How do blog gardens grow? Language community correlates with network diffusion and adoption of blogging systems

Lei Gu[+] Paul Johns* Thomas M. Lento[#] Marc A. Smith*

[+]University of Washington, Department of Sociology
*Microsoft Research, Community Technologies Group
[#]Cornell University, Department of Sociology
[+]*One Microsoft Way, Redmond, WA 98052
[#]323 Uris Hall, Cornell University, Ithaca, NY 14853
gulei@u.washington.edu, paul.johns@microsoft.com, masmith@microsoft.com, thomas.lento@cornell.edu

Abstract

Different language communities present in a multimedia blog and social network hosting system are contrasted. The disproportionate population of Chinese language speaking users is explored. Four stages of population growth and activity are examined: invitation, adoption, retention, and contribution. Using blog system logs we find that the Chinese language user population grew from just two invitations and demonstrates higher levels of invitation and retention which was highly correlated with contribution. Additional survey results confirm the difference between language communities on their attitudes and approaches to blogging systems. This suggests visible differences in the ways different cultures and language communities embrace and enact relationships through computer mediated publication.

Introduction

Do different language communities build different kinds of computer-mediated social relationships? The Wallop blogging system created a platform for exploring the ways groups of people can create and share content via the Internet, providing a light-weight way for users to express their social connections and foster the maintenance and development of social relationships (Farnham et al 2004). In the process it generated a rich data set and several interesting social phenomena that bear investigation. A central question is how the Wallop blogging system came to be populated by a majority of Chinese language speakers, despite the fact the interface was exclusively in English. The population distribution is all the more intriguing given the snowball social network mechanism which restricted access to the system to those invited by existing users. The result is an opportunity to explore a contrast: how did the Chinese language population compare to the rest of the user population in terms of adoption, contribution and invitation of other users? How did the diffusion of invitation and acceptance of Wallop lead to this population distribution? How do other patterns of activity compare among the different language populations?

The rapid and widespread adoption of blogging software has made it increasingly easy to post information and link to others via the Internet (Blood 2004). The result is a densely connected web of blogs built from individual messages, links to other web sites and blogs, and patterns of readership. Treated as data, collections of blogs can offer valuable opportunities to study social processes and patterns of difference across large populations. Some sites host multiple blogs and provide additional features like limiting access to content to select others, readership tracking, and membership through network propagation rather than open registration. These additional features have the secondary effect of generating yet more data about patterns of blog use otherwise invisible to a study of more decentralized blog content.

Blog systems have been adopted across many nationalities and cultures. A large number of world languages are present in the form of blogs and other web content despite the current (if waning) dominance of the English language on the Internet (Global Reach 2005). Recent work suggests that use of the Net in developing countries and the number of non-English websites are growing rapidly (Warschauer, 2000), and the proportion of English language use in computer-mediated communication was expected to fall from 80 percent to 40 percent within ten years, starting from 1996 (Graddol, 1997). As of June 30, 2005, 103 million users in China surf the Internet for at least one hour per week (CNNIC, 2005). 10.5% of Chinese users reported that they often blogged, and 9.6% hoped that blogs would be added as a complementary feature to text-based instant messenger services. Wallop data adds rich information about the mechanism of changes in the online activities of language communities, especially the Chinese group.

Given the observed variations in the ways age and intention are related to the pattern and content of blogs (teenagers blog differently than the CEOs of technology companies – see Nardi et al 2004 for examples), different cultures may adopt and enact relationships through blogging software differently. Language may offer a proxy for culture to a certain extent, and the mechanical recognition of differences in language can be exploited to automate the

analysis of large data sets. By hosting a blogging system with a closed invitation model, a snowball sample of potential Wallop users was recruited to the system. Users' resulting behavior, contributions and interactions with the system and one another created logs that can be analyzed to explore contrasts among different language populations.

Background

Although related work has explored the link structure and patterns of content production within collections of blogs (Marlow 2004; Kumar et al. 2003), the similarities and differences in blogging practices across cultures and communities are less well understood. Variations of user behavior that correlate with attributes like language use can offer some insights into the ways technology and culture interrelate. Why do some networks of relationships grow dense ties and thrive while others generate sparse ties and fade away? What are the social processes driving adoption and development of blogging systems? Do cultural and linguistic differences between groups of potential users affect the growth and spread of these blog communities?

Network studies of linking behavior provide an excellent means of examining structural patterns within blogging groups, and can be used to relate these patterns to a group's potential for growth and productivity. However, in order to draw any meaningful conclusions about why group dynamics differ, such studies are dependent on additional information (Kumar et al. 2003). Content analysis, survey, and ethnographic data can provide some of the missing context for these network dynamics, but such data is difficult to collect automatically and is often unavailable to researchers analyzing publicly uploaded blog content. When it is available, data about the referral, adoption, and contribution patterns of the users provides useful information about the mechanism for diffusion of a blog group. This data can provide a foundation for research assessing the importance of both structure and content in blog community dynamics.

Diffusion and Adoption in Blogspace

Ethnographic studies of bloggers provide detailed information about how people find blog communities and why people participate in them (Nardi et al 2004; Schiano et al 2004). Other work has contrasted the adoption of the Live Journal blogging system between the US and Russian user base. Russian users of Live Journal are typically in their late 20s, are businessmen or members of Russian literary society, and use Live Journal as an outlet to discuss a variety of topics ranging from art to politics. US Live Journal users, by contrast, are predominantly teenagers who use Live Journal as an online diary (Gorny 2004). Although some of these differences are related to the demographics of Internet users in the US and Russia, this study suggests a cultural difference in adoption and use of Live Journal.

Survey-based cross-cultural studies of online behavior also indicate differences in Internet usage. For example, Chau et al (2002) show that Asian consumers favor applications that facilitate social communication, while US consumers welcome improvements in information search, but like many cross-cultural studies they are focused on e-commerce rather than informal association. In one of the few studies of blog usage, Su et al. found that Internet users in China, Japan, and Taiwan are less socially connected to others through their blogs than those in Western Europe, Southeast Asia and Australia (Su et al., 2005). Although they did not find many other significant differences across cultures, they acknowledge existing research suggesting that cultural differences affect adoption and usage of communications technologies (Bell 2001; Bell 2003). Furthermore, the measure of culture used in Su et al. is nationality, which does not account for users outside of East Asia who come from similar cultural roots, and are closely connected with their nation and culture of origin.

A Diffusion Model for Adoption. We propose an approach to examining membership diffusion in different blog communities. We first consider the process individuals use when faced with the decision to adopt a blog community. This is typically modeled as a five-step innovation-decision process of: 1) knowledge, where an individual becomes aware of an innovation; 2) persuasion, where she develops a positive or negative attitude towards an innovation; 3) decision, where an individual takes some action, which in turn leads her to choose whether or not to adopt an innovation; 4) implementation, where she begins using an innovation; and 5) confirmation, where an individual evaluates the results of her decision (Rogers 2003).

This classic diffusion model provides a useful framework for comparing the adoption processes of two groups. However, in the case of online social spaces, it can be difficult to differentiate between the various steps in the innovation-decision process. Once they become aware of an online group, many users form their opinion of it and decide whether or not to use it by visiting the group and even joining or logging in. Many of these users in turn decide not to implement the innovation, and leave after just a few visits. The participation distribution for many online social systems is a power law function with a steep fall-off in which two-thirds of users contribute as little as necessary to qualify for inclusion in the data set (Smith 2004). The persuasion, decision, and implementation phases therefore appear to be encompassed by the same set of behaviors – e.g. logging in to the system – and are nearly impossible to assess without detailed ethnographic data.

In order to address this issue, we propose a four-step model based on the classic innovation-decision process. This model captures key points in the adoption process for an online social group that can be assessed through machine collectable data. The four steps – invitation, adoption, retention, and contribution – illustrate the differential population behavior that may start to define the range of variation of blog related behavior.

People discover the Wallop blogging system through an **invitation** from an existing user. This is similar to receiv-

ing an email referral, which is one of the most common ways users are attracted to other blog sites and distributed systems (Nardi et al. 2004). Users then decide whether or not to accept a referral. After **adopting** a group by logging into it at least once, an individual may drop out fairly quickly, or she may remain involved and be **retained** in the system. At any point, she may refer others to the site, and their likelihood of adopting and being retained as a user of the system is correlated with how much the referrer **contributes** and how well-connected she is to the group.

About Wallop

Wallop was designed as an invitation-based blogging and online interaction application that supports the selective publication of text, audio, and images and the expression of links between people (http://mywallop.com). Where many existing blog hosting sites focus on individual publishing, the Wallop system provides additional explicit support for sharing content among a selective population. Wallop differs from many other sharing tools, such as picture and music sharing sites, by placing additional emphasis on social interaction and relationship networks. Wallop encouraged connecting personal publishing to the formation, maintenance, and extension of social networks via the invitation system and the display of maps of users' Wallop social network.

Users first come into contact with Wallop via an invitation sent by an existing Wallop user. After registration and login, they are guided to the main interface (Figure 1). Users may provide their profile or contact information in the right pane of the interface. At any time after first login, users can invite other people to join Wallop by entering their email addresses into the system. Users were initially given five invitations, but could request more if they wished. Inactive users' accounts were deactivated after 60 days. Users were able to re-invite inactive members by selecting their names through the same invitation process.

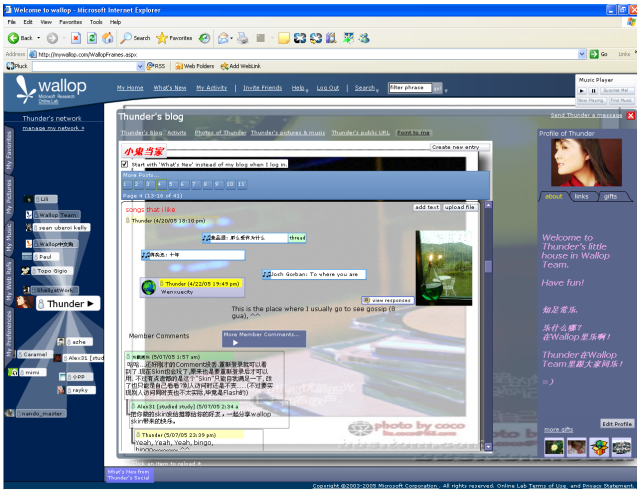

Figure 1. Sample Wallop user home screen

The left column presents a user's explicit network based on person-to-host interaction and the strength of their relationship, not exceeding a maximum of 15 people. "Manage my network" on the top links users to an expanded network beyond the top 15 (Figure 2). To show the invitation relationship, the inviter is listed in the invitee's explicit network upon the latter's first login, and the Wallop agent automatically assigns some of the inviter's active network friends to the invitee to help her build her initial network.

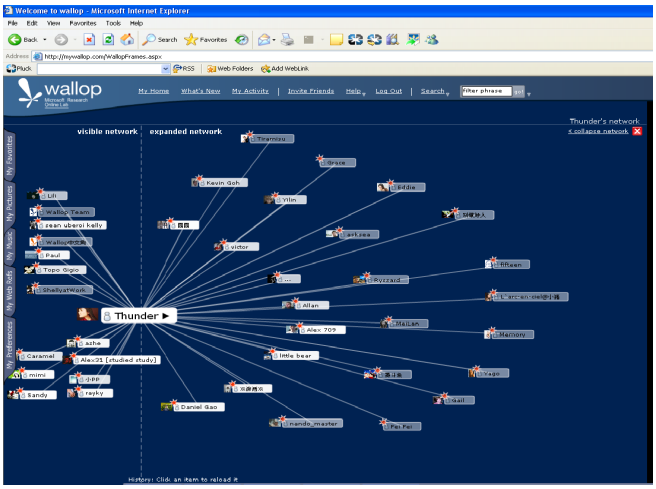

Figure 2. The Wallop extended social network interface

Users can change the members of their social network at any time by dragging name tags into and between the two types of network maps. The background color of the name tag indicates the degree of activity; the whiter the more active the contributor. By clicking on each person's name, people can easily navigate to this person's blog which is then displayed in the area in the middle main column, and blog entries are displayed in reverse chronological order.

Wallop Deployment. Wallop was developed with an exclusively English language user interface, and after a testing phase started as a small-scale experimental public deployment in October 2004. Six months later, 65% of the users who logged in at least once came from the time zone from which many Chinese network users originate. With so many Chinese users actively uploading content to Wallop, it seemed the "official" language of Wallop might become Chinese. Despite its roots in an English-language user population the users' most frequent request was for a Chinese interface.

Given that the system was designed for English language users, and that the initial users were predominantly English speakers, why did the Chinese language user population grow so quickly? At what point in the diffusion process did the Chinese group overtake the English users? What are the implications of this for other blog communities?

Testing the Model

The initial population of Wallop users were all English speakers, the user interface was entirely in English, and no special effort was made to attract foreign users. Therefore, the rapid growth and spread of the Chinese language user group must have come from some difference in the behaviors of the Chinese and non-Chinese language groups. These differences in behavior must necessarily manifest themselves in Wallop's diffusion process. But which steps in the process are critical to the rapid growth of the Chinese language user group? We examine each stage of four-step model described above by addressing the following:

Invitation. Do Chinese language users send more invitations than non-Chinese language users?

Adoption. Do Chinese language users accept invitations at a higher rate than non-Chinese language users?

Retention. Do Chinese language users remain active in the system for a longer period of time?

Contribution. Do Chinese language users produce more content in the system?

Addressing the Questions. Most available data is not well suited for a full analysis of the adoption mechanism that drives the growth and spread of blog communities. Data that can be readily downloaded from distributed blogs provides excellent opportunities to study linking behavior, but lacks information about referrals. Survey and ethnographic data are difficult to collect, and even in combination with network studies it is very difficult to compare the adoption patterns of different groups of users.

Wallop data solves many of these problems by including information about referral behavior, acceptance of referrals, and longevity within the system, which we use as a proxy for diffusion. Furthermore, the system records linkages between users, making it possible to explore the relationship between social cohesion and group activity.

In the following section we present data and an examination of this four step diffusion model. We focus on a comparison of the rapid growth and spread of the population of Chinese language users with the slower growth of the non-Chinese population. We identify **invitation** and **retention** as the crucial steps of the diffusion process in growing the Wallop blog community.

Data and Analysis

The Wallop blog system hosted 87,339 users over a one-year period from late 2004 through late 2005. Wallop users invited one another to join the system; no access was available to those who simply "walked-up" to the web site. Starting with the research team that developed the system, invitations propagated through each user's social network as users sent invitations to other potential users. The user base grew from a group of 7 initial inviters,through as many as 40 generations. Not every user invited others. Not every invitation was accepted. Not every acceptance resulted in an active user who contributed content and comments on the content of others. Not every adopter went on to invite others to join the system, which could start the process of diffusion over again.

Wallop User Population Breakdown	**Total Users**	**% of Total Population**
Chinese language users	33,204	38%
Non-Chinese Language users	17,060	20%
No comments/Not categorized	37,075	42%
Total Population	87,339	

Table 1. Wallop users classified by language use

Our model of the Wallop system's population dynamics captures and contrasts the relative rates of invitation, adoption, retention, and contribution among different populations of users based on patterns of language use. Of the users of the Wallop system who logged in at least once between October 2004 and September 2005, 33,204 posted comments containing Chinese language characters (Table 1). All user's comments were scanned to match against the range of character sets defined by the CJK range (Chinese, Japanese, and Korean). Users whose comments contained five percent or more characters in that range of the character set were assigned to the "Chinese language user population". Alternatively, if three or more comments contained a single Chinese language character the user was also assigned to the "Chinese language user population". All other users were assigned to the "Non-Chinese language user population," which was further sub-divided by those users who generated no comments at all and thus were impossible to categorize.[1] Random samples from each groups were inspected manually to estimate the effectiveness of the classification system. Less than 1% of users classified as Chinese language users were not using Chinese, including those using Japanese and Korean. Less than 5% of non-Chinese language users were using Chinese.

To compare invitation rates, we compared the frequency distributions of the number of invitations sent by each group, normalized to account for variance in group sizes. We also examined data on the number of invitation generations grouping each group. This measure provides some information about the relative fertility of the two groups. The more generations within the group, the more fertile, as invitees not only adopt but recruit others to the system.

We compared the average adoption rates, defined as the number of invitations accepted divided by the number of invitations sent, for each group, clustered by the number of invitations sent. The number of unique login days was used as a measure of retention. This is a better measure than total logins, because it includes information about how many days the user spent active in the system. It is also better than elapsed time between first and last login for

[1] These users most likely represent a mix of Chinese and Non-Chinese language users, and were relatively inactive. They were therefore omitted from the analyses.

each user, because it accounts for highly active users who have been in the system for a relatively short time. Activity was measured by the number of comments each user uploaded into the system, and the number of total files (pictures, audio, etc.) each user uploaded into the system. We compared the frequency distributions of both measures across the different groups.

Finally, in order to better visualize the invitation and activity behavior of the groups, we generated a treemap showing the invitation hierarchy and the number of comments posted into the system. The treemap allows us to see how invitation behavior is distributed across the two language groups, and gives us a sense of how many comments the users in each group are posting on average.

In addition to the data taken from the Wallop system logs, we examined an earlier survey of Wallop users. This data shows the proportion of each group who have more friends in their networks whom they have never met face to face. It corroborates the results of our activity log analysis.

Results

Analysis of the data suggests that Chinese language users are more likely to send invitations, have invitations accepted at a higher rate, stay in the system longer, and contribute more content to the system (Table 2).

Wallop User Behavior by Language Group	**Chinese Language**	**Non-Chinese Language**
Sent Invitations	51.76%	37.63%
Overall Adoption Rate	65.98%	60.80%
Average # Unique Login Days	16.60	6.23
Average # of Files Uploaded	13.02	6.20

Table 2: Summary of differences in user behavior

Throughout all of the contrasts between the Chinese language user population and the other populations we find that the biggest difference is in the percentage of moderate users – the groups are otherwise similar at their extremes, although the Chinese language user group does have a larger core of heavily active users.

Figure 3 shows the frequency distribution of invitations sent by percentage of Chinese language users and non-Chinese language users. The Chinese language user base has a higher proportion of moderately active users, and a larger heavily active base of users. The non-Chinese language user group has a substantially higher proportion of users who did not send invitations. The higher numbers of users who send 5, 10, 15, and 20 invitations is an artifact of the limits imposed by the invitation system. After every 5 invitations sent, the user had to request more before she could invite more users. This limit had a profound effect on invitation behavior, as only those users motivated enough to take the extra step required to get more invitations could move to the next invitation level.

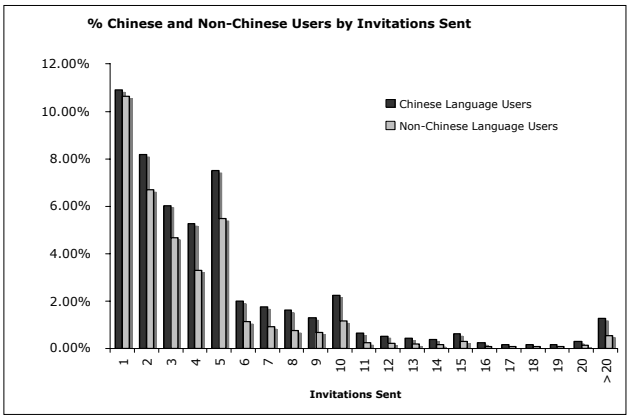

Figure 3. Normalized frequency of invitations sent as a percentage of two user populations: Chinese Language Users and non-Chinese language users[2]

Figure 4 shows the distribution of each user base across generations of Wallop invitations. Two users issued invitations initially to just 12 people who collectively had 33,204 descendant "child" Chinese Language users. In contrast, Non-Chinese language users issued a first generation of invitations to 209 people who ultimately generated a total of 17,060 descendant users.

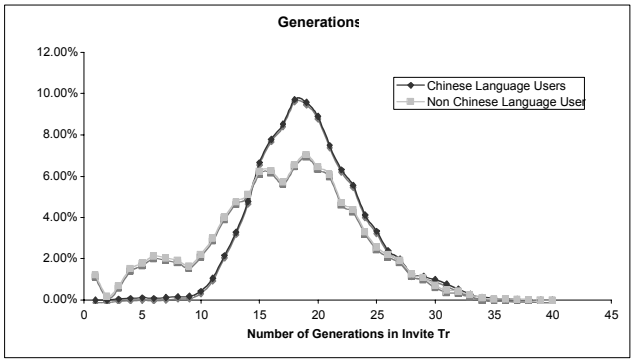

Figure 4. Distribution of maximum generations by language population

The average number of generations in each branch of the invitation tree varied by language population. Chinese language users had on average 19.39 generations from their initial invitations, in contrast to the 17.01 generations on average for Non-Chinese language users.

[2] In Figures 3, 5, 6, and 7 the bin sizes were determined by taking 20% of the standard deviation from the mean, calculated for the combined activity of the two groups. The minimum value for the final bin is 3 standard deviations above the mean.

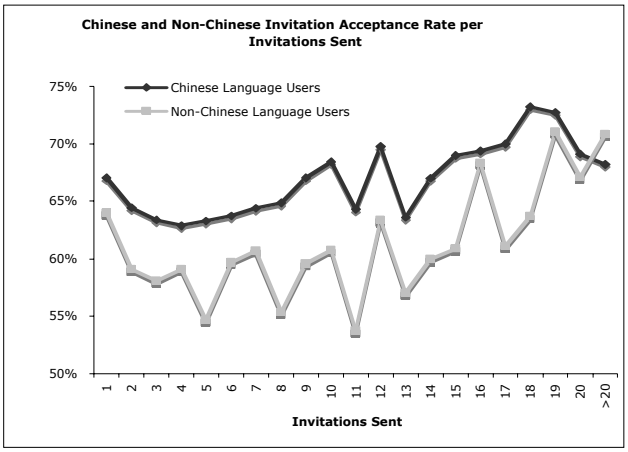

Figure 5. Percentage of invitations accepted per number of invitations offered for Chinese and non-Chinese language populations

Figure 5 shows the success rate of invitations sent. The patterns of the curves are puzzling, and although some of the peaks and troughs might be a result of the bias towards users who send five invitations, or an artifact of small sample sizes on the right side of the graph, it is not clear why the shapes of the curves are so different. Overall, the rates of adoption for Chinese language users are consistently higher than those for non-Chinese language users, but the overall difference was not as stark as might be expected. This suggests that the difference in the rate of diffusion of Wallop between the two groups was driven by higher rates of invitation and retention.

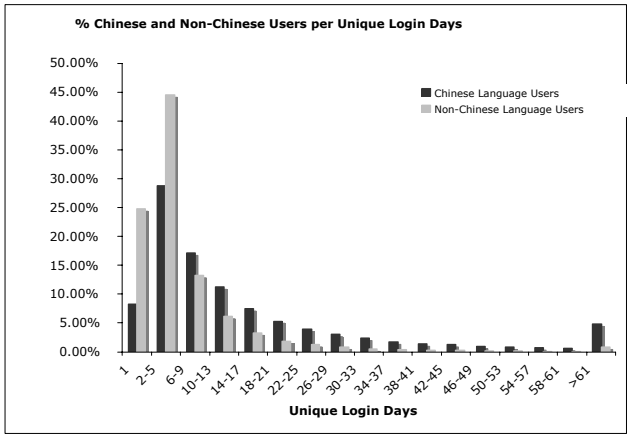

Figure 6. Normalized retention of Wallop users for Chinese and non-Chinese language populations

Figure 6 illustrates a striking difference in retention within the system between Chinese language users and non-Chinese language users. Non-Chinese language users tend to remain active in the system for just a few days, while Chinese language users are much more likely to remain for several weeks. Commenting behavior appears to be a strong predictor of retention, as users who make comments are far more likely to remain in the system for over a month than users who do not make comment.

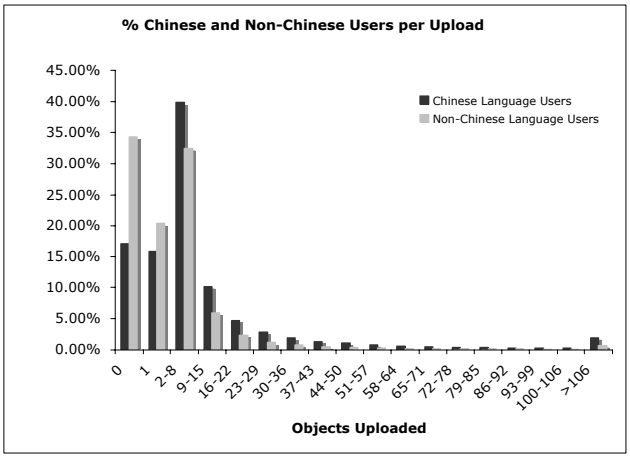

Figure 7. Normalized frequency of non-comment contributions (photos and audio files) for Chinese and non-Chinese language populations

As with invitation behavior, the Chinese language user population has a larger proportion of moderate contributors and a smaller percentage of marginal contributors than the non-Chinese language user population. Frequency distributions of comments (not shown) follow a similar pattern.

The driving factors behind the growth of the Chinese population are retention and invitation, which present the starkest contrasts of any of the data.

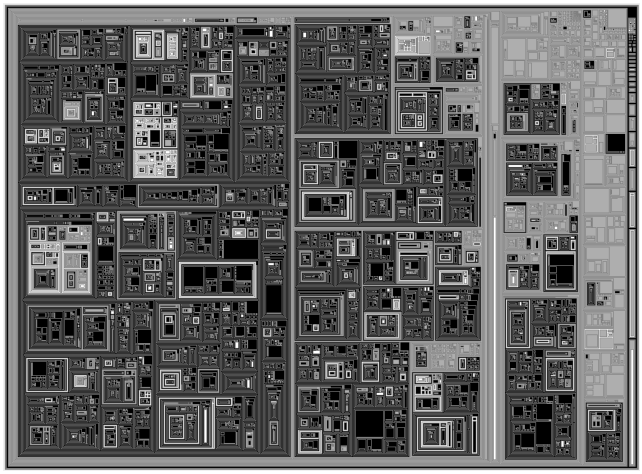

Figure 8. Treemap of Wallop invitation tree, Black: Chinese Language Users, Gray: Non-Chinese Language Users, Area: Number of Total Comments per User

The treemap visualization of the Wallop invitation tree provides some interesting insights into the growth and composition of the Wallop population. The Wallop system recorded the number of invitations each user sent and the

rate of success in terms of accepted invitations and later conversations to regular contributors was logged. This allows for the review of the process of social diffusion of access to the Wallop system.

The obvious preponderance of dark colored boxes in Figure 8 indicates the overwhelming dominance of Chinese language users in the system. Further, the nesting of boxes highlights the ways that the Chinese population had more generations of successive invitations and acceptances than the non-Chinese language population.

Wallop User Survey

To expand on the data from Wallop log files, we explored data from interviews, and survey research to provide a triangulated view around blogging practices. In March 2005, we conducted a survey on Wallop experience among the users who had logged-in at least once and had been in the system for at least two weeks by the survey time. Our survey segmented the Wallop population in terms of Chinese Language Users, non-Chinese but North American Wallop users, and finally non-Chinese and non-North American users. This model differed slightly from our categorization of log data which focused on use of language rather than self-description. Among the 220 responses, 135 were native Chinese speakers, 47 were English language users in North America and 38 come from the rest of the world.

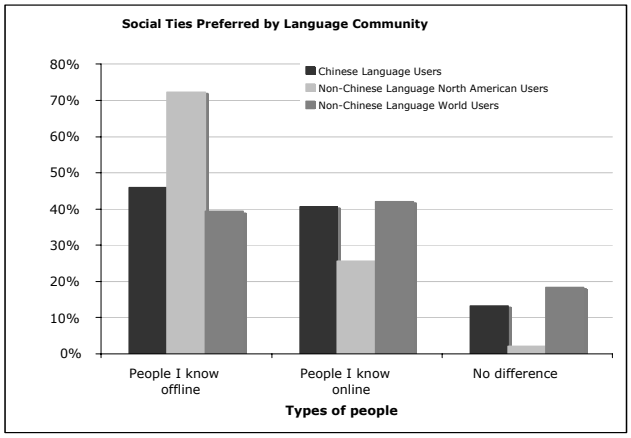

Figure 9. Types of social ties preferred within Wallop, grouped by language community

In our survey we found that the motivations for blogging were significantly different between language communities. In Figure 9. the three language groupings are plotted in terms of their main focus for blogging: staying in touch with already existing contacts, making new contacts and balancing new and old contacts. We find that Chinese language users are significantly more likely to agree that they used Wallop to build new social connections with people they did not already know face-to-face. In contrast, North American users reported significantly higher levels of Wallop use in support of maintaining existing social relationships. This is corroborated by the log activity analysis. Chinese language users were actively seeking ways to expand their online social interactions, while North American users tend to form a closed online social network.

Discussion

Blogging systems host a range of uses and are valuable sites for the study of behavior and its variation across various kinds of population sub-divisions. Because language communities can act as proxies for national and social membership we can leverage properties of the content contributed to a blogging system (which character sets are used, in this case) to provide a mechanically recognizable division in the data. This allows us to automate observations about the variation in behavior between major populations in the user community.

In Wallop the Chinese language population demonstrated significant differences from the other language communities present in each of the four major points of social process that was documented by the data logs. Chinese language users invited more other users to join the system and those users accepted those invitations at a higher rate than non-Chinese language users. Chinese language users further converted to active users and even to active contributors at a much higher rate as well.

These patterns of difference show that the Chinese language users and non-Chinese language users are accessing and approaching the Wallop system in clearly distinct ways. The survey data indicates that Chinese language users were seeking to expand their social networks, while non-Chinese users were looking to maintain their networks. Although the data presented here does not prove that these differences in behavior are directly related to cultural differences between the groups, the results of this analysis pinpoint the critical sections of the diffusion cycle. The rapid growth of the Chinese language user population is primarily a result of their tendency to remain active in the system longer than the non-Chinese users. Building on this insight, we can further our understanding of how these different populations use the Wallop system, and how important cultural factors are in explaining the differences between these groups.

Future Work

A major direction for this work is to more directly explore the reasons behind the different patterns of behavior observed in the two groups. Do Chinese users have a more active social network within Wallop? Is this why they remain active in the system longer? How can we explain the more rapid growth of the Chinese language user population? Is it a result of random chance, where the initial Chinese language user population happened to include a cluster of highly motivated users? Or is the successful

spread of the Chinese language user population the result of more effectively leveraging pre-existing network ties?

Social network analysis of the emergent structure of relationships within Wallop can provide some answers to these questions. When they leave comments on content uploaded by another user, Wallop users create social ties, which are catalogued in the system. Using this data, it is possible to derive social network data on the Wallop users. By comparing patterns of clustering and other network statistics, it is possible to establish the relationship between network density, centrality, and the activity of network alters on the likelihood of an individual remaining active in the system. Comparing the Chinese language user network with the non-Chinese language user network might highlight some of the factors that lead to the Chinese language users' more effective recruiting and greater longevity. Comparing both of these groups to randomly generated graphs will demonstrate the likelihood that this phenomenon is the result of random chance. Finally, combining data from this network of interactions with the invitation hierarchy data makes it possible to infer the extent to which users have a pre-existing social tie. This makes it possible to determine the whether or not pre-existing social networks affected the growth and spread of the two user populations.

Although the results of this study highlight some of the ways culture and language intersect with technology, there is much more research to be done on the relationship between culture and usage behaviors. This study, and further analysis of the Chinese and non-Chinese language user networks, can help direct more in-depth analysis. It is difficult to effectively examine cultural differences using computational approaches, but using these methods we can uncover interesting phenomena and the behaviors that give rise to them. This might provide a clue for other researchers, particularly those who might not be able to make progress in a system as large and varied as Wallop.

Acknowledgements

We would like to thank the Wallop development team and user populations (Nihao!), members of the Social Computing and Community Technologies Group at Microsoft Research, Lili Cheng, Shelly Farnham, Tony Capone, Danyel Fisher, and the folks who now do the Wallop operations.

References

1. Blood, R. 2004. How Blogging Software Reshapes the Online Community. *Communications of the ACM* 47:53-55.
2. Bell, G. 2001. Looking Across the Atlantic: Using Ethnographic Methods to Make Sense of Europe. *Intel Technical Journal* 5:1-10.
3. Bell, G. 2003. Other Homes: Alternate visions of culturally situated technologies for the home. *CHI 2003* April 5-10, Ft. Lauderdale, FL.
4. Chau, P.; Cole, M.; Massey, A.; Montoya-Weiss, M.; O'Keefe, R. 2002. Cultural Difference in the Online Behavior of Consumers. *Communications of the ACM* 45:138-143.
5. China Internet Network Information Center (CNNIC). 2005. 2005 China Internet Development Report.
6. Farnham, S.; Kelly, S.U.; Portnoy, W.; Schwartz, J.L.K..2004. Wallop: Designing Social Software for Co-located Social Networks. In *Proceedings on the 37th Hawaii International Conference on System Science*.
7. Global Reach, Internet Statistics by Language http://global-reach.biz/globstats/index.php3
8. Gorny, Eugene. 2004. Russian LiveJournal: National specifics in the development of a virtual community. Version 1.0 of 13 May 2004. http://www.ruhr-uni-bochum.de/russ-cyb/library/texts/en/gorny_rlj.htm
9. Graddol, D. 1997. The Future of English. London: The British Council.
10. Kumar, R.; Novak, J.; Raghavan, P.; Tomkins, A. 2003. On the Bursty Evolution of Blogspace. In *Proceedings of the WWW 2003*, 477-495.
11. LiveJournal, http://www.livejournal.com
12. Marlow, Cameron. 2004. "Audience, Structure, and Authority in the Weblog Community." Presented at the International Communication Association Conference.
13. Nardi, B.; Schiano, D.; Gumbrecht, M.; Swartz, L. 2004. Why We Blog. *Communications of the ACM* 47:41-46.
14. Netscan: http://netscan.research.microsoft.com
15. Rogers, E. 2003. *Diffusion of Innovations*. New York, New York: Free Press.
16. Schiano, D.J.; Nardi, B.A.; Gumbrecht, M.; Swartz, L. 2004. Blogging by the Rest of Us. *Ext. Abstracts of CHI '04*, 1143-1146.
17. Shneiderman, B. 1992. Tree Visualization with Tree-maps: 2-d space-filling approach. *ACM Transactions on Graphics* 11:92-99.
18. Smith, M. Measures and Maps of Usenet. In From Usenet to CoWebs, C. Lueg, D. Fisher, 2004. Springer.
19. Su, N.M.; Wang, Y.; Mark, G.; Aiyelokun, T.; Nakano, T. 2005. A Bosom Buddy Afar Brings a Distant Land Near: Are Bloggers a Global Community? To appear in *Proceedings on the 2nd International Conference on Communities and Technologies*.
20. Wallop : http://www.mywallop.com
21. Warschauer, M. 2000. "Language, Identity, and the Internet". In B. Kolko, L. Nakamura & G. Rodman (Eds.) Race in Cyberspace. 151-170. New York: Routledge.