

Ownership & Evolution of Local Process Representations

Thomas P. Moran¹, Tara L. Matthews¹, Laurian Vega²,
Barton Smith¹, James Lin¹ and Stephen Dill¹

¹ IBM Almaden Research Center, San Jose, California, USA
{tpmoran, tmatthe, n6hdn, jameslin, dill}@us.ibm.com

² Center for Computer-Human Interaction, Virginia Tech, Blacksburg, Virginia, USA
laurian@vt.edu

Abstract. Knowledge workers tailor collaborative business processes to local conditions. They own (i.e., create and maintain) representations of these local processes (such as checklists) to guide the work. Our goal is to design tools to support the ownership of collaborative local processes by enabling workers to flexibly adapt process representations to work situations. This paper focuses on how workers evolve representations for collaborative, locally-owned processes by updating them from situated experiences to keep up with changing business conditions. To understand this, we conducted a field study and a lab study. From the field study, we describe how factors like group roles and documentation purposes affect the evolution of process representations. Based on these observations, we propose a model of the practice of evolving local process representations that provides a framework for understanding activity documentation needs. The lab study then provides behavioral details on the ways people carried out the evolution practice. These studies yield design implications for collaborative activity support tools.

Keywords: Coordination, collaboration, business processes, activities, work practices, process evolution.

1 Introduction

Business processes are often treated as centrally-owned and mandated work [1,2]. However, various ethnographic studies of knowledge workers have shown that they formulate their own objectives for business processes and work out strategies and methods for attaining them [3,4,5,6]. We call these *locally-owned processes*. They are local in that they are adapted to local work situations and needs, even when workers are following centrally defined and mandated organizational processes, which we call *central processes*. Local processes are owned by the workers involved and responsible for doing the work. A number of studies show that people adapt central processes to local situations and needs, e.g., [2,7,8,9]. However, little is known about

how workers actually go about evolving local process representations over time, an issue investigated in the studies presented in this paper.

1.1 Local Process Representations (LPRs)

In everyday, informal work practices, processes exist as verbal or even tacit knowledge. However, prior work has found that knowledge workers also create representational artifacts, or documentation, to help manage these processes – checklists, diagrams, charts, how-to's, instructions, etc. [3]. These local process representations (LPRs) help workers: (1) carry out activities by providing to-do reminders, critical how-to information, key resources, etc.; (2) manage and coordinate activities by representing the dynamic state of the processes; and (3) support adapting *activities* (i.e. instances of local processes) to specific situations by being flexibly changeable.

The key to the vitality of LPRs is that people not only create them, but also *maintain* them. LPRs are never perfect, and even if they were, they would degrade over time as conditions change [10]. We are investigating whether workers will adapt and evolve LPRs, under what circumstances they will do so, and what types of LPRs they see fit to evolve over extended periods of time. Throughout the paper, we use the terms *LPR* and *documentation* interchangeably.

1.2 Studying the Practice of Evolving LPRs

To study practices of evolving LPRs, we distinguish two aspects: (1) the organizational context and (2) cognitive capabilities and working styles of the workers involved. To explore the effects of the both aspects on practices of evolving LPRs, we conducted field interviews. Based on observations from the field, we propose a model of the practice of evolving LPRs that provides a vision and hypothesis for the lab study. To explore the second aspect in more detail, we conducted a lab study to gather behavioral data on the ways people carried out the practice of evolving LPRs. These studies yielded design implications for collaborative activity support tools.

2 Related Work

Process representations and their coordination and use are research topics that have an extensive and varied literature. Since we cannot cover all this research in detail, we will show how our work is positioned along with some key examples.

The first broad research area, *activity theory* [11], examines how people work communally with socially constructed tools to create situated outcomes. In this framework, central processes serve as a tool to help people negotiate their roles and actions as mandated by the community's rules. Local processes and their documentation is seen as fashioned from the local needs of in situ activity. Sachs [4]

articulates this dichotomy as an “organizational, explicit” view of work (central process representations) versus an “activity-oriented, tacit” view (LPRs).

A second area of research, *distributed cognition*, examines how tools such as LPRs help mediate cooperation, learning, and error recovery. In a study of how groups coordinate with regular member attrition and on-the-job novices, Seifert and Hutchins showed how people organized the work of navigating a naval vessel with minimal errors [12]. The work on distributed cognition highlights the need to use tools, such as processes and LPRs, for shared understanding and cooperation.

A large body of *ethnographic work* examines how knowledge management systems – which can be considered as collections of process representations – came into being [13] and were reused [14]. Along with studies that were in-depth analyses of discrete yet complex work (e.g., call centers [15,16], ticket booking [5], and banking [17]), this research emphasizes that work is seen to follow a rote method, but is actually quite messy and locally adapted: “Study after study have demonstrated, unambiguously and beyond any doubt, that the status of these formal organizational constructs in the actual course of work is problematic in that these constructs are impoverished idealizations when taken as representations of actually unfolding activities” [5, p.166].

Another body of work studies *organizational routines* [18,19], which are “repeated patterns of behavior that are bound by rules and customs” that continuously change [18]. This literature also observes the “endless variations” of routines (e.g., in different departments) and the range of artifacts that enable and constrain them [19].

The literature on *articulation work* is another form of process reification. Articulation work is the detailing of who will do what and when. Articulation work is distinguished from “cooperative work,” which is the business of doing the work. (See [5] for theoretical postulates on the balance between the two.) This research reflects our findings about who, how, and when people go about revising documentation about both articulation work and cooperative work.

Suchman notes that central process technologies serve as “a basis for centralized assessments of the efficiency and correctness of the local operations in which the technologies are embedded” [2]. Similarly, Dourish argues that workflow technologies (which he describes as centrally-owned) act as “organizational accounting devices,” rendering employees’ work “observable and reportable” [1]. In contrast, the goal of our work is to inform the design of technologies to represent and coordinate locally-owned processes. While prior work shows that locally adapting processes is widespread, there are few studies of the systematic practices of evolving LPRs, the issue we explore here.

3 Field Study: The Practice of Evolving LPRs

Our field study’s goal was to explore and discover the processes for which LPRs are created, the roles people had in updating LPRs, and the practices for revising LPRs.

3.1 Participants, Method, and Data

In order to focus on LPR evolution, we selected 14 participants who coordinated processes, since they would more likely own the documentation. Our participants came from diverse work contexts: they worked across two states, in three different organizations (a large business organization, a university, and a home owners association), in seven positions (administrator, director of development, program manager, HR, department manager, graduate student, professor, and fiscal technician) and in eight different departments. Ten participants were female.

All interviews were conducted in person, in the participant's working context. Interviews lasted 40-90 minutes and were audio-recorded. The interviews were semi-structured, based on questions designed to gather data about the processes they carried out, factors affecting how they carried out the processes, and how they and their teams used and evolved process representations. We also probed topics that arose during interviews, and we collected copies of process representations.

Because participants coordinated processes, our data does not represent a general office population. However, this enabled us to gather concentrated data on practices related to evolving LPRs. We focused on gathering in-depth data about processes that had some aspect of *local ownership*, were *collaborative*, and were *repeated*.

3.2 Results

We discussed various repeated, collaborative processes with participants (26 different processes in total), including on-boarding new hires, disclosing inventions, planning various events and programs, evaluating employees, creating products, moving offices, and teaching classes. For all of these processes, participants had 40 LPRs they and/or their teams maintained through evolving business conditions and local contingencies. We report three main findings in this section: (1) people spent effort to maintain LPRs for four specific purposes; (2) LPRs were updated at different times depending on the way they were used; and (3) roles emerged in teams for practices of evolving LPRs.

Purposes for Maintaining LPRs. We observed that participants took time to maintain and reuse LPRs for four different purposes. Six LPRs were coded as serving two purposes.

- *Explicit and complete how-to information for a complex activity* (8 of 40 LPRs). For example, two administrative assistants at the same global company maintained and used how-to checklists for only two processes (on-boarding new hires and organizing a colloquium), even though they carried out many other processes regularly. The difference in these two processes was that they were particularly complex or long, warranting the effort spent to maintain LPRs. For their remaining, simpler processes, they had memorized the steps. These two checklists included details missing from central documentation and local customizations to the way the process was carried out (e.g., adding new hires to departmental email lists).

- *Maintaining status information for a complex activity* (8 LPRs). Complex activities often require participants to develop LPRs to track status. Both the documentation that organizes status information and the status information itself can evolve over time and participants found these valuable to maintain. For example, one participant coordinated a patent disclosure process by placing post-its, each representing a disclosure, in the various boxes of a state diagram drawn on her whiteboard. The position of a post-it told her the status of the corresponding disclosure. The disclosure process had gone on continuously for years, and the diagram had evolved over that time as she learned more about the process (e.g., she added boxes upon learning of new disclosure states).
- *Informational documentation used during an activity* (22 LPRs). Informational documentation was referenced or shared during a process to help users learn or remember what to do and how to do it. Email archives were most commonly used for this purpose (14 LPRs). For example, one participant who organized volunteer programs at a large company, archived her informational emails sent to employees about the programs each year and reused them next year, both to remind her about how to do the process and to revise and resend to employees.
- *Final product of a process* (8 LPRs). Products created as part of a process, such as papers or presentations, were commonly maintained and reused. For example, when planning and teaching a course a professor reused the previous year's syllabus, editing it for the current semester, and posting it for students.

Timing for LPR Updates. Participants updated LPRs at varying times in process life cycles, depending on the way the LPR was used and the process itself.

- *Just before reusing the documentation.* This was observed in cases when a participant was about to reuse old documentation, but knew that it was out-of-date or not specific and thus revised it first (7 processes). For example, a participant at a large global company kept archives of all the email she sent about the regular employee volunteer programs she coordinated. When a volunteer program was approaching, she would revise and then reuse her old emails about the program.
- *After completing the process.* This was observed in cases that the participant (1) discovered the existing documentation was inadequate or incorrect after using it and updated it for themselves (3 processes) or (2) wanted to keep a record of changes made to the process for people who would do the process in the future (1 process). For example, a participant who was working as a creator and coordinator knew that in the following year someone else in his club would have to coordinate a large yearly event that involved approximately twenty people. To help that future coordinator, he updated his LPR, an online to-do list with associated notes, a week after finishing while the details were still fresh in his mind.
- *While carrying out the process.* This was observed in two of the processes for which the documentation became a product of the process. For example, a wiki, used by a conference planning team of eight people at a large global company, was revised as part of the process: originally, the wiki was used to coordinate their activities that were discussed in meetings (e.g., who to invite as speakers). After the

meetings people would keep track of the progress by looking at the artifacts people were working on that they had posted on the wiki. Some of these artifacts were eventually molded into final products, like the conference agenda and speaker list.

- *As new process information arose.* This was observed in four processes where changes to some aspect (planning, coordination, status) happened continuously or at particular times. For example, one participant created a state diagram for disclosing inventions, which she constantly updated whenever the legal department informed her of a new patent status. She revised the diagram as she learned about the process, regardless of the process' status as starting, completed, or in-progress.

Roles in the Practice of Evolving LPRs. Related studies on documentation practices have noted that experts tend to create documentation, because they have the knowledge needed to do so [12,14,15]. However, our results suggest that additional roles are important in LPR evolution: *creator*, *coordinator* (of various kinds: managerial, administrative, partial, and overall), *team member*, and *advisor*. Of course, not all processes had all of these roles.

Process *creators* usually generated the initial process documentation. For example, a participant was the first president of his homeowners association and created a format for organizational meetings. This format was documented as the meeting agenda, which was discussed in meetings, annotated, and republished with notes monthly to the other four association members through email.

If there was a process *coordinator*, this person usually did the bulk of the LPR updates. If there was a *managerial* and an *administrative* coordinator, the latter tended to be in charge of documentation. For example, SN co-coordinates a yearly conference, but he has an administrative assistant who helps him keep the documentation for the process up-to-date. Coordinators maintained LPRs since they were responsible for making sure the process was accomplished.

In the absence of an explicit coordinator, one or more *team members* will sometimes maintain documentation for collaborative activities. For example, SN describes a repeated process in which he collaboratively creates presentations with team members. The presentation slides are created by whoever takes responsibility first, and they are edited by all the team members.

An *advisor* was a person who instructed others how to complete a process, but was not directly involved in doing so. For example, TS keeps a set of how-to LPRs for managers in her local population. Whenever corporate HR sends an email instructing managers to do an HR process, TS sends a follow-up email with helpful instructions.

4 A Model for the Practice of Evolving LPRs

We articulate a general model for the practice of evolving LPRs to clarify what we mean and to serve as a hypothesis for a lab study exploring individual workers' capabilities and tendencies for evolving LPRs (presented in the next section).

A *locally-owned process* is an organization of the work required to accomplish a business objective, by some more or less specific method, with various *actors* and

resources. The actors play different *roles* to carry out the process collaboratively.

Example:

Patenting at an engineering firm is initiated with a central Invention Disclosure Process, which is the first legal documentation of an invention that could be patented. The actors in this process are an Inventor, the inventor's Manager, an Attorney, and an Expediter. The Expediter coordinates the others to make sure inventions are disclosed correctly and efficiently.

A specific instance of a business process is called an *activity*. Every activity is carried out in the context of a specific work *situation*. A situation includes physical conditions, task constraints and requirements, and actor availability. A situation also includes the actors' knowledge of, experience with, and attitudes towards the activity being done. Situations vary because every work context is different. More importantly, there are trends in the variations. Situations *drift* over time as people, resources, organizations, and business conditions change and evolve. Local processes must also evolve to keep up with this drift, in order to remain relevant and effective.

Example (continued):

Joe and Jane each invented a new technique that might be patentable. Though their disclosures are happening at the same time, their situations are different—the inventions are different (one is time critical), they have different Managers and Expeditors (one is new to the role), and the shared files have been moved to a new server. These are local variations. The organizational context has also evolved: two months earlier, Managers were not part of the process, but they were given a role to reduce the burden on Inventors.

To help manage a process, actors can create a *local process representation (LPR)* or *documentation*, such as a checklist of steps describing the substantive work and coordination work needed to carry it out. Process representations can vary in detail, accuracy, and completeness, depending on how well actors maintain them. Example (continued):

The first Expeditors worked together to create a 16-step set of instructions to describe the invention disclosure process, which they put in a shared file directory. The steps in the instructions are assigned to the different roles, and they are in the approximate order in which they should be carried out.

Documentation can be imperfect for a particular situation. There are discrepancies between the instructions and the situation. A *discrepancy* is a specific inconsistency between what the instructions say and what the situation demands. Two types of discrepancies involve people (actors failing in various ways and roles being unassigned or misassigned) and steps (being too difficult, vague, or just wrong).

Example (continued):

The instructions assign Jane's Manager too many steps, but this Manager always delegates these to his senior engineer. When the Expediter encounters this discrepancy, he finds and works with the senior engineer to get the information needed.

In this model, we characterize a situation by its *delta*, the number of discrepancies between the situation and its documentation. Delta is the measure of the inaccuracy of the process representation in context.

As situations drift and processes evolve, the *documentation must evolve* to continue to be useful guidance for the process. In other words, documentation must be updated to keep the deltas manageable. It is the actors who must update them, since they own them. Thus the practice of evolving LPRs requires that actors not only carry out activities, but also update the process representations. Example (continued):

When the Managers assumed a role in the process, the instructions had to be altered to assign some of the Inventor's steps to the Manager, as well as to change some of the steps.

Actors must carry out activities in situations with discrepancies in the LPRs. When discrepancies are encountered, which can happen before or during a performance of the activity, actors must *reflect* on the discrepancy, *decide* whether it is useful to document their reflections, and actually *document* them by fixing the instructions. In some situations, discrepancies will cause a *breakdown* in the activity, with which actors must first *cope* (find a way around) and then *reflect* upon.

The model does not specify when revisions happen (since this varied based on field study results) or the roles of actors within the process (since these depend on the process). The focus of the model is how potential updates can arise from the experiences of carrying out processes in specific situations.

5 Lab Study: Individual Behaviors in Evolving LPRs

We conducted a lab study to explore people's capability to carry out a practice of evolving LPRs consistent with the model: Can they cope with the discrepancies? Can they fix documentation on the fly? What kinds of fixes do they make or not make?

Participants in the study carried out a typical local business process – an invention disclosure process (explained above). The participants used Google Talk to communicate and followed instructions on a single-page shared Google Document, which served as the LPR.

Each participant played a coordinator role (which we called the Expediter) in a single session, coordinating one complete invention disclosure activity. As Expediter, the participant's job was to gather information from the others, draft the disclosure, and coordinate the others in assessing, enhancing, and approving the disclosure. The researcher, using IM, played the roles of the collaborators – Attorney, Inventor, and Manager – according to predefined scripts that enacted various situations.

We varied deltas in different sessions by planting role discrepancies (unassigned and incorrectly assigned roles) and step discrepancies (incorrect, irrelevant, missing, and out-of-order steps). Participants encountered discrepancies while interacting with their collaborators. For example, if steps were listed in an incorrect order, the participant would notice this either when a collaborator did them in the correct order or when a collaborator would warn the participant that he did a step at the wrong time. Given these discrepancies, the participant was instructed to make sure the LPR (the instructions) were accurate for the next group doing this process.

5.1 Session Groups

Session situations were constructed to enact specific discrepancies to control the delta. We designed the situations as four related groups. Two groups had *stable situations*, where the situation was the same in each session; and two groups had *drifting situations*, where each session's situation had increasingly more discrepancies than the preceding session's. Two groups had *static instructions*, where the same

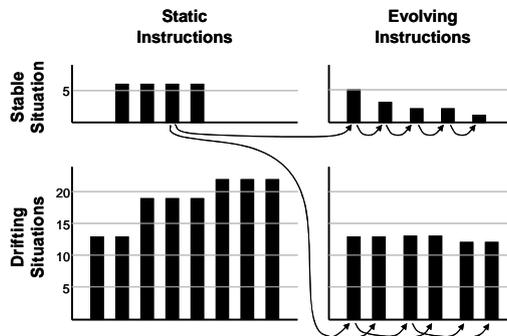


Fig. 1. The quadrants depict deltas for the four groups: the height of a bar represents the number of deltas a participant encountered during that session (only 23 out of 27 participants are portrayed, because data from the four participants who did not follow the practice of evolving LPRs were removed). Arrows show how updated instructions were passed between sessions in the two evolving-instructions groups.

instructions were used in each session; and two groups had *evolving instructions*, where the updated instructions from each session were passed on to the next session. See Figure 1. Our research questions were:

- *Stable-Static Group* (the situation and instructions are the same for each participant). How much variation is there in coping with and fixing discrepancies?
- *Drifting-Static Group* (subsequent situations have higher deltas, but the instructions are the same). What limits do people have for coping with discrepancies?
- *Stable-Evolving Group* (the situation stays the same, but evolved instructions are passed from one session to the next). Will the instructions become optimized to the specific situation, that is, will the deltas approach zero?
- *Drifting-Evolving Group* (subsequent situations have higher deltas, but evolved instructions are passed from one session to the next). Will the instructions evolve to keep up with drifting situations?

Figure 1 shows the four groups: the height of a bar represents the number of deltas a participant encountered in that session (only 23 of 27 participants are portrayed, since data from 4 participants were removed for reasons described below). For example, in the *Stable-Static Group* (upper-right quadrant of Figure 1) the four participants started with the same instructions and faced the same discrepancies. The number of deltas participants encountered in the two evolving-instructions groups (right quadrants of Figure 1) was determined by how many discrepancies the previous participants actually fixed (and introduced) in the instructions, since instructions were passed from one participant to the next. Our prediction was that the deltas in these groups would be lower than those of the static-instructions groups.

5.2 Participants and Procedure

We recruited 27 study participants from a global company: 14 were employees, 13 were interns. Interns had little or no experience in patent disclosures. Employees had

participated in some part of the patent process, either writing or reviewing patent disclosures. The process and LPR used in the study were significantly different from the process and LPR some participants were familiar with, so all participants were relative novices. Participants ranged in age from 23 to 60 years; 7 were female.

Each session was conducted as follows: The participant was (1) introduced to the tools; (2) read and discussed a document explaining the task, roles, and instructions; (3) began the invention disclosure activity with his collaborators, for which there was a one-hour limit; (4) took up to ten minutes to update the instructions for use by subsequent participants; (5) completed a survey about how he felt the activity went; and (6) was interviewed by the researcher to better understand how he felt the activity went, what changes he made to the instructions, and why. Participants were compensated for their time with lunch vouchers. The revised instructions and qualitative feedback were the raw data from the study.

A critical part of the procedure is how we conveyed the basic idea of the practice of evolving LPRs to the participants. In part (2) of the session procedure they read:

Edit the instructions to reflect how your group actually did the work. You can change, add, or delete; and you can add advice or notes as bullets under the steps—whatever you need to best reflect your work. These instructions are not perfect. People have worked on them before you. The previous groups could have added information that is useful. But they could also have added information that is not helpful or even wrong.

Further, they were told their revised instructions would be used by participants in the next session. We motivated participants to improve the instructions by offering them an extra lunch voucher if their revisions helped the next participant.

5.3 Results

The study verified that most participants could follow the practice of evolving LPRs. Specifically, participants experienced minor breakdowns when they faced discrepancies; but all of them, even those in high delta situations, were able to cope with the breakdowns and complete the activity correctly. Participants were able to reflect on breakdowns and document about half of them by fixing the instructions. Revised instructions had an impact on subsequent participants. However, we also observed that not all participants followed the practice of evolving LPRs.

Participants Not Following the Practice of Evolving LPRs. Our most basic question is whether people understand and agree to follow the LPR practice. This includes knowing that they own and can change the instructions, viewing the instructions as a valuable tool for completing the activity, and making changes that document their experience. Four out of 27 participants did not follow the LPR practice. In the lab study situation, these four participants did not update the LPR either because they did not find it was a useful practice, they did not feel they should change the LPR since they did not feel they owned it, or they felt that the LPR was ineffective and completely re-engineered the process by creating a new LPR. Overall,

this finding indicates that not everyone will understand or agree with the practice of evolving LPRs.

Participants Following the LPR Practice. Of 27 participants, 23 followed the practice of evolving LPRs. Figure 1 depicts the number of deltas encountered during these 23 sessions. We focused our analysis on how well participants evolved the instructions, measured by the number of discrepancies they *fixed*. A *fix* removes a discrepancy from the instructions, and thus is a positive change. The grouping of sessions also enabled us to examine the effects of evolving instructions on the number of deltas, in both stable and drifting situations. Evolved instructions had the predicted effect: deltas for sessions using evolved instructions were lower than the corresponding static-instruction sessions (see Figure 1).

First, consider the two groups with stable situations. The Stable-Static group had a delta of 6 (same instructions, same situation). But the Stable-Evolving group showed that continually passing on revised instructions allowed participants to optimize the instructions to delta of 1. The final participant in this group said, “Everything was straightforward and really easy to follow, so I made no changes” to the instructions. The instructions could be continually optimized since the situation was a stable target.

Next, consider the groups with drifting situations, where the situations were designed to become more discrepant with the initial situation. The Drifting-Static group shows the continually increasing deltas in succeeding sessions. But the Drifting-Evolving group was able to hold the deltas down by passing on revised instructions. These situations were a distinct improvement compared to the deteriorating situations in the Drifting-Static group. (The fact that the deltas were held constant is an accidental artifact of our having made the situations drift at the same rate as the participants could fix discrepancies.) Regardless of delta, all participants could cope with all discrepancies to correctly complete the disclosure process, even if they did not fix all of them. As delta increased, the number of fixes also increased, but the percentage of discrepancies fixed remained constant, averaging 41%.

These results may seem obvious, but we think it is important to empirically demonstrate that most people can follow an effective practice of evolving LPRs. For example, participants might not have been able to cope or to revise instructions in a way that benefits succeeding participants.

How Participants Followed the Practice of Evolving LPRs. The lab study also revealed *how* people follow the practice of evolving LPRs.

First, participants did more than fix discrepancies. Participants had no way to distinguish our planted discrepancies from any other parts of the instructions that they felt could be improved. We use the term *enhancements* to label any changes participants made to the instructions that were not fixes of planted discrepancies. While we reserved the term *fix* for changes in the LPR that were positive remedies to planted discrepancies, enhancements could have negative effects: they could actually introduce discrepancies into the instructions (25% of enhancements were *incorrect*, thus increasing the deltas for the following participants) or could lead to verbose, confusing instructions (one participant in the Stable-Evolving group spoke about her instructions, which had been revised by three others: “Some of the notes ... were lengthy and confusing....”). Overall, participants made six enhancements per session

on average. We found that the number of enhancements did not depend on delta and did not vary across the different session groups. This is surprising, because intuitively one might expect fewer enhancements in sessions with more discrepancies to fix or with instructions that had already been enhanced.

Second, participants fixed different *types* of discrepancies with differing frequency. We planted two types of discrepancies into the study situations: role and step discrepancies. For sessions with deltas 6–22, participants fixed 60% of the role discrepancies and 40% of the step discrepancies. (In lower delta sessions, participants had fixed all role discrepancies.) The participants' enhancements can be classified into step-oriented (74% of all enhancements), role-oriented (21%), and other (5%). Editing a step was by far the most common (55% of all enhancements), and most of these were to add clarifying content to the step (43% of all enhancements).

Participants could revise the instructions while they worked on the activity (up to 60 minutes) or in a 10-minute period after the activity ended. Participants fixed discrepancies *during* the activity more often than *after* in the higher delta sessions (deltas 13–22): there was an average of 6 fixes during the activity and only 1 fix afterwards. Enhancements were made equally during and after the activity.

Participants fixed 41% of discrepancies on average. The most common reason for this (revealed in interviews) was that participants simply did not notice they had coped with a discrepancy (10 participants). Other reasons included: feeling like a novice in the disclosure task (5 participants); being confused by the instructions, which could have reduced their ability to document discrepancies (2); forgetting to fix a discrepancy (1); having limited time to make revisions (1); and believing that some discrepancies were so minor they were “not worth recording” (1). These results indicate that iterations over several sessions may be needed to optimize instructions, since some discrepancies will often be overlooked in each session.

6 Discussion and Design Implications

Our field interviews focused on practices and roles affecting the practice of evolving LPRs, while our lab study focused on the capabilities of individuals. Using two empirical methods to explore the same issue proved useful since results from each study provided data and design implications the other could not. One lab study finding was that there was a limit to the number of problems people could recognize and fix in the documentation (a fairly constant rate of 41% for the experimental situation). The lab study also helped us to identify a distinction between different kinds of revisions: *problem fixes* (which were more objective) and *enhancements* (which were more subjective) – a distinction designers may want to consider to help users maintain usable documentation.

The field study helped us explore contextual factors that were hard to simulate in the lab. We found that roles emerged for document evolution. Processes had a primary documenter (usually the coordinator), indicating that tools may need different functionality for documenters and other participants. We also identified four specific purposes motivating people to spend effort maintaining LPRs: how-to and status information for complex activities, informational documentation using during an

activity, and final products of processes. A design implication is that tools should support creating and maintaining these types of LPRs. Specifically, *how-tos* should be findable and the audience should be explicit since LPRs are tailored for specific groups. To track *status* information, designers should enable the creation of custom status views, like state diagrams. Since email was so commonly used as *informational documentation*, it will be important for tools to tightly integrate with email. Finally, to help teams collaborate on *products*, tools need to support iterating and commenting on drafts and transitioning content to final product format.

Other results from the lab and field studies complement each other. Lab study participants were told they were in a coordinating role (the Expeditor) and were in charge of documentation. Even so, not all participants adopted this role (3 made no changes to the instructions). However, *all* participants edited the patent abstract. Similarly, we saw in our field study that more team members collaborated on deliverable documentation. Such documentation does provide some information about process status, though it is primarily a product of the process. For design, this implies that mixing general process documentation with the product of a process in tools may increase team participation in maintaining both.

In the lab, most people revised instructions *while doing* the activity, but we saw more varied behavior in the field. This difference may have been due to differences in processes and situations. For example, lab participants had no way of discovering problems before doing the activity, and revising. The high variability of processes in real contexts shows it is important that documentation can be easy to revise at any time.

7 Conclusion

Our goal is to inform the design of tools supporting ownership of local processes. We conducted a field study and learned how participants evolved LPRs: factors like group roles and purposes of documentation affect the evolution of process representation. Based on our findings, we developed a model of the practice of evolving LPRs to explore a specific vision of bottom-up ownership. We used a lab study to explore the cognitive issues of whether and how people understand and carry out the practice. The model and studies revealed implications for designing tools that can not only support coordination through process representations, but also make evolution of LPRs integral to such support.

References

1. Dourish, P., *Process descriptions as organisational accounting devices: the dual use of workflow technologies*, in *Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting Group Work*. 2001, ACM: Boulder, Colorado, USA. p. 52-60.
2. Suchman, L., *Making work visible*, in *Communications of the ACM*. 1995, ACM: New York, NY, USA. p. 56-ff.

3. Kogan, S.L. and M.J. Muller, *Ethnographic study of collaborative knowledge work*. IBM Systems Journal, 2006. **45**(4).
4. Sachs, P., *Transforming work: collaboration, learning, and design*, in *Communications of the ACM*. 1995, ACM. p. 36-44.
5. Schmidt, K. and C. Simone, *Coordination Mechanisms: Towards a Conceptual Foundation of CSCW Systems Design*. Computer Supported Cooperative Work (CSCW), 1996. **5**: p. 155-200.
6. Gerson, E.M. and S.L. Star, *Analyzing due process in the workplace*, in *ACM Trans. Inf. Syst.* 1986. p. 257-270.
7. Suchman, L.A., *Office procedure as practical action: models of work and system design*. ACM Trans. Inf. Syst., 1983. **1**(4): p. 320-328.
8. Fischer, G. and J. Ostwald, *Knowledge Management: Problems, Promises, Realities, and Challenges*. IEEE Intelligent Systems 2001. **16**(1): p. 60-72.
9. Hardstone, G., L. d'Adderio, and R. Williams, *Standardization, Trust and Dependability*, in *Trust in Technology: A Socio-technical Perspective*, K. Clark, et al., Editors. 2006, Springer: Dordrecht, The Netherlands. p. 69-103.
10. Bertelsen, O.W., *Organisational learning is crystallised into artefacts*. SIGOIS Bull., 1996. **17**(3): p. 37-39.
11. Kuutti, K., *Activity theory as a potential framework for human-computer interaction research*, in *Context and Consciousness: Activity Theory and Human-Computer Interaction*, B.A. Nardi, Editor. 1995, Massachusetts Institute of Technology: Cambridge, MA. p. 17-44.
12. Seifert, C.M. and E.L. Hutchins, *Error as opportunity: learning in a cooperative task*. Human-Computer Interaction, 1992. **7**(4): p. 409-435.
13. Halverson, C.A., T. Erickson, and M.S. Ackerman, *Behind the help desk: evolution of a knowledge management system in a large organization*, in *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. 2004, ACM: Chicago, Illinois, USA. p. 304-313.
14. Markus, M.L., *Toward a Theory of Knowledge Reuse: Types of Knowledge Reuse Situations and Factors in Reuse Success*. Journal of Management Information Systems, 2001. **18**(1): p. 57-93.
15. Ackerman, M.S. and C. Halverson, *Considering an organization's memory*, in *Proceedings of the 1998 ACM conference on Computer supported cooperative work*. 1998, ACM: Seattle, Washington, United States. p. 39-48.
16. Clancey, W.J., et al., *Brahms: simulating practice for work systems design*. Int. J. Hum.-Comput. Stud., 1998. **49**(6): p. 831-865.
17. Cluts, M.M., *The evolution of artifacts in cooperative work: constructing meaning through activity*, in *Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work*. 2003, ACM: Sanibel Island, Florida, USA. p. 144-152.
18. Feldman, M.S., *Organizational Routines as a Source of Continuous Change*. Organization Science, 2000. **11**(6): p. 611-629.
19. Pentland, B.T. and M.S. Feldman, *Organizational routines as a unit of analysis* Industrial and Corporate Change, 2005. **14**(5): p. 793-815.
20. Wasserman, A.I. and D.T. Shewmake. *Rapid prototyping of interactive information systems*. in *In Proceedings of the Workshop on Rapid Prototyping*. 1982. Columbia, Maryland: ACM Press, New York, NY.