Importance

Returning to the basic decision people face – is this an important contact? – we sought to find out how our informants dealt with this issue. Informants responded with a surprising consensus. Since they could not make this decision at first exposure, they relied largely on the *history of their prior interactions*. Further individual factors, such as communication style seemed to affect the number and type of contacts selected.

In our interviews, we probed informants to identify specific aspects of interaction history and communication style that were critical in determining important contacts. We asked people to walk us through their contact management tools and explain why particular contacts had been included.

Communication History

Eric, an investment banker, characterized communication between himself and valued contacts in terms of interaction *frequency* and *recency*:

“Important people to me are people that I talk with a lot. And people I’ve talked with a lot in the last week or month, then they are really important to me.”

Other informants noted that frequency and recency alone were insufficient criteria, instead pointing out that *longevity* of communication is critical. Barney, a researcher, put it this way:

“Often you get intense interaction with people associated with a specific event like organizing a trip or a meeting. But that doesn’t mean in the long-term that person is going to be important. There needs to be serious sustained interaction.”

Simon, a corporate lawyer, describes the complexity of the process of deciding when to create a detailed record for a contact.

“The people who end up in my Outlook [address book] are people who I think that I’ll have some reason to contact at some future time. Some people I talk to I have longstanding relationships with, but most of them I don’t. I just have to call them back. Now if they call me, and there’s no reason for me to call them back, then I don’t waste my time putting them into Outlook. So if people are trying to sell me a service that I don’t want or they are trying to form a relationship with Company-X that the company doesn’t need, then they don’t get in there. Of course the problem is that a lot of the time its hard to tell at the outset in some cases, so I tend to be a bit conservative about putting people in.”

Here, Simon echoes the view that important contacts are those with whom he has long-term interactions. He does not create detailed records when returning a one-shot call. He also points out the need for *reciprocity*. Some people make unsolicited efforts to get in touch with him, but unless he sees a reason to respond, there is no need to record their information. Receiving unsolicited communication from someone does not make them a valued contact. Finally, he points out that his selection process is error-prone, because of the difficulty of predicting long-term relationships on the basis of initial interactions. So he is cautious about when he records detailed information about a contact.

Individual Differences of Communication Style

We also noted differences between users in their contact management behaviors. Different communication styles led to different contact management behaviors. Some intense communicators working across organizations ended up creating and maintaining extremely large sets of contact records, whereas others working alone or in small teams got by with much smaller sets.

People were also influenced by contact affiliations. Although most of our informants had contacts outside their own organization, informants reported they were more likely to work with, and hence have reason to keep in touch with, contacts from within their organization or workgroup.

Summary

These interview data pinpoint two related problems with contact management. These are identifying important contacts and data entry. The data also suggest several criteria people use for identifying important contacts based on communication history, namely frequency, reciprocity, and temporal aspects such as the recency and longevity of communication. There also seem to be individual differences in the number and extent of contact records that stem from different communication styles.

The next section describes an experiment in which we tested these observations to further investigate how people identify important contacts. We examined the relationship between these aspects of (a) *communication history*; and (b) *communication style*; on (c) *user’s perceptions of contact importance*, using contacts extracted from email. Using experimental techniques we were able to test explicit hypotheses about these criteria.

EXPERIMENTAL STUDY OF FACTORS UNDERLYING PERCEPTIONS OF CONTACT IMPORTANCE

The experiment examined the criteria underlying user’s judgments of contact importance. We presented people with sets of contacts automatically extracted from their email archive. The archives included messages sent by the user to others. They excluded messages that users had received but deleted, as we had no way of accessing these. For each extracted contact we asked users whether they wanted to include that contact in their contact management system in order to keep in touch with them.

For each email contact, we also recorded header information from the email archive about each message involving the contact. From this data, we can compute quantitative characteristics of that contact’s communication history involving the user, including the *frequency*, *recency*, *reciprocity* and *longevity* of their exchanges. We can therefore determine the extent to which the decision to select a particular contact correlated with these aspects of communication history. The second part of the study examines individual differences in *communication style* on contact selection.

We considered an alternative experimental design, where instead of having users select contacts extracted from email, we simply looked at the contacts already in their email address books. However our field study research suggested using existing contact entries was problematic: address books often contained out-of-date contacts who had never been removed, or contacts who had been added in anticipation of interactions that never materialized. We wanted instead to collect information about *currently* important contacts.

Hypotheses

The hypotheses are derived from the criteria suggested by our users in the field study.

Communication History Hypotheses

The communication hypotheses concern frequency, reciprocity, recency and longevity.

First we expected important contacts to interact more frequently with the user. Frequency is defined as the total number of messages exchanged between contact and user, divided by the longevity of their relationship.

Frequency: Important contacts should have more frequent exchanges with the user than unimportant ones.

We also expected important contacts to show greater reciprocity, so that messages exchanged with important contacts should contain roughly equal numbers of sent and received messages. Reciprocity is defined as $(\text{number of messages sent})/(\text{number of messages sent} + \text{number of messages received})$. This definition gives a high reciprocity score to a user sending multiple messages to a contact, but receiving few in return. This situation demonstrates a high investment on the part of the user in maintaining the communication, which we would expect to be reflected in a high perceived value for the contact. Other possible definitions of reciprocity involve the use of message replies (re:). However the header logs that we collected did not contain the message subject lines needed to extract this information.

Reciprocity: Important contacts should demonstrate greater communication reciprocity than unimportant ones

We also made a more specific prediction about *unsolicited* communication, which is a specific instance of (lack of) reciprocity. We define unsolicited communication as cases where a contact sends messages to the user, but there is never any communication from the user to the contact. While this definition is simple, it may however, overestimate unsolicited communication by including people who have sent messages that the user intends to respond to.

Unsolicited communication: Contacts who send messages to the user, but never receive any communication from the user should be more likely to be classified as unimportant.

The next two hypotheses concern the *temporal* aspects of the communication history, longevity and recency. *Longevity* is defined as the total number of days between the dates of the first and last messages exchanged by contact and user. *Recency* is the number of days since the last message exchanged between user and contact.

Longevity: Important contacts should interact over longer periods than unimportant ones.

Recency: Important contacts should have interacted with the user more recently than unimportant ones.

Individual differences in communication style

The next hypothesis concerned individual differences between users in terms of their communication *style*. We classified all users into high and low frequency

communicators based on whether they exchanged more messages with contacts than the overall sample mean. We expected more intense communicators to select more contacts because of the greater effort they invested in communication.

Communication style: People who are high frequency communicators should select a greater proportion of the contacts they are exposed to.

Our final hypothesis was about the *affiliation* of the contact. Our interviewees said that they were more likely to judge as important, contacts from within their company. We therefore examined the domain name of the contact's email, to determine whether it was the same as the user's.

Affiliation: Important contacts should be more likely than unimportant contacts to come from within the users' own organization.

Method

Users

Seventeen users from a large corporate research laboratory took part in the experiment. They included researchers, managers, secretaries, computer support staff and marketing managers. Participants had been using their email system for an average of 3.0 years (standard deviation 1.8 years), and so all had substantial numbers of messages in their archives.

Task

We presented users with an on-line list of extracted contacts. For each contact we showed contact name (e.g. Jane Smith), email name (e.g. jsmith), domain name (yahoo.com), the number of messages received by the user from that contact, the number of message sent by the user to that contact, the date of the first message exchanged by user and contact, the date of the last message exchanged.

This information was presented in a spreadsheet-like table. The columns could be sorted, making it possible to order contacts by the number of messages they sent to the user, or by the domain name of the contact, and so on. This allowed users to examine and order the extracted contacts in multiple ways, while making their choices. One concern is that the columns in the table may have biased users to focus on particular contact characteristics. However, pilot studies showed that without techniques to systematically sort and view data, users quickly became overwhelmed by the task of judging hundreds of contacts.

We asked users to select important contacts for inclusion in their contact management system. They were told to choose contacts based on whether 'you might want to be in contact with them again'. Users could make three possible judgments. They could decide that contacts: (a) should be added to their contact management system, i.e. that they were worth keeping in touch with; (b) should be excluded from the system, i.e. they were not worth keeping in touch with; and (c) that they were unsure of the

status of the contact. We automatically recorded all decisions. We classified contacts selected for inclusion in the system as *important*. Excluded contacts and those that people were unsure about were classified as *unimportant*. We gave users as long as they liked for the task. They took between 15 and 75 minutes to do this, with a mean of 50 minutes, and most users said that they enjoyed the process. We also had them rate on a scale of 1-100 how satisfied they were that they had identified all important contacts, where 100 was ‘completely satisfied’ and 1 was ‘completely dissatisfied’. Overall satisfaction was 82%.

To make certain that users were genuinely satisfied with their judgments, we allowed them an optional second session one or two days later, where they could re-examine their choices and update them. Seven users chose to do this, taking an average of 12 minutes for this session.

Validating our Procedure Using Subjective Judgments

To test the validity of this procedure as a way of evaluating the value of potential contacts, we administered a brief questionnaire after the judgment task. We presented each user with a stratified random sample of 12 contacts of three types: important (i.e. those included in the contact management system), unimportant (i.e. excluded from the system), and those contacts that the user was unsure about. For each contact, we asked users questions based on [4], about various aspects of their relationship with the contact – how frequently they: (a) socialized with; and (b) collaborated on work projects. Each frequency judgment was made on a 6-point scale ranging from ‘very infrequently’ to ‘very frequently’. We also had users rate how close they were to the contact on a 6-point scale from ‘very close’ to ‘distant’. Frequency judgments were converted to numeric values with ‘very frequent’ being allocated a score of 6 and ‘very infrequent’ a score of 1. Closeness judgments were mapped in a similar way. If our procedure allows users to reliably distinguish important and unimportant contacts, then work, socialization and closeness judgments should all be greater for the important contacts.

To test this, we conducted three one-way ANOVAs with contact type (important, unimportant) as independent variable and the various user judgments (i.e. closeness, socialization and collaboration frequency) as dependent variables. Important contacts (i.e. those that the users added to the contact management system), were thought to be closer ($F_{(1,110)}=84.48, p<0.00001$), socialized with more ($F_{(1,110)}=28.80, p<0.0001$) and collaborated with more frequently than unimportant contacts ($F_{(1,110)}=9.01, p<0.003$). These results indicate that our procedure is a valid way to distinguish important and unimportant contacts.

Results

Characteristics of Extracted Contacts and the Selection Process

Before testing our hypotheses, we present some general observations about the characteristics of the original archives and the set of contacts our users rated as important. We also present some observations about the selection process.

	Mean	Standard Deviation
Number of Contacts Extracted from Archive	859	775
Number of Contacts Rated as by Users as Important	119	79
Percentage of Extracted Contacts Rated by Users as Important	19%	13%
Percentage of Extracted Contacts that Engaged in Unsolicited Communication	30%	12%
Percentage of Extracted Contacts Only Appearing as ccs or bccs	37%	21%
Percentage of Extracted Contacts From External Organizations	72%	15%

Table 1: The Character of Extracted Contacts

Table 1 shows overall statistics for the extracted contacts. A number of features of the table are worth noting:

- Consistent with our interview results, users are exposed to a large number of contacts. We extracted an average of 859 contacts from each person’s archive.
- Again consistent with our interviews, users rated as important only a small percentage (19%) of the contacts they were exposed to.
- A substantial proportion of contacts (30%) extracted from user’s archives engaged in unsolicited communication, i.e. they sent messages to the user, who never responded to them.
- A substantial proportion of contacts (37%) extracted from user’s archives never communicated directly with the user. They only appeared on the cc or bcc lines of messages.
- Another somewhat surprising statistic was the large proportion of contacts (72%) we extracted from the archive that were from outside the user’s organization, as indicated by their email address.
- There were large differences between users for all measures as reflected by the large standard deviation scores.

We also recorded users’ spontaneous comments as they made their classifications, and these bear out the statistics. Most users were surprised by the large number of contacts

we extracted from their archives, i.e. the total set of contacts that they had been exposed to. Their comments indicated that they were unfamiliar with many of these. One user with an archive containing 3229 contacts commented: “I only recognize about a hundred of these names.” The fact that users don’t seem to recognize many extracted contacts is consistent with the view that they only want to record information about a small proportion of possible contacts. There would seem to be little point in recording information about unknown people.

Users also made comments about the fact that the status of various contacts was subject to change. They pointed out that there were various people in the archive who they had previously interacted with a great deal, but who for organizational or social reasons were now unimportant. “I used to talk to Martin all the time about the reorganization but now that I’m not working on that any more we don’t need to talk”.

Testing importance hypotheses

We tested the hypotheses using logistical regression, where the dependent measure was whether the contact was judged to be important or unimportant. The analysis involved a total of 14,598 extracted contacts. The independent measures were frequency, reciprocity, unsolicited communication, longevity, recency of interactions between user and contact. We also included users as an independent variable to control for individual variability. An additional independent variable was affiliation - whether the contact was from the same organization as the user. All data were normalized. The model was significant (McFadden’s $Rho^2=0.27$, $p<0.0001$). The model along with significance calculations for each parameter (excluding users) is shown in Table 2.

Overall each hypothesis was confirmed. The coefficients in Table 2 are normalized so they show the strength of

each variable in the regression equation. As the regression coefficients in Table 2 indicate, longevity, recency and reciprocity were strong predictors of contact importance, with affiliation and unsolicited communication being weaker predictors, and frequency being a very weak (though significant) predictor.

The results for each hypothesis were the following:

Frequency was confirmed. People interacted more frequently with important than unimportant contacts (0.54 vs 0.34 messages/day).

This was a relatively weak effect as the regression coefficient reveals. Users pointed out one potential explanation for this: there are some important contacts whom they want to keep in touch with, but with whom they nevertheless interact only infrequently. Furthermore, users observed that they had previously interacted frequently with people who they now judged unimportant. Often this occurred because organizational changes meant that those contacts were no longer relevant to them.

Reciprocity was confirmed. Important contacts showed greater reciprocity than unimportant ones (0.23 vs 0.04).

As we had expected, and as the regression coefficient indicates, reciprocity was a strong predictor of contact importance. People were more likely to both send and receive communications from important contacts.

Unsolicited communication was confirmed. Important contacts were less likely to be unsolicited communicators than unimportant contacts (24% vs 27%).

Although the hypothesis was supported, the effect was somewhat weaker than expected. One surprising observation was the number of important contacts who were unsolicited communicators. One possible explanation is that our definition of unsolicited communication was too stringent. Some of the contacts we classified as unsolicited communicators may in actuality have been people whom the user intended to communicate

Parameter	Important Contacts (mean)	Unimportant Contacts (mean)	Regression Coefficient	Standard Error	t-ratio	p-value	Hypothesis Confirmed?
Frequency (# messages exchanged/day)	0.54	0.34	0.070	0.022	3.179	0.001	Yes
Reciprocity (# messages sent /total messages exchanged)	0.23	0.04	0.565	0.026	21.728	0.0001	Yes
Unsolicited Communicators (Contacts who sent but never received messages from the user)	0.24	0.27	0.202	0.032	6.230	0.0001	Yes
Longevity (# of days of correspondence)	269.0	85.00	0.594	0.028	21.497	0.0001	Yes
Recency (# days since last message)	12.4	337.5	0.475	0.044	10.809	0.0001	Yes
Affiliation (Contacts from within same organization)	0.25	0.19	0.301	0.070	4.310	0.0001	Yes

Table 2: Logistical Regression showing Effects of Communication History on Contact Importance

with at some future time.

Longevity was confirmed. Users interacted over much longer overall intervals with important than unimportant contacts (269 vs 85 days).

Longevity was a strong predictor of contact importance. The somewhat surprising fact that users communicated with unimportant contacts for about 85 days may again be the result of changing relationships. People may have had significant interactions with these contacts in the past, but social or work changes mean that they have now become unimportant.

Recency was confirmed. Users were more likely to have communicated recently with important than unimportant contacts (13 vs 340 days).

Again recency was a strong predictor of importance. Important contacts were much more likely to have been in touch with the user within the last two weeks, whereas unimportant contacts had most recently communicated about a year ago.

Affiliation: Important contacts were more likely to come from within the user's own organization (25% vs 19%).

Affiliation was a fair predictor of importance as the regression coefficient shows. However we were surprised at the number of external contacts selected as important. This may reflect a trend observed elsewhere – that modern work takes places across organizational boundaries [7].

Communication style was confirmed. People who are high frequency communicators judge a greater proportion of extracted contacts to be important (23% vs 9%),

We analyzed the effect of *communication style* on the contacts users selected, using a one-way ANOVA with communication style (frequent/infrequent) as independent variable, and the likelihood of rating a contact as important as the dependent variable. Recall that frequent communicators were those who exchanged more than the mean number of messages/day for the entire user sample. This difference shown in Table 3, is significant ($F_{(1,14518)} = 548.9, p < 0.0001$).

	Unimportant Contacts (%)	Important Contacts (%)
High Frequency Users	77.2	22.8
Low Frequency Users	91.0	9.0

Table 3: The Effect of Communication Style on Judgments of Contact Importance

Both user comments and our experimental findings suggested that there were frequent changes in the set of people who users considered important. Our experimental design does not allow us to investigate such changes directly, as we only asked users to make importance judgments at one point in time. Nevertheless, we can

partially investigate change in contact selection by examining the effects of longevity and recency.

Longevity of Communication (Years)	Proportion of All Important Contacts Chosen	Recency (Years)	Proportion of All Important Contacts Chosen
Less than 1	0.66	Less than 1	0.65
1-2 years	0.24	1-2 years	0.26
2-3 years	0.08	2-3 years	0.07
3-4 years	0.01	3-4 years	0.02
4-5 years	0.00	4-5 years	0.00

Table 4: The Effect of Longevity and Recency on Judgments of Contact Importance

Table 4 shows that important contacts tend to have recent, short duration relationships with users. Most important contacts have been in touch for less than two years: 66% of important contacts have relationships lasting a less than a year, and a further 24% between one and two years. The picture is similar for recency -- 65% of important contacts have been in touch in the last year, and 26% in the last one to two years. This prevalence of recent, short duration relationships with important contacts suggests that users' important contacts change frequently.

COMPARISON OF FIELD AND LAB STUDIES

Our field data suggested a significant, but currently under-researched problem, that of contact management. People are exposed to large numbers of potential contacts, but the onerous nature of data entry means that they end up being conservative about who they add to their contact management systems. Despite this, people have a large number of contacts that they have to manage, but end up using a variety of ad hoc tools for this purpose.

Our experimental results confirm the interview data in two important respects. First, consistent with the interview data, people are exposed to a large number of contacts (859 on average), only 19% of whom they judge as important. This supports the idea that people are exposed to many more contacts than they want to keep in touch with. This in turn suggests that *contact selection* is an important process.

Second, the experiment confirmed the *criteria* that our interviewees suggested for identifying important contacts. We found evidence that a contact's *communication history*, and *communication style* were important determinants of whether a contact was selected. Frequency, reciprocity, longevity, and recency predicted subjective importance, as did contact affiliation and the style of the user's communication.

However there were a number of ways in which the experimental results diverge from our interview data.

First, the number of contacts extracted, while large overall, is small when one considers the age of people's archives (on average 3.0 years). People have an average

of 859 distinct contacts in their archive. This means that people are exposed to, and have to make a decision about a mean of 859/1095 contacts/day, i.e. fewer than one novel contact each day. Ignoring deleted messages, which were obviously not in their archives, users may therefore have to make fewer decisions about contact importance than our interviews led us to believe, making contact selection a more straightforward process.

Another way in which the contact selection process may be easier than anticipated, is the large proportion of extracted contacts (67%) who are either unsolicited communicators, or people who only appear in cc or bcc lists. While all of these contacts cannot be immediately rejected, it may be easy for users to downgrade the importance of many of these contacts, as they never engage in direct, reciprocal conversations.

Secondly, users selected a relatively small number of contacts as important (119). This contrasts with the large numbers of contacts and complex systems described by our field study informants. One possible reason for the discrepancy is that our experiment did not allow participants full opportunity to identify all their important contacts. However this seems unlikely, as there were no time limits on the experimental procedure and users were 82% satisfied that they had identified all relevant contacts. They were also allowed an optional session to revisit their judgments.

A more likely reason for the discrepancy between studies may be that the long-term systems that we explored in the interviews may contain 'stale' or out of date contacts. In our experiment, we had users set up a new system from scratch, whereas the field study looked at systems acquired over many years. One significant finding from the lab study is that contacts change frequently - with the majority of important contacts being encountered in the last two years. If users do not 'clean up' their contact archives over the years, this may mean that older address books or PDAs are filled with no longer relevant contacts. Other research has documented the infrequency with which users 'clean up' various other types of email and voicemail archives to remove outdated information [13,14,15]. Contact management systems may be no different in this regard.

A different possibility for the discrepancy between studies is that the contact management systems we observed in our field study contain many contacts who never emailed our informants. It may be that users have important contacts who they often interact with by phone or face-to-face. Logging these phone and face-to-face interactions would have been an extremely complex undertaking, however, although phone logs are now becoming available [5]. In future research we intend to investigate contact selection in these other media.

A final limitation of the lab study is that our archives excluded deleted messages, although collecting large complete archives would have taken many years. The absence of deleted messages does introduce some potential confounds, however. There are large differences between users in the likelihood and frequency that they 'clean up' other archives [13,14,15]. This may have affected the nature of individuals' archives and their subsequent contact selections.

DESIGN AND THEORY IMPLICATIONS

Several design suggestions follow from these results. First, our regression analysis is a model for identifying important contacts in email, and this could be implemented directly as an algorithm. The ability to automatically identify important contacts from communication archives might be used in a number of applications, allowing us to improve messaging applications, support reminding and provide social recommendation. Messaging applications are currently poorly integrated with contact management tools, but future systems could exploit information about important contacts in a variety of ways. These might include alerting, filtering and prioritization of incoming email or voicemail messages based on the sender's importance. Tighter integration of contact information with messaging logs could be also used to manage relationships with contacts, e.g. reminding the user when they haven't talked to an important contact in a long time. We have implemented contact-based alerting and reminding in a social network based user interface to communication and information [7]. Finally social recommendation systems might be able to exploit information about a register of important contacts to either direct a user query or guide information access.

Other design implications concern contact management tools directly. We could improve address book utility by using our algorithm to automatically recommend that a potentially important contact should be added to the address book, based on their communication history. But even if we provide ways to better identify significant contacts, data entry is still a major problem. One possible way to address this would be to identify contact information from other sources, such as Internet home pages containing addresses. We may also be able to mine other types of records such as phone and voicemail logs, or use reverse lookup to provide detailed addresses for contacts. Having general techniques for populating address books is clearly important. One unexpected finding from our research was that 72% of important contacts came from outside the user's organization. While this may depend on the specific user population, it suggests that corporate address books or intranets have limited utility as a way to provide detailed addresses for contacts.

We also identified problems with contact management tools containing 'stale contacts', suggesting that current systems would benefit from better tools to 'clean up' contact archives. We could use our algorithm to identify 'stale' contacts, possibly relegating them to a secondary interface view. By reducing the set of immediately visible contacts, such categorization might allow users to identify important contacts more easily while browsing contact records.

Other theoretical questions concern the process of contact selection itself. There are several ways to refine our hypotheses about what determines contact importance, and hence improve the algorithm. More sensitive measures of reciprocity (e.g., using messages replies or "re:" tags) might increase its predictive capability, as would improved definitions of unsolicited communication. We also plan to investigate other factors that affect importance, such as user job type. Research on social networks shows the effect of job type on both communication and social networks and this merits further study [4]. We might also examine the content of messages to see how this could be used for determining importance. Message content could also be used to provide addressing information from signature files. Another major theoretical issue is the role of change. We need to better understand how one's set of contacts is modified over time and also what causes these changes. Are they the result of corporate reorganizations, changes in work, or changes in social interest?

There are also links between this work and social network research [3,4,12], which examines people's communication patterns and social relationships. However we diverge from that research program by focusing on user-centric representations of contact's importance - using such information to provide tools that help people to manage their communications with others. Our work has greater overlaps with social data mining research [1,2,9,11] that attempts to represent and analyze textual conversations, so that these can be used as an informational resource.

Finally our results reinforce the need to refine our theories of asynchronous communication, to include important background outeraction [6] processes that make long-term communication possible. Elsewhere, we have outlined the critical role that contact management processes play in making asynchronous communication possible [14]. We hope that such work can contribute to new CMC theories that focus on critical aspects of asynchronous communication, such as long-term communication and persistence, rather than relying on models that are derivative of synchronous face-to-face interaction.

ACKNOWLEDGMENTS

We thank everyone who participated in the experiment and the interviews for their help, and especially to Bonnie Nardi for many discussions of these ideas.

REFERENCES

1. Donath, J., Karahalios, K., & Viegas, F. (1999). Visualizing conversations. In *HICSS-32*, IEEE, 583-592.
2. Erickson, T. (1999). Persistent Conversation: Discourse as Document. In *HICSS-32*, IEEE Press, 5-14.
3. Granovetter, M. (1973). The Strength of Weak Ties. *American Journal of Sociology* 78, 6, 1360-1380.
4. Haythornthwaite, C., Wellman, B., & Mantel, M. (1995). Work Relationships and Media Use: A Social Network Analysis. *Group Decision and Negotiation*, 4, 193-211.
5. Millen, D. and Henderson, D. (2000). PhoneMan – The Benefits of Personal Call Histories, In *CHI2000*, New York, ACM, 155-156.
6. Nardi, B., Whittaker, S., & Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. In *CSCW2000*. New York, ACM, 79-88.
7. Nardi, B., Whittaker, S., Isaacs, E., Creech, M., Johnson, J., & Hainsworth, J. (2002). Integrating Communication and Information Through ContactMap. *Communications of the ACM*, 45, 89-95.
8. Rice, R. E. (1988). Issues and concepts in research on computer-mediated communication systems. *Communication Yearbook*, 12, 436-476.
9. Smith, M. and Fiore, A. (2000). Visualization components for persistent conversations. In *CHI2001*. New York, ACM, 136-143.
10. Terveen, L. and Hill, W. (2001). Human Computer Collaboration in Recommender Systems. In Carroll, J. (Ed.), *HCI in the New Millennium*, Addison Wesley, 223-242.
11. Viegas, F.B. and Donath, J.S. (1999). Chat Circles, In *CHI99*. New York, ACM, 9-16.
12. Wellman, B. (2001). Physical Place and Cyber Place: The Rise of Networked Individualism. *International Journal of Urban and Regional Research* 25, 227-252.
13. Whittaker, S., Davis, R., Hirschberg, J., & Muller, U. (2000). Jotmail: a voicemail interface that enables you to see what was said. In *CHI2000*. New York, ACM, 89-96.
14. Whittaker, S., Jones, Q., & Terveen, L. (2002). Persistence and Conversation Stream Management. In *HICSS-35*. IEEE Press, 335-344.
15. Whittaker, S. and Sidner, C. (1996). Email overload: exploring personal information management of email. In *CHI'96*. New York, ACM, 276-283.
16. Whittaker, S. Terveen, L., Hill, W., & Cherny, L. (1998). The dynamics of mass interaction, In *CSCW98*. New York: ACM, 257-264.

