

© 2006 by Ankur Kalra. All rights reserved.

SPICE - EMPOWERING BLOGGERS AND THEIR COMMUNITIES

BY

ANKUR KALRA

B.S., Georgia Institute of Technology, 2004

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Computer Science
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2006

Urbana, Illinois

Abstract

The ease with which blogs allow bloggers to communicate to their community has led to their adoption in a variety of contexts, including classrooms, newspapers, political campaigns and corporations. However, the ability for the members of a community to communicate with each other and with the blogger is constrained by the limited affordances of commenting, thereby preventing the blog from fully fostering community-driven content. In *spice*, a blog platform designed to explore social enhancements for blogs, we apply real-world metaphors to blogs to better address this need of a blog's community.

Prior work has also been invaluable in determining the composition of the blogging community, and in identifying common trends in blog usage. Viegas and others, for instance, have identified key disconnects between the characteristics of the medium and most blogger's perceptions of these characteristics. Bloggers tend to appropriate an informal, conversational tone in blogs, which conflicts with the inherently persistent and indexable nature of blogs. Additionally, bloggers often tend to only recognize the presence of their core audience, forgetting that their posts are persistently available to a much larger and wider audience. Disconnects like these have led to rude awakenings - for example, 36% of all bloggers have gotten into trouble with their friends, family or employers for something that they have said on their blog. In *spice*, we explore a combination of features that attempt to better address these needs of bloggers.

We released *spice* online and studied the usage patterns of eleven users around the world. The users we chose to follow represented a cross-section of bloggers, as identified by others in previous research, and blogged from across Asia, Europe, the Middle East and different parts of the United States. Additionally, we conducted informal focus groups with local teenagers, as this demographic was difficult to represent in our extended field study. The results of these evaluations are described in detail.

om

Acknowledgements

This project would have remained a collection of vague ideas if it weren't for the support of very many people. My advisor, Professor Karrie Karahalios, was invaluable in her role as intellectual guide, close supporter and good friend. I thank her for the opportunity to work with her, for her support in completing this work and for our conversations.

I'd like to thank my colleagues and office-mates at the Social Computing group – the discussions we had over many late-night cups of IHOP coffee shaped spice into what it is today.

I'd also like to thank our numerous users and participants for volunteering to experiment with spice, despite its shortcomings and 'technical issues'. In much the same vein, I'd like to thank our reviewers whose feedback continually drove me to improve and refine the work.

Perhaps most of all, I'd like to thank my friends and family – I do realize that a year's worth of conversations about my research can get a tad repetitive, and when any of you choose to write a thesis, I promise to nod along to your descriptions of things I barely understand.

Lastly, I'd like to thank the baristas of the Green Street Coffee House, whose patience and never-ending supply of tea were perhaps two of the most critical ingredients in the completion of this thesis.

Table of Contents

List of Figures	viii
List of Tables	ix
1 Introduction to Spice	1
1.1 Community-Driven Content	2
1.2 Transience and Privacy	2
1.3 Roadmap	3
2 Background	4
3 Community-Driven Content	8
3.1 Design Motivators	8
3.2 Design Concept	11
3.2.1 Annotation	11
3.2.2 Interaction History	13
3.2.3 Spatial filtering	13
3.2.4 Flexible vocabulary	13
3.3 Implementation	16
3.4 Evaluation	18
3.4.1 Focus Groups and Poster Session	19
3.5 Future Work	22
4 Transience and Privacy	25
4.1 Transience	26
4.1.1 Design Concept	27
4.1.2 Design Motivators	29
4.1.3 Implementation	30
4.1.4 Evaluation	31
4.1.5 Future Refinements	32
4.2 Privacy	34
4.2.1 Design Concept	34
4.2.2 Design Motivators	35
4.2.3 Implementation	39
4.2.4 Evaluation	39
4.2.5 Future Refinements	40

5	Evaluation	42
5.1	Design	43
5.2	Demographics	45
5.2.1	Respondents	45
5.2.2	Users	45
5.3	Overall Feedback	46
5.4	Limitations	47
5.5	Feedback from Focus Groups and Poster Sessions	48
5.5.1	High School Focus Groups	49
5.5.2	IDCSA 2006	50
6	Key Lessons and Conclusion	52
6.1	Lessons learnt about alternative models of social interaction	52
6.2	Lessons learnt about the difficulties in evaluating a blogging platform	55
6.3	Conclusion	56
	References	58

List of Figures

1.1	The spice logo	1
2.1	A typical blog	4
2.2	Rate of growth of blogs	5
2.3	Blogs as a mean of expression in otherwise repressive societies	6
3.1	Annotations and margin notes	9
3.2	Annotation in spice	11
3.3	A spice post with four comments	12
3.4	Interaction history and spatial filtering in spice	14
3.5	Flexible vocabulary examples	17
3.6	Commenting in spice	18
3.7	User-chosen posts for spice	20
3.8	Feedback from users	21
3.9	Refinement to the commenting system of spice	23
4.1	Transience in spice	27
4.2	Activity in spice	28
4.3	Existing privacy options	35
4.4	Multifaceted online identity	36
4.5	Privacy options in spice	37
5.1	Countries represented by the users in our study	43
5.2	States represented by the users in our study	44
5.3	Poster session materials	50

List of Tables

5.1	The composition of the focus groups, drawn from six high schools from the state of Illinois	49
-----	---	----

1 Introduction to Spice

First introduced in 1994, blogs have now been around for over a decade. Over the last twelve years, they have grown from an experimental form of communication for techies to a medium appropriated by celebrities, politicians, entrepreneurs, homeless people and everyone in between. Indeed, the flexibility of blogs and the ease with which they can be maintained are part of what makes a blog a blog.

The popularity of blogs has led to considerable attention, both in academic circles and the mainstream media. As a form of computer-mediated communication, the raw potential for blogs is still being appreciated, but considerable work has been done in trying to understand the psyche of bloggers. Who are they? Why do they blog? Why don't they blog more? What fascinates them about the medium? What upsets them? What impact does their blog have on the rest of their lives? Like any other form of computer-mediated communication, blogs connect people, and have both a technical and a social impact on their lives.

Nonetheless, the enhancements made to blogs in recent years, important as they are, have been mostly technical. The social implications of blogs are just beginning to be understood, and not much work has yet been done in experimenting with enhancements that have a social impact on bloggers and their community. To explore this design space, we built SPICE – a Social Platform for Interacting and Communicating with Everyone¹. As a platform, spice helped us address the needs of bloggers and their communities, as identified by an increasing body of prior work.

It is important to mention that we do not consider spice to be the solution to all the identified shortcomings

¹Admittedly, this is a bit of a backronym.



Figure 1.1: The spice logo

of blogs. Our goal with spice has been to explore the design space of solutions and to thereby better aid future work in the area.

1.1 Community-Driven Content

Even before blogs were prevalent, some have observed [19] that the network fosters an almost hive-like capacity towards content creation. This has been evidenced in numerous topical contexts, ranging from the successes of Wikipedia to those of the open source software movement. Although they have received less formal attention, the same effect can also be observed in more individual contexts. Blogs, for instance, are especially interesting – although blogs can be very personal, they are essentially a collaboration between a blogger and their community.

Indeed, it is the community of the blog that transforms the contents of a blog from the opinion of one person to the thoughts of many. The word ‘Slashdot’ refers as much to the community of ‘Slashdotters’ as it does to the website itself [6]. The blog is the seed around which the community crystallizes their opinion; the blogger is the moderator of a discussion larger than their blog. Unfortunately blogs, as designed today, poorly support this model of interaction.

Blogs today are blogger-centric – the comments of the community are secondary to the contents of the post, and are divorced from the context it provides. We wanted to explore the effects of a system that was explicitly designed to *empower the community* by supporting community-driven content in what is traditionally a very individual medium.

1.2 Transience and Privacy

In trying to understand the impact of blogs on the rest of the blogger’s life, Viegas’ identified that 36% of bloggers had gotten in trouble with family, friends or employers for things they had written on their blog [28]. 34% of bloggers knew of other bloggers who had gotten in trouble for things they had written online, and 75% of the sampled bloggers admitted to having gone back to edit previously published posts. Viegas’ work pointed towards a fundamental disconnect between the characteristics of the medium – persistent,

indexable, searchable, universally accessible – and the transient, conversational tone appropriated by bloggers in their usage of the medium.

In spice, we designed in features to better align the medium to the ways in which blogger's used the medium. We explored metaphors to convey and enforce transience, and built in sophisticated access-control mechanisms so that bloggers could preserve their privacy across multiple dimensions without having to censor themselves. In these ways, we tried to *empower the blogger* in ways that had little technological impact but considerable social impact.

1.3 Roadmap

This thesis is broken into six chapters. Chapter 2 discusses the world of blogs and the state of prior work in the area. Chapter 3 describes the mechanisms we built into spice that were designed to support the community. Similarly, Chapter 4 describes the features that we designed to support the needs of the blogger. Chapter 5 discusses the various approaches we took to evaluate the features that we had built and leads up to a discussion of key lessons learnt, covered in Chapter 6.

2 Background

Blogs, are websites maintained by one or more authors where content is displayed in reverse chronological order. They are notable for the ease with which they allow their authors, called bloggers, to maintain and update content. First introduced in 1994, blogs have typically been text-based, although in recent years photo-blogs, audio-blogs (typically called ‘podcasts’), mobile-blogs (also called ‘moblogs’) and video-blogs (‘vlogs’) have also become common. The word ‘blog’ is itself a contraction of the term ‘web log’.

Blogs have come a long way in terms of adoption in the past decade. Xanga, a popular blog site that only had a hundred or so blogs in 1997, had over 20 million blogs by 2005 [7]. In the time since, blogs have received considerable attention in the popular media [3, 31], and have mushroomed in popularity – Technorati, one of several blog indexers, tracks nearly 60 million blogs alone [7].

As powerful a cultural phenomenon as blogs are, their evolution has been mostly technical. In terms of the interaction models they support, blogs today are very similar to the blogs of almost a decade ago. Since their introduction in 1996, the interaction model of blogs have only undergone two significant changes – the introduction of commenting in 1997 gave the com-



Figure 2.1: A screenshot of a typical blog




Figure 2.2: In recent years, blogs have exploded into the mainstream, resulting in almost 60 million blogs tracked by Technorati as of Oct. 2006

munity of a blog it's voice, and the introduction of RSS syndication in 2001 redefined the dynamic between bloggers and their audience. Although many other notable enhancements have been made over the years, they have been primarily technical enhancements that have served to increase the accessibility and feature-set of blogs. These enhancements have been critical to the success of blogs as a mainstream cultural phenomenon, but have only begun to explore the potential for complex social dynamics between bloggers and their audiences.

Several organisations have appropriated the medium as their primary means of communication [12, 18, 20], and many people have leveraged the unique characteristics of the medium to document events as they happen [3, 5, 29], a common challenge to more traditional journalistic channels. People have even turned to blogs to express themselves in an otherwise censored society [9].

Blogs have also received attention in academic circles – Herring, et al's genre analysis of weblogs has laid the way for much work in the area [16]. Viegas and others have identified key limitations of blogs today [13, 28], and Gumbrecht has done ethnographic work to determine the boundaries of what bloggers consider their 'protected space' [13].



Figure 2.3: Due to a strongly censored society, blogs like khabgard.com are an extremely popular means of expression in Iran.

Unfortunately, bloggers' perceptions notwithstanding, blogs are usually very public spaces, and this conflict has resulted in more than a third of all bloggers getting in trouble for things they have written on their blog [28], often with quite serious consequences [1, 7, 24, 26].

Although the implications of privacy and control online are just beginning to impact blogs, these effects have been well-studied in the past in other contexts. Most notably, there was considerable backlash when Usenet first transitioned from an essentially transient medium into a searchable, indexable medium [8, 14]. However, there is a crucial difference between Usenet and blogs today – Usenet existed as a public resource that users specifically had to seek out to participate in, an online public commons of sorts; blogs, on the other hand, are typically considered the private space of the blogger and are much more personal and individual, as Gumbrecht's work indicates [13]. This has considerable implications for the approach taken to address the privacy needs of the blogger. To our knowledge, Viegas' Collections work for personal photographs has been the only other major effort to design socially-aware access control mechanisms for personal content [27].

Similarly, much work has been done validating the need for better visual cues for consumers of digital content. Nielsen indicated that users typically skim through online content and often rely on prominent visual landmarks

[22]. Hill, et al and Wexelblat have long since demonstrated the successes of such a system for digital content in general [17, 30]. Much more recently, Golder and others have experimented with annotation mechanisms for online content [4, 11, 32], although to our knowledge, no such mechanisms have yet been explored for online content in social contexts.

Blogs have been around for more than ten years, yet the big changes introduced in recent releases of major blog platforms have been a more responsive interface, the ability to upload pictures and movies and WYSIWYG (What You See Is What You Get) text widgets [21]. While they undoubtedly contribute towards the overall blogging experience, they address absolutely none of the identified social shortcomings of blogs [28].

Our work with spice leverages the body of prior work in the research community to explore ways to address the needs of both bloggers and their communities better.

3 Community-Driven Content

Although the vast majority of blogs have a single author, most blogs can not be fairly described as the product of a single individual. The power of a blog lies as much in the community that it fosters and the connections that it facilitates as it does in the content of the individual posts.


For a form of online interaction that is inherently social in nature, blogs are surprisingly limited in their ability to foster community driven content. Even the commenting system available to a blog's community sets a very deliberate tone in the styles of interaction it supports. Comments on a blogger's post are both visually and spatially secondary to the post; comments often need to be approved before they are shown – indeed, since most platforms do not support threading, even discussions that the community chooses to actively participate in are crippled in their ability to scale.

Although there are examples of successful community-driven content online, such as Wikipedia, these examples are a lot rarer in the blog world. Sites like Digg and Slashdot attempt to tackle these issues at a macro level by enabling the community to blog by proxy, but these blogs are highly specialized and their approaches only make sense for blogs of that style [16]. An individual's personal blog can not appropriate these styles of interaction without significantly redefining 'individual' and 'personal'.

With these limitations in mind, we designed the commenting features of spice to better leverage the communities that bloggers thrive in. In the subsequent sections of this chapter, we describe the motivators, design and implementation of this system, concluding with a discussion of its evaluation and plans for future refinement.

3.1 Design Motivators

One of our primary sources of inspiration for the concept of community-driven content comes from the richness of commentary found in many used



(a) Typical Used Book



(b) The Talmud

Figure 3.1: Annotations and margin notes serve to provide context in books

books. Indeed, it is not uncommon for people to prefer used varieties of certain books, due to the content enhancers one can often find in them [30]. In textbooks, for instance, underlined and highlighted text serves to indicate key concepts, and margin notes routinely help explain tricky points. Even in literary works, margin notes and dog-ears help provide insight into other people’s perception of a written work, perhaps by providing amusing commentary or by referencing other similar works.

Considering the richness of commentary found in a book previously read by a single reader, an obvious next question is – what happens when content has been read by very many people before? In the world of blogs, for instance, it is not uncommon for single blog posts to have several dozen commenters, and hundreds of readers. Nonetheless, existing commenting systems do not allow for nearly the same level of richness that a used book affords a single reader. This is especially surprising considering that these concepts have been previously explored in the context of offline digital content [17, 32], and usability studies have called for these kinds of cues [22, 30].

However, existing comments do serve an important function. As a form of annotation, their existence extends the blogger’s static, single-person post into content that represents the opinions and insights of many; by presenting supporting or opposing viewpoints, by providing references to related content or to relevant background information and even, on occasion, by providing comic relief. As comments are currently the primary means for a community

to enhance the content of a blog post, for our exploration of the design space, we chose to address some of their characteristics that limited their ability to truly foster community-driven content.

- *Delayed annotation* – The annotation functions of comments, while useful, are limited in their potential. Since these annotations are spatially subjugated to the content of the post, it is only after the reader has read the entire post and decided to read through the comments as well that he or she discovers that there is a flaw in the original argument, or an error in a presented statistic.
- *Comment dispersion* – Only a reader who is diligent enough (and has enough time) to review all the comments on a post can take advantage of their insight. Even if only a small set of the comments relate to the parts of the post that the reader cares about, the reader still has to process each comment to be able to identify the ones that are relevant.
- *Limited affordances* – Although the text of a comment can be as expressive as the literary skill of the commenter permits, the form of the comment constrains the ways it lets commenters interact with other members of the community. The vocabulary of interactions that a commenter has at his or her disposal is limited to threading, and with threading, commenters can indicate no more than a relationship between two comments. Commenters have no ability to indicate that their comment is challenging the previous commenter, as opposed to supporting it; or that while they agree with the previous commenter’s conclusion, they do not agree with the previous commenter’s reasoning. Solving this problem in a meaningful way involves redefining the model of interaction to provide a platform for expression that commenters can adapt and appropriate in novel and unexpected ways.
- *Repeat processing* – Comments represent the reaction of a member of the community to the content of the post. Regardless, the forty-second reader of a post, upon first viewing it, has no indication that forty-one people have already processed and reacted to this same post. A statement made by the blogger that has, over time, organically evolved and been enhanced by the community, is still presented in its original, raw




Figure 3.2: An example of annotation in spice – the readers have annotated an image the blogger uploaded with their own caption

form to any subsequent reader. At first glance, the reader has no visual cues to indicate how controversial the post is, how much discussion it has generated, or even how often it has been viewed. Affordances for these kinds of subtle cues are often taken for granted with printed works – wrinkles, notes, dog-ears and even the condition of the binding are all cues that affect one’s perception of such work.

3.2 Design Concept

When a post is loaded in spice, comments are rendered as small comment-shadows. When moused over, these comment-shadows intelligently expand to include the comment text, taking into account nearby obstacles, the length of the comment, and the preferences of both the blogger and the reader. Part of the process of leaving a comment involves ‘dropping it off’, so that the system knows what part of the post the commenter meant to comment on.

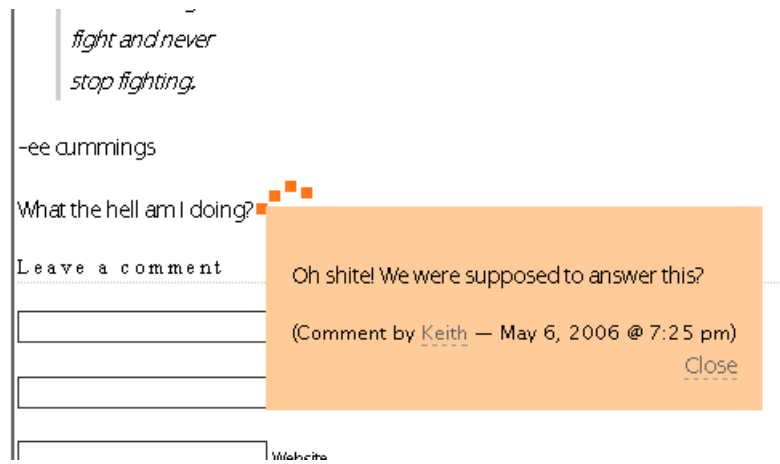
Comments are not second-class citizens in the world of the spice blog post – they occupy the same space as the content of the post and are no longer divorced from the context it provides. By virtue of their co-location, they indicate to future readers parts of the post that have generated much discussion, even before the readers have started reading the post. Their spatiality naturally groups together related threads of discussion. The *form* of the comments allows for very flexible usage, enabling commenters to use it as another channel of interaction.

3.2.1 Annotation

In spice comments, as annotations, help provide external context to the post *while it is being read*. Consequently, readers can now identify that there is a



(a) Comment-shadows



(b) An expanded comment

Figure 3.3: A spice post with four comments

subtle flaw in an argument, or look at background material for some statistic, while the argument is being made or the statistic is being presented.

3.2.2 Interaction History

While each comment adds value individually as an annotation, many comments together add value by providing traces of *interaction history* as well [30]. The dense presence of very many comment-shadows in a certain region give readers processing cues to better understand the content of the post even before they start reading. Cues like these have been shown to help consumers of digital content – for example, by enabling readers to get a better high-level understanding of the community’s reaction to a post at a glance.

3.2.3 Spatial filtering

By associating comments with a spatial location, commenters are essentially tagging their comments with a specific semantic part of the post. This is a crucial step from existing blog systems because for the first time, we can group together semantically related comments. This also allows readers to ignore comments that are related to a part of the post that they are not interested in – even if a reader chooses to read every comment made, this system allows them to ignore comments that are unrelated to the discussion they are currently following.

An important difference between this system and a threaded comment section is that even though a threaded comment section allows for semantic grouping, it does not enable the reader to identify the topic being discussed in a specific thread without actually reading the comments. With this system, not only are comments grouped, they are also linked to a specific topic by virtue of their co-location with the blogger’s post.


A fortunate side-effect of this is that unwanted posts – such as those automatically generated by comment spammers – algorithmically end up grouped together as well, thereby allowing readers to ignore them en-masse.

3.2.4 Flexible vocabulary

Due to the diverse nature of personal blogs and their communities, we wanted our design to provide a flexible platform for interaction. We wanted to build

3 AM Epiphany

Filed under: Me, myself and Hem — hemlyn @ 3:10 am

*To be nobody but yourself
in a world which is doing
its best day and night to
make you everybody else
means to fight the hardest 
battle which any
human being can
fight and never
stop fighting.*

-ee cummings


What the hell am I doing? 

Figure 3.4: An example of both interaction history (interesting regions are apparent even before the text is read), and spatial filtering (multiple threads of discussion can be followed independantly)

in a versatile vocabulary of interactions that each community could appropriate in their own ways. In spice, the spatiality of the comments helps us accomplish this.

With traditional comments, commenters could at most indicate the specific comment that they were replying to – thereby forcing the interaction between the members of the community of the blog to be explicit. With this commenting system, the placement of comment-shadows gives rise to a whole new implicit and visual vocabulary of interaction. Although placing a comment next to a previous comment indicates a relationship between the two comments, information is also encoded in the spatiality of the two comments – if a comment-shadow is placed on top of an existing comment-shadow, for instance, the interaction is very different than if they were placed side-by-side. A few possible interactions that can be exhibited by comment-shadows are¹:

- *Conflict* – by choosing to place your comment-shadows on top of another’s, it is possible to try to drown out your opponent.
- *Disagreement* – by choosing to place your comment-shadow near, but distinct from previously placed comments, it is possible to indicate a differing viewpoint on the same topic.
- *Cooperation* – by choosing to place your comment-shadow closer to the viewpoint (of multiple opposing viewpoints) that you agree with, it is possible to indicate your support. This has the effect that it is visually apparent which of multiple groups has the most support in the community.
- *Play* – by choosing, as some of our users did during the study, to interpret your comment-shadow less as an indicator of a comment and more as a really large pixel, it is possible to draw simple pixel-art (such as a smiley-face). The comment-shadows allow for play *with* the comment-shadows, allowing for playful interactions that are not just about the discussion. As a simplistic art form, this pixel-art can transform a comment into both visual- and content- graffiti.

¹It should be noted that these are merely examples mentioned to motivate the discussion, and that this list is by no means exhaustive. None of the design was influenced by these specific use-cases – rather, the design was intended to allow for flexibility in usage, and this list is included merely to illustrate the diversity of interactions supported.

3.3 Implementation

For us to be able to conduct a meaningful evaluation, we needed to be able to support a wide variety of platforms for an extended period of time. Any impact on the user experience would affect the communities around the spice blog directly, and result in users unwilling to complete the study. Like most web-based applications, our system had to be robust to screen resolutions and browser rendering quirks. Unlike most applications, however, we needed to render consistently across browsers as well. For the commenting in spice to work, it is not enough for the Firefox experience to be consistent and the Safari experience to be consistent – it is also important for comments made in Firefox to render in the same spatial location when viewed in Safari.

The comment form in a spice blog has an additional link, usually entitled ‘Commenting on?’. Clicking on this link activates an embedded Javascript based mouse tracker that renders a preview comment-shadow at the location of the mouse. This mouse tracker ensures that the comment-shadow follows the movement of the mouse, until the user has clicked on the part of the post that they want to leave their comment on. At this point, the mouse tracker deactivates and saves the final co-ordinates – the comment-shadow persists in that location as a visual indicator of the saved co-ordinates. If the user wishes to move the comment-shadow to a new location, clicking on the ‘Commenting on?’ link will reactivate the tracker and restart the process. When the user is satisfied with the location of the comment-shadow, and has submitted their comment, the location of the comment-shadow is saved by the system along with the comment.

In an earlier iteration, this location was represented by a two-dimensional co-ordinate, with the origin fixed at the upper-left corner of the page. Saving this information, however, was meaningless – laying out the page is the responsibility of the browser and unless the screen-size, resolution and browser versions matched exactly, it would be unlikely that the stored co-ordinate specified by the server would point to the same semantic part of the post as before. Furthermore, if the blogger were to ever change the layout of their blog, all the previously stored co-ordinates would be horribly wrong.

In our final iteration, we addressed this problem by introducing a simple transformation, which was found to considerably increase the robustness of the system. The body of the post was encapsulated by an HTML div element.

... time, consider
... Isn't it amaz



... than some of tl
... it is really old, i
... So even thou

(a) Conflict

... enthused about grow
... wish I didn't, and I'm
... ns of rust.



Sheer brilliance! I'm

(b) Disagreement

demo @ 6:49 pm

... ole wisdom tooth #2



red)

(c) Support

... the comments
... ve discussions,



(d) Play


Figure 3.5: Examples of the flexible vocabulary that spice provides – in each case, the yellow box indicates the comment currently being placed in the context of previous comments

Leave a comment

Name (required)

Mail (will not be published) (required)

Website

[Commenting on?](#) 

When the yellow box appears, point your mouse at the part of the post you want to comment on and click again. If you make a mistake, click the "Commenting on?" link again to start over.

Figure 3.6: The form used to leave comments in spice – the orange box follows the mouse to indicate where the comment will be placed when the user clicks

The location of this wrapper div on a rendered page is trivial to compute with basic, cross-browser Javascript. Before submitting the co-ordinates of a comment, this location is computed, and a transformation is performed so that the origin is now the upper-left corner of this wrapper div. The reverse transformation is applied when the page has finished loading or the window has been resized. As long as the content within this div – usually just plain text – is rendered consistently across browsers, the comments will render accurately.

3.4 Evaluation

When the system was described to them, all of our users were enthusiastic about the new commenting system. Eight of our users had readers that were enthusiastic with the new commenting system from the very start of the study. Of our three remaining users, one user's readers gradually became comfortable with it, one user chose to maintain a private blog that took advantage of other aspects of spice, and one of our users discouraged

commenting on her blog throughout the study.

One of our users had readers that were almost excessively enthusiastic in their usage of the commenting system – their usage of comments spanned all of the uses we’d thought of while designing the system, sometimes in a single post. For example, the post illustrated in the figure exhibits annotation, interaction history, spatial filtering and playful interaction, all in one post.

Although there was a lot of positive feedback to the commenting system, there was also a few areas that were identified as needing some refinement. Notably, some of our users remarked that this system *appeared* too similar to existing commenting systems, which occasionally confused readers who were expecting it to behave the same way. This led toward a prototype (described in the next section) with different affordances that met with better initial feedback. Additionally, some users felt that commenters had too much freedom in their placement of comments, and asked for functionality that would prevent commenters from leaving comments over one another. It is interesting to note that this behaviour only ever occurred accidentally during the course of the study, which suggests a disconnect between the levels of conflict within a community, and a blogger’s perception of these levels. Lastly, our users suggested several enhancements, such as custom comment shadows and landmarks, which are discussed in the next section.

3.4.1 Focus Groups and Poster Session

These sessions were intentionally kept informal, in order to encourage feedback, and every participants was provided with convenient access to a spice blog.

The focus groups were conducted in three sessions of roughly ten members each. In each of these sessions, the motivation behind the commenting system of spice was presented, after which the members of the focus group were given a few minutes to explore the commenting system at their individual stations. The members of these focus groups were high-school girls from various schools in Illinois. They spent a weekend at the Department of Computer Science, being introduced to and exposed to different styles of research, to better expose them to what majoring in Computer Science entailed.

Of the high-schoolers present, no more than two identified themselves as non-bloggers – each of the other members described themselves as either

Rhapsodic

April 5, 2006

Men and weirdness
 Filed under: [uncategorized](#) — hemlyn @ 9:37 am

So, Keith and I have do this thing where we exchange mp3s all the time. Sometimes by email, I think once over MSN and sometimes via removable disks. Anyway, so he sent me a song by the Bloodhound Gang and while I was editing my mp3 tag I discovered the name of the album; "Hooray for Boobies!"

Men. Honestly. Not hooray for candy! Or Yipee for new shoes! But Hooray for Boobies. Of all the things to hooray about...

On a side not, HAPPY BIRTHDAY, TASI! We're all so *bluetoday*. Come cheer us up? 😊

Leave a comment

Name

Mail (will not be published)

Website

Commenting on?

When the yellow box appears, point your mouse at the part of the post you want to comment on and click again. click the "Commenting on?" link again to start over.

blogroll

- [my other blog](#)
- [the world keeps spinning around](#)
- [spice](#)
- [british born confused desi](#)
- [the compulsive confessor](#)
- [a walk in the clouds...](#)
- [postsecret](#)
- [why have you forsaken me?](#)
- [going up in flames](#)
- [scratchpad](#)
- [back from hell](#)
- [think tank](#)
- [random thoughts on just about anything](#)
- [yacketyakking screaming vomiting whispering why me?](#)
- [through my eyes](#)
- [the epic abyss of incredulity](#)
- [do not go gentle into the goodnight!](#)
- [lunacy unplugged](#)
- [seven days from nowhere, camping on the horizon](#)
- [time is in a hurry, it cant wait to meet its final destination...](#)
- [the restaurant at the end of the universe](#)
- [epileptic raptures](#)
- [clarity? coherence? logic? here? indbloggers - index of indian bloggers](#)
- [blah...](#)
- [life is good yeah right !!!](#)
- [ronald's road to perdition](#)
- [r i e v i](#)
- [miscalculated mischief of a misguided mind](#)
- [falon](#)
- [mute](#)

categories:

- [uncategorized](#)
- [me, myself and hem](#)
- [spice](#)

search:

archives:

- [april 2006](#)
- [march 2006](#)

meta:

- [login](#)
- [rss](#)
- [comments rss](#)
- [valid xhtml](#)
- [xfn](#)
- [wp](#)

Figure 3.7: This post, by one of our study users, exhibits in itself annotation (lower-left), interaction history (lower-right), spatial filtering as well as playful interaction (upper-left)

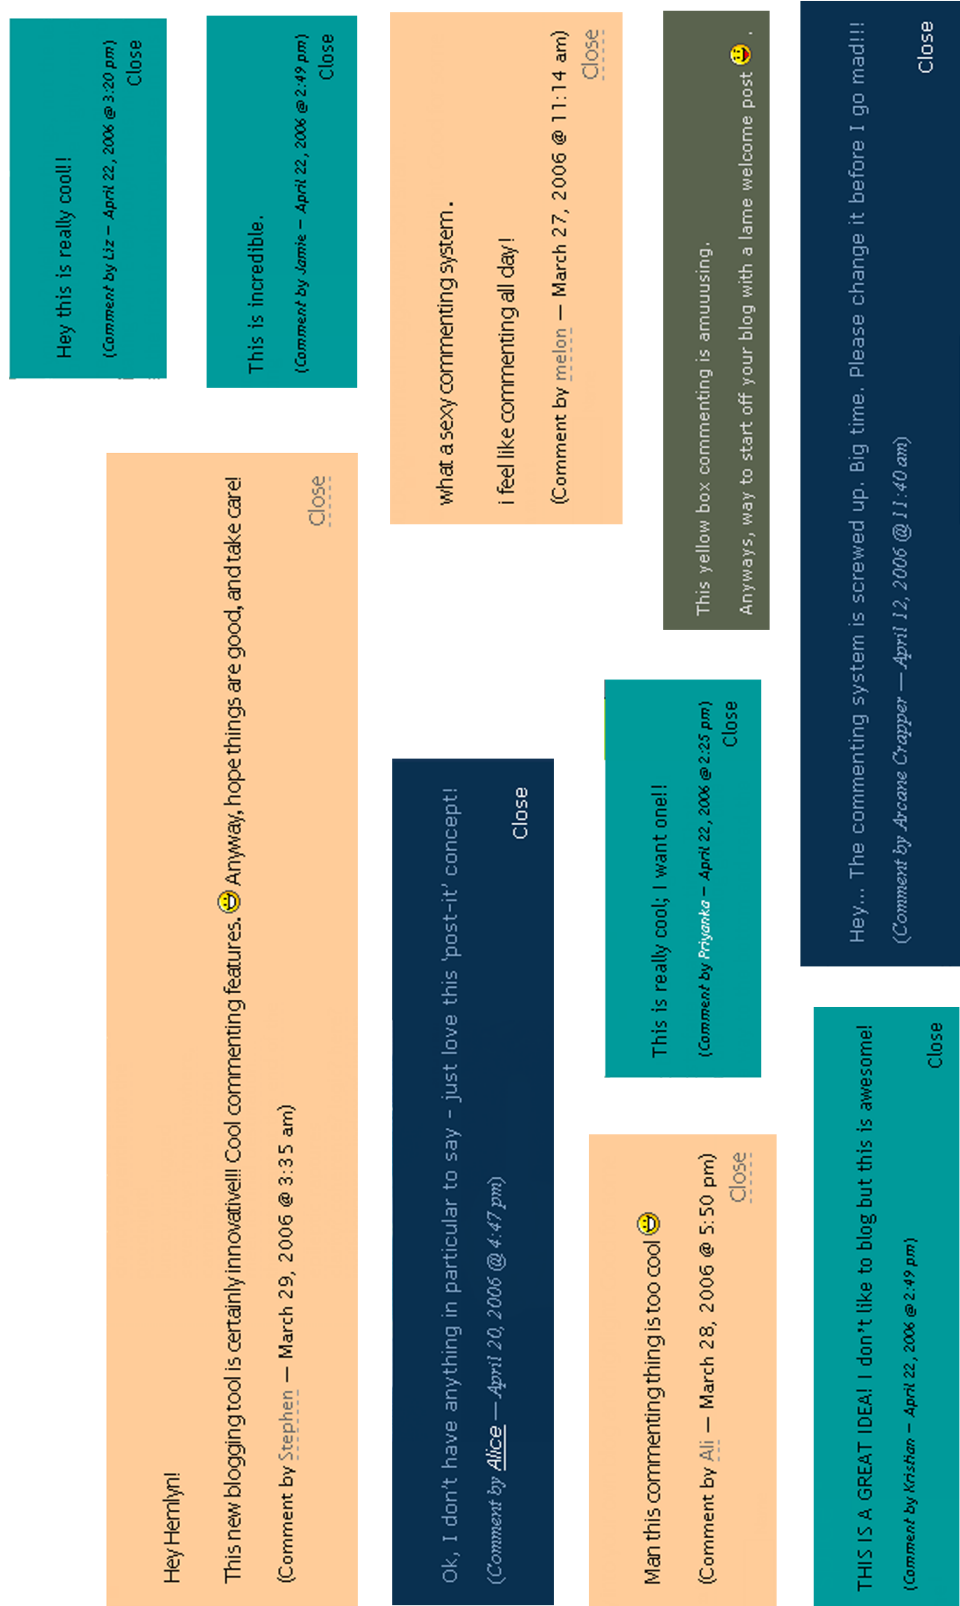


Figure 3.8: These are a few of the comments that have been left on spice blogs that are online and available to the public

avid or casual bloggers. Although a few had their own custom blogs, most of the members of the focus groups blogged on one of the major platforms (LiveJournal, MySpace, Xanga, or something similar). Each of the groups had at least three members request their own spice blog before the end of the thirty-minute session. There were also many requests to ‘make it work for Xanga’.

During the poster session, participants were introduced to the larger spice system through a poster, while a neighbouring interactive station demonstrated the commenting features of spice. The participants in the poster session were mostly people affiliated with technology in some way, either as students, researchers, faculty members or in industry. Some of the participants had already begun to indulge in the playful behaviour we had noticed during the diary study, which we found to be an encouraging sign.

3.5 Future Work

Although we do feel that the design of this novel commenting system was a success, we also believe that the idea would be better served by continued refinement. Some of the things we plan to try out in the future, based off of the feedback we received, include:

- *More streamlined commenting* – One of the biggest challenges we faced with this implementation was that it *looked* very similar to existing commenting systems – thereby not providing enough cues to readers and commenters that it was a novel system that worked differently. A refinement we’re working on provides for a more streamlined commenting experience, and better integrates the commenting and the marking steps – a preview is included in the figure.
- *Customizable comment-shadows* – Currently, each comment is rendered the same way - as a small blue (or other theme-appropriate colour) box. To mirror corresponding developments in traditional commenting systems, a future refinement will allow for the blogger’s comments to be rendered in a different colour, or with a small icon. Interfacing with third-party systems like Gravatar would allow for similar support for all commenters.

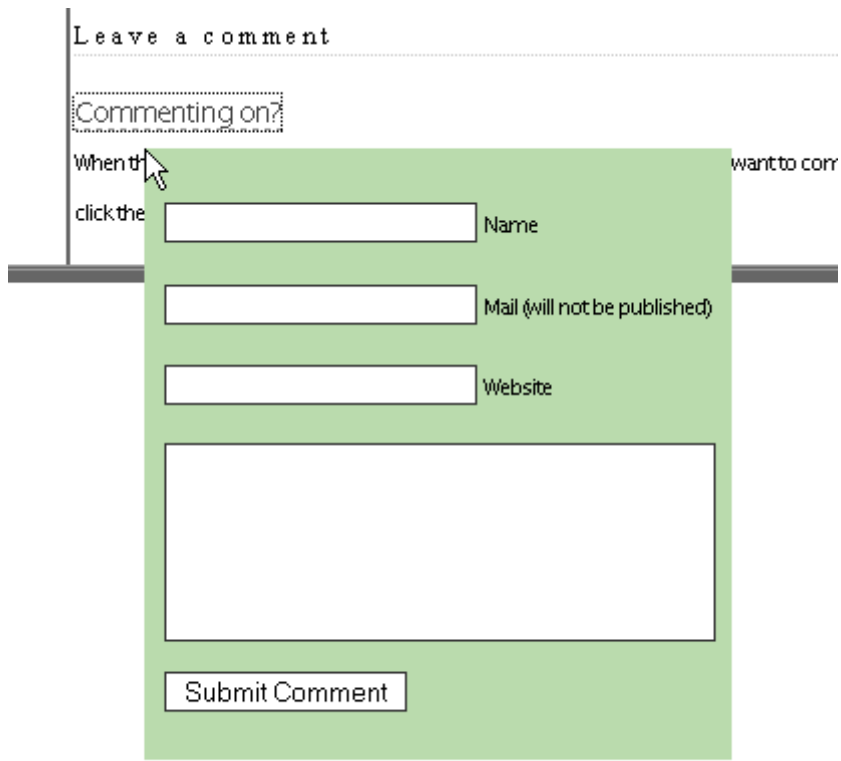


Figure 3.9: In this version, the commenting form follows the mouse, and has to be dropped in the appropriate place before it can be filled out, which better maintains the metaphor of leaving comments ‘in-place’

- *Custom-designed themes* – While we met with positive feedback, we still felt that the success of the system was severely limited by the fact that the themes and layouts used were mere adaptations of existing themes designed for other systems. We’re currently exploring design options that take into account this novel style of commenting – one could imagine, for instance, a theme that had ‘landmarks’ that commenters could leave comments on. An immediately obvious example scenario for this would be for polls, where the number of comments left near each option would immediately indicate its relative popularity.

4 Transience and Privacy

Bloggers today are in a unique position – the ease and power of blogs have given the average blogger access to an audience that has typically been the privilege of writers and journalists.

The blade cuts both ways, however – such unprecedented potential has led to a poor understanding of the true nature of the medium. Bloggers often learn the hard way that their medium is a little too powerful. Once it’s been written, with a few exceptions, almost anyone in the world can read your blog. Once it’s been written, powerful indexing and archiving systems in place ensure that almost anyone in the world can read your blog *anytime* – in two minutes, in two weeks, in two years. Most blogs today, however, are of a much more personal nature, and thereby appropriate more personal modes of interaction – bloggers tend to treat their blog as a web-based diary instead of as a newspaper [28]. The tools they have at their disposal, however, are ill-designed for the needs of these kinds of bloggers.

To better empower individual personal bloggers, we looked at their usage of blogs in the wild, so that we could identify the dissonance between what they expected from the medium and what it provided them. Viegas’ work in this area was invaluable. She identified 36% of bloggers as having gotten in trouble with family, friends or employers for things they had written on their blog. 34% of bloggers knew of other bloggers who had gotten in trouble for things they had written online, and 75% of the sampled bloggers admitted to having gone back to edit previously published posts [28]. Viegas’ work pointed towards a fundamental disconnect between the characteristics of the medium – persistent, indexable, searchable, universally accessible – and the transient, conversational tone appropriated by bloggers in their usage of the medium.

Although many bloggers have come to terms with the realities of the technology the hard way, the question begs to be asked – what if blogs were not persistent and universally accessible? What if the content was transient

and temporal, much like conversation and dialogue? What if some blog posts were only accessible to long-time readers of your blog, even if you couldn't name them? What if you could post an entry for that one reader whom you've never met and couldn't identify, but whom your site logs tell you has been checking your blog thrice a week from Australia?

For our exploration in this space, we built two sets of related features. The first feature attempts to introduce a temporal metaphor to blogs by making content transient. It borrows metaphors from the real world to subtly convey age and make posts, in some ways, perishable. The second feature, on the other hand, explores alternative access-control mechanisms that allow bloggers to tailor content to readers based off of their behaviour instead of their identity. This is a crucial step, as prior work has identified that bloggers can typically identify only 30% of their audience [28]. By providing bloggers with tools to personally interact with the remainder of their audience, we enable them to build meaningful connections with a much larger section of their audience. The remainder of this chapter describes our approach towards exploring these needs, the results of our evaluation, and some directions for future exploration that we have identified.

4.1 Transience

The online persistence of content is not a new problem – web-decay and link-rot are problems nearly as old as the Internet itself [2] – however, the social aspects of the blogging domain magnify the impact of these problems. Similar social issues were also identified when Usenet discussions were first archived and made publicly available [8, 14]. However, since a blog is generally perceived to be a blogger's personal space, the social implications of this persistence are much more complex; people are less likely to censor themselves on their blogs (which are typically considered personal space) than on Usenet (which is usually considered a public commons).

It is with these challenges in mind that we designed the temporal system for spice. This system borrows and extends visual cues that users are already familiar with from real-world metaphors, and adapts them for online usage. Although they were developed and tested with spice, the design was intended to be flexible across a variety of web-based media (blogs, wikis, websites, et



Figure 4.1: Over time, content in spice visually fades away to represent age (and other factors) with little to no modification – this flexibility allows the cues to be reinforced across the multiple facets of a user’s online experience.

4.1.1 Design Concept

Posts in spice visually fade away over time, until they are eventually no longer publicly accessible. The reduced visual contrast is aligned with existing real-world metaphors, such as old books and thermal facsimile messages, to help convey age.

In practice, this is implemented by tying the distance of the foreground

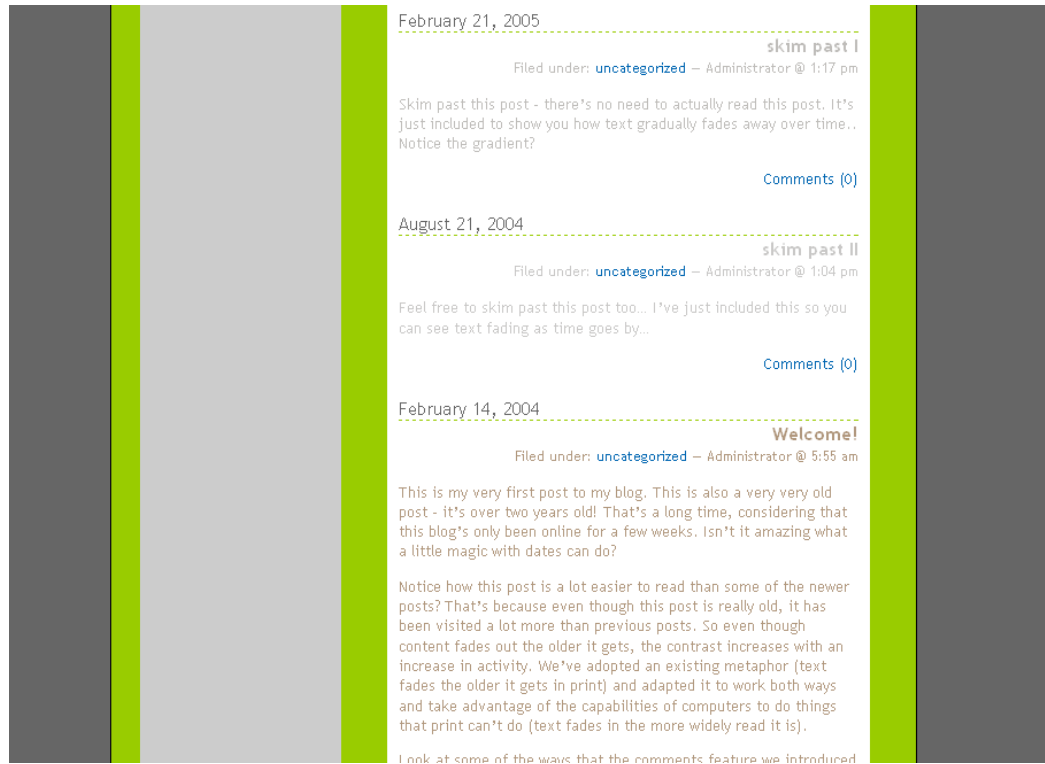


Figure 4.2: More active posts fade back in to emphasize their relevance

colour from the background colour to the time since the post was created.

Older posts that are visited disproportionately, however, have a corresponding increase in visual contrast applied to their content. This allows us to step beyond existing metaphors, and extend contrast to be a metaphor of *relevance* instead of age¹. This is done by mixing in a third colour which complements and is visually distinct from both the background and foreground colours, and is of a similar saturation as the foreground colour. This colour is typically theme-dependant, and is therefore usually chosen by theme-designers.

These, however, are merely meant to serve as visual cues to the reader – in both of these cases, it is still possible for the reader to access the full-text of the post. However, the transience model of spice has an additional feature that goes a step further. In spice, posts only fade away for a certain amount of time. Each post has an associated expiration date, and once a post has

¹The term ‘relevance’ is used her in the most generic of senses – this approach does not attempt to gauge the relevance of a post to an individual reader. However, this approach does reward both recency and activity, thereby approximating relevance to both the blogger and the larger community.

passed its expiration date, it is no longer publicly accessible. The blogger, naturally, always has access to the post.

These features are designed to complement each other – a post fades away smoothly until its expiration date, at which point it is no longer publicly accessible. Posts that expire in weeks will fade much quicker, while posts that won't expire for years 'keep better' and fade much slower.

When a blogger sets up their spice blog, they are asked to estimate the typical life of their posts. This is used to suggest a default expiration date for the post when it is being made. The blogger, however, can alter a post's expiration date at any time, including after it has already expired.

4.1.2 Design Motivators

Recency

The time-based depiction of content is one of the defining characteristics of blogs [31]. In fact, the ease with which blogs are updated to provide relevant, timely content is one of their core strengths – making them to the preferred medium for communicating events as they happen [3, 5, 29].

This model of content organization implicitly values more recent content as more relevant. Visitors, however, often reach a blog through alternative entrances – search engine results, direct links from other blogs, et cetera – and have no immediate visual indicators of either age or activity. Although most blogs do include a timestamp, finding, interpreting and processing these timestamps is a conscious exercise – we were interested in exploring a subtler, visual indicator that conveyed age without the reader having to actively seek it out.

Independent of the visual cues that we wanted to explore with spice, we wanted bloggers to have control over the transience of their content. Prior work [28] had identified that the persistent nature of blogs had gotten bloggers in trouble with their friends, family and employers, and we wanted to explore mechanisms that gave bloggers better tools in these situations. The concept of an expiry date was a natural extension of the visual cues we had developed, and allowed bloggers to limit the exposure of content in a flexible way. Leveraging a system-wide default that could be overridden on a per-post basis enabled bloggers to set different expiration dates for different kinds of posts – perhaps a distant expiration date for a carefully crafted post,

and an expiration date two days away for a particularly vitriolic rant – in a lightweight way.

Activity

While associating contrast with age implicitly values recency, it does so in a manner that is inherently blind to usage patterns. Arguably, a blog post that is a year old but is still heavily visited is more relevant today than a blog post that is four months old but has only been visited a half-dozen times.

This motivated us to determine ways to encode activity into our visual representation of relevance. Although traditional hit-counters could have been used, these would have diluted our contrast-based metaphor of relevance. Since relevance is intrinsically relative, we developed our representation to increase contrast proportional to the *relative* activity on a specific post. This required us to build in a model for the blog’s activity over its lifetime so that an appropriate measure of relative activity could be efficiently calculated, the details of which are discussed in a later section.

This approach allowed us to extend the metaphor to behave in ways that the real-world analogue itself couldn’t – indeed, we would handicap the potential of online content if we limit it to behaviour that mirrors printed content.

4.1.3 Implementation

To encode age into contrast, the following ratio was first calculated:

$$ageRatio = \frac{currentDate - postDate}{expirationDate - postDate}$$

Similarly, for activity, the following ratio was calculated:

$$activityRatio = \frac{currentActivity}{2 * averageActivityOfAllPriorPosts}$$

Both ratios were clamped to the interval [0, 1].

The three colours – the background colour, the foreground colour, and the alternative foreground colour – define a two-dimensional space of colour that

can be navigated in many different ways. However, we found that the simplest approach also gave the best results – we interpreted each ratio as a co-ordinate along each of the foreground basis vectors, and merely averaged the resulting colours.

In order to be standards compliant, spice’s presentation layer is almost exclusively driven by Cascading Style Sheets. Although a separate CSS file could be generated for each post in the blog (to reflect its accurately interpolated foreground colour), this would significantly affect performance because of the inability of the browser to cache an almost entirely static CSS file. For these performance reasons, the CSS file was unchanged, with the colours computed on the server – the actual changing of the colour of the text was accomplished by minor post-specific code. This allowed us to take advantage of performance gains provided by existing infrastructure without sacrificing the cues we had chosen to use. This also has the additional benefit that with very little tweaking, the code can be adapted to work on other existing blog platforms as well.

4.1.4 Evaluation

Although the transience system was evaluated along with the rest of spice, we expected very preliminary feedback – it would be difficult for our users to provide significant anecdotal feedback in one month for a system with subtle effects for the first (by default) 18 months. However, we structured our design so that both bloggers and readers would be introduced to the concept and exposed to the system, and prepared our evaluation to be primarily conceptual – we tried to ensure that our users understood the concept of the system and how it behaved, and then we sought their opinion on the usefulness of the system. Additionally, the effects in the first month would be noticeable but not yet helpful, so we were somewhat prepared for some of our bloggers turning off the system (i.e., setting the default expiration date to ‘Infinity’).

When our users were first introduced to the concept, their responses were considerably more positive than we expected. All of our users mentioned that this was something that they thought would be useful and that would add value to existing blogs. Surprisingly, none of our users turned off the feature, although they were all given the option to do so.

Although we anticipated the benefits of the visual cues to be primarily for the reader, some bloggers felt that it would benefit them as well by serving as a forcing function – they felt that the embarrassment of having a front page that was considerably low-contrast and faded out would force them to post more frequently.

Additionally, some noted that the benefit of this system was best exhibited when scanning through archives – it became a lot easier to skip over old regular, day-to-day entries because they had faded out, and consequently, the old heavily-visited entries stood out much more strongly, because of their increased contrast relative to their temporal peers.

The biggest criticism we received was in the choice of terminology that we used – by associating each post with an ‘expiration date’, a few of our users were initially concerned that on that date, their posts would be deleted or somehow disappear. We are exploring other metaphors that we can use that are not quite so terminal – our next iteration is employing a ‘retirement’ metaphor.

Another drawback of our implementation was caused by the fact that we relied on a client-side script to change the foreground colour to a correctly interpolated colour. This had the unfortunate side-effect that on slow network connections, readers would be able to see the post and would start reading it while the page finished downloading, only to have the page finish downloading and then have the browser execute the script – consequently, the colour of the text that the reader was reading would sometimes change while they were reading it, providing for quite an unnerving experience. We are currently exploring alternative means to achieve the same effect.

4.1.5 Future Refinements

Apart from the refinements listed above, there are several other ideas we are currently exploring that could shape the future direction of this work:

- *Piecewise suppression* – This is a completely different approach from the contrast metaphor we’ve taken, and instead of merely providing cues, this approach actively makes it harder to access content over time. In this approach, as a post gets older, its content is incrementally suppressed – after three months, you will no longer be able to read the comments made on the post, after nine months, you can only read the

title of the post and and after a year you can only realize that a post was made on that day, and not be privy to any of its contents.

- *Noun substitution* – An idea we’ve toyed with and prototyped is the substitution of nouns as time passes. So a rant today might be very specific, but a year from now, names in the rant may have been replaced by other, more generic, names. Consequently, an angry post written today will hopefully not come back to haunt you years later when it is no longer relevant. Doing this consistently (by ensuring that Jennifer always gets replaced by Jane and Susan always gets replaced by Mary, so that stories that span posts still make sense) would result in automating a process that many bloggers already participate in manually.
- *Semantic spaces* – A similar approach that is somewhat trickier to implement, and with less clear implications, is doing something similar in a semantic space, instead of with nouns. Over time, sentences would be reworded to be increasingly vaguer and more semantically ambiguous over time. For example, ‘I had lunch with John – the roommate I hate – today’ would become ‘I had lunch with John – the roommate I don’t like – today’ after three months, and ‘I had lunch with someone I don’t like today’ after nine months. This would require very sophisticated linguistic techniques, and it is not clear how users would react to their posts being reworded. Nonetheless, this approach best mimics what people do in real life when recounting old stories.

4.2 Privacy

Although research concerning blogs is still relatively preliminary, work by Viegas, Nardi, Gumbrecht and others have identified privacy as a significant concern for bloggers [13, 28]. With the notable exception of LiveJournal, access-control and privacy mechanisms built into most blog platforms are extremely simplistic. WordPress, one of the more commonly used blog platforms on the Internet, allows its bloggers to give each post a public or private status only, thereby forcing its users to divide their readers into binary trust zones. Trust, however, is a complex function that is both fluid and multi-dimensional [10, 27]. The Collections photo-sharing system, which has been around for more than half a decade, recognized and leveraged the multiplicity of these social contexts that users find themselves in, resulting in a much more sophisticated access-control system [27].

While considerable innovation has taken place in other aspects of the blogging experience, privacy and access-control mechanisms have been largely ignored, despite significant indication in both academia and the media that these are of urgent concern to bloggers [13, 28]. We designed the privacy features of spice to be relatively simple changes to technology that would provide for significant changes from a blogger’s perspective and would result in very flexible usage in a wide variety of scenarios. These features are not considered to be mutually exclusive with systems like Collections – indeed, we feel that they are complementary and that careful integration could add value to both systems. Our intention is merely to illustrate that reasonable access-control systems are not technically infeasible, and that with a little thought, considerable sophistication can be attained – which will directly impact the further adoption of blogs.

4.2.1 Design Concept

WordPress, the underlying system upon which spice is based, allows bloggers to give a post either a draft status or one of two ‘finished’ statuses – public or private.

In spice, bloggers are provided with several more shades of privacy. Absolute privacy, which is what the traditional interpretation of ‘private’ posts in WordPress implies, is maintained and renamed to ‘Only I can see this post’.




Figure 4.3: Existing blogs allow relatively simple privacy options

Other options available to the blogger include ‘Readers who have been reading since (date) can read this post’, ‘Readers who have been reading as often as (once/twice/thrice) a (week/month/year) can read this post’ and ‘Readers who have read post number (number) can read this post’.

On the reader’s end, this system is completely transparent – for all intents and purposes, a post that has been made some form of ‘private’ does not exist for them, unless they meet the criterion specified by the blogger. It is important to note that the reader must satisfy the requirements put forward by the blogger *at the time that the post is made* – the reasons for this are explained in detail in a later section.

4.2.2 Design Motivators

Behaviour versus Identity

These options allow bloggers to indicate trust by describing the *behaviour* of the readers they trust, instead of by describing the *identity* of the readers they trust (in the form of usernames or the like). This is a crucial difference because, it is not uncommon for bloggers to only be able to identify a few core members of their audience (who return frequently and comment) [28]. It is conceivable that a blogger would be inclined to trust someone whom they can not identify, but who has been reading their blog semi-regularly for a number of years, more than they trust someone whom they can identify but has only been reading for a week (such as a prospective employer). Since in most online social settings, lurkers constitute the vast majority [23, 25], this type of access-control allows the blogger to interact with, and indicate their trust of, a significantly larger segment of their audience.

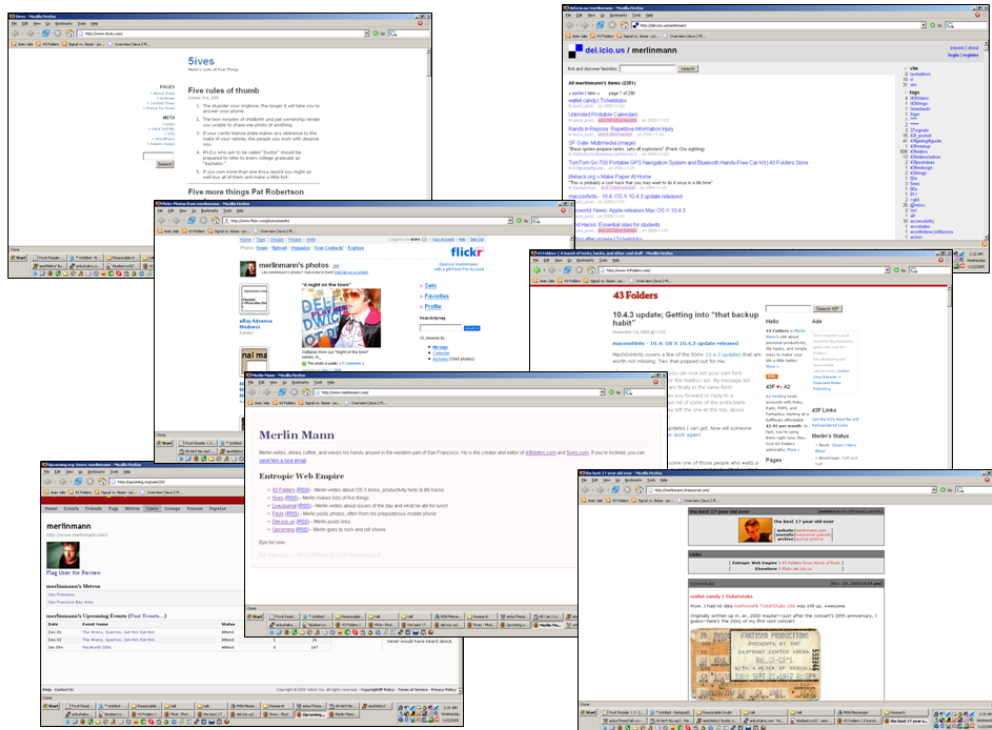


Figure 4.4: Some people, such as the pictured Merlin Mann, maintain multiple online presences for different aspects of their life

Multiplicity of Social Contexts

One person, despite the singularity of their existence, is usually the composition of very many different personas. For example, I am simultaneously a researcher, a food enthusiast, a jazz lover, a brother, a son, a student, a teacher, and very many other things besides. While it is unlikely that I would wish to blog about all these sides to my personality, it is true that people exist in multiple social contexts and it is not uncommon for people to maintain separate blogs for the separate hats that they wear in public. Although little formal research has been done to identify how common this practice is, much anecdotal evidence has surfaced, even in prior research work [27].

Although technically, it is possible to blog about these separate roles in one place, and separate them by using categories or tags, such an approach is much more likely to cause ‘spillage’ between social contexts. Sometimes this is acceptable (my boss finds out about my fascination with Portugese jazz), but it is just as likely to be catastrophic (my boss finds out about my

Advanced

Post Status: Published
 Draft
 Private

Only I can see this post

Readers who have been reading since can read this post.

Readers who have been reading as often as a can read this post.

Readers who have read post number can read this post.

Send trackbacks to:

Separate multiple URIs with spaces

Post slug:

Post author:

Edit timestamp

Figure 4.5: The privacy options in spice allow bloggers to interact with their audience based off of their behaviour, instead of their identity

rant about the incompetence of management at work) [1].

Interaction Audiences

Although the privacy options provided to the blogger in spice are still very simple, they allow the blogger to segment their audience temporally. This is important because it gives bloggers a means of grouping together readers that are likely to be similar to each other in terms of their relationship to the blogger.

For instance, one would be able to use the temporal segmenting to group their audience into people who started reading while they were an undergrad (likely to be people they met during their undergrad years or earlier), people who started reading while they were a graduate student (likely people they met while in grad school), and people who started reading after they graduated (likely people they met at work).

Much more powerful, however, is the final option – it accepts readers who have read a specific previous post. This lets bloggers build together complex and interesting interactions, examples of which include:

- If I post an update to a previous post, I can ensure that only people who have the proper context can be shown the update
- If I make a post criticizing someone or something, I can limit its access to people who have read previous positive posts, to ensure that only people who know that I'm not usually like this see my negative posts
- If I blog about my experiences at a specific restaurant, I can make a subsequent post organizing another trip to the same place

- If I blog about a dinner party I am hosting, a subsequent post linked to it can include directions or a change of plans.
- If I make a somewhat controversial post that generates a heated discussion, I can create another post linked to it to provide a summary or a conclusion (or a retraction).
- If I link a new post to an old already-expired post, I can ensure that only people who have been reading for a considerable amount of time can read my new post.

This kind of interaction radically changes the blogger-reader interaction from one that is linear and homogeneous to one that is non-linear and tailored to the interests of the reader. The reader effectively self-selects (by choosing to read or not read posts) themselves into a set of groups that reflects their interests, and since the blogger tailors content to each of these groups, the aggregation of the reader's interests results in tailored content for the reader.

Reader's Privacy

Unfortunately, the ability to provide additional privacy for bloggers requires sacrificing some of the privacy of the readers – to be able to tailor content to a reader, we have to have some way of keeping track of who the reader is. Although we can certainly (and do) take every possible step to ensure that we merely keep track of behaviour and not identity, we do need to be able to string together multiple visits to a blog by the same person.

Realizing that this is a significant departure from how readers *perceive* their visits to blogs, we felt the need to require the readers to participate in some way before they were tracked by our system. In this way, readers would have the option of completely preserving their anonymity, if they chose to do so. Although the participation level required is trivial to tweak, it is currently set to 'commenting' – a reader has to make atleast one comment before they are included in the tracking system.

This has the fortunate side-effect of filtering out webcrawlers and one-off visitors that would have otherwise inundated our tracking system, although at the cost of depriving some otherwise-willing lurkers the ability to take advantage of the system.

4.2.3 Implementation

When a reader first makes a comment on a blog, most blog platforms typically set a cookie on the commenter’s system, with basic information such as the commenter’s name and email address. As a result of this, the commenter does not have to fill that information out the next time they want to leave a comment on that blog. In spice, the cookie also contains a reader identification number that is generated on the fly on the server. The reader identification number is essentially a random number with some tweaks to make for more efficient lookups, and depends solely on the order of issuance (as opposed to on the reader the number is being issued to). All subsequent activity on the blog by that reader is tracked through the reader identification number.

When a blogger decides to make a specific post private to some section of their audience, spice interprets the rule that the blogger has set and determines which of the readers currently in the system satisfy the rule. It then generates a list of these reader identification numbers, and stores this list along with a representation of the rule in the database with the post. The permissible readers need to be determined when the post is made because otherwise it becomes possible for the reader to game the system – a reader could write a script to sequentially read every post on the blog after a protected post has been made, thereby giving them access to the protected post. With the approach we’ve taken, the reader has to decide to do this before there is any protected post to speak of – by the time a protected post is made, it is too late for the reader to do anything to gain access to it.

When spice retrieves the list of posts to display to a visitor, it first checks if the visitor has a reader identification number set. It skips over all protected and private posts, unless the protected post allows the specific reader identification number (if any) that spice found in the visitor’s cookie.

4.2.4 Evaluation

Although some of our users had blogged on LiveJournal and Xanga before, both of which allow for some amount of access-control, most of our users were quite unprepared for a privacy system of this type. Although our users did express an interest in using the system, and came up with some particularly interesting ways to use it, usage throughout the month-long study remained

low. While a large part of this can be attributed to the fact that meaningful temporal segmentation of a blogger’s audience requires considerably longer than a month, it seems that the users also had a hard time internalizing the capabilities of a system as different from what they were used to as this was.

This limitation is understandable – the model of privacy employed in spice is considerably different from anything in common usage today, and is in direct contradiction to the prevailing norm of identifying trusted people by username. Indeed, the examples that our users came up with to describe how they would use this system tended to use this new machinery in much the same way as they had used the old. For example, one of our users suggested making a temporary post, telling all her closest friends to go see it, and then taking it down. All subsequent posts that she wanted only her closest friends to see would be linked to that post, providing for a seamless segmentation without the need for accounts and usernames that are disruptive to the experience of engaging with a blog. In theory, this could be repeated multiple times with multiple groups of friends as well, but despite her excitement in describing this to us, neither she nor any other user actually used the system in this way.

Although encouraging, and with the undoubtable value of doing away with needless accounts and usernames for readers, it is nonetheless true that this user still appropriated this system in a manner that allowed her to engage only with people she could identify. However, we feel that over time, as users become more comfortable with this alternative way of thinking about their audience, considerably more novel and interesting interactions will arise.

4.2.5 Future Refinements

Although this system does provide users with considerably more flexibility in access-control than the norm, there are very many things that can be done to make it more powerful still. Some of the things we’re currently evaluating include:

- *Better Feedback* – For the blogger, this would ideally consist of a live display near the post status field, so that the blogger can see how the rule they choose and the parameters that they choose for it impact the size of the audience that is privy to that specific post. Something along

the lines of ‘This rule gives access to **27** people, who have been reading for an average of **4.3** months’.

- *Better Tracking* – Currently, our tracking system is susceptible to all the same concerns that cookies are susceptible to, such as the same reader using multiple browsers or even multiple computers to access the blog. Using a more robust tracking mechanism, such as employing a user-agent string, an IP address and a cookie as a tuple for identification, and considering two out of three matches as a successful match, would give better results. Although we have not yet implemented this level of functionality, the tracking problem is well-studied in industry and quite sophisticated techniques have evolved.
- *More Complex Rules* – Some of our readers suggested that we incorporate more complex rules into the privacy screening. While certainly feasible, the trade-offs involved include designing an interface that lets bloggers express these rules in a way that is not computationally ambiguous, as well as ensuring that doing the necessary tracking is tractable. One rule that has been often requested is the ability to use IP addresses as screening mechanisms – while they are unreliable as unique identifiers, they tend to be quite effective in blocking out whole geographic regions – thereby enabling bloggers to talk about experiences in different cities without having to censor themselves.

5 Evaluation

Although the motivation for spice lies primarily in concerns raised by the blogging community, it is nonetheless an exploratory project and certainly the first of its kind that we are aware of. Our design process was specifically guided both by interactions with bloggers and by prior preliminary work done in the field.

For these reasons, we feel that a study of spice would be incomplete without a meaningful evaluation. Although a token laboratory evaluation is certainly possible, the issues we were trying to address are not tied to specific interface concerns. Our goal was not to streamline the blogging experience by interface enhancements, it was to *radically redefine* the blogging experience to better understand the social effects of blogging as it exists today, and thereby better envision what blogging would ideally look like tomorrow from a social computer-mediated communication perspective.

For this reason, we felt that the best approach would be to conduct a field study – although the specifics are described in the following sections, we wanted to find real bloggers and ask them to use spice for an extended period of time. We wanted to reach as diverse a set of users as possible, and thereby reach as diverse a set of existing communities as possible, and then evaluate their reaction to the differences we'd introduced.

For our study, we recruited bloggers from around the world and with varying levels of expertise. Our bloggers represented various parts of North America, Europe, the Middle East and South Asia; and while some of our bloggers had multiple blogs with around three hundred regular daily readers, some had never blogged before. Our bloggers included students, journalists, gallery managers, software engineers, architects, production designers and even several people who weren't sure how to define what they did. In every manner, we tried to maximise diversity and to encourage differences, with one notable exception – although we did encourage different styles of blogs, we tried to restrict our users to personal blogs. Although it is likely that

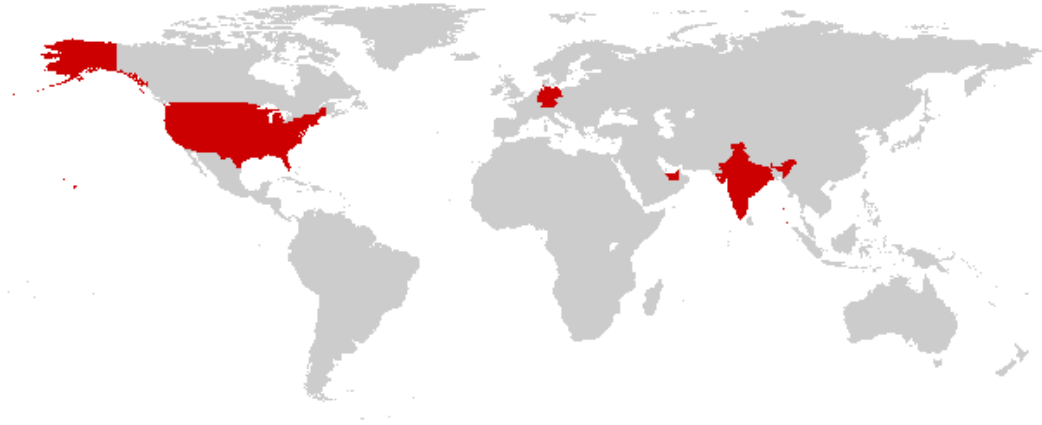


Figure 5.1: Countries represented by the users in our study

some of these ideas would apply equally well to news sites or to corporate blogs, their social dynamics are unique and our design was not intended to be appropriate for them. Within the realm of personal blogs, however, maximal variability was encouraged – indeed, photo blogs, work-related blogs, private blogs and even anonymous blogs were all represented in our sample.

5.1 Design

For our study, both existing and potential bloggers were recruited. This recruitment was extensive and indiscriminate – direct emails, messages to mailing lists and public signs were all employed in an attempt to recruit widely. Each of these initial respondents was asked to fill out a basic pre-experiment questionnaire that asked for some basic demographic question and tried to ascertain the respondent’s blogging experience.

Of these respondents, we selected ten to move forward with. These respondents (henceforth referred to as *users*, to distinguish them from the larger set of respondents), were chosen with two primary goals in mind:

- To maximize variability and diversity in the sample, and to avoid too many ‘similar’ users
- To best mirror documented reports on the distribution of the blogging world. Although we recognize that some of these reports are known




Figure 5.2: States represented by the users in our study

to be skewed, as an artefact of the sampling process, we did feel that there was value in ensuring that the distribution of our sample did not vary too widely from prior reports.

- To ensure that the natural blogging frequency of the blogger, when it could be determined, would allow for a reasonable amount of data to be collected over a one month period

Each of the users were then contacted with instructions on moving forward. Each user was asked to select a domain name (which would be registered in their name for a period of one year, and would count as their remuneration for the study), as well as to setup an online appointment with the researchers to discuss the technical and design aspects of their custom spice installation. At this point, the researchers went through and set up custom installations for each of the users, and then went over the features of spice and the functionality of the platform with each of the users, addressing any concerns and questions they may have had.

For the next month, users were encouraged to blog as naturally as possible. At periodic intervals, researchers would conduct diagnostic checks to ensure the installations were working smoothly. At all points, users were encouraged to communicate any difficulties or desired functionality to the researchers. The researchers ensured that they checked-in with each user at least once a week.

At the end of the month, users were debriefed and their qualitative feedback on the different aspects of spice was solicited. As much as possible, researchers tried to separate concerns about the implementation from concerns about the underlying ideas. This feedback has generally been presented inline in the preceding chapters, although overall concerns are outlined in the following sections as well.

5.2 Demographics

5.2.1 Respondents

Of our 20 respondents, 6 were male and the average age was 23.05 (median: 22, oldest: 40). 3 of our 20 respondents had not actively blogged for an extended period of time before. Of the 17 who had previously blogged, 2 respondents maintained three or more blogs, 6 maintained two blogs, and 9 maintained a single blog.

5.2.2 Users

Although we had originally planned to move forward with ten of our twenty respondents, we decided to do the study with eleven users. Our eleventh user was not able to start the study in the same general timeframe as our other ten users, but seemed like they would provide a unique and interesting viewpoint, and we felt that the additional perspective would be well worth any logistical difficulties.

Of our 11 users, 3 were male and the average age was 22.36 (median: 22, oldest: 25). 3 of our 11 users had not actively blogged for an extended period of time before – thereby allowing us to test the effectiveness of spice amongst those who did not have a pre-conceived notion of what blogging entails. Of the 8 who had previously blogged, 1 user maintained three or more blogs, 3 maintained two blogs each, and 4 had just one blog. These 8 bloggers averaged almost a hundred regular readers (average: 84.5), although the distribution is not normal and is heavily skewed towards two bloggers who individually accounted for 250-300 regular readers each. Without these two bloggers, the average lowers to a much more reasonable 21 regular readers for each of the remaining six bloggers.

5.3 Overall Feedback

Of the three features, the commenting system received the most discussion and consequently, the most feedback. The most visible change on the blog, it impacted everyone who visited a spice blog, whether they chose to leave a comment or not. All but one of our users seemed in general to support the idea, and were cognizant of the ways in which it added value to the discussions within the community. All of our users, however, also had varying levels of discomfort with the way that the idea was realized in code, and many suggested additional refinements. Needless to say, as an exploratory project, this is a tremendously positive result – our implementation was not ideal, but it was successful in the sense that it conveyed the underlying idea effectively enough for our users to suggest meaningful refinements. Consequently, our own thinking refined as the study progressed, enabling us to spot subtle conceptual flaws in our implementation that we have since been able to correct for the next iteration. The details of these refinements have been covered in the relevant chapter. The community that has had the most success with this feature, notably, were very vocal in communicating the shortcomings of the system to us (through the blogger) early on in the process, enabling us to refine the implementation multiple times for them within the first week of the study alone. The sole user who had a negative experience with the commenting system went to great lengths to discourage commenting on her blog to begin with, requiring her users to complete a multi-step process before they would be allowed to post a comment, and requiring it to be moderated even then. Although she finally did get rid of all these requirements, she did so three days before the study ended, and these facts taken together suggest that her unhappiness with the system be treated as an outlier.

Although some of our bloggers mentioned uses that they had for it as well, the fading system was mostly designed to be taken advantage of by the readers of a blog. Coupled with the fact that the effects of the fading system were designed to be subtle across short time periods and ‘normal’ traffic levels, we were prepared to receive little to no direct feedback on the implementation. Choosing to focus more directly on the idea, in both the introduction of the system to the users and the eventual debriefing sessions, revealed positive results – users indicated that this was something they had felt the need for, and that the approach we had taken seemed like a reason-

able and lightweight way to address that need. This feedback by our users validates both our work and prior reports that have indicated this to be a need felt by the blogging community. Admittedly, more study is needed in determining the implementation that best addresses this need.

The privacy features of spice also received a lot of attention, although usage remained low throughout the study. Although both existing bloggers and first-time bloggers seemed to understand the capabilities of the system equally well when it was being introduced to them, we do feel that the difference in *style* of privacy afforded by spice relative to existing systems contributed to its low usage. It is important to note that the prevailing style – that of making content private by username or group – is ubiquitous not just in other blogs, but also in existing photo-sharing websites and social networking websites. All of our users, however, expressed an interest in using access-control mechanisms of this sort in some manner, and some even suggested more intricate and complicated sets of parameters and queries to drive it.

5.4 Limitations

Like any evaluation, ours is not without several limitations. Several limitations, as well as other things to keep in mind while analyzing our results, include:

- *Limited userbase* – Although most of our bloggers had prior blogging experience, all of these prior bloggers had experience on managed systems such as Xanga, LiveJournal and Blogger. None of our respondents had experience with self-managed systems, which are of the sort that WordPress users tend to be. Consequently, the complexity necessitated by a self-managed system like WordPress (which is what spice is based on) adversely influenced user’s opinions about spice. Many of the complaints we received were rightly user-interface issues of the underlying system, and although we did our best to address those where possible, they were not related to the research we were doing, despite having an obvious impact on it.
- *Restricted userbase* – Some populations of bloggers were intentionally avoided – although news sites and corporate blogs are certainly valid

blogs and of interest to many in academia, our focus is on the social impact of blogs. The social dynamics of personal blogs are of primary interest to us, and it is with these blogs in mind that we designed spice. Consequently, we restricted ourselves to personal blogs for the evaluation of spice, and although some ideas may have wider applicability, our evaluation results should be considered indicative of personal blogs only.

- *Technical limitations* – Although every attempt was made to ensure a smooth blogging experience, each of our user’s blogs were hosted on a UIUC DCS TSG¹ server. These resources were very well-managed by the TSG staff, but spice did experience occasional downtime. Although relatively insignificant, our users were comparing us to their prior blogging experience – which, in the case of Blogger users, was powered by Google’s servers. Many of the requests we received in the early stages of the study were related to design and technical support for extremely esoteric cases (such as WAP support for spice) that we, as researchers, were simply not equipped to provide. Unfortunately, as an artefact of the way the study was designed and the blogging landscape, we – as researchers – were being compared to commercial offerings.
- *Too realistic bloggers* – One of our bloggers abandoned her blog prior to the completion of the study. Although this has some impact on the study, this is a very common practice in the larger blogging world, and was expected to some degree. Indeed, abandoned blogs are extremely common and are frequently a concern amongst commercial providers of blogging services. Future field work with blogs should take this form of user attrition into account when designing their study.

5.5 Feedback from Focus Groups and Poster Sessions

Although our formal study provided for several key insights, we also presented our work both internally and externally, and received valuable feedback from

¹University of Illinois at Urbana-Champaign Department of Computer Science Technical Support Group

High school class	Frequency	Age	Frequency
Freshmen	8	14	4
Sophomore	11	15	17
Junior	8	16	3
Senior	3	17	5
		18	1

Table 5.1: The composition of the focus groups, drawn from six high schools from the state of Illinois

these sources as well. Two specific sessions that were particularly fruitful are detailed in the following subsections.

5.5.1 High School Focus Groups

Focus groups were held with thirty teenage girls from six high schools in the Illinois region. Their ages ranged from 14 to 18 years old, and freshmen to seniors were represented.

Each group had roughly ten teenage girls from across multiple high schools, as well as various college-level counsellors. Each group was introduced to spice and the underlying goals behind it, and the commenting system was described. Each member of the group was then given fifteen minutes to play with the commenting system on a demo spice blog, while a discussion simultaneously took place. Although a quick overview of the other features of spice was also provided, due to logistical and resource constraints, the discussions were primarily focused on the commenting system.

Roughly 80% of the girls presented to described themselves as active bloggers. Although platforms varied, most used either MySpace, Xanga or LiveJournal. Of the non-bloggers, more than half indicated that their reasons for not blogging primarily included privacy concerns. The perception that they lacked anything interesting to say was also often cited.

The responses we received were overwhelmingly positive. Even some of the college-level counsellors expressed a strong interest in using a spice blog. Although a few of the girls (< 5) had initial trouble understanding how to use the system, once they saw others do it they picked it up quickly and had fun with it. We had multiple requests to ‘build one for Xanga’, and very many more requests to give people custom spice blogs. Indeed, weeks after the focus groups were held, we were still receiving email requests from some



Figure 5.3: The poster presented at IDCSA 2006

of the girls asking for spice blogs, even on our private email addresses!

5.5.2 IDCSA 2006

Along with 84 other projects, spice was presented as a poster at IDCSA, Illinois in April 2006. The audience consisted of a cross-section of industrial affiliates, most of whom were not bloggers themselves. Nonetheless, all but one of the people presented to (out of over a hundred) personally knew at least one blogger and were familiar with the concept of blogging. Unlike the high school focus groups, the presentation at IDCSA equally covered all sides of spice, including the evaluation. Additionally, although this was not originally planned for the poster session, a demo of the commenting system was arranged for some participants upon request.

The ideas behind spice were warmly received, and several people requested their own spice blogs. Although the audience being presented to was incredibly eclectic – including thespians, managers, dancers, executives and programmers alike – the discussions were insightful with participants

relating the ideas behind spice to concepts they were already familiar with, despite the diversity of the sample. Indeed, the level of support for spice was perhaps best demonstrated by its being awarded the Best Poster Award, through a combination of metrics that reflected the opinions of both the judges and the general audience.

6 Key Lessons and Conclusion

Over the course of its design, implementation and evaluation, spice has been an interesting project. From the very start, spice was not meant to be the definitive answer to problems other researchers had identified with blogs; our goal has always been to use spice as an exploratory tool to begin discussions in areas that had been under represented in the development of blogs, and as a way to experiment within the design space.

During the project, we learnt a lot about the nature of bloggers and their communities. In many ways, however, what we learnt about the challenges uniquely associated with doing an evaluation of this sort has had a larger impact in the direction we have chosen to take for spice. Indeed, we feel that these contributions of the work are complementary to and as important as the specific social questions that spice raises.

6.1 Lessons learnt about alternative models of social interaction

The evaluation of the mechanisms we had built to foster community-driven content was overwhelmingly positive. Although there was one exception, both bloggers and their communities reacted positively to it.

Some of the comments left on our user's blogs by their communities include:

- 'This new blogging tool is certainly innovative!!! Cool commenting features :D'
- 'Ok, I don't have anything in particular to say – just love this 'post-it' concept!'
- 'what a sexy commenting system. i feel like commenting all day!'
- 'THIS IS A GREAT IDEA! I don't like to blog but this is awesome!'

- ‘The commenting system is screwed up. Big time. Please change it before I go mad!!!’
- ‘This is incredible.’

From our discussions with our bloggers, the following aspects of the commenting system were commonly cited as key to its success:

- *Lightweight and simple* – the work involved in commenting and annotating was minimal, with only one additional step from existing systems, and the rendering was subtle and didn’t detract from the content. One user summarized it as ‘..I’m loving it! I think the idea is really good – just takes people a while to get used to it, but it is so good afterwards... you’d want it on all blogs. Just makes more sense.’
- *Flexible and expressive* – the system could be appropriated in different ways to do different things. For example, one of our user’s commenters left comments next to images, providing her own humorous captions for each of the images to counter the blogger’s captions. Another user’s commenters indulged in a game of hide-and-seek, trying to find locations in posted images that were of a similar shade as the comment-shadow. Although it didn’t actually happen, one of our bloggers mentioned that they expected their commenters to start using comment-shadows as a means of ‘covering up’ controversial parts.

There were a few suggestions for improvement as well. The most notable ones mentioned were:

- *Intelligent defaults* – as implemented, the system would assume a default posting location for all comments that didn’t explicitly pick one. This had the negative effect that some comments ended up being rendered on top of one another, thereby effectively drowning out the older comment. More than 75% of the suggestions for improvement concerned this issue, although the specific suggestions varied – some felt that it would be better to pick more intelligent defaults, others felt that it was more appropriate to force commenters to always explicitly pick a location.

- *Non-topical comments* – bloggers felt the need for a means for commenters to leave comments without relating it to a specific part of the post. This need was anticipated for more general comments that merely indicated support or were only tangentially related to the contents of the post. Even though only two of our bloggers and none of their commenters specifically asked for this, when prompted, other bloggers expressed that they felt this would add value.
- *Loss of time* – the chosen representation of comments loses the encoding of time, which made it challenging to determine in what order comments got added. Although this was not asked for by any of our bloggers, from observing several comments threads, we felt that encoding time in some manner would add value to spice. It is interesting to note that in the absence of any such system, the commenters on some of the blogs developed their own conventions to encode time – in one blog, each comment in a specific area was posted higher up on the page and closer to the right edge of the page than all the comments in the same area that preceded it.

The improvements we made to empower the blogger, on the other hand, were harder to evaluate due to their inherent complexity and realm of influence. Although, they were very useful in helping us identify difficulties in traditional evaluation methodologies for blogs, there were nonetheless a few interesting takeaways:

- A model of transience does address some of the concerns of bloggers about the persistence of their content. Only one user discussed turning this off, and none actually did. At least two users took repeated advantage of the visual contrast metaphor to find heavily-visited posts. Both bloggers and their communities did ask for clearer expectations about the impact of this transience model in the context of webcaches and search engines.
- Access control mechanisms that are seamless and behaviour-based appeal to users. When introduced to it, almost all of our users immediately understood and extended the model by asking for more complex behaviour-based rules than we were able to support. Nonetheless, none

of our users actually used this system. At this point, it is difficult to ascertain whether or not this is due to the short nature of the study and the fundamental disconnect between this approach and more traditional forms of access-control.

6.2 Lessons learnt about the difficulties in evaluating a blogging platform

While conducting our evaluation, we ran into a few unique challenges that we weren't expecting. We've documented these below so that other researchers can keep them in mind earlier in the process when they design similar studies:

- *Continuity* – Bloggers value their blog for its ability to archive thoughts as much as they do for the ease of expression it provides. Temporarily using another system impacts the continuity in their archives, and affects their willingness to participate.
- *Attrition Rate* – The larger blog community has a remarkably high attrition rate, sometimes estimated to be as high as 66%, depending on the sample [15]. This reduces the likelihood that most of the users of a blog study will actually complete the study, which increases the number of users needed.
- *Natural posting frequencies* – Different bloggers post at different frequencies, ranging from multiple times a day to a few times a year. Accounting for these differences makes the design of a study considerably more complex.
- *Risk to the blogger* – By participating in a study for an extended period of time, bloggers risk alienating members of their community that don't react well to the experimental treatment.
- *Long-term impact* – Some enhancements, such as the transience model built in spice, or some aspects of the access control mechanisms in spice, require considerable periods of time to manifest themselves.
- *Sophisticated alternatives* – The prevalence of blogs in the mainstream means that there are a considerable number of well-built and free blog-

ging services. By constantly increasing blogger’s expectations, these services make it difficult to provide an experimental blog platform that can compete with the feature set and release cycle of a commercial product.

6.3 Conclusion

Through a review of prior work, we identified key shortcomings in blogs that limited their abilities for both bloggers and communities. The limitations of comments constrained the kinds of interactions that the community could participate in, both with each other and with the blogger. Additionally, the characteristics of the medium and the limitations of access-control mechanisms limited the freedom of the bloggers to express themselves.

We then developed a collection of features that addressed these needs of both bloggers and their communities, and incorporated them into a custom-built blogging engine. By developing a novel commenting system, we hoped to address some of the characteristics of comments that limited their ability to foster community-driven content. By building in a transience model, we tried to design around the persistent nature of blogs that prior work had indicated was counter-intuitive to bloggers. Our behaviour-based access control was designed to allow bloggers to interact meaningfully with members of their audience that they couldn’t explicitly name, which prior work has shown to be the vast majority of their audience.

We evaluated these alternative interaction models with a diverse set of users around the world, and solicited their feedback as well as that of their communities. We used this feedback to evaluate our design and to suggest directions for future work. We also evaluated the system with focus groups drawn from local high schools and at poster sessions.

Our study revealed a few lessons about the models of interaction that we had developed. Over the course of the study, we also determined a set of challenges in conducting evaluations of this sort that may aid future researchers in the area. The qualitative feedback we received about the transience model and access-control mechanisms we’d built for bloggers was promising, but of a preliminary nature due to the inherently long-term, subtle effects of these features. Initial feedback, however, suggests that these avenues merit further

exploration. Although there were a few suggestions for refinement, the novel commenting system that we proposed met with considerable positive feedback from both bloggers and their communities, and met our goal of better fostering community-driven content.

References

- [1] Heather Armstrong. Dooce.com [<http://www.dooce.com/>].
- [2] Ziv Bar-Yossef, Andrei Z. Broder, Ravi Kumar, and Andrew Tomkins. Sic transit gloria telae: towards an understanding of the web's decay. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 328–337, New York, NY, USA, 2004. ACM Press.
- [3] Paul Boutin. Live coverage of bill gates ces keynote [<http://www.engadget.com/2006/01/04/live-coverage-of-bill-gates-ces-keynote/>].
- [4] A. J. Bernheim Brush, David Barger, Anoop Gupta, and J. J. Cadiz. Robust annotation positioning in digital documents. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 285–292, 2001.
- [5] Deadly Katrina Contributors. Deadly katrina: Tracking the aftermath of hurricane katrina [<http://www.deadlykatrina.com>].
- [6] Various Contributors. Slashdot.org [<http://www.slashdot.org/>].
- [7] Wikipedia Contributors. Blogs [<http://en.wikipedia.org/wiki/blogs>].
- [8] Wikipedia Contributors. Deja news [<http://en.wikipedia.org/wiki/dejanews>].
- [9] Alireza Doostdar. "the vulgar spirit of blogging": On language, culture, and power in persian weblogestan. *American Anthropologist*, 106(4):651–662, 2004.
- [10] Erving Goffman. *The Presentation of Self in Everyday Life*. Anchor, June 1959.
- [11] Scott A. Golder. Webbed footnotes: Collaborative annotation on the web. Master's thesis, Massachusetts Institute of Technology, 2005.
- [12] Google.com. Official google blog [<http://googleblog.blogspot.com/>].
- [13] Michelle Gumbrecht. Blogs as 'protected space'. In *WWW 2004 Workshop on the Weblogging Ecosystem: Aggregation, Analysis and Dynamics*, New York, NY, USA, May 2004.

- [14] Michael Hauben and Ronda Hauben. *Netizens: On the History and Impact of Usenet and the Internet (Perspectives)*. Wiley-IEEE Computer Society Pr, 1997.
- [15] Jeffrey Henning. Perseus blog survey: The blogging iceberg – of 4.12 million hosted weblogs, most little seen, quickly abandoned [<http://www.perseus.com/blogsurvey/thebloggingiceberg.html>].
- [16] Susan C. Herring, Lois Ann Scheidt, Sabrina Bonus, and Elijah Wright. Bridging the gap: A genre analysis of weblogs. In *HICSS '04: Proceedings of the Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) - Track 4*, page 40101.2, Washington, DC, USA, 2004. IEEE Computer Society.
- [17] William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. Edit wear and read wear. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 3–9, New York, NY, USA, 1992. ACM Press.
- [18] Weblogs Inc. Engadget.com [<http://www.engadget.com/>].
- [19] Kevin Kelly. *Out of Control*. Addison-Wesley, Reading, MA, 1994.
- [20] Gawker Media. Gizmodo.com [<http://www.gizmodo.com/>].
- [21] Matthew Mullenweg. Wordpress 2 [<http://wordpress.org/development/2005/12/wp2/>].
- [22] Jakob Nielsen. How users read on the web [<http://www.useit.com/alertbox/9710a.html>].
- [23] Blair Nonnecke and Jenny Preece. Lurker demographics: counting the silent. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 73–80, New York, NY, USA, 2000. ACM Press.
- [24] John Oates. Queen of the sky gets marching orders [http://www.theregister.co.uk/2004/11/03/airline_blogger_sacked/].
- [25] Elizabeth Reid. Electropolis: Communication and community on internet relay chat. Master's thesis, University of Melbourne, 1992.
- [26] Gaurav Sabnis. Vantage point [<http://gauravsabnis.blogspot.com/2005/10/update.html>].
- [27] F. B. Viégas. Collections - adapting the display of personal objects for different audiences. Master's thesis, Massachusetts Institute of Technology, 2000.

- [28] F. B. Viégas. Bloggers' expectations of privacy and accountability: An initial survey. *Journal of Computer-Mediated Communication*, 10(3), 2005.
- [29] Melissa Wall. blogs of war: weblogs as news. *Journalism: Theory, Practice & Criticism*, 6(2):153–72, 2005.
- [30] Alan Wexelblat and Pattie Maes. Footprints: history-rich tools for information foraging. In *CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 270–277, New York, NY, USA, 1999. ACM Press.
- [31] Dave Winer. What makes a weblog a weblog? [<http://blogs.law.harvard.edu/whatmakesaweblogaweblog>].
- [32] Polle T. Zellweger, Susan Harkness Regli, Jock D. Mackinlay, and Bay-Wei Chang. The impact of fluid documents on reading and browsing: an observational study. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 249–256, New York, NY, USA, 2000. ACM Press.