

Software Development at Microsoft Observed: It's about *people ... working together*

Gina Venolia	User Interface Architect	gina.venolia@microsoft.com
Rob DeLine	Researcher	rob.deline@microsoft.com
Thomas LaToza	Intern (<i>Summer 2005</i>)	tlatoya@cs.cmu.edu

Microsoft Research
Human Interactions in Programming team
<http://research.microsoft.com/hip/>

© 2005 Microsoft Corporation. All Rights Reserved.

ABSTRACT

To understand Microsoft developers' typical tools and work habits and their level of satisfaction with these, we performed two surveys and eleven interviews with developers across all business divisions. This report provides a summary of the resulting data. From the set of potential problems we gave them, the top three that Microsoft developers agree they have are: understanding the rationale behind a piece of code (66%); having to switch tasks often because of requests from teammates or managers (62%); and being aware of changes to code elsewhere that impact their own code (61%). The most notable take-away from the data is that developers go to great lengths to create and maintain rich mental models of code and don't rely on external representations. The mental nature of these models requires frequent, disruptive, face-to-face meetings to keep individuals' models in sync, which greatly slow the rate at which a newcomer to a team can become productive. These interruptions also burden more senior development team members, as they have to recover what they were doing in the code following the interruption from team members.

INTRODUCTION

The Human Interactions in Programming (HIP) team in Microsoft Research applies human-centered research techniques to build tools that improve the software development process. The joke goes, “we build tools as if software were made by *people ... working together.*”

As a human-centered effort, we draw from various research fields including human-computer interaction, information visualization, computer-supported cooperative work, and social computing. The central tenet of these fields is that one needs to understand the user in order to design tools to support them. To this end we have initiated a series of investigations to understand software development at Microsoft. Our research builds on a rich history of research into professional software development practices [1, 2, 4, 5, 6, and 7]. This document describes the process we used, what we learned, and directions for future user research.

METHODOLOGY

We performed two surveys and several face-to-face interviews of developers at Microsoft during the summer of 2005. The first survey contained 205 questions asking how developers spend their time, what tools they use, and the severity of various problems they face. We deployed it to 1000 architects, software developers and software test developers randomly selected from the Microsoft address book by job title. We received 157 responses, though the data presented here includes only the 104 responses from the developers.

Next we performed semi-structured interviews with six software development leads and five software developers drawn from our survey respondents. Each was done by two interviewers, who took copious notes. Each interview lasted about an hour. Most were recorded on audio. To find the themes latent in the notes we transcribed them onto ~800 3x5” cards and did a massive card sort exercise (see Figure 1).

The interviews and survey lead us to several preliminary hypotheses. We developed a second survey to test the hypotheses and investigate them more deeply. We deployed the 187-question follow-up survey to 1000 randomly-selected software developers and software development leads, excluding those who received the first survey. We received 187 responses.



Figure 1: We analyzed notes from the interviews using a card sort, some of which is captured in this photo.

OBSERVATIONS

Obviously we have a tremendous amount of data about software development at Microsoft. Rather than recount it in detail, this whitepaper presents the key observations that we drew from our research, which fall into three themes:

- Developers go to great lengths to create and maintain rich mental models of code and don't rely on external representations.
- Understanding the rationale behind code is the biggest problem for developers. When trying to understand a piece of code developers turn first to the code itself and, when that fails, to their social network.
- Developers and development managers use a variety of tools and work practices, and are actively looking for better solutions.

The next sections present details of and evidence for these observations.

Mental Models

Developers create and maintain intricate mental models of the code. Through our interviews we know that developers can talk spontaneously about the architecture, how it's implemented, who owns what parts, the history of the code, to-dos, wish-lists, and meta-information about the code. For the most part this knowledge is never written down, except transient forms such as sketches on a whiteboard. One of our interviewees summed it up well, saying, "Lots of design information is kept in peoples' heads."

Personal Code Ownership

Mental models are expensive to create and maintain. Developers have a strong notion of *personal code ownership*, which constrains the amount of code they have to understand in detail. In our follow-up survey, 77% of respondents agreed¹ with the statement, "There is a clear distinction between code that I own and the code owned by my teammates." (On the other hand some teams have a policy to avoid personal code ownership because it makes individuals too indispensable and promotes, in the words of one of our interviewees, "too much passion around the code.") Code ownership is a long-term proposition, reducing the number of times that a developer has to learn a new code base. In our first survey, the average time on the current code base was 2.6 years, with 32% reporting 6 years or more. Personal code ownership is usually tacit, i.e. part of the mental model; written records of ownership, when present, are often out-of-date and distrusted.

We received conflicting information about design documents for issues within a team. In the interviews design docs were described almost as write-only media, serving to structure the developer's thinking and as an artifact to design-review, but seldom read later and almost never kept up-to-date. On the other hand our follow-up survey respondents reported a different picture of design documents for issues within the team: their feature teams wrote an average of 8.0 documents in the prior year, and kept 51% of them up-to-date. We were surprised with these numbers and can't reconcile them with the results of the interviews.

Team Code Ownership and the "Moat"

Even stronger than personal code ownership is a notion of *team code ownership*. A whopping 92% agreed with the statement, "There is a clear distinction between the code my feature team owns and the code owned by other teams." (For this survey the term *feature team* was explicitly defined as "the core group of developers you work with.") Feature teams are small – an

¹ Throughout this document, the word *agree* means that the participant selected either "Somewhat agree", "Agree", or "Strongly agree" from a seven-point Likert scale.

overwhelming 93% stated that their feature team consisted of 2-4 people (including the respondent). There seems to be a sweet-spot at three-person feature teams, reported by 49%. Feature teams are almost always collocated, facilitating informal knowledge sharing.

One of the ways that developers keep their mental model of the team's code up-to-date is by subscribing to check-in messages by email, though several interviewees disliked the lack of detail provided by teammates.

Small feature teams and their strong code ownership form a kind of *moat* around the team, isolating them from outside perturbations. The moat is defined by design documents, which specify the interface across the moat. Design docs for cross-team issues were less common than those relevant to issues within the team. The average number of design docs written by the respondent's team in the last year for cross-team issues was 4.8, vs. 8.0 for within-team issues. On the other hand cross-team design docs are slightly more likely to be design reviewed (87% vs. 83%) and kept up-to-date (63% vs. 51%). The greater care taken with cross-team design docs reflects their important role in defining the moat.

Unit tests, used by 79% of our respondents, are an important part of the development process for many reasons. One of their surprising functions is to help to defend the moat against outside perturbations – 54% of respondents agreed that an important benefit of unit testing is that “they isolate dependencies between teams.”

Dependencies that cross the moat are anathema to development teams. One of the ways they have to avoid them is what we call *clone-and-own*, where a development team takes a snapshot of another team's code and integrates it into their own code. This might be done initially because the donor team will not ship at the required time or will not add a feature needed by the cloning team. Despite this practice's obvious long-term costs, it is often the most rational decision in the short term. It's not generally perceived to be a problem – only 29% agreed that their team's code is made more difficult to manage by “a large body of code imported from another team into my team.”

Almost all teams have a *team historian* who is the go-to person for questions about the code. Often this person is the development lead and has been with the code base the longest.

Newcomers

Creating a mental model from scratch requires a lot of energy for the new team member and the team as a whole. Often the newcomer is assigned a mentor, often the team historian, designated as the first point of contact for questions about the code. The mentor helps to jumpstart the newcomer's mental model and social network. Newcomers are much more likely to read the team's design documents than seasoned team members. Some teams maintain documents

Problem	% Agree
Code Understanding	
Understanding the rationale behind a piece of code	66%
Understanding code that someone else wrote	56%
Understanding the history of a piece of code	51%
Understanding code that I wrote a while ago	17%
Task Switching	
Having to switch tasks often because of requests from my teammates or manager	62%
Having to switch tasks because my current task gets blocked	50%
Modularity	
Being aware of changes to code elsewhere that impact my code	61%
Understanding the impact of changes I make on code elsewhere	55%
Links between Artifacts	
Finding all the places code has been duplicated	59%
Understanding who “owns” a piece of code	50%
Finding the bugs related to a piece of code	41%
Finding code related to a bug	28%
Finding out who is currently modifying a piece of code	16%
Team	
Convincing managers that I should spend time rearchitecting, refactoring, or rewriting code	43%
Convincing developers on other teams within Microsoft to make changes to code I depend on	42%
Getting enough time with senior developers more knowledgeable about parts of code I'm working on	34%
Expertise Finding	
Finding the right person to talk to about a piece of code	39%
Finding the right person to talk to about a bug	38%
Finding the right person to review a change before check-in	19%

Table 1: *Percent of developers who agreed that “This is a serious problem for me.” (The survey participants were presented the questions in a different order, and without categorization.)*

specifically for newcomers. Unguided exploration of the code is rare; more commonly the newcomer is assigned bugs specifically to introduce them to the code while minimizing risk.

Rationale behind the Code

Understanding the rationale behind code is the most serious problem that developers face. In our initial survey, 66% of the respondents agreed that “understanding the rationale behind a piece of code” was a serious problem (see Table 1 for the complete list of problems we asked about). There are many facets to the rationale problem: 82% agree that it takes a lot of effort to understand “why the code is implemented the way it is,” 73% “whether the code was written as a temporary workaround,” 69% “how it works,” and 62% “what it’s trying to accomplish.”

Investigating Code Rationale

When investigating a piece of code the developer first turns to the code itself: on average respondents to the survey spent 42% of the time examining the source code, 20% using the debugger, 16% examining check-in comments or version diffs, 9% examining the results, 8% using debug or trace statements, and 5% using other means. In other words, the code itself is the best source of information about the code. However it is not flawless. It is common for developers to become disoriented in source code, and it is often difficult to discern the relationships between observed behavior and the source code.

When the code itself does not give the answers the developer needs, one might expect them to turn next to the vast amount of information that's written about it – the bug reports, the specs, the design documents, the emails, etc. This however is emphatically not the case.

The second recourse for investigating the rationale behind code is in fact the social network. If the developer thinks that someone within his team might be able to provide the needed information (or the name of the person who might), he will walk down the hall to talk with the teammate. Unplanned, face-to-face meetings happen frequently with teammates, averaging 8.4 in the prior week, and much less frequently with non-teammates, averaging 2.6. Email is used instead of face-to-face meetings when the issue is low priority, involves multiple people, or crosses the team moat, averaging 16.1 sent to teammates in the prior week and 5.9 to non-teammates. (Note that these data show that communication within the team is much more common than communication across teams, indicating that the culture of informal communication works well and that the team boundaries are typically drawn in the right places.)

Once the developer has the desired information, he returns to his office, applies the newfound information, and gets on with his work. This information is precious: it is demonstrably useful, demonstrably hard to ascertain from the code, and was obtained at a high cost. Yet it is exceedingly rare for the developer to write this morsel down anywhere. The next person who encounters the same information need has to go through the same laborious discovery process. There are plenty of reasons that a developer would choose to not record the information – the overhead of checking the code out, editing it, and checking it back in (possibly triggering check-in review processes, merge conflicts, test suite runs, etc.) is enough to dissuade the developer from recording the information as a comment in the code. But the loss of this precious knowledge is, from an organizational perspective, a crying shame.

Interruptions

Each of these unplanned, face-to-face meetings represents an interruption of at least one person. Recovering from these interruptions is a substantial problem, ranking second with 62% of developers agreeing that this is the case (Table 1). People have adopted various strategies to mitigate the effects of interruptions on themselves, such as using a closed office door, “do not

disturb” sign, or closed relight blinds to deflect interruptions, working on complicated tasks at times of the day when interruptions are infrequent, staving off an interruption for a moment while finishing a thought, or scheduling “office hours.” Other strategies mitigate the effects of interruptions on the *other* person, such as using email instead of face-to-face for a low-priority issue or emailing a warning 10 minutes before the office visit. While many (though not all) interviewees indicated that they received too many interruptions, all acknowledged that interruptions were an important part of the way work happens at Microsoft. Interestingly, two interviewees indicated that interruptions were more of a problem since their teams had adopted agile processes.

An Environment of Experimentation

Developers and development teams are constantly trying new tools and work practices to try to optimize their work. Developers use a variety of tools to do their job. When writing code, 49% use two or more tools, and 19% use three or more.

Visual Studio has become the dominant development environment within Microsoft. An average of 53% of the time developers spend writing or editing code is in Visual Studio. Source Insight² is a distant second, accounting for 17% of the time.

In our interviews we found several development teams that were experimenting with “agile practices,” a collection of behaviors intended to make the software development process more efficient³. Some teams were gingerly dipping a toe into the agile water, while others were jumping in with both feet. In our follow-up survey we found little overall use of agile practices (see Table 2). On the other hand 48% of respondents to the follow-up survey agreed that their team was using two or more of the eight practices, 32% to three or more, and 20% to four or more. Six respondents (3%) reported that their teams use seven or all eight of the practices. Most developers want to continue the adoption of agile practices (53% agreed that they thought their team “should adopt agile software development methodologies more aggressively”) while a few were skeptical (14% agreed that their team should adopt *less* aggressively).

The degree of experimentation suggests that there is ample opportunity for our team to successfully deploy tools to Microsoft developers.

² <http://www.sourceinsight.com/>

³ <http://agilemanifesto.org/>

Agile Practice	% Agree
Collective code ownership within the team	49%
“Sprints,” i.e. a development cycle that last four (or so) weeks	42%
An intentional policy to involve customers (internal or external) deeply into design and planning	33%
“Scrum meetings,” i.e. a brief daily status meeting including all stakeholders	25%
“Burndown” estimate or chart, i.e. a measure of the time remaining in the sprint	24%
An intentional policy of preferring face-to-face over electronic communications	16%
Pair programming, i.e. developers working together, shoulder-to-shoulder on a problem	16%
A “bullpen” or other open-floorplan space for the team	10%

Table 2: *Percent of developers who agreed that their teams use these agile development practices. Adoption of agile methodologies is happening piecemeal and in islands around the company.*

FUTURE WORK

Surveys and interviews are good tools but they are limited because they rely on what the participants *perceive* and *say* about their own actions. Both methods are notoriously biased, and are limited in the types of questions they can answer. There is much we still don’t know. Here are just a few of the things we can’t learn from surveys and interviews but we think are important to know:

- We know that when trying to understand existing code developers mostly just read it. What are the patterns and strategies that developers use when trying to understand code?
- Developers have told us that they can get disoriented when browsing through code. What are the strategies and tools that developers use to avoid disorientation, become aware that they are disoriented, and recover from disorientation?
- We know that the “hallway culture” is a primary resource – when a developer can’t figure out code by examining it, he turns to his social network. What transpires in those spontaneous, face-to-face conversations?
- We know that developers have rich mental models of the code, and that these models are rarely externalized. What happens when two developers talk about code? How do they align their mental models?

More user research needs to be done to address these questions.

BIBLIOGRAPHY

1. Ko, A., Aung, H., and Myers, B., “Eliciting Design Requirements for Maintenance-Oriented IDEs – a Detailed Study of Corrective and Perfective Maintenance Tasks,” in *Proc. ICSE’05*.
2. Kraut, R., and Streeter, L., “Coordination in Software Development,” in *CACM* 38(3), pp. 69-81, 1995.
3. Hill, W.C., Hollan, J. D., Wroblewski, D., and McCandless, T., “Edit Wear and Read Wear,” in *Proc. CHI’92*.
4. Perry, D., Staudenmayer, N., and Votta, L. G., “People, Organizations, and Process Improvement,” in *IEEE Software* 11(4), pp. 36-45, 1994.
5. Singer, J., “Practices in Software Maintenance,” in *Proc. ICSM’98*.
6. Singer, J., Lethbridge, T., Vinson, N., and Anquetil, N., “An Examination of Software Engineering Work Practices,” in *Proc. CASCON’97*, Toronto, pp. 209-223.
7. Walz, D., Elam, J., and Curtis, W., “Inside a Software Design Team: Knowledge Acquisition, Sharing, and Integration,” in *CACM* 36(10), pp. 63-77, 1993.