

Ethnographic study of collaborative knowledge work



S. L. Kogan
M. J. Muller

We present an ethnographic study in which we examine the ways collaborative knowledge work gets done in a process-oriented environment. The purpose of the study is to identify the kinds of support that knowledge workers could benefit from and to make recommendations for tools that might provide such support. The participants in this study, knowledge workers in various business domains, work in a collaborative environment; their skills are in their areas of expertise rather than computer science and programming. The data we collected are based on field interviews, on observation sessions, and on validation sessions using prototypes. We analyzed the field data using selected principles from grounded theory, and the results of each cycle were used to guide the research in subsequent cycles. In our findings we describe how knowledge workers develop their own strategies and techniques for getting their work done in complex, dynamic environments in which prescribed work processes serve only as reference models. By presenting instances of such environments from our study data, we illustrate how such individualized work processes are created and demonstrate the need for new supporting technologies and tools.

INTRODUCTION

According to Tom Davenport, a knowledge worker is “someone with high degrees of expertise, education or experience and the primary purpose of their jobs involves the creation, distribution, or application of knowledge.”¹ The term was coined by Peter Drucker in 1969 to describe someone who adds value in the workplace by processing existing information to create new information which can be used to define and solve problems.² Examples of knowledge workers include managers, salespeople, nurses, doctors, lawyers, judges, and analysts. To get their job done knowledge workers rely heavily

on tacit knowledge, the kind of knowledge that cannot be codified, but only gained through training or personal experience.

Companies consider knowledge workers among their top talent and are looking for ways to improve their effectiveness. These workers rely on the ability

©Copyright 2006 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the Journal reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to republish any other portion of the paper must be obtained from the Editor. 0018-8670/06/\$5.00 © 2006 IBM

to work collaboratively, leverage relationship capital, and deliver new solutions.^{3,4} Understanding how they work and what their needs are is a critical step toward creating tools that enable them to perform more efficiently. If we can improve technologies and work practices for knowledge workers, we may impact the knowledge work component of many jobs.⁵

We describe an ethnographic study whose goal is to better understand the ways knowledge workers get their jobs done, to identify the kinds of support they could benefit from, and to make recommendations for tools that might provide such support. We conducted this study as part of a requirements gathering initiative for future workflow products for business users (in this paper the terms “workflow” and “process” are used interchangeably).

The knowledge workers in our study have no special computer skills—we refer to them as nontechnical business users of information technology (IT). We focus on knowledge work that involves collaboration and business processes (we use collaboration in the sense that at least two people are involved in the given process). The data we collected are based on field interviews, on observation sessions, and on validation sessions using prototypes. We analyzed the field data using selected principles from grounded theory and used the results of each cycle to guide the collection of data in subsequent cycles.

In our findings we describe how knowledge workers develop their own strategies and techniques for getting their work done in complex, dynamic environments in which prescribed work processes serve only as reference models. By presenting instances of such environments from our study data, we illustrate how such individualized work processes are created and demonstrate the need for new supporting technologies and tools.

The rest of the paper is organized as follows. In the next section we describe our methods and study design, including approach, tools, study participants, and procedures followed. In the following section we present the results of the study. Because the study was realized as a three-cycle process, the results are presented by cycle. The last section consists of a discussion of the results and related work.

METHODS AND STUDY DESIGN

We describe in this section our approach to carrying out this study, the study participants, and the methods and procedures we followed. The field research component of the study was conducted by an ethnographer (the first author) and was reviewed by a project team whose 12 members included IT specialists in design, development, marketing, product management, and research. For confidentiality reasons the names of people and businesses are fictitious.

Approach and tools

We conducted this study over a six-month period in 2003 and 2004 at five different business sites in the Boston area. At each site we conducted interview sessions with a number of study participants. Each session included a questionnaire-based interview, an observation period, a task analysis segment, and a validation segment using prototypes. The tools we used included semistructured questionnaires, a low-fidelity storyboard, and two high-fidelity storyboards.

We performed the data analysis using selected principles from grounded theory (GT).⁶ Grounded theory is a qualitative analysis method used in the social sciences to find relationships and distill patterns from loosely connected data. The collected data are analyzed and this analysis guides the collection of additional data. The process can be summarized as:

Collect data → Define concepts → Build relationships between concepts → Discover patterns in data

Consistent with the GT approach, the study consisted of three cycles, whereby the results of each cycle affected the course of the following cycles.

Participants

The 52 participants in this study (all three cycles) are knowledge workers who are business users of IT and have no specialized computer skills. They are domain experts in the following areas: biotechnology, high technology, medicine, health care, professional services, retail, manufacturing, and law. *Table 1* shows the grouping of participants by job title and the number of participants in each group.

Procedures

The study was conducted in three cycles, and the results of each cycle were used to direct the data

collection in the subsequent cycles. At the end of each cycle, the results were checked in a validation segment that involved the participants, and then reviewed by the project team. The project team included people from design, development, marketing, product management, and research. Content for the storyboards was developed iteratively with the help of the study participants.

First cycle: Observation and task analysis

Observation sessions lasting approximately three hours were conducted with a total of 10 participants, two at each of the five sites. Participants allowed us to observe them while they worked, and also provided tours of their work environment. The tour included the participants' personal work areas, meeting rooms, reading areas, service areas such as mail rooms and lunch rooms, and areas where people gathered for informal breaks. Time was set aside for questions at the beginning and end of each session, when we completed questionnaires, collected work-related artifacts, and took photographs if permitted.

At the end of the visit, we asked the participants to describe the work processes they either used or intended to create. We made sure we understood the job-related tasks, the strategies used to perform these tasks, the tools used, and the problems encountered. We also solicited suggestions for new tools and strategies for managing the work.

Empirical data collected from this cycle were coded and then organized into distinct groups (known as "open coding" in the GT approach). Concepts that would account for perceived patterns in the data were developed for each group. Notes (in the form of memos, as per the GT approach) that captured and compared relationships between the concepts were created. As more data were collected, additional comparisons were made to further refine the concepts. Data collection continued until a point of diminishing returns was reached, when no additional insights were generated from the data analysis (known as saturation in the GT methodology).

The results were reviewed with the project team and the study goals for the next cycle were identified. A prototype, in the form of a low-fidelity storyboard, was created for validation work. Using a GT approach, dimensions/categories of importance were identified. Testing and validating the low-

Table 1 Grouping of study participants by job title

Job Title	Number of Participants
Administrators/Support staff	5
Biostatisticians	2
Clinical research managers	2
Consultants	6
Credit managers	3
Executive administrators	8
Financial clerks	3
Lawyers	4
Managers	5
Physicians/Researchers	2
Research assistants	4
Sales people	4
Store managers	4

fidelity prototype was the goal of the second cycle of the study.

Second cycle: Low-fidelity prototyping

One of the processes encountered in our study involved scheduling meetings for groups of people. A low-fidelity prototype based on this process was used to validate results obtained in the first cycle. Validation sessions were conducted with seven participants. The sites were the same as the ones in the first cycle, but the participants were different. Two to three site visits were conducted with each participant, and a few of these were followed up by phone sessions.

The prototype consisted of a fictional storyboard with a narrative involving several people attempting to schedule meetings for a number of different purposes: meetings with customers and staff meetings for updates and for performance reviews. The storyboard consisted of rough, paper-and-pencil sketches.

We began the session by walking the participants through the storyboard and then encouraged them to comment. The participants suggested changes, and the storyboard evolved with each session, as each additional participant offered new details or

new insight based on personal experience. We designed the storyboard to contain attributes common to all sites, but in addition, also included aspects specific to only certain sites. The latter helped stimulate discussion and revealed how decisions were made as well as nuances across domains.

This study cycle enabled us to collect both confirmatory and contrasting use cases. Consistent with

■ The data we collected are based on field interviews, on observation sessions, and on validation sessions using prototypes ■

the GT approach, we coded the results and refined the concepts and categories defined in the first cycle. The results from this cycle were used to design two high-fidelity storyboards to be tested in the third cycle.

Third cycle: High-fidelity prototyping

Two high-fidelity storyboards were created in PowerPoint** and Flash** and reviewed with the 35 participants in this cycle. The storyboards were based on processes involving document review and approval, buying a car, and repairing a car. The goal was to elicit methods of how processes were completed and how exceptions were handled. To validate the results the participants were interviewed using a semistructured questionnaire. Some of the interviews took place in face-to-face meetings—the rest were conducted by Web conference.

RESULTS

In this section we describe the results of the study by cycle.

First-cycle results

We asked the participants about the business processes and supporting tools they used to carry out their tasks. The participants asserted that aside from occasional use of some formal company-wide business processes such as expense reimbursement and procurement, they did not use any formal business processes. Most participants thought that process tools were inflexible and not easy enough to use. They also thought that the tools were too complicated to use or adjust for use in their own

processes for their daily tasks. The participants described to us how they collaborated with others to get their work done, the situations in which they depended on others for completing work projects, and the ways in which they were bound by company or client policies or procedures.

While observing participants in the process of performing their tasks, we noticed that the wall spaces and boards were plastered with pieces of paper containing checklists, diagrams, how-tos, company policies, flow charts, and various instructions. In some cases, there were several copies of the same list with slight modifications. Whereas the participants were ready to discuss these lists in detail and to give copies away, they made sure to get the notes back and return them to their place on the wall or desk. These notes were described as “vital” or “critical” to their job and to getting work done. In the typical case the note was created by the participant, although some were modified versions of notes created by a colleague. Some of these were heavily annotated with information regarding status, reminders, and next steps.

After hours of observation, it became evident that these checklists, diagrams, and how-tos were manifestations of aspects of processes, the full form of the process consisting of a combination of the artifacts and each participant’s knowledge of how to use the artifacts. The participants presented this situation simply as “their work,” how they get things done, or how they work with others; that is, how they deal with the dependencies that arise from working with others.

These are the participants’ work processes that they streamlined and personalized for their own use. Some processes were simple, such as “how to reserve the LCD projector for a meeting” or “how to submit a book purchase order.” Others were more complex, such as “how to track and report adverse events in a clinical trial.”

Some processes were carried out exclusively by e-mail or instant messaging. Recreating the message trail involved searching and filtering the e-mail many different ways to ensure that the most recent information was captured. It was up to the participant to find all the relevant e-mail, attachments, and supporting documents and place the artifacts in the proper sequence. The challenge was

Table 2 Examples of human-centric processes and the associated form

• Processing a new employee (photocopied checklist) • Updating an organizational chart (diagram) • Reserving the LCD projector (list and photocopied form) • Making travel arrangements, reimbursement of travel expenses (checklist) • Managing documents (workflow) • Managing time (Gantt chart on board) • Updating a calendar, scheduling meetings (diagrams and lists) • Organizing meetings (flowchart) • Creating compliance-based processes (diagrams and text doc) • Registering for an online course (list) • Completing a request for proposal (RFP) or request for information (RFI) (checklist) • Conducting a merger and acquisition (diagram and checklist) • Updating and completing a clinical trials monitoring report (text file) • Analyzing clinical trials data (text file)	• Preparing a data analysis plan (text file) • Processing a work order request (photocopied form) • Following standard operating procedures (text files) • Developing protocols (text file) • Monitoring safety in a clinical trial (flowchart) • Reviewing papers or grants (list and e-mail) • Creating reminders to complete formal company processes (text files and photocopied forms) • Buying books/procurement (list) • Tracking inventory (printed spreadsheet) • Creating a work action plan (table) • Forecasting inventory and sales requirements (printed spreadsheet) • Managing 'shrink' (printed spreadsheet) • Tracking adverse events (checklist and text file) • Conducting a design review (e-mail)
---	--

to ensure that the most up-to-date information was available.

Because we differentiate between a formal business process (referred to by the participants as the company's process) and the personal work process of the participants (the one they use day-to-day to get their work done), we refer to the latter as the *human-centric process* (or human-facing process). The knowledge worker is the owner of the personal work process, whereas the company process is perceived to be "automated" and without an owner. Unlike the company process, which would send computer-generated messages with the injunction "do not reply to this e-mail," the human-centric process has a clear owner. It is always the knowledge worker who makes the decision to proceed and who evaluates the quality of the information and the integrity of the data before going ahead. **Table 2** shows examples of human-centric processes we encountered and the form associated with each.

A small set of detailed use cases were developed based on the task analysis. One of the more popular use cases was the New Employee Checklist—an example is shown in **Figure 1**. All of the companies had some variation of this, and all participants were

able to relate to it. This was a photocopied form used throughout one of the company sites—no electronic version existed.

Because there are many steps and people involved in the new employee process, participants explained that the process gets complicated very easily. The checklist involves five people representing five different departments. Because a department may not always be represented by the same person, no single point of contact exists, which may lead to additional work and delays.

Analysis of the checklist shows that with the 19 steps in the checklist assigned to five different roles, there are many pitfalls. Some of the steps, but not all, need to be performed sequentially, but this is not evident from the checklist. The checklist serves merely as a guide to what needs to be done, following a one-size-fits-all approach. Taking into account exception-handling cases, there are in fact over 50 possible steps involved in this checklist. For the successful completion of the tasks in this checklist, a large number of e-mails, phone messages, forms to be filled, and face-to-face meetings are required. Some participants maintain several

New Employee Checklist		Dept#:
Name:		Start Date:
Position:		Manager:

Task	Owner
1 Offer Letter signed and returned to HR	HR/Mgr
2 Forms - medical, dental, taxes, etc.	HR/Mgr
3 Hardware - need to order in advance?	Mgr/IT
4 Software - our tools installed from net, e-mail, etc.	Mgr/IT
5 Meet people, explain roles	Mgr
6 Nameplate	RS/Fac
7 Office space selected	Mgr/RS
8 Network access and home directory	Mgr/IT
9 User name added to defect and other repositories	Mgr/IT
10 Connections to proper network volumes	Mgr/IT
11 Introductions to personnel	Mgr
12 Tour of the office (kitchen, rest rooms, supply closet, etc.)	Mgr/RS
13 Office supplies, cube set up	RS/Fac
14 Access cards/Picture	HR/Mgr
15 Added to CTG Group e-mail aliases	RS
16 Update organization chart, personal information, and floor plan	RS
17 Expense report procedure explained	RS
18 E-mail address	HR/Mgr
19 Phone and extension working number	RS/IS

Key:

Fac = Facilities	IT = Information Technology	RS = Executive Administrator
HR = Human Resources	Mgr = Manager	

Figure 1
New Employee Checklist from one site in our study

versions of this checklist to cover special cases, such as non-U.S. employees, new employees by department, and new employees by type (e.g., university recruits). RS related to us the way in which the steps in the New Employee Checklist are executed (the following steps of the hiring process take place after a candidate has been chosen).

The first few steps in the process begin with the hiring manager sending out the offer letter with the necessary forms and informing RS that a new employee is starting. The expected process flow is as

follows: **[1]** Hiring manager initiates formal letter offering the position; **[2]** standard HR forms are sent along with the letter; **[3]** hiring manager informs new employee’s supervisor of the hiring decision; **[4]** supervisor orders hardware and software for new employee; **[5]** supervisor requests user ID and accounts for new employee; **[6]** RS requests nameplate, office space, and telephone service for new employee. In practice, however, exceptions occur so often that they become their own informal work processes. We use boldface process step numbers in brackets to indicate the formal process and italic

process step numbers in parentheses to indicate the informal exception-handling improvisations and workarounds. For example, if the new employee is not a U.S. citizen, then (1-*visa*) the hiring manager initiates a visa-approval process, with a series of e-mails, phone calls, and often months of delay. These delays can propagate into further informal processes, such as (4-*visa*) postponed hardware and software orders and (6-*visa*) postponed office space and telephone service. In such cases, the simple [1]-[2]-[3]-[4]-[5]-[6] process envisioned by workflow designers becomes a tangle of formal and informal processes: [1]-[2]-(1-*visa*)-[3]-[4]-(4-*visa*)-[5]-[6]-(6-*visa*).

Restarting a delayed process is often a confusing and time-consuming process, involving additional informal steps that depend on how far the formal process had gone before it was interrupted. In many cases, restarting a delayed or deferred formal process may be more time-consuming than simply canceling and then reinitiating it, incurring additional costs and delays as a result of executing most of the steps in the formal process twice: [1]-[2]-(1-*visa*)-[3]-[4]-(4-*visa*)-[5]-[6]-(6-*visa*)-[2]-[3]-[4]-[5]-[6]. There are many other potential complications in the process, but this illustrates how the first few steps play out in the case of an exception.

RS is an executive administrator at Company A and supports two company vice-presidents and several directors. She is involved in all hiring decisions and collaborates with the Human Resources department. On average there are five candidates being considered for employment at Company A at any given time. As RS described the details and nuances of the new employee process, she reminded us that this process is an added responsibility on top of everything else she is responsible for in her job and her daily life.

RS has several versions of this checklist, in different colors, augmented with sticky notes and annotations, to remind her of what to do in special cases. She said that she feels like each case is a special case, and there are always tweaks to the process. She updates her own checklist and does not share it for fear of complicating things further. She knows what she needs to do and does it. Because she likes to track what others are doing, there is always a flurry of e-mail and many voice mails traded back and forth. During busy hiring times, meetings are

held twice weekly to keep track of status and who is doing what. She would like a technology-based solution but does not know what to use or how to implement it. She thought of building a fancy calendar widget at one point that would help track things, and even submitted a request to the IT department, but only heard back from them three months later (she showed us the ticket number and description). Because it was too hard to explain to the IT department, she decided to continue with her current method.

RS expressed the desire to have a simple tool that would support the collaborative tasks she performs daily. This tool would create a work process, support the assignment of tasks to workers, provide

■ We analyzed the field data using selected principles from grounded theory, and the results of each cycle were used to guide the research in subsequent cycles ■

status information about these tasks, and support task-related notifications. For example, the tool would allow RS to find out if someone were working on a task without asking directly to see if and when the task was completed. She even made the attempt to get the IT department involved in developing such a capability, which was labeled as “create new calendar widget,” although the effort did not come to realization.

Types of users

The data collected through the user profile questionnaire helped us group participants in terms of their business role and IT skill level. We identified two major types of users in our study: business users and power users. For comparison purposes we also list here the profile of the IT professional (developer).

- Business user
 - Goals: aligned with the business goals
 - Focus: getting the job done
 - IT: computers are simply a means to an end; there is little business value to learning a new tool or technology.

- Skills: business expert who uses computers (e.g., browses the Web) but has no development skills, not even HTML, and has no desire to develop these skills

- Power User

- Goals: mix of business and technical goals
- Focus: sometimes known as the local IT expert (guru); provides help to others in installing or upgrading of software and some troubleshooting; usually the first person called upon for help before approaching the IT department
- IT: usually the early adopter in their group; uses computers to get the job done and has interests in new tools and technologies
- Skills: business expert who, although not a developer, understands technology and is willing to learn to use a new tool if the benefits justify the effort

We know much about the developer from previous studies. The developer role is different from the business user and power user described above.

- Developer

- Goal: mix of business and IT goals but predominantly IT-focused
- Focus: provide IT support, work closely with business users
- IT: use IT solutions to improve business processes; solutions may be either custom built or based on commercially available products; able to deliver quick turnaround for customers
- Skills: software development, Web technologies

Second-cycle results

For the participants involved in storyboard sessions (validation sessions using the low-fidelity storyboard), these sessions helped them think about changes to their work processes. They came up with ideas for changes not only for the business process depicted in the storyboard, scheduling a meeting, but other processes as well, such as design reviews, requests for proposal, and clinical trial protocol development. Using paper and pencil, participants sketched out how they thought these processes should work and where and when they might encounter exceptions.

The major weaknesses in the current work process were identified as the inability to track and monitor processes, the inability to repeat routine steps or tasks as needed, the difficulty to customize, and the inability to capture and reuse processes. Not able to reuse artifacts and routines from previous processes, some participants felt that each new process required starting over from scratch. In spite of the repetitive nature of many process elements, participants approach each process as highly personalized and often in need of customization. Most of the communication among knowledge workers happens asynchronously and is partially paper-based. The paper-based artifacts are heavily annotated, and there are multiple copies everywhere. In an effort to collaborate and share work product and status, these annotated artifacts are sometimes faxed or photocopied and sent around to collaborators.

We found several largely transactional tasks or steps that workers would like to automate. Participants report that the transactional work is time-consuming and distracts from what they refer to as the “real issues.” They see a need for easy-to-use and flexible tools to improve productivity and streamline some of the work processes.

The ways in which a job is done within the same company, or even in the same department, varies widely. The need for personalization applies not only to the work process, but also to the process instance. Participants confirmed that they customize processes to suit their work style and needs, and also to adapt to a specific customer or a given situation. For example, two participants at the same consulting firm and the same group not only modified the RFP process to match their work styles, but also adapted it for certain clients. Work practices vary for the same client depending on the goals and project scope. One participant had seven different versions of the same RFP process. A previous study confirms that when knowledge workers are faced with a similar situation, each comes up with a different solution, and it is this diversity of ideas that brings benefit to the company.⁷ Reuse of processes or process elements can lead to savings and thus increased business profits.

Because participants expect that each work process is unique to some degree and needs to be tweaked, they require tools that are smart enough to help them get started and flexible enough to accommo-

date the uniqueness of the work process instance. In their words, they “live in the exception rather than the rule.” Participants told us that often there is significant similarity among their work processes. It is hoped that with better tools, the 10 percent that is different could be customized and the 90 percent that is shared could be reused. Bardram shows that exceptional situations are not unusual in the course of work activities and offer the opportunity for learning and for enhancing existing plans for readiness for future action.⁸ The comments from the study participants suggest we need a better way to support the handling of exceptional situations.

We collected data on a diverse set of work processes. We defined process attributes and applied these to each individual process. As **Table 3** illustrates, each attribute is expressed as a continuum between two opposing values. A given process can be structured (e.g., expense reimbursement), unstructured (e.g., conducting research for a proposal), or somewhere in between (e.g., organizing a meeting). A given process can have one or more of these attributes.

Processing a new employee using the New Employee Checklist (see Figure 1) can be characterized as semistructured, fairly static, predefined, multiperson, repeatable, close to business critical, and not automated.

In exceptional cases or for trouble shooting, an instance of this process could have different values for its attributes.

Third-cycle results

As a result of the validation sessions using the high-fidelity storyboards and the semistructured questionnaires, we produced a final set of concepts and categories and started defining requirements. We developed scenarios and use cases and shared them with the project team. We also developed a set of prioritized future requirements and presented them to the project team.

We concluded that a tool is needed to support collaborative work processes for knowledge workers. This tool has to be flexible and easily configurable in order to accommodate changing work requirements and meet the needs of the different user types. Indeed, processes are tweaked for various reasons—customer requirements, different circumstances, new information, and most impor-

Table 3 Attributes associated with work processes

Unstructured	←————→	Structured
Static	←————→	Dynamic
Ad hoc	←————→	Predefined
One person	←————→	Multiperson
Single use	←————→	Repeatable
Business critical	←————→	Not business critical
Automated	←————→	Not automated

tant, the need to improve the process. Participants said that they would benefit greatly from being able to reuse an already existing process—it would not only expedite the start-up, but they could tweak the weak points and build on the successes.

In this cycle we collected a set of general requirements which included:

- Making specialist knowledge accessible to non-specialists
- Capturing best practices
- Reusing self-customized processes
- Customizing a process instance with ease
- Transferring and sharing knowledge
- Expediting process start-up-time
- Improving personal productivity
- Making their work more visible, documentable, and transparent

We also found that the participants were using many different tools to get their work done. Many of the tools were computer-based (desktop or laptop), some were paper-based, and some were based on mobile devices. Currently there is no easy way to consolidate all tools. Additionally, there is redundancy among systems. For example, the participants may receive multiple messages associated with the same event through mail, e-mail, voice mail, FAX, and so on. Most participants find this distracting. Inboxes are quickly overloaded, and voice-mail messages and instant messages keep coming fast and furious. Participants struggle at times, trying to determine their next step. Through their participation in the study, the participants were able to view their work processes in a new light and were

surprised to discover that their focus and time allocation could be improved.

Another requirement that emerged from the study related to a tool or mechanism that would help the knowledge workers focus on what they should attend to next. A tool is needed to help them get to the *right* things, the *right* people and the *right* events.

DISCUSSION AND CONCLUSION

The study participants made a clear distinction between their own work processes and the company processes (e.g., travel arrangements, expense reimbursement, HR requests). Knowledge workers make decisions and perform their work by drawing

■ Knowledge workers develop their own strategies for getting their work done in complex, dynamic environments in which prescribed work processes serve only as reference models ■

on a particular set of skills, knowledge, experiences, and intuitions.⁹ Although the work is dependent on the individual, process plays a major role.^{10,11} The work processes we observed differ from the enterprise-level processes in several ways:

- These processes are rarely duplicated; some are ad hoc, some are semistructured, and all are in a state of change.
- They are defined and owned by the knowledge worker.
- The process cannot be codified; decisions are made by people and cannot be automated.

Thus, the work processes we observed are not governed by explicit rules. These types of processes cannot be automated, but can be supported or enhanced. We refer to them as human-centric processes (or human-facing processes). They depend on human skills, judgment, discernment, and decisions. This means that the human acts as the “process engine” by handling the rules and moving the work along to the next step in the process. The burden of completing the process rests with the

person in charge—the human. It is the person’s responsibility to integrate the information at each step, evaluate the circumstances, decide whether they are adequate, and move on until completion or resolution.

As we examined the life cycle of some of these human-centric processes, we observed, for example, that they may start out as ad hoc and informal processes but can quickly grow into best practices. This finding is corroborated by studies of activity-centric computing. Muller et al. found that work that begins as ad hoc and unstructured evolves over time into a semiformal process or a reusable pattern exemplifying best practices.¹² Knowledge workers need a simple way to change unstructured, informal processes into more formal, structured processes. They also need flexible tools to customize and reuse existing processes. The essential point is that business users are looking for ways to simplify their interactions with co-workers and for more efficient ways of collaborating and tracking status.

Much time is being spent in e-mail, chat sessions, and on the phone enacting protocols that could be formalized to some degree. In-boxes are quickly overloaded, and these tools are not integrated with process methods. When a process occurs by e-mail, it is time consuming to sort and filter the thread and check voice mail and chat history when a decision needs to be made. Users need a fast, simple way to mix, match, and link disparate sources of information related to a process.

Information overload and attention management were also problematic for the participants because there were so many competing factions and no easy way to organize and consolidate information. Each participant had his or her own *modus operandi* for managing their work but wanted tools and strategies to improve this. Especially frustrating for the user was information discovery. So much time was spent sorting and searching that it was hard to see when something was done in advance, and it was hard to find answers to upcoming questions.

Although business processes (workflows) and other formalized models of work have been proposed as means of structuring and managing knowledge work, there is a growing body of evidence that actual work is more situational and contingent than can be accommodated by formal models.^{13–16} Guindon

shows that the actual practices of individual software engineers engaged in software problem solving do not follow the prescribed formal descriptions of that process, but instead are filled with opportunistic process violations.¹⁷ Olson et al. show similar results for teams of software engineers.¹⁸ Suchman directly challenges the claimed adequacy of workflow and other formal approaches.¹⁹ Muller et al. show that telephone operators (who had been described as performing routine, invariant actions in response to a limited number of work scenarios²⁰) actually spend much of their time performing an under-recognized form of knowledge work.²¹ Dourish provides an integrative analysis of the need to make systematic violations of ordained workflow processes in order to work better, faster, and with greater customer satisfaction.²² We found in our study sites that an overly formal process may contain important ambiguities and may constrain work in ways that reduce both productivity and quality.

Whereas the extent of knowledge work varies from one production function to another (see, e.g., the discussion by Taylor et al.²³), there appears to be a consensus that many jobs are composed of both knowledge work and routine operations. From that perspective, we suggest that our findings may be of value in supporting knowledge-oriented aspects of many different types of jobs.

Our research contributes to this literature by showing the complexity and diversity of seemingly straightforward business processes (Table 2), such as hiring a new employee (Figure 1). Whereas transactional, procedural descriptions of these processes are important, we tentatively agree with Guindon, who argues that formal versions of work may function more as models than as actualities, and may provide their principal value as reference versions of what must be done by the conclusion of an activity, rather than as maps of how a business activity should actually progress.¹⁷ In contrast to Dustdar's proposal to reconcile knowledge management and workflow through linkages between human activities and artifacts²⁴ (similar approaches are explored in other papers in this issue of the *IBM Systems Journal*), we argue for a careful calibration of those human activities. We do not advocate any specific outcome, such as "routinize everything" or "all human activity should be treated as knowledge work." Rather, we advocate that analysts, designers, and implementors should take a more critical stance

and should carefully examine the current and future forms of the work to determine which aspects are likely to benefit from routinization and which aspects are likely to benefit from a more knowledge-intensive approach.

Our study confirmed that knowledge work is a delicate combination of both tacit and transactional work, and that it is largely idiosyncratic. Consistent with other studies, we observed that each task or process uses a different proportion of each activity type.^{3,12,25} From a business perspective, there are major opportunities in helping knowledge workers to get their work done and to focus on the tacit work that can result in true innovation. The challenge of building a new set of tools for knowledge workers is to help extend the breadth and impact of their tacit activities and provide their companies with a business advantage as well. The ability to capture best practices and share and reuse them presents a unique business opportunity and also could raise the productivity of a company's most valuable employees.

Our proposed calibration calls for determining how much structure is needed at each stage of the human activity, and how much variation is permissible and even desirable around those structured stages. There is a need to enable people to mix, match, and link together disparate applications in support of their human-centric processes. Currently, collaboration tools are serving as a diverse set of proxies for a process tool, and much of the benefit is being lost. We saw that processes were happening outside of any system and were not being captured in a way that made the data accessible or reusable.

This study yields a deeper understanding of how processes are created, used, and reused and emphasizes the need for organizational models that accommodate a range of process formalisms, from highly structured procedures to the more flexible and tacit processes explored in this paper. This study also identifies important and as yet unfilled gaps between the domains of tacit and transactional work. Future work will build on our new understandings of the complexities of procedural and tacit work and will examine how people bring the scattered and disconnected resources that our participants accessed into a functional whole which makes sense to people and is one of the principal goals of activity-centric collaboration.

ACKNOWLEDGMENTS

We thank all the people who participated in our study for generously giving their time and providing feedback throughout the course of this study. We also thank Chris Reckling for helping with the data interpretation, Lori Small for help with the validation of the user types, and Ralf Heindoerfer for sharing with us his expertise on workflows. Finally, we thank Charlie Hill for encouraging us to write this paper and for his guidance and support throughout this research.

**Trademark, service mark, or registered trademark of Microsoft Corporation or Adobe Systems Incorporated in the United States, other countries, or both.

CITED REFERENCES

1. T. H. Davenport, *Thinking for a Living*, Harvard Business School Press, Boston (2005).
2. P. Drucker, *The Age of Discontinuity*, Harper and Row, New York (1969).
3. K. Groth, "A Framework Supporting Knowledge Sharing in Organizations," *Proceedings of the International Conference on Knowledge Management (I-KNOW '05)*, Graz, Austria (2005), pp. 28–35.
4. B. A. Nardi, S. Whittaker, and H. Schwarz, "It's Not What You Know, It's Who You Know: Work in the Information Age," *First Monday* 5, No. 5 (May 2000), http://www.firstmonday.org/issues/issue5_5/nardi/.
5. E. Davenport, "Mundane Knowledge Management and Microlevel Organizational Learning: An Ethological Approach," *Journal of the American Society for Information Science (JASIS)* 53, No. 12, 1038–1046 (2002).
6. A. Strauss and J. Corbin, *Basics of Qualitative Research: Grounded Theory Procedures and Techniques*, Sage, London (1990).
7. A. Kidd, "The Marks are on the Knowledge Worker," *Proceedings of Conference on Human-Computer Interaction (CHI)*, ACM, New York (1994), pp. 186–191.
8. J. Bardram, "Plans as Situated Actions: An Activity Theory Approach to Workflow Systems," *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work Conference (ECSCW '97)*, September 7–11, 1997, Lancaster, UK, pp. 17–32.
9. S. C. Beardsley, J. M. Manyika, and R. P. Roberts, "An Introductory Note: The Next Revolution in Interactions," *The McKinsey Quarterly*, Number 4 (2005), http://www.mckinseyquarterly.com/article_page.aspx?ar=1690&L2=18&L3=30&srId=9&gp=1 (requires registration).
10. P. Dourish, "Process Descriptions as Organisational Accounting Devices: the Dual Use of Workflow Technologies," *Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting Group Work*, September 30–October 03, 2001, Boulder, CO, pp. 52–60.
11. T. P. Moran, A. Cozzi, and S. P. Farrell, "Unified Activity Management: Supporting People in e-business," *Communications of the ACM* 48, No. 12, 67–70 (2005).
12. M. J. Muller, W. Geyer, B. Brownholtz, E. Wilcox, and D. R. Millen, "One-Hundred Days in an Activity-Centric Collaboration Environment Based on Shared Objects," *Proceedings of Conference on Human Factors in Computing Systems, CHI 2004*, ACM, New York (2004), pp. 375–382.
13. A. Agostini, G. De Michelis, M. A. Grasso, and S. Patricia, "Re-Engineering a Business Process with an Innovative Workflow Management System," *Collaborative Computing* 1, No. 3, pp. 163–190 (1994).
14. R. Medina-Mora, T. Winograd, R. Flores, and F. Flores, "The Action Workflow Approach to Workflow Management Technology," *Proceedings of the Conference on Computer Supported Cooperative Work, CSCW 92*, ACM, New York (1992), pp. 281–288.
15. T. Winograd and F. Flores, *Understanding Computers and Cognition: A New Foundation for Design*, Ablex Publishing, Norwood, NJ (1986).
16. M. Hammer and J. Champy, *Reengineering the Corporation: A Manifesto for Business Revolution*, Harper Collins, New York (1994).
17. R. Guindon, "Designing the Design Process: Exploiting Opportunistic Thoughts," *Human Computer Interaction* 5, 305–344 (1990).
18. G. M. Olson, J. S. Olson, M. Carter, and M. Storøsten, "Small Group Design Meetings: An Analysis of Collaboration," *Human Computer Interaction* 7, 347–374 (1992).
19. L. Suchman, "Do Categories Have Politics? The Language/Action Perspective Reconsidered," *Computer-Supported Cooperative Work* 2, No. 3, 177–190 (1994).
20. W. D. Gray, B. E. John, and M. E. Atwood, "Project Ernestine: Validating a GOMS Analysis for Predicting and Explaining Real-World Task Performance," *Human-Computer Interaction* 8, 237–304 (1993).
21. M. J. Muller, R. Carr, C. A. Ashworth, B. Diekmann, C. Wharton, C. Eickstaedt, and J. Clonts, "Telephone Operators As Knowledge Workers: Consultants Who Meet Customer Needs," *Proceedings of Human Factors in Computing Systems, CHI 95*, ACM, New York (1995), pp. 130–137.
22. P. Dourish, "Process Descriptions as Organisational Accounting Devices: The Dual Use of Workflow Technologies," *Proceedings of the ACM International Conference on Supporting Group Work, Group '01*, September 30–October 03, 2001, Boulder, Colorado, pp. 52–60.
23. P. Taylor, G. Mulvey, J. Hyman, and P. Bain, "Work Organization, Control and the Experience of Work in Call Centres," *Work, Employment and Society* 16, No. 1, 133–150 (2002).
24. S. Dustdar, "Reconciling Knowledge Management and Workflow Management Systems: The Activity-Based Knowledge Management Approach," *Journal of Universal Computer Science* 11, No. 4, 589–604 (2005).
25. P. Dourish, *Where the Action Is: The Foundations of Embodied Interaction*, MIT Press, Cambridge (2004).

Accepted for publication June 16, 2006.

Published online October 16, 2006.

Sandra L. Kogan

IBM Software Group, One Rogers Street, Cambridge, MA 02142 (sandra_kogan@us.ibm.com). Sandra Kogan is a user researcher and user-experience designer in Workplace, Portals, and Collaboration Software. She has a B.S. degree from Concordia University and an M.S. degree in computer

science and human computer interaction from the University of Massachusetts at Lowell. Before joining IBM, she worked as a human factors engineer in medical informatics at the Brigham and Women's Hospital, Harvard University, Boston, Massachusetts; she also has extensive clinical research experience in neurology, geriatrics, and gerontology. Since joining IBM in 2001, she has led several design strategy initiatives for Lotus Development. Today, her focus is on empirically derived user models and quantitative analysis.

Michael J. Muller

IBM Research Division, One Rogers Street, Cambridge, MA 02142 (michael_muller@us.ibm.com). Dr. Muller is a research staff member in the Collaborative User Experiences group. He received a Ph.D. in cognitive psychology from Rutgers University in 1983. A member of Computer Professionals for Social Responsibility and the Association for Computing Machinery, Dr. Muller is an internationally recognized expert in participatory design. ■