

In Praise of Inefficiency

Notes from the Ethnography Underground

Marijke Rijsberman
November 19, 2000

This case study of the development of a collaboration tool celebrates inefficiency at two levels:

- organizational inefficiency in the development of internal systems
- inefficiency in the internal systems themselves

This is not to argue that inefficiency is always—or even ever—a value to strive for. But inefficiencies can be associated with significant invisible benefits. Because of its recognition of complexity in human social activities and structures, because of its commitment to holism, ethnography is crucial in uncovering such incidental benefits—*if* our methodologies, rapid or otherwise, do in fact bring out the full range of complexity and interconnectedness of human experience. This is the first challenge.

On the other hand, your users are right there, and the general inefficiency of the development cycle in many IT departments¹ means that you may have six months rather than two weeks to learn about them. Of course, you're unlikely to have a budget for ethnography—the second challenge. Rather than the domain of an ethnographer, ethnography must become a “distributed” expertise in such environments, a skill to be developed by all team members and particularly project managers.

Both the challenges and the opportunities are illustrated in the unfortunate history of our collaboration tool.

Project Background

The project, code named Mercury, was carried out in the network consulting division of one of the largest technology companies in this country. In the early '90s, when the division was small and the number of engineers still very modest, the company had set up an email list that enabled the engineers to ask each other for help with

challenges they ran into at the client site. For purposes of knowledge reuse, all email was also kept in a searchable archive.

As the engineering group grew, the original list was carved up into different lists, each devoted to a different subject. Engineers were still obliged to receive every question, but they were given the option of not receiving the rest of the discussion for any or all of the smaller lists. Even so, by early 1999 the user community numbered about 3,000 engineers, and the amount of email generated was overwhelming both to them and to the company's Exchange servers.

A Cautionary Tale

That's when the company undertook to build a new system that would offer users a choice of email, newsgroups, and the web to participate in the collaborative problem solving. Each subject would become a “forum,” and it would be up to the engineers to decide how to interact with it. The web interface was meant primarily for archive searches but the design was to permit participation in problem resolution as well. Unfortunately, the team badly underestimated the design challenges involved, as well as the development effort required. In fairly short order, a traditional requirements process was completed. The sum total of user research was an informal survey asking engineers how likely they would be to use newsgroups if the new system offered them.²

The developers began building a prototype in September 1999 but suspended the work when it was discovered that the news server interfered with the threading of messages in the archive. In late February 2000 development resumed. A different news server had been found. An estimated launch date was set for May 31. The “prototype” was two weeks away from being

¹ See Beyer and Holzblatt, *Contextual Design* (1998) pp. 33-36, for some of the reasons for the persistent difficulties that beset IT departments.

² As an interesting side note underlining the divide between saying and doing, 80% of the engineers said they would use newsgroups. It became apparent later that the majority had never used newsgroups and had only a vague sense of how they worked.

finished. That's when I came on board—not as an ethnographer, but as the hapless project manager.

The following months were a textbook example of the consequences of poor planning and a design process that ignores users. An elaborate review cycle involving the entire end user community resulted in the systematic dismissal of any feedback that didn't fit the client's preconceptions about design and functionality.³ Late-stage testing turned up serious design flaws. Launch had to be delayed five times until September 2000, for a total of three-and-a-half months. Now, three months later, there still are serious bugs that have not been fixed.

There are some very happy users, but they are in the minority, and there's been considerable grumbling. Many of the complaints are about broken functionality, particularly the fact that content is not identical across access channels, a state of affairs that undermines trust in the system. But there are also serious usability problems in the web interface. Heavy use of the option to send email to all system users, designed as a last-ditch resort, is creating unwanted email traffic and discontent among engineers.

Initial measurements indicate that the new tool is by no means a failure, but there is some cause for concern. There appears to be less activity in the new system than there was in the old, with fewer threads started each day.⁴ The number of responses per thread is also down, but that may indicate more effective problem resolution rather than less enthusiastic participation. Only about 2200 out of 3450 engineers, or 64%, have bothered to register to use the system—a disappointing rate considering that participation is more or less mandatory and is a significant factor in the engineers' performance reviews.

³ The most telling feedback was about usability problems in the web interface, which was found to be generally easy to learn and easy to use, except for the display of discussion threads. The client dismissed feedback with the argument that the engineers wouldn't use the website to participate in the discussion anyway. That's what the newsgroups were for. The survey encouraged this mindset—another example that bad data are worse than no data.

⁴ The data are tough to interpret, because there was considerable variation in the old system and because upheaval in the organization may also be causing (or contributing to) the decrease, which appears to be around 15%.

Ethnography to the Rescue

It's easy enough to conclude that ethnography would have made the difference in the design of this tool—after all, I've told the story so as to lead you there, carefully sticking to the “morphology” of the user research folk tale. But if by ethnography we mean a one- or two-week task-focused inquiry, an investigation into the nature of work à la Beyer and Holzblatt, then the answer is a resounding no.⁵ Sure, the web interface would most likely have been much better designed, and I'd hope that the fallacies of that horrible survey would have been exposed. But the most significant and fundamental issues—those to do with community—would most likely not have been identified.

So let me try to tell the story a little differently. The engineers are stationed at client sites, sometimes for weeks, months, or even years. They call themselves “road warriors,” a term that, in truth, applies to them even more than to salespeople: the engineers don't even have a spot in the local home office. They can be reached by pager only. They see each other for occasional trainings and for after-hours office parties. They are occasionally sent on an engagement outside their own territory and meet other engineers that way. The only events that pull engineers together from all over the world are new hire training and another training in their second year on the job.

Much as everyone hated the old system, it helped sustain a very strong sense of the larger community, allowing engineers to get to know each other and work together. It also provided some degree of “awareness”—a feeling for who's out there at any given time—in a primitive fashion.⁶ Not that anyone would have, or even *could* have, told you so if you had known to ask. This underlying reality was entirely “covered” by the tasks involved in asking for help with technical issues and offering it to others. It qualifies as a truly invisible incidental benefit of the inefficiency of the old system. And it became apparent only as a consequence of careful, long-term observation.

⁵ I choose Beyer and Holzblatt as a paradigmatic example because their work, which I hugely admire, is so often referred to as a model of ethnography.

⁶ On awareness, see for example Thomas Erickson and Wendy A. Kellogg, “Social Translucence: An Approach to Designing Systems that Support Social Processes,” at http://www.pliant.org/personal/Tom_Erickson/st_TOCHI.html.

The Social Economy

Finding #1: The lists were “self-policed.” The knowledge management group formulated a policy and depended on the entire community to enforce it. One thing that happened with some regularity was the accidental sending of personal email to the group.⁷ Someone would post a message and then another engineer would respond, “Hey dude. Haven’t heard from you in a long time. How ya doing?” The collaboration must have sparked personal exchanges far more often than we actually got to see.

Finding #2: The day the new system went live, traffic on a company-wide non-technical email list more than tripled. Topics on that list range from “Monday humor” to discussions about personal buying decisions, to griping about the stock price, to questions about lap top configurations. Traffic on that list has continued at the new high. It is tempting to conclude that the informal list meets the engineers’ need for a sense of community. However, not nearly all engineers are on that list. And because it is divorced from the technical collaboration tool, it does not have the same ability to support the technical discussion.

Finding #3: Overuse of the emergency feature that enables the sending of global email is motivated by a drive to reach particular people—whether these are people the engineer knows or people who have a certain characteristic (for example, engineers who have installed and configured mirrored Cisco LocalDirector routers to provide fault tolerance for ecommerce systems). In the new system, it is not clear who will see your posting and how soon if you post your message to one of the forums, so the global email feature is resorted to far more often than it was meant to be.

Initial conversations confirm that

- to *some* engineers the community aspect of the old system was enormously important
- the new system breaks that sense of the larger community
- the technical discussions are supported by the sense of community and the virtual relationships it fosters

⁷ There are technical solutions to this problem of course, but I’m not much interested in them. I’m just grateful the improperly addressed email was sent, giving me some insight into how the social and professional were intertwined in the community.

After all, time spent in collaboration is non-billable, and the willingness to help is based in part on the social relationships that the old system facilitated. Future enhancements must focus not only on the usability issues in the web interface, but, more importantly, find a way to support a sense of community, encourage interrelationship among individual engineers, and provide for a better sense of awareness within the forums.

Questioning Productivity

Much as I admire *Contextual Design*, the fact is that Beyer and Holzblatt accept the extremely constricted definition of productivity that has held sway in business environments since Frederick W. Taylor’s scientific management attempted to circumvent the full humanity of workers in the name of efficiency in production. There is more and more evidence that productivity is linked to social and emotional engagement.⁸ Collaboration is becoming more and more important to many forms of work, which also brings questions of meaning and relationships into the forefront.

The challenge is to achieve an understanding of the bearing of social and emotional experience on well-designed systems, not just in entertainment and consumer services, but in the work environment as well. Ethnography offers a perspective that can take us beyond task-based investigation to the entire range of experience of users. That means drawing on ethnography’s traditional attention to social and emotional factors in the construction of meaning. It means questioning task-focused ideals of productivity, taking a good look at inefficiencies to see what social and emotional economies they may hide. It means paying attention not only to patterns, but also to the points where they break down and become incoherent. As Renato Rosaldo remarks, “Although the classic vision of unique cultural patterns has proven merit, it also has serious limitations. It emphasizes shared patterns at the expense of processes of change and internal inconsistencies, conflicts, and contradictions. By defining culture as a set of shared meanings, classic norms of analysis make it difficult to study zones of difference within and between cultures.”⁹ Applicable not only to larger cultural

⁸ See for instance, Jon R. Katzenbach, *Peak Performance* (2000) on emotion and productivity.

⁹ Renato Rosaldo, *Culture and Truth: The Remaking of Social Analysis* (1989).

systems but also to the smaller cultures that dominate workplaces, this focus on inconsistency and conflict is invaluable in breaking through accepted stories about productivity and efficiency and other workplace myths that stand in the way of truly user-centered systems.

The Ethnography Underground

Finally, it is my conviction that applying ethnography to internal systems design also means applying at least some of the traditional ethnographer's patience. Although you can do a pretty good job on the exploration of tasks, it is unlikely that you can get a feel for the diversity of an end user community and for its social and emotional needs in one or two weeks. Combined with a widespread reluctance to fund ethnography studies for internal systems, this simple fact means that ethnography must become a distributed skill and may even have to go underground. My own experience on this project is instructive.

I was endlessly frustrated by the inefficiencies and immaturities of the process and the design failures in this project. But the fact is that, because of its protracted nature, I "lived with the natives" for about six months. I got to know many engineers, I had ready access to most relevant data, I had an opportunity to observe the collaboration, the formal and informal interactions. I worked with the committee of engineers and with the volunteer group. I facilitated discussions between the engineers and the client about the pros and cons, the whys and wherefores of the design. I did presentations at the new hire training meetings. I was subscribed to the old system and was able to see how users interacted with it and with each other first-hand over an extended period of time. I was subscribed to the informal, company-wide mailing list. I trained the Helpdesk on the new system design and learned about the connectivity challenges that engineers struggle with from them.

It is true that much of the knowledge I gained in the course of the project could not be put to use right away. But it is not lost. It will inform future enhancements at the least, and it may form the basis for further study. A proposal to study the tool and the community that uses it in greater

depth is currently under consideration. It is worth remembering that this is a client who would never have considered ethnography as a matter of principle, of basic methodology. The proposal is being taken seriously because my data and interpretations gave sufficient evidence that there is indeed more to worry about than the usual client-centered requirements.

Entire teams can be involved in paying attention, keeping records of anything odd or against expectation, and discussing their observations. But project managers are crucial, both in that they have the power to make it happen and in that their position facilitates a role as observer. While there were moments when it was tempting to take a partisan position, to think that I already knew the answer to a particular question, by and large the role of project manager is that of a facilitator, a mediator between different constituencies—a role that makes it possible to explore the structures and relationships that characterize the project environment and the end user community.

Conclusion

Organizational inefficiency in internal systems development gives iterative design a whole new meaning, basically requiring that you resign yourself to getting it wrong the first time around. But it also offers the opportunity to do an extraordinarily good job the second time around, building systems that are based on a deep understanding of the end user community and what helps them to do their work. It also offers the opportunity to do a great deal of careful if surreptitious fieldwork.