Cogn Tech Work
DOI 10.1007/s10111-008-0126-z

ORIGINAL ARTICLE

Information technologies and intuitive expertise: a method for implementing complex organizational change among New York City Transit Authority's Bus Maintainers

Lia DiBello · Whit Missildine

Received: 1 July 2008 / Accepted: 16 October 2008

Abstract This paper describes an attempt to implement a complex information technology system with the New York City Transit Authority's (NYCTA) Bus Maintainers intended to help better track and coordinate bus maintenance schedules. IT implementation is notorious for high failure rates among so-called "low level" workers. We believe that many IT implementation efforts make erroneous assumptions about front line worker's expertise, which creates tension between the IT implementation effort and the "cultures of practice" among the front line workers. We designed an aggressive "learning intervention" to address this issue and called "Operational Simulation". Rather than requiring the expected 12 months for implementation, the hourly staff reached independence with the new system in 2 weeks and line supervisors (who do more) managed in 6 weeks. Additionally, the NYCTA shifted from a reactive to a proactive maintenance approach, reduced cycle times, and increased the "mean distance between failure", resulting in a estimated $40 million cost savings. Implications for cognition, expertise, and training are discussed.

Keywords Organizational change · Information technology · Intuitive expertise · Simulation-based training

L. DiBello · W. Missildine
WTRG, CUNY Graduate Center, New York, NY, USA
e-mail: lia@wtri.com

W. Missildine (✉)
3333 Camino del Rio South, Suite 110, San Diego, CA 92108, USA
e-mail: wm@wtri.com

1 Introduction

The Bus Maintainers on the shop floor of the New York City Transit Authority (NYCTA) were in the middle of a complex organizational restructuring. At the time, the NYCTA was about to experience a 10% increase in ridership and had been informed of a $300 Million budget cut in their operating and maintenance divisions. They had also bought a new fleet of buses from a new vendor (a procurement process that literally takes up to 5 years) but when the buses arrived, they did not meet standards and had to be taken out of service. A prior vendor was consulted, but could not replenish the fleet to replace the buses that were due to be taken out of service. Ultimately, the new fleet never arrived, so the maintainers had to deal with the 10% increase in ridership with no new fleet. Increased ridership meant that the maintainers had to do more with less. In other words, they were dealing with an increase in wear and tear on the buses, while charged with the task of increasing what is called the "mean distance between failure" (MDBF)—a key metric in assessing the "health" of the bus and the efficiency of repairs or proactive maintenance plans. Higher MDBFs also translates into fewer days in the shop and more days in revenue-earning service, translating to a better bottom line for the NYCTA.

In addition to all of this, NYCTA management wanted to implement a "cycle based" preventative maintenance system involving a new complex IT system, replacing the standard "reactive" system already in place. Given the disappointing track record of IT implementation among so-called "low level" workers, such as the Bus Maintainers of the NYCTA, management was nervous. The Workplace Technologies Research Group (WTRG) had been called upon to address this particular anxiety, which is common among IT implementation efforts: How do we design an efficient, yet complex IT system that is useful for workers with little computer skills training and whose existing culture of practice is generally resistant to such efforts? Further, how do we deploy it work with their considerable skills, which may be critical to the success of maintenance activity but which are not captured in systems per se? While conducting field research in shops in preparation for our "learning intervention", we met "Ed", and we asked Ed to explain the process of a periodic inspection of the buses, based on the existing system:

"Ed": Well, there is not much to this. We just go down the checklist. Nothing to it really.
Lia: So we start at the top and just go down
Ed: No, I don't do that. I mean, I skip around the list.
Lia: Why is that?
Ed: Well, the order doesn't make sense. See that guy back there (points to rear of bus), I'll be in his way if I start back there. And if I follow the list exactly, I'll be running around the bus all day, literally. So I begin with the things in front. And since I have it up on the lift [for reasons unrelated to the inspection] I begin with the things underneath first.
Lia: Okay.
Ed: (looking at steering arm bushing under bus). Here, hold this flashlight for me (picks at dirt and rust around bushing).
Lia: What's that?
Ed: That's the bushing. What's bothering me here is that it looks like some rust here. That's not good. Shows me there's a problem. Let's look and see when this is due back in (looks at schedule of inspections and picks more at the dirt and rust around bushing).
Lia: What's up?
Ed: Well see this bushing over here. Shine the light right here. This is good. See, no rust mixed in with the dirt. Now look at this one. There is some rust in here. But not too much. Not very red. See that?
Lia: (researcher sees no difference).
Ed: That bushing really needs to be changed. But given that this is coming in in 3000 miles for an A inspection, we can take care of it then. It's got at least that much time on it left. And they need this bus this afternoon. It's gotta wait. So we will make a note of it.
Lia: How do you know it has at least another 3000 miles left on it?
Ed: Well, it's obvious. By the color of the dirt. The amount of rust in there.

Ed's explanation of this standard bus inspection provides a window into the complex intersection of everyday expertise among front line workers ("well, it's obvious, by the color of the dirt…") and the tools they interact with that organize their labor ("going down the checklist"). And, ultimately, this intersection is at the center of a complex of factors contributing to IT implementation failure.

As stated in a statement issued by the National Science Foundation, "The achievement of business success in most areas associated with innovative technology not only will require the development of innovative tools and techniques, but also will require a comprehensive understanding of their applications". For our work, the significance of this statement is not in the problems associated with developing "innovative tools and techniques", though these are important issues. Rather, we need to be critical about our assumptions of how workers achieve a "comprehensive understanding of their applications." At first glance, it did not seem that these workers could use a technology such as this; it was a very "unfriendly" text based system and the majority of these workers had poor English reading skills (80% had English as a second language), poor writing and spelling skills and virtually no experience with computers. Thus, as we elaborate, it turns out that the training and education for these systems is not a straightforward issue. In the end, while the problems at first appear to reside in the complexity of the system, at root we needed to address assumptions about workplace learning.

For this study, by way of illustration, we describe a project in which shop floor workers managed to use an advanced technology with surprising effectiveness (as measured by the site's financial performance) after experiencing an intervention we call the "Operational Simulation" (OpSim), a method that helped elicit and reorganize expertise among front workers. We discuss this in greater detail below. We then discuss what it means that this method worked, in terms of training, the impact of complex technology, and adult workplace learning. We focus in particular on the ways in which training efforts may be misguided for certain populations of so-called low skill workers.

1.1 The problem of information technologies from a business point of view

Since the 1970s large information technologies have been fundamentally changing many industries. Specifically, they involve large, highly integrated information systems that capture relatively "live" data for analyses and decisions about rapid changes in business strategy. Two examples we have studied in depth are Enterprise Resource Planning (ERP) and Computerized Maintenance Management Systems (CMMS). We have focused on these because we think they exemplify the ways in which highly-integrated information technologies have changed the face of workplaces and ultimately, workplace skills. That is, they exemplify a general trend in business systems in that they are highly integrated, cross over into several portions of the business and require very detailed and up-to-date information in order to work.

Enterprise level information technologies such as ERP and CMMS can lead to a large-scale and widespread re-orientation of the way business is done, while costs associated with redundant effort can be eliminated. For example, planning can be automated by computer and integrated directly with purchasing and inventory movement. In the case of manufacturing, the communication of product specifications can occur almost instantly between customers and potential suppliers of raw material. In transportation, down time of vehicles can be predicted or planned, allowing spare factors to be greatly reduced (e.g., the number of "extra" trucks or trains required to substitute for those out of service) along with the associated costs of caring for the spare assets. This results in reduced costs to the customer with better and more reliable service.

Important as they are, technologies such as these have enjoyed only modest success in workplaces. Some of the literature on their failure (e.g., Boldt 1994, 2000) indicates that the more highly integrated information technologies are costly and hard to implement. Typical implementation times are on the order of 12–18 months (per site) and success rates have been low. Additionally, ever since the often-cited "CHAOS Chronicles" were published in 1995 and 1998 (The Standish Report 1994, 1998), IT implementation and integration have been widely considered among the most failure prone industries. Based on their criteria, the CHAOS authors found that in 1998, only 26% of IT implementation projects were successful in meeting their original objectives, on time and on budget. In that same year, 46% of those projects were considered "challenged", meaning that the project was forced to downscope while the budget and timeline were expanded, or the original objectives were significantly modified. Furthermore, and perhaps most disconcerting, 26% of these projects were cancelled altogether. Despite these failures, in 1998 corporations spent $275 billion on IT implementation, and over the last decade, by 2007, the top 500 IT implementation companies reported combined revenues of $379 billion.

Some have criticized the vague definitions of "success" and "failure" (Glass 2005), and by 2003, IT implementation efforts saw some improvement. Project success rates had risen to 34%, and project failure rates declined to 15%. Of the remaining 51% of projects were considered "challenged", half of them had cost overruns of less than 20%. On the downside, of these "challenged" projects, time overruns had actually increased since the original reports to 82% (The Standish Report 2003). In part, it may not make sense to talk about IT implementation failure in the same terms as other kinds of large-scale capital intensive projects. One of the primary sources of implementation failure is flexible design, which creates uncertainties in terms of project objectives, necessary resources, and a fixed timeline. Software must adapt and respond rapidly to changing business environments, whereas other large-scale projects, like building bridges, often have a frozen design, with easily identifiable goals and few intervening conditions.

Another difficulty is the perceived disconnect between the skills of front-line workers and complex technologies. New IT platforms are probably best used and deployed by front line workers (such as mechanics, assemblers, or service workers interacting directly with customers) who are in contact with the details of day-to-day operations or have detailed knowledge of equipment. However, this group is often the least likely to possess an understanding of computerization, the role of data in the analysis, and the abstract business goals driving these technologies. Many organizations have also noted that they exhibit significant resistance and have not responded well to classroom instruction on these systems. For this reason, some firms try implementing technologies only among more "qualified" computer savvy staff. However, employees with a background in computer systems (typically MIS staff and programming staff) lack the necessary contact with the front line of the business, or lack the breadth of knowledge about the work itself needed to render these efforts effective. Perhaps more importantly, as technologies increase in their connective power and rely more on up to the minute data, the front line, low level workers' role increases in importance and the consequences of their labor magnify. This presents organizations with a kind of IT catch-22: the users least likely to understand the technology are the most likely to realize the IT benefits and those most likely to understand the technology are least likely to realize the IT benefits.

1.2 Attempts to solve the problem through design and training

While a comprehensive review of research addressing these issues with training (see Capelli et al. 1997) and design considerations (see for example, *Human Factors*: The Journal of the Human Factors and Ergonomics Society; Ergonomics Abstracts and Human Factors in Manufacturing) is beyond the scope of this paper, three key factors emerge for our purposes here. First, in our work, we have seen that workers often use of these technologies for purposes other than they were intended, but these other purposes may be more interesting in the long run. Thus, complex technologies are actually more powerful when they somehow (usually not intentionally) enable workers to look differently at workplace processes, communications and other actions. Second, many standard training methods are proving inadequate for determining the actual impact of training efforts on the company's bottom line. Evaluations that link training to either learning or work performance are not often done or do not draw an unambiguous line between learning and performance (DiBello 1998).

Lastly, after many years of failed attempts to teach skills in ERP-like systems, both firms and employees usually rely on professional societies for training in ERP as well as certification. The emphasis has been on training and certifying master users with the hope that a small group of managers can run these systems successfully. These efforts have not been a straightforward success by any means (see DiBello and Glick 1993 more recent ref available). As systems become more complex and more highly integrated, firms are slowly recognizing that nearly all employees have an impact on the system's data and, therefore, on its ability to analyze business scenarios correctly. Strategies for dealing with this have ranged from requiring the technology vendor to provide training for all users to restricting what users can do (e.g., making systems "idiot proof"). Neither of these solutions has increased the success rate of complex technologies.

1.3 The problem of information technology from a cognitive point of view: failure to see the "user" as an "intuitive expert"

It is estimated that all manufacturing companies have at least tried to implement an ERP system of some kind since the late 1970s. In transportation, power generation and distribution industries, CMMS systems are turning out to be critical to competitive survival in that they can allow companies to re-allocate resources and capital in cost reducing ways while at the same time enhancing performance and service delivery. For large-scale public services, CMMS systems are required.

Both ERP and CMMS systems have caught our attention for three reasons: (1) they are affecting a great many industries, and indirectly, many jobs, (2) they are data intensive and highly integrated, forcing workers at all levels to have a comprehensive understanding of the system's logic and the functions of other users, and (3) they often have a built in logic that is counter-intuitive to the existing expertise of front line workers. As a result, workers are not only compelled to understand the whole business better and participate with that knowledge, but the model of the business itself is changing as technologies make unforeseen things possible. Thus, these new technologies can fundamentally alter the way they represent the work processes they control, analyze and make recommendations about, introducing a deep cognitive impact among front line intuitive experts (see Dreyfus and Dreyfus 1986; Zsambok and Klein 1997; Klein 2004, for more on intuitive expertise).

In light of these considerations, there is a key overlooked element in this complex equation: The "change" made new technologies may not be a function of the technology itself, but rather, may be a function of what the technology makes available to someone who is already an intuitive expert in a given domain. That is, much like a biologist with a more powerful microscope, the "more powerful" is only function of what the biologist already brings to the equation. Therefore, what is important about these technologies is not the ways that they introduce change to workplaces, but rather the ways that they introduce change *into workers*. Specifically, we think they change what kinds of problems and solutions can now be thought about and entertained by those with significant experience. Instead of designing new technologies to replace the kind of "intuitive expertise" that comes with experience, designing for the user as intuitive expert may be the best way to the benefits of large-scale IT implementation. We think this potential is largely unrecognized as evidenced by the way technologies are implemented:

1. Experienced workers are seen an impediment to change.
2. New hires are targeted for technology training.
3. Schools and researchers operate with the assumption that is a generic set of new basic skills that enable workers to use technologies.
4. People with rich hands-on knowledge and poor "school skills" are passed over for advanced technology training or are the first to be eliminated.
5. When training does occur it emphasizes not what workers know, but what they do not know, such as how to use a mouse or basic PC training and data entry drills.

Many of these questions have been addressed from a cognitive point of view under a number of headings, such as "novice-expert shift" (e.g., Chi et al. 1988) "situated cognition" (e.g., Rogoff and Lave 1984), or "naturalistic decision making" (e.g., Zsambok and Klein 1997) and our work has been influenced by the methods and theoretical models from all of these various approaches. However, since the focus of our inquiry concerns the *development* of different ways of thinking in different domains, the research has been most influenced by the theories and methods of developmental psychology and particularly the developmental theories of Vygotsky (1987) and Scribner's application of them to workplaces and workplace cognition (e.g., as summarized in Scribner 1988).

We see cognition and skills as developing in the service and support of activities at work (DiBello 1997, 1998). This is the principle difference between "school learning" and ongoing learning at work. As one participates in a particular industry or occupation, particular strategies and ways of understanding the business at hand are selected and reinforced as they prove over time to have direct bearing on accomplishing important goals (DiBello and Kindred 1992; Scribner et al. 1992). For example, in our first studies of workplaces undergoing technology changes we made a small discovery at a plant north of New York City that influenced a great deal about our subsequent work. In a study of workers using ERP (Scribner et al. 1992) in two different factories—one with a successful implementation and one with an unsuccessful implementation—classroom instruction was shown to be an ineffective strategy for preparing workers to effectively use ERP at either plant (DiBello and Glick 1993; Scribner et al. 1991, 1992). Despite this, at one plant, many individuals managed to master ERP and reduce their inventory by 72%. It turned out that on-the-job activity proved to be critical to developing the necessary skills, and yet it was certain kinds of activity—not activity in general—that made the difference.

When looking at workplace "culture" as really being about skills developed in service of accomplishing goals, it is not surprising, that embedded practices developed over time in a particular workplace are often seen as "resistance" to the process of change as business goals are realigned to adjust to market changes. In fact, when a change is being introduced, change agents (i.e., new management, consultants or a process improvement team) will often disregard any of usefulness that previous strategies may add in the process of change. They are often unaware of the important role that prior knowledge can play in the "new" vision (Chamberlain and DiBello 1997). Their strategy is often to replace all legacy practices and the skills associated with them. Many times they see this involving "selling" the change or eliminating key resisters. This usually does *not* increase the chances of a productive transition. For one thing, it does not acknowledge the importance of content knowledge employees have accumulated over the years. We think the process of integrating useful aspects of legacy skills with practices, which support new and changing business goals, is required for any positive change (DiBello 1996).

In many ways the argument we are proposing here is not new and is not considered radical for technologies or tools that have been designed for "official" experts, such as CNC consoles for machinists or simulators for physicists. What is new is the idea that this may be true for any worker with significant front line contact. Illustrating the point in the context of this study, we turn to CMMS to better understand the relationship between embedded work practices and complex IT implementation efforts.

1.4 CMMS: it's deeper principles—"reactive" vs. "cycle based" maintenance

The approach represented in CMMS is very much opposed to traditional methods for asset maintenance, which are—at least formally—highly reactive. Reactive methods, as opposed to preventative, or "cycle based" methods, of maintenance are an outgrowth of post WWII scarcity. Reactive maintenance methods assume that life cycles of equipment are unpredictable and that the most cost effective approach is to milk an asset for all it is worth by running it to failure. Further, the tracking necessary to determine life cycles (and then verify them) was costly and very laborious.

Now that both computing power for analysis and components for replacement can be procured easily, reactive methods of maintenance are considered to be unnecessarily costly. In fact data provided by the Society of Automotive Engineers indicates that in the 1960s the parts/labor ratio was 2/1. By the 1990s the ratio completed a full reversal (parts/labor = 1/2). Maintenance practices that emphasize reactive repairs also require redundant systems, large spare factors and significant loss of revenue opportunity when equipment is down for repair. Further, with large fleets of buses or trains that support the economic functioning of large metropolitan areas—where reliable service is expected—running to failure introduces unacceptable uncertainty.

With increased emphasis on service combined with the recognition that down time costs money (and unplanned down time costs even more) there was renewed interest in exploring "cycle based" preventive maintenance through component tracking and preventive replacement. However, this approach represented a shift in thinking among workers already expert in functioning well in reactive environments. The basic difference is a shift in the cost of time vs. the value of an operating asset. The idea that the service life of an asset might be cut short in the interest of saving (more expensive) time is a radical shift for most maintenance workers at all levels. The chart below represents the key differences in the underlying assumptions (Table 1).

To summarize, the ERP and CCMS systems are actually implemented to solve a business problem or to affect a shift in business model or practice. Failure to acknowledge this overall purpose is at root a matter of problematic deployment and ineffective training. To be effective, complex IT systems require two co-existing properties on the part of the user: (1) an understanding of the underlying logic, and (2) detailed content knowledge of the work. Detailed content knowledge only comes from years of experience,

Table 1 Key differences between cycle based scheduled maintenance and traditional "reactive" maintenance

Cycle based scheduled maintenance	Traditional "reactive" maintenance
The life spans of individual components are predictable and cyclical in nature, at least within the same environment, as per the physics of wear and tear	The life spans of individual components are unpredictable or factors affecting wear and tear are not constant or predicable
Tracking components' symptoms and repairs can reveal the cycles in time or mileage terms	Tracking component repair reveals only "repeaters", chronic problems or bad repairs
Preventive replacement according to a cycle is more cost effective as the asset's downtime can be planned	Preventive replacement wastes the potentially extended life of a component and it is better to run to failure

but is often embedded in a traditional, reactive maintenance paradigm. On the other hand, someone who understands of the underlying logic but lacks experience is likely to pick the wrong system for tracking or miss important components entirely during the identification process. Most firms and vendors will deploy these systems among those who seem to understand the logic and are computer savvy, but who have little front line contact and knowledge. This rarely works, because the logic assumed by these systems is at odds with the logic assumed by the traditional structure of the workplace.

1.5 The maintenance information diagnostics analysis system project: pilot study for the OpSim intervention

1.5.1 NYCTA and scheduled maintenance

For many years, NYCTA management wanted to implement a centralized "cycle based" maintenance system. Manual systems proved unwieldy, given the size of the fleet (over 4,000 buses and 8,000 subway cars), and it was widely acknowledged that early information technologies failed for many of same reasons cited earlier in this paper.

1. The information needed to make them work had to be extremely accurate at the right level of detail. Ideally the information should be inputted by the mechanic him or herself.
2. Efforts to train mechanics on computer use had not been successful historically. In general, it was widely acknowledged that this population did not contain, by and large, "classroom learners". At New York City Transit, many of the workers did not speak English as a first language (about 80%) and had virtually no keyboard training.
3. System sabotage. Front line workers are usually threatened by information systems on the shop floor, often seeing them as "time and motion" studies in PC form. Such attitudes and the relative vulnerability of computer systems led to widespread system sabotage or damage to expensive computer equipment.

At the onset of the design process senior management of the Department of Buses made a decision to have the hourly Bus Maintainers, Class B (BMB's) enter their own repair data into the system without clerical assistance. There were two reasons for this decision. First, NYCTA—being a public utility—had not always been cost effective in its labor practices. The budget cuts were forcing them re-examine "redundant" work in particular. Asking mechanics to record information in long hand then ask clerks to type the same information into a computer represented a particularly costly redundant practice. Secondly, there was considerable evidence that the hand-written records were much more accurate than what the clerks eventually entered. In fact, this was an industry wide problem and many transit properties were attempting to deal with it. Therefore, the decision was made by senior management to attempt moving all data entry responsibility directly to the shop floor. The decision was an occasion for considerable nervousness among middle management ranks. In general, this approach had never succeeded in anything but private transportation companies (such as UPS), where workers are carefully screened before hiring.

Since this had never been tried before on such a large scale with a "management" technology, the maintenance information diagnostics analysis system (MIDAS) team believed that any traditional education approach based on this very different kind of user would be inappropriate at best. In a small study sponsored by the Spencer Foundation, our group had helped mechanics in the compressor shop at NYCTA subway department overcome the typical long learning curve for ERP systems. We did this by designing a very simple (in retrospect) exercise simulating the mechanics workplace and inventory concerns. These worked very well and seemed to bypass the need for "pre-requisite" knowledge of computers.

Our relatively minor success with training mechanics on a complex computer system using manipulative simulations was seen as a way of making an important large scale, front line computer systems a reality at New York City Transit Bus. Specifically, the senior management saw our project as a success in getting mechanic acceptance and mitigating against system sabotage. At the point of our first conversations, they did not recognize that the mechanics accepted the system because they had learned its business purpose and were using it as a tool for their work. Nor did they agree that user knowledge of the buses might be critical to the successful use of the system from a management perspective. That is, they did not recognize that understanding the reason for the system might affect the quality and nature of the data entered by mechanics, and that this level of quality would, in turn affect the analytic results of the system's pattern analysis capability. In short, the front line mechanic was not seen as a person with knowledge of the buses that could be critical to cycle identification.

Rather than attempt to convince the management that these factors were important, we proposed to do a "training pilot" at one location, ostensibly to increase "user acceptance" and prevent system sabotage. The actual project we designed and eventually rolled out to 19 locations actually addressed user understanding of the reasons for the system and we designed measurements to examine the relationship between user knowledge, data quality, and financial impact. These measurements turned out to be critical for two purposes. First, we were able to show a marked financial improvement when worker knowledge is brought to bear on a business issue in a systematic way. Secondly, we were able to deepen our own understanding of the relationships between workplace performance, workplace "culture" and the individual's ways of thinking and doing. With a better understanding of these critical levels of a functioning workplace, we are better able to understand how it is that the individual can have such a widespread impact—especially with the presence of a highly integrated technology—and at the same time better know what exactly each individual worker needs to know or understand in order to have a positive impact.

1.6 Learning about the culture of fixing buses

As we have indicated above, our entire method rests upon the assumption that efforts to change a workplace culture most often fail because there is an already functioning, cohesive culture that is actively competing with the change. In order to effect change, we must know as much as possible about the competing culture of practice. In particular, we need to know the situations in which the same goal or task is understood and handled differently or even in opposing ways.

We have also found repeatedly, that the official "narrative" of the management tends to reflect where the organization wants to go (although this does not mean the management understands fully what that means) and can also offer some insight into where the organization "has been", but it does not reflect too much about where the organization is now. For example, NYCTA had been explicitly a "reactive maintenance" operation, with an explicit policy to allow components to "run to failure" with little attempt to predict lifecycles, preventatively replace components or sacrifice any of an assets service time. Part of the rationale for this was component unavailability; part was due to the impossibility of verifying life cycles due to a lack of record keeping. Periodic inspections were an attempt to catch dangerous levels of wear and tear on components such as brakes or steering arms. On the surface it seemed that our "competition" was reactive maintenance and the attendant belief that parts do not have natural life cycles. However, this still did not tell us what actual practices instantiate these beliefs. Also, from our experience, we find that on the front line of the business (usually the shop floor level) the picture is more complicated. Usually, the senior management is not moving toward a new paradigm unless there is some tension or inefficiency at the front line of the business. In these situations, legacy methods are already under challenge and new things are being tried. This is what we call the "informal" domain of practice. This is usually our real source of competition and the real source of culture change failure.

Many ask how we "get at" the legacy domain of practice. Very few people in a given workplace are explicitly aware of the dominant "domain of practice", but most are aware when they are operating effectively within its parameters. That is, they know who is effective; who knows what is "going on" and they are able to assess the meaning and significance of situations that are baffling to outsiders. The trick is to tap into the ways that these workers understand their workplace and its business. There is ample reason to believe that people who have implicit expertise in a given area are not the best at narrating their processes of working and making decisions (Dreyfus 2000), especially in dynamic settings such as vehicle maintenance. Further, workers may not share our assumptions about the purpose of self-narrative and interviews in general. Simply put: workers may not necessarily see this as an opportunity to tell what they know, but rather to accomplish some other goal, such as appearing not to be a braggart. In any case, most of the mechanics interviewed off-line greatly underestimated their knowledge of planned maintenance.

Therefore, in order to understand more about how mechanics actually think about the business of fleet maintenance we felt we needed to begin by observing them on the job, but in such a way to understand what it is like to do the job, from their point of view. In order to make this a natural and comfortable observation while still allowing us to ask questions as they worked, we did our fieldwork in the role of "quasi-apprentices". In this role, it is normal to ask questions, want explanations for decisions and be curious about the underlying reasons for doing things. Also, it puts the experienced worker in the role of "master" or "teacher", which is a role they have had to assume many times when breaking in new workers. In all we conducted about 100 h of observation on all the major functions of bus maintenance, which included planned government mandated inspections, proactive maintenance and "planned upgrades", and troubleshooting buses that had been reported defective by drivers.

Revisiting the quote from the mechanic that opened this paper as an example, it is clear that the mechanic believes, as reported on an earlier occasion, that he does not "think", but rather does what he is told to do. Despite his belief around "not thinking", there is a significant amount of situation assessment, analysis and information coordination and (cycle based) maintenance being done here. What this and other observations tell us is that experienced mechanics do have an intuitive understanding of the life-cycles and the coordination of life-cycles among components, within one piece of equipment. This and other observations like it also show us that the depot is probably able to reliably make service (outfit the routes with the right number of working buses) due to decision situations such as this. In other words, there is already an informal culture of preventive, coordinated maintenance operating when the formal practices of reactive maintenance threaten the depot's ability to make service requirements. However, it is not yet systematic or consistent and has a "plan B" status as a practice.

We also were able to observe how the mechanics learn during the course of doing their work. None of the workers we observed considered themselves to be strong classroom learners or "read and write" types. However, as was clear from their performance and their organization's dependence on their ability to diagnose problems and repair complex vehicles virtually unsupervised, they do learn well in complex domains. Most contributed to their own ongoing learning by "puzzle solving", and when stumped, drew on the opinions and observations of peers to help them understand the equipment through systematic group experimentation.

These two observations—the existence of an implicit scheduled maintenance "domain of practice" and the mechanics' evolved method of learning—greatly influenced the next design decisions of the project. Namely, the cognitive probes designed to tap into the individuals' ways of thinking about maintenance and the training exercise to move them into a new way of thinking.

1.7 Field work and the construction of the cognitive probes

In a sense, cognitive probes are conducted as a way of looking at the domain of practice at a different level. Rather than a measure of individual capability, the probes are actually a way of examining how the workplace culture has influenced the thinking of the individual's that comprise it.

Once we completed the field work at NYCTA Bus, we had a pretty good idea of the business model that is desired, the reasons for doing it, and its "competition". Our next task was to design a way to assess the quantitative influence of each domain in the daily business of doing work. We have found that the best way to do this is at the level of the individual using cognitive probes. These probes are very similar to those originally used by Klein (1999) and resemble in spirit his critical decision method (CDM) (Klein et al. 1989) for interviewing experts. However, there are some important differences. Klein's method is a retrospective interview method that employs a set of cognitive probes to non-routine incidents. The CDM attempts to get at an expert's implicit knowledge and situation assessment skills by telling a story from their work history and exploring the methods by which he or she reasoned it through. By the time we are conducting cognitive probes, we wish to constrain the problem-solving context and see how our interviewees view and handle the constraints we have defined. This involves setting up the problem and the tools available to solving it in a uniform way, while at the same time having a situation that "invites" the interviewees' implicit skills and situation assessment biases. Our method involved the following steps:

1. Identify the strategies and practices associated with each domain that make sense only within the "world view" of that domain. E.g., most workplaces have more than one theory of the work being done, and occasionally these compete. Our field work showed us that maintenance is both reactive and proactive, depending on the goals considered most important and the available resources. In NYCTA, the lack of good data for proactive planning and the "make service at all costs" emphasis in the organization tended to favor reactive maintenance. However, proactive practices were also in play, when time and information permitted.
2. Identify behaviors associated with these strategies in the workplace in which we are doing the research. Simply put, we looked at how the different kinds of thinking manifested itself in day-to-day decisions.
3. Design a meaningful "problem" situation into a "cognitive task" that can be solved using the strategies and behaviors from either domain, or a mix of both during a short interview.
4. Design a problem situation that is similar to that in #3, but which is more abstract and generic than the site specific version.

5. Develop a scoring form that permits a coder to check off the strategies/behaviors easily and calculate the proportion of the strategies used from each domain.

For NYCTA we constructed two basic tasks, each of which had three variations. The first task was an "active" task; given a pile of work orders, we asked the interviewee to look them over and then make five piles for each day of the workweek. In other words, schedule the work. Below is a small sample of the strategies for solving this task, by domain of practice (Table 2).

The same task was given in two forms, another piece of equipment that is commonly known (bicycles) and "Machines" listed as "Machine A" through "Machine N", purely made up items that do not really exist with only meaningless codes as defect or component indicators (such as defect Mu8). The second task required the interviewee to interpret information in bus repair histories. Again a bicycle repair history and one for "machine T" was also included. A similar set of strategies (only for interpreting data) as those shown above was used to code the protocol. Photographs were taken of the interviewees' piles and any drawings or writing and all talking and "thinking aloud" was audio taped. As researchers we were sensitive to having an "intrusive" presence among the interviewees. As such, we adopted what we call the "expert apprentice role". That is, while we shadow the workers on site, we ask them to explain how they do their job as though we were going to be doing it. We try to dress in such a way that reflects the kind of dress the workers wear. And we attempt to build a sense of "we are all in this together" insofar as we are helping them solve a problem, and they are helping us to understand the problem. After such measures are taken, audio taping, picture taking, and eliciting descriptions of the work in the cognitive probes are rarely seen as intrusive.

Figure 1 shows the proportional balance of strategies from each domain, before our training exercise was delivered. As can be seen, over 60% of the interviewees' strategies were reactive and less than 40% were proactive, indicating that some proactive planning skills had developed in the workplace. There was striking homogeneity among interviewees in the pilot depot, suggesting a strong workplace cultural effect.

1.8 The "Operational Simulation": a "learning intervention" that would enable the front line Bus Maintainers to do data entry

Based on our fieldwork and the cognitive battery results, we decided to construct a three-part manipulative simulation of a miniature depot, constraining the goals and resources in such a way that in order to "win" (i.e., make service requirements and stay within budget) the participants had to use proactive strategies. From what we could tell, "constructive" activities in real workplaces lead to learning because they elicit the implicit knowledge that the worker has to bring to the problem and at the same time select against non-workable strategies (through experiences of failure). Therefore, the first part of the exercise was designed to "engage the default". Simply put, we gave them the miniature depot, a set of non-

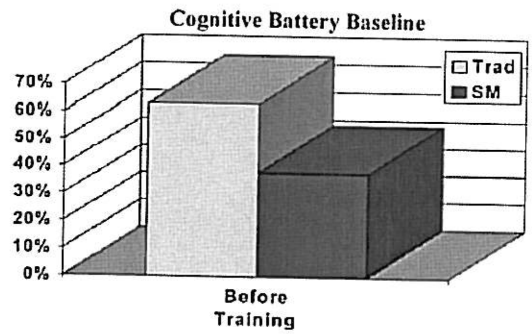

Fig. 1 Proportional balance of maintenance strategies before training. "Reactive" traditional maintenance compared to "proactive" scheduled maintenance

Table 2 Differences in strategies for solving work order problems between a cycle based scheduled approach and a "traditional" maintenance approach

Cycle based scheduled maintenance	"Traditional" maintenance
Interviewee sorts work-order cards first by equipment ID number. Or asks: can the same asset be taken out of service only once to satisfy multiple problems?	Interview does initial sort of work-order cards by type of job or type of trade needed to do job, regardless of equipment ID number
Interviewee compares number of assets coming in shop with number needed for service. Assigns work accordingly. Brings in asset twice on two different days only as necessary to make service	Interviewee distributes work-orders evenly among the days, regardless of the type of work needed to be done
Looks at symptom remarks on the work-orders in order to ascertain nature of problem and if the repair can be coordinated with another repair or scheduled inspection on same vehicle	Looks at the symptom remarks on the work-orders to ascertain if problem is repeater and if part will need replacing

negotiable goals (meet service requirements, stay cash flow positive, make money, keep the buses in good repair) and a number of tools for solving the problem, but little direction.

Specifically, teams of eight participants each were asked to "run a depot" of 40 plastic buses with relatively complex interior components. Ultimately, 3,000 people over 13 months participated in the exercise across the NYTCA bus depots. Two locations would run the exercise at a time for 2–3 weeks and then we would move on to the next two locations. The process lasted 13 months. Each depot had between 200 and 300 people. For the OpSim, participants were broken up into teams of eight according to their shift. For example, if the day shift had 80 people, there were 5 sessions on the day shift with 16 participants running through the exercise at a time, set up in two competing teams of 8 participants each.

The goals were to maintain 32 buses in service at all times (limiting the number out of service to eight), order all the materials (within a budget) needed for doing so and evaluate daily operator reports (each "day" being 20 min) that might indicate potential problems (e.g., noisy engine). The activity was "rigged" so that the only way to meet these goals was to predict what was due to break next. The breakdown patterns of all components followed time/mileage cycle rules and were pre-calculated using a computer. The toys were actually "broken" according to this pattern. The participants were given adequate tools to predict and calculate this breakdown, (printouts of every bus' repair history among other things) but were given other tools as well, including those similar to those used to do "reactive" maintenance.

Our trainers also played a role. One acted as "dispatcher" regularly demanding buses to satisfy routes while the other acted as a parts vendor and an FTA (Federal Transit Administration) inspector, looking for safety violations or "abuses of public funding" such as overspending or cannibalizing. Close examination of the use the participants made of the tools offered pretty much approximated their work history. People tended to construct a solution to even a novel problem that fit with their experience, even when explicitly instructed to avoid doing so. In fact, the participants were rarely aware they are replicating their normal methods.

Rather than interfere with this tendency, the trainers allowed the participants to "wing it", while carefully documenting the cash flow, labor flow, inventory acquisitions and the number and type of on -the -road failures that result from failing to predict problems. Meanwhile, heavy fines are levied for expensive "reactive" problem solving strategies, such as "cannibalizing" an entire bus for a few cheap parts that will get other buses back on the road. As the activity progresses, participants are continually shown the financial consequences of their decision making patterns and asked "what they were thinking" by the vendors/inspectors and dispatchers. By the end of the first day, the "depot" is in crisis and the participants are realizing their budget is being expended to react to mounting problems. The activities are stopped and the team is sent back to work or to lunch.

On the second day of the exercise, participants reflect on what they did, as recorded by the trainers. The participants discuss among themselves what thinking led to various decisions and begin to identify practices that lead to bad outcomes vs. practices that are preventative. It is only at this point that the participants were truly open to new ideas about how to solve the problems of vehicle maintenance. They also began to understand in detail the ways that their "gut feel" decisions reveal how they have actually misunderstood preventative maintenance. In the last hour of part 2, the trainers facilitate participants in building a manual scheduled maintenance system. The participants identify cyclical patterns from histories (which were available from the first but which now take on new meaning) and set up the predictive data structures, identifying true cycles and—most importantly—coordinating cycles so that their "system" is bringing in a bus only once to satisfy several cycles at once. For example, in the MIDAS workshop, the participants quickly realize that a 15,000 mile cycle and a 30,000 mile cycle can be coordinated so that at least half the time the 15,000 mile cycle co-occurs with a 30,000 mile component cycle. The participants construct a maintenance allocation chart for the whole fleet over a number of months and evaluate the stress this will put on the shop. After doing this, they then enter these data on an actual test region in MIDAS and create and assign the work orders according to this schedule.

During part 3, the participants complete their data entry and print out their work assignment sheets and work orders. They run their miniature depot again using MIDAS and see the difference in profits and ease of workflow. Usually only after 5 "days" the team can afford to buy an additional bus to add to the fleet and thereby increase their fare-box revenue.

The last activity of the workshop involves entering the data on work orders (paying attention to detailing the components, defects and symptoms involved) and closing out both work orders and work assignment sheets. At this point, participants also learn how to get various reports that they now realize they will want, such as a 30-day history on a bus. After operating *as* MIDAS and then *with* MIDAS, participants navigate through the actual system more easily, know what to look for and ask informed questions. Even computer illiterate individuals show little hesitation when exploring the system.

2 Results of the OpSim learning intervention

This same exercise was conducted with over 3,000 people in a period of 13 months. As indicated above, about 80% were not native English speakers and fewer than 20% were computer literate. Many mid-career individuals had not completed high school. None wanted to attend the training and most were resistant to the idea of having to do their own data entry.

Despite these features of the trainees, they mastered the system at record speed: Rather than requiring the expected 12 months for implementation, the hourly staff reached independence with the system in 2 weeks and line supervisors (who do more) managed in 6 weeks. The one exception was a location that received classroom training but no OpSim intervention. After 8 months, the implementation was being declared a failure.

2.1 Evaluation of the outcomes

Because the participants in our exercises were depending on us to orient them to the system, we could not arrange for a control group. Rather, those who were out of work on the first day of their scheduled exercise (due to illness, personal days or other reasons) were scheduled to go through the exercise at some much later time and were measured as a kind of control group before they could be trained. There were 12 of these individuals out of about 150 pilot participants. After 6 months, we conducted two measures of the project and developed two others that would require more than 6 months worth of data. These latter two measures were used at all the remaining locations when the pilot program was rolled out to all 3,400 maintenance department employees. We will return to a more complete discussion of these below.

After 6 months, the maintenance personnel produced a mirror image of their former strategy. Rather than solving the scheduling and data interpretation tasks with a "reactive" dominant approach, they exhibited about 70% of their strategies in the "proactive" domain. Further, when asked about how this compared to their prior performance, most did not remember doing it another way, and several could not replicate their former solution to the problem.

2.2 Data entry patterns

2.2.1 Depth of detail

Two measures used to determine the quality of the system use concerned the data inputted by the mechanics. We know from industry standards that coding the location of an equipment defect at a sub-system level is required for trend analysis. When users do not understand the level of detail required, they code at too general a level for the data to be useful. After six months, our trainees were coding at the 4th level (greatest level of detail in the bill of material—the most detailed component level within the assembly) of the equipment template most of the time, and about the same number of times that the repair is probably due to a life-cycle ending. This is an unprecedented result in the transit industry. It indicates that the users were knowledgeable of the uses the system makes of the data and were coding appropriately.

2.2.2 Code variation measures

Downloads of workers' navigation through the system and data entry practices were analyzed for component code variation and homogeneity. For the first we measured the frequency with which any component code was chosen from a finite universe of about 2000. In general, the data from systems like MIDAS have been considered poor quality, or inaccurate when the same symptom, defect and component codes are chosen over and over because they are both general and easy to remember. An example would be inputting a code for "malfunctioning" (symptom) "broken" (defect) and "fuel system" (component). When analyzed by the system for patterns that indicate life cycles, these kinds of data are basically useless. When users understand this, they tend to code more specific kinds of information. A more detail example of the above would be: "sporadic power surges" (symptom) "cracked" (defect) injector valve-aft (component). When users are inputting correctly, a certain level of variation should naturally occur in the choices they make; for example, if an individual uses only three different symptom codes all month and logs in work two to three times a day, there is clearly little variation. In that case, each code used would have a high hit rate. Therefore, we looked for low hit rates per code, per user. In general, we found that a hit rate of 1.2 per month, per code per person was the level of variation indicating good use of the system when the employee was logging in about once a day. After system use became more widespread (and the universe of codes did not expand) users were logging in about 100 times a month or more. The hit rate then went to about 2.5.

Figure 2 below shows the average frequencies per code, per person in two groups: our trainees, and the 12 "controls" (we did not use the site with no training as a control in this analysis because those users were not even logging in after 2 months!). As can be seen below in the lowest frequency was among the trainee group while controls were higher in both frequency and variance. We continued analysis on the codes and the MDBF for 3 years formally. When we went back and visited the chief maintenance

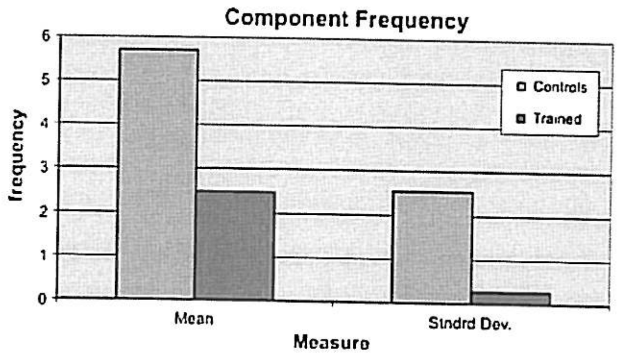

Fig. 2 Average frequency per code, person in the "trained" group and the "control" group, who had not received the training

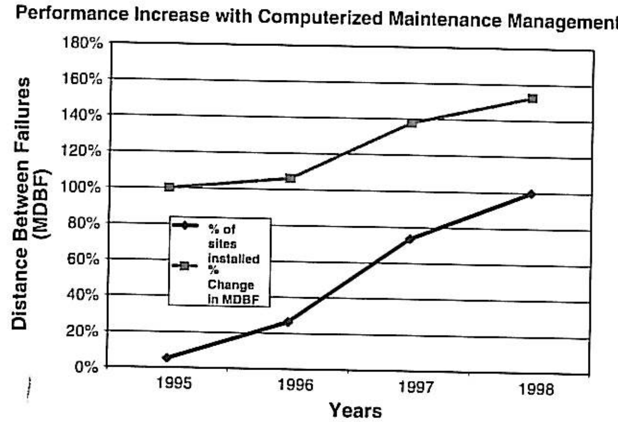

Fig. 3 Rise in "Mean Distance Between Failure" as the implementation progressed

officer, there was no degradation in the performance of the fleet and the new buses still had never arrived.

For homogeneity we used a Scheffe test of the standard deviation. Experts tend to be more homogenous than novices, thus by doing an ANOVA and seeing a smaller standard deviation around the mean is a sign of expertise. The Scheffe test was used because there were unequal sample sizes due to the unequal sizes of the jobs to be done. The test is more sensitive to incremental differences in homogeneity of variance.

The low standard deviation among the trainees indicates a homogeneity effect. That is, there was considerable within group consistency in the type of detail entered even though the input was much more complex. In contrast, the control group entered much less detailed component information and showed significantly greater idiosyncrasy in their choices of what to enter.

2.2.3 *Mean distance between failure*

Mean distance between failure is calculated by taking the number of in-service vehicle failures divided by the in-service distance traveled. At NYCTA it is calculated both system wide and for each depot. With the introduction of MIDAS it could be calculated on a by vehicle basis. In general, unanticipated equipment failure indicates that a proactive maintenance plan is failing (e.g., lifecycles are unknown, incorrectly estimated, or that wrong repairs were done the first time). Thus, elevations in MDBF indicate that proactive maintenance efforts are working. From a business point of view it also means that the revenue-earning asset is out earning money and is not incurring maintenance cost from repair labor.

Therefore, simply put, the higher the better. As can be seen in Fig. 3, the MDBF rose system wide at the same rate that the MIDAS mechanics participated in our exercise. The savings from the increased MDBF is an estimated $40 Million. At the time of the OpSim, the NYTCA had bought a new fleet of buses from a new vendor, but they were not good enough, and the old vendor could not replenish the fleet to replace the buses that were due to be taken out of service. The new fleet never arrived, so the maintainers had to deal with a 10% increase in ridership with no new fleet. They had to take old buses and increase the MDBF and they had to extend it beyond the manufacturer's prediction. The MDBF is the overall system-wide average MDBF for the whole fleet. The overall average rose with the number of depots that had completed their MIDAS implementation and hence were having a greater first pass yield on repairs. The savings in field supervisor time (handling the return of broken down buses) is estimated to be 208,000 h times a fully loaded hourly rate of $70, or $14,560,000. These numbers represent the financial benefits that were incurred even before there were enough data collected to do the kind of trend analysis needed for true preventive replacement based on life cycles. That analysis was just beginning about 2 years after the system was fully implemented.

For the transit industry, the MIDAS project has been the first successful front line deployment of a CMMS system and, one of most successful and enduring implementation of a CMMS in general. The interesting portion of the project from our point of view is that when the technology was implemented as a tool for "experts", it extended the existing content knowledge of the workforce into a form that could be used for an aggressive change in the way business was done. The educational process involved not a content driven course, but rather a preparation for people who were already experts of one kind, who needed to be experts in a different way. The "preparation" was an opportunity to reconstruct an implicit framework of the business into one that was more appropriate to current goals. Further, and perhaps more importantly, it allowed the workers to re-locate themselves in the larger picture. This is critical when a technology of such high connectivity is involved.

3 Conclusions: what does our work tell us about users, technologies and the relationship between them?

We have made a number of points in this paper about technologies, their properties, their impact on users, and ultimately, their impact on the performance of a workplace. In many ways this discussion has been a struggle of languages and domains. Our main point is simple: the way that workers understand their work and their role in the workplace acts as a kind of operating theory that affects how they do their work and what actions they choose at various decision points. It also affects the kind of skills they seek to develop. The "operating theory" of work and its attendant skills develop or are learned as a result of pressures from the workplace culture itself. Workplace culture manifests as a set of activities, practices and procedures that have evolved historically in response to having to accomplish important goals with specific resources. When the resources change, an opportunity is provided for new means. Sometimes, as with complex technologies, a new set of business goals or "theory" are also often possible.

Getting workers on board with new complex technologies requires thinking very differently about the problem than many have done so far. Our work with the bus mechanics suggests that "learning" is not additive, in that they can simply read and study new ways of operating. Rather, "learning" in this context requires a radical shift in the basic framework, or "operating theory" that organizes their rich content knowledge. This would only be possible, however, with people who are already exhibiting intuitive expertise in a particular domain. However, it may also be that these workers benefited not only from their experiential knowledge, but also from *ongoing contact* with the front line of work. If the MIDAS system was responsible for the decrease in un-needed repairs, they would be the first to experience both the results and the exact ways it contributed.

Once the MIDAS system was put into place, MDBF rose dramatically. We believe that the power of the OpSim exercise lies in the first two parts: (1) when the participants actively encounter the basis for their resistance (existing expertise and "default" ways of doing things) and (2) when they re-tailor it to fit new demands and priorities. We think that explanation, simplification and instruction have not worked because each individual has a different prior perspective that must be reckoned with. In this context, "resistance" to learning may actually be the assertion of existing expertise. Whenever a teacher "simplifies" material for his or her students, he or she is really anticipating the "entry point" of the learners. This method often fails with experienced workers because the entry point is not always predictable or universal (for example, "simplified" is often not helpful for those experienced in thinking through vast amounts of detail). The learners—when allowed—actually do better at breaking it down for themselves in a way that is useful to them. In this sense, it seems that learning through "constructive activity" actually involves reorganizing the way that one already understands something, not simply adding new knowledge to existing expertise.

Further, mechanics increased their use of an "optional" feature of the system, the free-form notes attached to each work record. Mechanics not only entered notes with increasing frequency, but read the notes of others as well. As time went on, these system notes became increasingly in the private language of mechanics. As the mechanics grew more comfortable with the system, it became harder for us to know what they were doing with it. In other words, they grew beyond us in their understanding of what the data were saying and the best way to enter it.

Below is an excerpt from the "notes" section of a work order at New York City Transit:

> *Worked on 7016, which came from ENY minus the following items: one entrance door partition, one station upwright and grabrail, one dome light partition cover and front dest sign lock. Remove dest compart locks from bus 7033—which is waiting for other parts—to meet req. All other items listed were obtain from spare buses at yard. Tap-out damage Riv-nuts installed new ones on same. Interior close to be continue.*

There are two striking features of this passage. The first is the admission of "cannibalism", (stealing parts from one bus to get another into service), a practice that could have led to dismissal before MIDAS was implemented. Using MIDAS, mechanics soon realized that indicating parts shortages in the components fields helped MIDAS correct parts ordering forecasts, making cannibalizing unnecessary. Telling other mechanics where the stolen parts came from helped them address missing parts problems in the cannibalized buses later. Other notes helped the mechanic on the next shift begin where the other left off. The other striking feature is that we cannot decipher very clearly what is going on. In other words, the notes are not useful to us non-mechanics. This trend became more pronounced as the system produced more profound financial benefits. What has happened here is that the system has become a tool for the mechanics, and perhaps this has been the problem all along with failed technologies.

3.1 The implications for low skill workers

In this paper we have described one of our projects in an attempt to make several other general points. One of the more important points is that the structure of workplaces is changing with advanced technology. There are a number of consequences to this. As with the MIDAS project, so called low-skill workers became central to the success of senior management's financial goals. One could argue that this is always the case to some extent. What is different here is that the directness of the role and the workers' awareness of it helped them hone their own skills toward actualizing transit wide goals. In the end, their exact method of increasing MDBF (increasing the first pass success of repairs by using data and "notes") was unanticipated by both senior leadership and the MIDAS software designers. In other words, once they understood the goal, they invented the best solution using the increased visibility provided by the software and their knowledge of the front line content. Given that most low skill or entry level workers are at the front line of work, we think that the ability to translate goals into action based on front line experience will be increasingly important. We suspect that the front line experience may not have to be extensive as much as ongoing. As we saw throughout the OpSim, operations managers who were once on the front line for many years were not able to formulate methods as well as those still on the front line.

Clearly the notion of "prerequisite" skill—as we now define it—seems to be outdated in this context. Preparation for inexperienced, low skill workers may need to focus on the changing context of their role in work and less on procedures involved or skill sets that map onto to school topics. On the other hand, preparation for experienced, skilled workers with intuitive expertise can no longer be viewed as an additive process that does not engage, and provide an opportunity for, reorganizing their rich content knowledge. At this point in our thinking, we are uncertain that any new "basic" skills will be ever be identified that clearly map onto the ways that front line work is changing. The mechanics in the study we have described may have done nothing more than realize what was needed and then reorganized and redeployed their existing expertise to meet situational demands. The OpSim intervention, instead of dictating the ways in which the mechanics needed to learn the new system, provided an activity space and a set of non-negotiable outcomes. It may be that our OpSim intervention was a success because, with minimal instruction, it gave them than a way to understand their workplace situation in a different way.

References

Boldt R (1994) TCRP synthesis #5, management information systems, a synthesis of transit practice. Sponsored by the Federal Transit Administration, Transportation Research Board, National Research Council publications, Washington, D.C

Boldt R (2000) Management information systems: a synthesis of transit practice. Sponsored by the Federal Transit Administration, Transportation Research Board, National Research Council publications, Washington, D.C

Capelli P, Bassie L, Katz H, Knoke D, Osterman P, Useem M (1997) Change at work: how American industry and workers are coping with corporate restructuring and what workers must do to take charge of their own careers. Oxford University Press, Oxford

Chamberlain ES, DiBello L (1997) Iterative design and implementation: a model of successful scheduled maintenance technology deployment. Transportation Research Record, August 1997, no. 1571. National Academy Press, Washington, D.C

Chi M, Glaser R, Farr M (1988) The nature of expertise. Lawrence Erlbaum and Associates, New York

DiBello L (1996) Providing multiple "ways in" to expertise for learners with different backgrounds: when it works and what it suggests about adult cognitive development. J Exp Theor Artif Intell 8:229–257

DiBello L (1997) Measuring success in non-trivial ways: How we can know that a DSS implementation has really worked. In: Proceedings of the 1997 IEEE International Conference on Systems, Man and Cybernetics. Society of the Institute of Electrical and Electronics Engineers, Inc

DiBello L (1998) Exploring the relationship between activity and the development of expertise; paradigm shifts and decision defaults. In: Zsambok C, Klein G (eds) Naturalistic decision making. Erlbaum Publisher's book series "Expertise: research and applications". Lawrence Erlbaum and Associates, New York, p 43

DiBello L, Kindred J (1992) Understanding MRPII systems: a comparison between two plants. Technical report for Laboratory for Cognitive Studies of Work, City University Graduate School, New York

DiBello L, Glick J (1993) Technology and minds in an uncertain world. Paper presented at the annual conference of the national society for performance instruction, Chicago, 12–16 April 1993

Dreyfus H (2000) Being in the world: a commentary on Heidegger's being and time, division 1. The MIT Press, Cambridge

Dreyfus HL, Dreyfus SE (1986) Mind over machine: the power of human intuitive expertise in the era of the computer. Free Press, New York

Glass RL (2005) IT failure rates—70% or 10 to 15%?. IEEE Software, Washington, pp 100–112

Klein G (1999) Sources of power. MIT Press, Cambridge

Klein G (2004) The Power of intuition: how to use your gut feelings to make better decisions at work. Doubleday, New York

Klein G, Calderwood R, MacGregor D (1989) Critical decision method for eliciting knowledge. IEEE Trans Syst Man Cybern 19:44

Rogoff B, Lave J (1984) Everyday cognition: its development in social context. Harvard University Press, Cambridge

Scribner S (1988) Head and hand: an action approach to thinking (Occasional Paper No.3). National Center on Education and Employment, New York

Scribner S, Sachs P, DiBello L, Kindred J (1991) Knowledge acquisition at work. Technical report no. 22. National Center on Education and Employment, Teacher's College Columbia University, New York

Scribner S, DiBello L, Kindred J, Zazanis E (1992) Coordinating knowledge systems: a case study. Research report prepared for The Spencer Foundation. Laboratory for Cognitive Studies of Work, City University of New York Graduate School and University Center, New York

The Standish Report (1994) Chaos chronicles. Technical report, The Standish Group, Boston
The Standish Report (1998) Chaos chronicles. Technical report, The Standish Group, Boston
The Standish Report (2003) Chaos chronicles v3.0. Technical report, The Standish Group, Boston
Vygotsky L (1987) Thinking and speech. In: Rieber RW, Carton AS (eds) The collected works of L.S. Vygotsky, Volume 1: problems of general psychology. Plenum Press, New York
Zsambok CE, Klein G (1997) Naturalistic decision making. Erlbaum, Mahwah