Process Descriptions as Organisational Accounting Devices: The Dual Use of Workflow Technologies

Paul Dourish
Dept. Information and Computer Science
University of California Irvine
Irvine CA 92697-3425 USA

jpd@ics.uci.edu

ABSTRACT

Workflow technologies present a problem for CSCW. On the one hand, they are perhaps the most successful form of groupware technology in current use; but on the other, they have been subject to sustained and cogent critiques, particularly from perspective of the analysis of everyday working activities. This leads inevitably to the question: in the face of these critiques, just why and how do workflow technologies prove effective? This paper suggests that part of the solution lies in the fact that workflow technologies play more than one role in organisations, and that, in fact, the success of workflow technologies may have little to do with the typical relationship of those technologies to the accomplishment of everyday work. On the basis of the notion of a dual role for workflow technologies, I lay out a framework for considering the design and analysis of workflow systems that may help to bridge between these two roles.

INTRODUCTION

For better or worse, workflow technologies have become one of the most successful genres of systems supporting cooperative working. Workflow technology emerged not least from early research in the then-nascent field of Computer-Supported Cooperative work, as represented by Winograd's research on the language/action perspective (Winograd, 1986; Winograd and Flores, 1986) and the family of products that it engendered. Subsequently, with the popularity first of Business Process Reengineering and later of the Internet (and "intranets"), workflow technology became the dominant paradigm for "groupware" in everyday commercial computing practice.

If workflow technology is CSCW's first child to reach adulthood, adolescence was anything but easy. From the earliest developments, workflow technologies have been subject to sustained theoretical and practical critiques. In particular, practitioners of the equally paradigmatic "situated action" perspective on the practical accomplishment of collaborative activity have been vocal critics of the workflow stance. They have argued that workflow technologies, rooted as they are in predefined formal descriptions of working processes, are fundamentally in opposition to their orientation to the sequential organisation of work as an improvised, moment-by-moment accomplishment (as well as being politically questionable, to boot). As the years have passed, and the arguments have raged back and forth, little compromise has been sought and little afforded.

It is not my intention, in this paper, to attempt to pour oil on those troubled waters, nor to chart a new channel through them. Elsewhere, with various colleagues, I have argued for a perspective on workflow technology which introduces a greater separation between process description and actual work, and which frames workflow systems as a coordinating technology that mediates between these two levels (Bentley and Dourish, 1995; Dourish et al., 1996). In this paper, however, my goal is to step back and consider more generally what workflow technology actually does.

This might seem to be a naive, or unnecessary, question. Surely we all know what workflow technology does? After all, we have plenty of examples of workflow systems in action. I will argue that our consideration of workflow technology in action is generally based on only a partial view of what the technology is doing. Drawing on an earlier ethnographic investigation of the use of a workflow system in a particular organisation, I will argue that workflow technologies play a dual role, and that, in considering them, we must take each role into account. On the basis of this, I will reflect on a framework to guide the assessment and design of workflow systems.

WORKFLOW TECHNOLOGIES

The term "workflow" covers a wide range of potential systems, including systems for email processing, forms management and document routing, in addition to more familiar systems from the CSCW literature. It is worth being clear about precisely the range of systems that will be of interest here.

For the purposes of the discussion presented here, I take the fundamental, constitutive component of a workflow system to be a *process representation* that drives the functioning of the system. A "process" is a set of tasks to be performed, perhaps entirely by people or at least with some human involvement, and some form of ordering amongst them. A process representation is an explicit computational representation of this process. So, I do not take a programming language compiler to be a workflow system; although it may have a representation of tasks (such as lexical analysis, optimisation and code generation), these are tasks for the system itself to perform, rather than user tasks to be coordinated. Similarly, I do not take, for instance, Sepia (Haake and Wilson, 1992) to be a workflow system; although it supports various phases of a writing process through its use of different tool spaces, it does not explicitly represent that process. Examples of workflow system, by this definition, include Milano (Agostini et al., 1997), Action Workflow (Medina-Mora et al., 1992), Regatta (Swenson et al., 1994), and Freeflow (Dourish et al., 1996)[1].

In this section, I will outline some of the current issues surrounding workflow technologies and their application. I will then go on to review an ethnographic study of a workflow system in action, which will help to locate these issues in a larger frame of workflow in use.

Arguments in the Case "For"

Proponents of workflow technology observe that common patterns of interaction in collaborative work, such as requests for information, commitments to perform tasks, approvals of activities and so forth, operate across a wide range of processes and constitute, essentially, a "grammar" of action. Specific processes, such as processing a travel request, fulfilling a product order or reviewing a paper for a technical conference, can be described according to this grammar. A specific workflow application, then, will encode a specific organisational process according to that grammar. The terms of the grammar vary from system to system. At its most primitive, the grammar might involve a notion of "tasks" (components of a larger process), the transitions from one task to the next, and the groups of people who can perform each task (often through an intermediate notion of "role"). A higher-level grammar might offer the ability to describe the task according to more abstract notions of, for example, commitments and negotiations, or might even encode these into the very notion of task.

Managing a collaborative process according to a process description, proponents argue, provides a number of benefits. Most particularly, it allows the computer to manage or take over the complex issues of coordination and negotiation between individuals; individual users can concentrate simply on the *performance* of the activity, and not on its coordination. Using roles, at the same time, allows the work to be evenly distributed across an organisation as appropriate, affording the efficient use of scarce human resources.

Workflow technologies became particularly popular when coupled with Business Process Reengineering (BPR). BPR is an approach to business management which seeks to improve organisational efficiency by taking a holistic view of the activities of the organisation from an explicitly process-based perspective. BPR advocates the use of this process perspective to "reengineer" or radically realign organisational practice; and it uses workflow systems as an enabling technology for introducing process-based approaches into practice. Thus, advocates of BPR have often become advocates of workflow, since workflow technologies are their means of effecting organisational change.

Arguments in the Case "Against"

For as long as they have been an active area of CSCW research and practice, workflow technologies and their underlying theoretical orientations have been subject to persistent critiques, originating particularly in sociological investigations of the structure of everyday action, and the "situated action" perspective common to CSCW research. A variety of problems have been raised with the workflow perspective; I will outline three here.

The first objection concerns the account of work that process-based systems offer. Workflow critics object that the process-based descriptions that workflow systems embody are based on a notion of uniform, repetitive, routine work, performed the same way time after time. The process that the workflow system embodies is like a program to be executed each time the work is to be performed. However, they argue, ethnographic studies of work, even the supposedly "routine" work of business processes, has revealed that the sequential organisation of activity is an ad hoc, improvised affair (Suchman, 1985; Button and Harper, 1993).[2] While formalised procedures may be useful as a resource in the management of activity, or even as a description of activity, they have no generative function in organising behaviour. The regimentation of process-based workflow systems stands in opposition to the observed fluidity of working activity.

A second criticism has attacked the theoretical foundations on which workflow systems are based. A number of current workflow technologies are theoretically founded in Speech

1. Although it will not be immediately relevant to the argument to be presented here, it is worth noting that these are all *generic* workflow systems, that is, frameworks within which specific workflow applications can be generated and/or executed. It will not be necessary, here, to distinguish between workflow "systems" and workflow "applications".

2. This is not to say, of course, that workflow systems are not based in observations of work, of their own sort. Randall et al. (1995) discuss the relationship of ethnographic investigations of work to those of BPR.

Act theory, and the notion of "Conversation for Action" as originally developed, in this application, by Winograd and Flores (1986). Critics such as Button (1995) have argued that Speech Act Theory is, itself, flawed on a number of counts as an appropriate theoretical foundation on which to build a model of conversation or linguistic behaviour. These critics typically focus their attention on the irreconcilability of speech act theory's explicit, stable and abstract categorisation of utterances with the situated, indexical and circumstantially contingent nature of linguistic action.

The third criticism is more political in nature. Suchman (1994) takes as her starting point the observation of the power wielded in categorisation and naming:

> *"systems of categorization are ordering devices, used to organize the persons, settings, events or activities by whom they are employed or to whom they refer." (p182)*

It follows, then, that the categorisation of actions that lie at the heart of a workflow design, in the creation of process representations, must be seen in the context of a set of people to whom they refer, and whose work is to be ordered by them, for some set of purposes. To the extent that a process representation presents a "cleaned up" view of the world that strips away irrelevant detail to reveal the essence of the work, the question must be asked, to whom are these details irrelevant? Who defines the essence?

The Conundrum

As I intimated above, it is not my goal, in this paper, to take a position in this debate. Instead, what motivates this paper is a simple question: why is workflow technology commercially successful? There can be no question that, however long the debate proceeds in conference presentations, panel sessions, and later over drinks, workflow technology continues to be deployed and used. It is, arguably, one of the most successful form of technology to have emerged from research in CSCW. Does it mean that the critiques are invalid? If so, how? If not, then how can that be?

A WORKFLOW SYSTEM IN USE

To look for answers, I will turn to an ethnographic investigation of the use of a workflow system in a commercial organisation. This is the study of Establishment Printers, reported first by Bowers, Button and Sharrock (1995). Since this study has been described before, I will give just a brief outline here, before turning to the particular feature of the situation on which I wish to concentrate.

Bowers, Button and Sharrock report on the use of a workflow system (called PF2) in an organisation called Establishment Printers, or EP. EP runs a number of large production printing facilities, as well as running print facilities located at directly at client sites, through an "outsourcing" arrangement.

In a number of their sites, alongside existing traditional print technologies such as offset lithography, EP is experimenting with the use of very high-end digital printing and reprographic devices. These machines are built around digital processing components and can be connected together on digital networks, accessed via networked computer systems, and so on. For EP, this offers the opportunity to electronically receive and distribute jobs and manage their printing service.

The computational nature of these devices implies that the printing devices can do more than simply print. Software can be loaded onto these machines to interact with system integration, job management, accounting and inventory software systems. PF2 is one such system. It was designed explicitly to manage print shop activities, and its hardware and software can interact directly with the digital reprographic machines. Print shop workers can interact with PF2 through "shopstations" available on the shop floor, and they use these stations to update the system with information about jobs in progress, such as which job it is, who is working on it, what its requirements are, what materials have been consumed or wasted, and so forth.

Bowers, Button and Sharrock report a number of problems encountered in the deployment of PF2 at EP. They detail a range of ways in which PF2 interfered with the ability of the shopfloor workers to manage the flow of work through the printshop. "Managing the flow", in this case, involves a careful regularisation of the working activities of the print shop so as to avoid both lulls and backlogs. For example, when the operators know that a regular job will soon come up, they may figure this into their plans for allocating jobs to machines, and begin the regular job ahead of time so as to take advantage of machine time that might otherwise be wasted. PF2, however, cannot allocate a job number before the job is actually submitted, and without a job number it cannot proceed. "Jumping the gun" may allow the print workers to manage their activities evenly, but it is not an activity that PF2 supports. Similarly, PF2 cannot accommodate individual workers performing two different jobs at the same time, or overlapping the different stages (such as printing and finishing) of a single job, even though these are common procedures that they employ to even out the flow of work through the print shop. Nor can PF2 represent the fact that multiple workers might be working together on a single job. PF2 will disallow such actions, and so interferes with the workers' ability to manage the smooth flow of work through the print shop.

At the different sites they studied, Bowers and his colleagues observed a variety of responses to these problems. These included technical workarounds (such as fictitious worker identifiers that could be used when one person worked on two jobs, or facilities in PF2 that allowed multiple jobs to be processed together as a "gang"), and administrative workarounds (such as maintaining records on paper, and then resolving these with the electronic records stored by PF2 later). Workarounds such as these were necessary to

address the persistent and serious problems that PF2 introduced if the print shop was to function at all.

A Second Role for PF2

As can be seen from this discussion, the use of workflow technologies was problematic at EP. PF2 interfered with the organisation of work on the print shop floor, and necessitated the development of a variety of dodges and workarounds in order to get on with the primary goal of getting the job done. We might therefore ask the question, in the face of all this adversity, why was PF2 still in use? The system was surely a problem for the print shop workers, since its impoverished model of the work of the print shop floor interfered with the ways in which they organised their work. At the same time, it was surely a problem for the print shop management, since it lowered productivity, resulted in more overtime and interfered with the organisational goal of getting the maximum value from the print machines. This being the case, why was it still in use?

The reason that PF2 was still in use at EP was, at heart, simple; the use of PF2 was *contractually mandated* by EP's customers. The use of PF2 to manage the process was a legal stipulation. In those sites discussed by Bowers, Button and Sharrock, EP was managing dedicated printing facilities run on behalf of (and often in the premises of) other organisations. Those organisations, as part of their contract with EP to "out-source" their printing needs, required EP to use PF2 to manage the process.

The reason for this stipulation was that, on the basis of information collected from the reprographic machines on which it was running, PF2 could produce monthly reports on the activities of the print shop, detailing the number of hours worked, the jobs performed, the resources consumed, and so forth. The host organisations who had out-sourced their document production work to EP wanted to be able to ensure that EP was managing the job efficiently and appropriately, and would do this through the reports that PF2 generated. So, while, on the one hand, the print shop managers would find PF2 troublesome and struggle with the ways in which it interfered with the efficient management of the print shop, they would find at the same time that PF2 was an absolute requirement in that, without it, they could not maintain their relationship with their customers. So, on those occasions when print shop workers would work around the features of PF2 so that they could get the job done, someone would often have to go back, later, and retroactively process the information in such a way that PF2's monthly report would provide an adequate and readable account of the work of the print shop, eliminating fictitious workers and those other elements of the "coping strategies" the print workers had developed.

Although this effort was necessary to meet their contractual obligations, the use of PF2 as a view of the work of the print shop also served EP's interests. They believed that their clients failed to understand what was their jobs entailed, and so valued the use of PF2 as a way to convey this.

In other words, PF2 served two roles for EP. Its first role was to manage the work of the print shop floor, ensuring that activities were carried out in the correct sequence and managing their coordination. The problems concerning the implementation of PF2 in actual everyday use that are detailed by Bowers and colleagues relate to the use of PF2 in this role. At the same time, PF2 also acted as a means by which the print shop could account for its own activities. It was a device whereby the multifarious activities of the print shop could be regularised, regimented and organised for inspection by the client organisation. In this second role, PF2 was considerably more successful.

THE DUAL USE OF WORKFLOW TECHNOLOGIES

So, in order to understand the success of PF2, or at least its persistence, we have to understand both the roles that it played in this organisation. It acts – at one and the same time – both as a *coordination technology* and as an *organisational accounting device*.

Traditional approaches to workflow technology have primarily emphasised this first role. Workflow systems, as technical artifacts, have largely been understood in terms of their coordination function, and their aim to relieve the user of the "burden" of coordination. As such, their design, deployment and evaluation have largely been conducted in those terms. However, particularly in the face of the "situated action" critiques, an analysis purely in terms of this coordinative function leaves us at something of a loss to explain the success of these systems in many settings.[3]

Organisational Accounting Devices

The second role of workflow systems, what I refer to here as as organisational accounting devices, is one that Suchman (1993) has classed under the umbrella of "technologies of accountability." Accountability, in this case, has a very specific meaning.[4]

The concept of "accountability" is central to the ethnomethodological perspective, as developed in the work of Harold Garfinkel (1967). Ethnomethodology is concerned with the common-sense methods and understandings by which practical social action is accomplished. Garfinkel's goal was to reorient sociological reasoning, moving away from traditional models which framed social action in terms of nebulous social rules that governed the social organisation of activity. Ethnomethodology's observation was that these

3. This is not to say that all critiques have been formulated in these terms; I shall return, later, to the relationship between Suchman's critique of workflow and the perspective of organisational accounting devices.

4. Suchman actually means "accountability" in both this technical sense and in the sense of fiscal accounting, but it is the technical meaning that is relevant here.

social rules were actually the outcomes of social action, not the motivation for it. Consequently, in this respecification of the subject matter of sociology, Garfinkel also turned ethnomethodology's attention away from abstractions such as class, capital and economy, and towards the mundane accomplishment of everyday practical social action. His concern was to find how, within the structure of this activity, one could find the genesis of stable social order.

One of Garfinkel's key insights is that the means by which sequential activity is organised and the means by which it is recognised are one and the same. What this means is that the sequential organisation of activity contains, to anyone who is able to see it, the means to understand just what that activity is. Garfinkel refers to this as the "accountability" of action, making it available to others:

> *"When I speak of accountable [...] I mean observable-and-reportable, i.e., available to members as situated practices of looking-and-telling. I mean, too, that such practices are an endless, ongoing, contingent accomplishment; that they are carried on under the auspices of, and are made to happen as events in, the same ordinary affairs that in organizing they describe" (Garfinkel, 1967:1).*

So, the accountability of action resides in the way that it renders that activity as "observable-and-reportable". We use the term "accounting device" to describe a scheme or mechanism by which such accountability is achieved. Bittner (1965) has discussed the application of ethnomethodological perspective to organisations; we use the term "organisational accounting device" to describe a means by which the activity of an organisation can be made accountable. We will explore Bittner's concerns in more detail shortly.

First, though, what does it mean, to say that process descriptions serve, as in the case of PF2, as organisational accounting devices? It means that the process description furnishes, to those who must offer an account of the organisation's activity, a *structure within which* that activity can be made sensible. It regularises and describes the work of "being" that organisation, so that the multifold activities of the organisation can be described as contributing to that organisation's function. Let me set out some specific features of this.

The first significant feature is the way in which the accounting device organises a view of the activity itself. It is not simply that the process description shows what is going on, although that is clearly important. What is particularly significant here is that the process description shows that what is going on *is* the work of being the organisation. That is, it says not simply "job A is being processed on machine M, and job B has been completed"; instead it says, "The job of being a print shop is making sure that jobs get printed, which involves these various steps; and, as you can see, the activity in which we are engaged is oriented towards accomplishing each of those steps in a coordinated manner". That is, at its most fundamental level, it does not simply describe the activity, but renders it observable-and-reportable *as being* the activity that print shops exhibit.

A second feature we might consider is the question of to whom, and for whom is this account manifested. Garfinkel's notion of accountability brings with it a notion of membership, or more generally, the notion of being a participant in a set of practices to a degree sufficient to be able to offer and to recognise an account. In other words, when action is rendered observable and reportable, we must ask, observable by whom? In the case of PF2 at EP, the process description rendered the activities of the print shop visible and reportably rational to EP's customers, whose concern was to ensure that their money was being well spent. In other words, the job of the process description, in this case, was not to show the print shop as being a place where people stopped for breaks, moved supplies from one place to another, monitored machines, or engaged in any of the other myriad activities that fill a working day. Instead, its job was to show, on behalf of the print shop, how the print shop was a place where client's needs were met promptly and efficiently.

The Relationship Between the Roles

Before we can move on to discussing some of the technical implications and opportunities presented by this dual-role view of workflow technologies, we need to consider the nature of the relationship between the roles.

The argument I have presented is that a significant feature of the use of workflow systems in real settings is the way in which they provide an organisation with the means to account for its activities. This role accompanies the more traditional technical conception of the workflow system as managing the coordination of individuals and tasks. In current systems, these two roles are by no means independent. It is through the very regularisation implied by the coordinative function that the account can be given. That is, the only way for the system to be able to account for the activity is to ensure that all and only activity which can be appropriately codified is performed. What the process description offers is a structure for describing the activity, but the only way to ensure that the activity can be *described* in this way is to introduce it as a mechanism for the *performance* of the work. The coordinative role, then, does not simply sit alongside the explanatory role, but, in traditional technological approaches, *enables* it. From this arises the regimentation and restrictions that workflow systems introduce into the management of activity; activities can only be undertaken that can be adequately accounted for by the process model.

This is a source of many of the problems that critics of workflow have identified. The process description is being used, here, to serve two purposes; that of an organisational account, and that of the management of action. The key observation that, although this may be a feature of many current workflow systems, it need not be. The conflation of these two purposes arises, I would argue, from the fact that

the dual purposes are not typically recognised. Once they are taken to be two separate but related functions achieved by the system, we can begin to investigate opportunities for relating the two goals in more creative ways than have traditionally been available to us.

ACCOUNTING FOR ORGANISATIONAL BEHAVIOUR

The preceding section introduced the notion of organisational accounting devices as a role that process descriptions play in the everyday use of workflow systems. In this section, I will explore relevant perspectives from other analyses of organisational activity.

The Concept of Organization

I alluded earlier to Bittner's exploration of an ethnomethodological view of organisation (Bittner, 1965). Bittner's goal in that paper is to reflect on the concept of organisation as it is used in sociological research, and to open up the topic of how organisational research is and should be conducted. In particular, he concentrates on the relationship of the formal and informal lives of organisations, and on how these notions feature as facets of sociological theorising. In the course of putting forth an ethnomethodological critique of the sociology of organisation, Bittner introduces some illustrations of features of organisational life as understood from the ethnomethodological perspective, which have some bearing on this discussion.

One of these is what he refers to as "the gambit of compliance." From the perspective of the concept of organisation as a "rule of conduct," he discusses how this rule might operate and relate to the activity that it "governs." The underdetermining nature of rules is a common ethnomethodological theme, and so this throws light on the question of how rules might be seen to operate in an organisational context:

> *"[A] good deal of the sense we make of things happening in our presence depends on our ability to assign them to the phenomenal sphere of influence of some rule. [...] Extending to the [organisational] rule the respect of compliance, while finding in the rule the means of doing whatever needs to be done, is the gambit that characterizes organizational acumen." (original emphasis)*

In other words, the rules constituted by an organisation are features by which activity can be seen and understood as rational. The compliance of activity with organisational life is the result of an active and explicit process.

Another of Bittner's examples is a reconsideration of the concept of formal organisation as "a model of stylistic unity", that is, a unifying whole by which the meaning of individual elements and actions can be found:

> *"The resulting coherence will be in evidence as outwardly proper conduct and appearance."*

These two elements of Bittner's analysis are, clearly, strongly related to the perspective developed here of organisational accounting devices. For Bittner, the function that formal organisational structures serve is to provide an interpretive frame for organisational activity that simultaneously gives it meaning and while reinforcing and reproducing the institutional reality of organisation. The role proposed here for process descriptions in workflow systems is very similar; they provide a framework for seeing and presenting the organisational activity as "outwardly proper conduct", by providing a set of rules to which organisational activity can be seen as being compliant.

Stakeholders and Boundary Objects

In presenting a notion of organisational accounting devices, we must also address the question, just who or what is the organisation? An important feature of the EP study that has not been addressed so far concerns the different groups involved.

The two roles of the workflow system involve three different groups; the print shop workers, the print shop management, and the organisation to which EP contracts. These groups have different concerns and different issues at stake in the deployment and use of a workflow system. The customers want to ensure that their money is being spent wisely, particularly since they have chosen to outsource what was previously an in-house print service. The print shop management want to ensure that the customer stays happy and will continue to support the service, and that the shop itself functions efficiently and profitably. The shop floor workers want to ensure that work keeps coming into the print shop but that superhuman efforts are not regularly required of them to get the job done. The local management and workers, together, want to defend their work in the face of technological developments (such as distributed printing) that might threaten it. In the centre of all this is a process description which, in theory, supports all these different needs. As we have observed, the practicalities are somewhat different.

Loosely, we might think of process descriptions as serving, then, as boundary objects (Star, 1989; Star and Griesemer, 1989). Boundary objects are:

> *"objects that are both plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to maintain a common identity across sites." (Star, 1989: 393).*

I give this full definition here because it reveals important consequences for the technical elaboration of the dual-use perspective. It seems clear that the process description, in the case of EP, is intended to act as a boundary object in the most general sense. It is a mediating representation by which the different goals and needs of the various stakeholders can be represented and communicated, although those needs may not, themselves, overlap or manifest themselves in the same ways. And yet, in implementation, it fails to meet the criteria that Star lists in her definition. Where the process descriptions fall down in their plasticity. The process description at the heart of PF2 is a rigid and unyielding structure, and all activity must be shoe-horned into it, resulting in precisely the sorts of failures that Bowers, Button and Sharrock describe,

whereby the process description undermines the very accomplishment of the work it seeks to describe and delineate. In its rigid inflexibility, it imposes, rather than reveals, an order to the activities of the print shop.

This brings us back, of course, to Suchman's concerns about the politics of categorisation and formalisation at work in the development of workflow technologies and applications. Suchman's discussion of categorisation as a form of discipline and social control draws attention to the fact that, of the various stakeholders, it is the joint organisational management (the managers of the print shop and accountants of the customer organisation) whose view of the process of shop floor work is to be enshrined in the technology. In response to discussion of her original article, Suchman (1995) is at pains to observe that her goal is not to "demonise" management, nor is that my point here.[5] Rather, I want simply to observe that the traditional workflow approach, in which a single representation is used as the basis for representing and categorising activity to quite different ends, inherently introduces problems for one perspective or another.

The question before us now, then, is whether the dual-use perspective on workflow offers any technical alternatives; and that is the next question to be addressed.

REIMAGINING WORKFLOW

This analysis suggest that workflow technologies, in their present form, play multiple roles in an organisation. In turn, this suggests that our analyses of the use of workflow systems and the roles of process descriptions need to be broader and to take the "dual use" into account. However, we can take this idea further. In this section, I want to turn my attention from the specific concerns of organisations using workflow technology to our conception of workflow technology in general, and consider what the dual role might imply for the design of workflow systems.

Workflow as a Visualisation Technology

The traditional model of workflow is as a *coordinative* technology. As has already been described, this manifests itself, in observational or analytic accounts of the use of workflow systems, in a preoccupation with the relationship of the system to the efficient management of tasks, the coordination of individuals and the concerted achievement of work. At the same time, this concern with coordination has also manifested itself in the ways in which the technology is conceived and designed. It places an emphasis upon such features as negotiation protocols, perspicuous task descriptions, on the transitions that can be taken from one task to another, and on the people who can be called upon to play different roles in the process.

The alternative view developed earlier has emphasised, by contrast, the explanatory role that workflow systems can play. This is not simply a description of what is going on, but an articulation of the work available to the system; one that presents it in an organisationally rational way. Technically, we could consider this as a form of *visualisation* system. Information Visualisation, as a field of HCI design, deals with the ways in which representations of data can be presented visually so as to allow the eye to pick out patterns and gain an at-a-glance understanding of a large body of information (Robertson et al., 1993). Our second perspective on workflow technology could be imagined similarly.

This would propose that workflow systems could *visualise* work rather than manage it. The end result is the same: that a given pattern of organisational activity can be rationalised and presented as an ordered, coherent whole. The difference is that, considered as a visualisation task, the role of the system is to *find and present* the order in the work, rather than, as traditionally conceived, to *prespecify and enforce* it.

In this model, then, the force of the process description is explanatory, not controlling. A workflow tool can be invoked in order to present a model of the world according to a particular process description. There are a number of implications. One is that the process description does not need to play any role at all in the actual *performance* of work. Unlike the usual model, where the workflow system gets involved in the delineation of one task from another and in the coordination between them, in this model the workflow tool is simply a way to explain what's going on. The actual work is decoupled entirely from the description of that work. In other words, to say that a job has reached a certain stage in print production does not necessarily imply that all the job reached that stage at the same time, and that some of it is not already in the next stage, or that a single person has been responsible for getting it that far.

This leads to a second implication, which is that, if the process description is not deeply intertwined with the work activity itself, there is no reason for there to be simply one, absolute process description. A variety of process descriptions might be used to give a variety of explications of the activity in question. In this way, a single action or resource could play a role in a number of different concurrent and interconnected processes. For instance, when I order a new PC for my project, I am simultaneously meeting a project goal, taking a step in the budget management process for my group, and initiating a process for the computer support organisation who will configure and install the machine. To have a single activity play a part in a number of concurrent processes in traditional workflow technologies would be difficult without defining a single enormous process representation that encodes the activities of all the different

5. Indeed, an analysis of the situation in terms of a customary opposition of the goals of the local print shop management and the print operators would be inappropriate. In fact, they worked together quite seamlessly in developing ways to work around the system and yet maintain an adequate account of their work to the client, in the face of emerging technologies that might endanger the print shop's existence.

groups, which clearly has its own problems. However, a visualisation perspective in which the activity and the process have been decoupled until they are actively combined in an occasion of "looking-and-telling" means that the multivalent nature of individual activities can be made available to the system and the user.

Process Descriptions as Organising Resources

We can build upon this visualisation approach to take another view of the role of workflow technologies and process descriptions, seeing them as organising resources for the conduct of work.

From one perspective, of course, this could be seen as a codification of Suchman's explorations of the role of plans not as scripts to be performed, but as guidelines around which the actual sequential organisation of activity is improvised (Suchman, 1987). This would certainly constitute the use of a process description as an organising resource, and particularly, a resource for the coordination of collaborative activity. Grinter (1996) provides evidence for the use of process descriptions in just such a role in the work of software engineers. Her study illustrates how the use of software development tools, designed to manage the software development process, are exploited by software engineers to understand the patterns of activity within a workgroup and the relationship of their activities to that of the group. In this way, then, the process description provides them with a view onto the activity of others (often quite literally, as a graphic representation of the activity over a module), and hence a resource by which to organise their own activity. One intriguing route for exploration, in the spirit of the visualisation perspective offered above, would be to integrate the dynamic role of software development tools as visualisations the work of the group, with tools that present visualisations of large code bases such as those of Ball and Eick (1996).

To go further, however, we can also consider using the process description to organise not only an individual's activity, and their understanding of the activity of others, but also the other resources that they may need as a part of this activity. For example, Dourish et al. (1999) describe a prototype system that uses a process description as an index into a historical record of activity over a document and other associated materials. In this case, then, the process description serves to organise not just activity, but also the objects of that activity, and not just in the present, but extended over time.

Process Descriptions Coordinating Documents

One issue of concern that arises when we consider these alternative models for workflow systems is that we may lose those elements of the workflow system that might actually be useful. That is, if we regard workflow as a visualisation tool rather than a coordination tool, does this mean that we also lose workflow's ability to, for instance, route documents through an organisation, alert people to relevant changes and activities, and so forth?

One example that demonstrates that this is not the case is that described by LaMarca et al. (1999). They describe a prototype workflow-like system, called Maui, based on a principle they call "document-centered collaboration." The Maui example is designed around a travel approval process; however, rather than being embodied in a monolithic workflow application, this process is handled through "active document properties." Document properties are objects that users can attach to documents to manage, organise or control them; *active* properties embed small pieces of code that can handle operations for that document, and respond to document events. Maui has been developed on top of the Placeless Documents system, an infrastructure for document management based on active properties (Dourish et al., 2000).

In Maui, the various coordinative functions of a workflow system, such as recording status, routing documents and alerting users of changes are performed by active properties, in response to user actions over the documents such as reading them and editing them. Since the workflow activities are associated with documents rather than with applications, users can use their existing tools to process the documents, and the workflow system similarly imposes no constraints on activities performed over the documents. In this way, then, the coordinative function of workflow systems can be retained at the same time as the restrictions on individual operation are removed. Freeflow (Dourish et al., 1996) took a similar approach, but Maui takes it further by providing a deeper level of integration with a wide variety of existing tools and the automatic execution of active property code.

BROADER CONSIDERATIONS

As the previous section outlined, the "dual use" model for workflow technologies has a range of technical consequences for the design of novel workflow systems and new ways to understand the role that workflow systems can play in organisations. More broadly, though, I want to relate these observations to some issues in CSCW research.

The primary CSCW critique of workflow technologies has concerned, in particular, with the way in which their conceptions of the sequential organisation of work do violence to the details of real practice. The observation here that workflow technologies play a second role is also not a unique one. Indeed, the PF2 example makes just this point, and elaborates Suchman's argument that, as I have described here, workflow technologies are "technologies of accountability" (Suchman, 1993). Similarly, Grinter (2000) and Schmidt (1996) have detailed some of the ways in which systems that represent work can be used this way. In their original paper, Bowers et al. (1995) use precisely this secondary role of PF2 to criticise the "equation of CSCW with informal, non-structural interaction". That is, they point out that the formality of

the description that the workflow system provides is an important feature of its organisational role.

However, the argument I want to present is somewhat different from Bowers et al's critique (much as I agree with it). My point is three-fold.

First, it is critical to recognise that the technologies of accountability in this sitting are precisely and identically the technologies of execution. It is not the case that one system is used to manage the activity, and another to "clean it up" and present it to the outside world as a rational and efficient achievement. If that were the case, then there would simply be no problem here at all. What creates the tension is that these are one and the same system, being put to two different uses. It serves the needs of both, and it is there that the tension arises.

Second, then, the misjustice that the representation does to the practice of the work is inevitable. It may even be a desirable characteristic of the representation, in as much as, if it serves to bring the world into alignment with some formalised model, then the greater the leeway between the representation and the practical work, the the greater are the opportunities to find creative matches between what goes on and the formalised account. It is inevitable that the representation will not do justice to the lived detail of the work; if it did, then it's other purposes would be undermined. Clearly, there are better and worse ways of doing this, but we need to recognise that no single account of working action can adequately serve both purposes.

Third, this mismatch is not automatically a source for concern. In fact, it represents a design opportunity, in the ways in which I elabored earlier. It opens up the possibility for collaborative technologies to move into the space between the work and its representation, and contribute to the procedures by which people will bridge that gap. This argues that a key role for technology in collaborative work is mediating between individual or collective activity and representations that we might want to make of that work for various purposes. This space should not be eliminated, but creatively colonised.

CONCLUSIONS

The question this paper set out to explore was why, in the face of a range of theoretical critiques, workflow technologies have proven to be commercially viable. If workflow systems interfere with the achievement of work, then one would expect that the concomitant drop in productivity would render them useless. This is clearly not the case.

I have argued that one explanation for this is that workflow systems perform multiple roles within an organisation. Critiques of workflow have typically focussed on how the process descriptions in workflow systems relate to the work they coordinate. Drawing on a field study of a workflow system in use, this paper has suggested that workflow systems not only coordinate work, but they also help to explain it. In this sense, process descriptions can constitute organisation accounting devices; a means by which to present the work of an organisation as externally rational. I have argued that the problems observed in studies of workflow technologies, and that have been the focus of both theoretical and practical critiques, are rooted, in part, in the way existing systems attempt to force a single process representation to support these different roles.

Suchman (1993) uses the term "technologies of accountability" to refer to technologies through which a rational order of activities is achieved over for the purposes of practical reasoning and action, and notes elsewhere (Suchman, 1994) that the accountability that these technologies introduce is, often, not only that of the ethnomethodologist, but also that of the accountant. So it is, certainly, in the use of the process description of print shop work at EP. What is particularly of interest, here, though, is the way in which the process description is used by the organisation as a way to *formulate* its activities, or made to them appear as features within the action (Garfinkel and Sacks, 1970). The process description gives to the print shop a way to demonstrate the organisation of its activity *as* the activity of a print shop.

This view of the process descriptions as a means not only to coordinate but also to re-represent the work of an organisation leads to a set of new opportunities for the development of workflow systems. I have outlined a number of them, taking as a starting point a new conception of workflow as a visualisation technology. These are, however, only first steps. The development of new document management platforms such as Placeless Documents offers the opportunity to create a new generation of workflow systems that respect the multiple roles that process representations can play.

Acknowledgments

Much of what is presented here is based on work on the Freeflow system conducted with Allan MacLean, Jim Holmes, Pernille Marqvardsen and Alex Zbyslaw at the Xerox Research Centre Europe (then called the Rank Xerox Research Centre). In particular, these ideas first emerged in a conversation in which both Jim Holmes and Theakston Brewers played significant parts. It also draws heavily on the EP fieldwork conducted by John Bowers, Graham Button and Wes Sharrock, and I am particularly indebted to Graham Button for early discussions of the EP study. Finally, I am grateful to Graham Button and Beki Grinter for valuable feedback on drafts of this paper.

REFERENCES

1. Agostini, A., De Michelis, G., and Grasso, A. (1997). Rethinking CSCW Systems: The Architecture of Milano. *Proceedings of the European Conference on Computer-Supported Cooperative Work ECSCW'97* (Lancaster, UK). Dordrecht: Kluwer.
2. Ball, T. and Eick, S. (1996). Software Visualization in the Large. *IEEE Computer*, 29(4), 33-43.

3. Bentley, R. and Dourish, P. (1995). Medium versus Mechanism: Supporting Collaboration through Customisation. *Proceedings of the European Conference on Computer-Supported Cooperative Work ECSCW'95* (Stockholm, Sweden). Dordrecht: Kluwer.
4. Bittner, E. (1965). The Concept of Organisation. *Social Research*, 32, 239-55 (reproduced in Turner, ed, *Ethnomethodology*. Harmondsworth: Penguin).
5. Bowers, J., Button, G. and Sharrock, W. (1995), Workflow from Within and Without: Technology and Cooperative Work on the Print Industry Shopfloor. *Proceedings of the European Conference on Computer-Supported Cooperative Work ECSCW'95* (Stockholm, Sweden). Dordrecht: Kluwer.
6. Button, G. (1995). What's Wrong with Speech Act Theory. *Computer-Supported Cooperative Work*, 3, 39-42.
7. Button, G. and Harper, R. (1993). Taking the organisation into accounts. In Button (ed.), *Technology in Working Order: Studies of Work, Interaction and Technology*. London: Routledge.
8. Dourish, P., Holmes, J., MacLean, A., Marqvardsen, P., and Zbylsaw, A., (1996). Freeflow: Mediating between Representation and Action in Workflow Systems. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work CSCW'96* (Boston, MA). New York: ACM.
9. Dourish, P., Bentley, R., Jones, R. and MacLean, A. (1999). A Matter of Perspective: Using Process Descriptions to Index Document History. *Proc. ACM Conf. Supporting Group Work GROUP'99* (Phoenix, AZ). New York: ACM.
10. Dourish, P., Edwards, K., LaMarca, A., Lamping, J., Petersen, K., Salisbury, M., Terry, D. and Thornton, J. (2000). Extending Document Management Systems with Active Properties. ACM Transactions on Information Systems, 18(2), 140-170.
11. Garfinkel, H. (1967). *Studies in Ethnomethodology*. New Jersey: Prentice Hall.
12. Garfinkel, H. and Sacks, H. (1970). On Formal Structures of Practical Action. In McKinney and Tiryakian (eds), *Theoretical Sociology*, 337-366. New York: Appleton Century Crofts.
13. Grinter, R. (1996). Supporting Articulation Work Using Configuration Management Systems. *Computer-Supported Cooperative Work*, 5(4), 447-465.
14. Grinter, R. (2000). Workflow Systems: Occasions for Success and Failure. Computer-Supported Cooperative Work, 9(2), 189-214.
15. Haake, J. and Wilson, B. (1992). Supporting Collaborative Writing of Hyperdocuments in SEPIA. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work CSCW'92* (Toronto, Ontario). New York: ACM.
16. Hammer, M. and Champy, J. (1994). *Reengineering the Corporation: A Manifesto for Business Revolution*. New York: HarperCollins.
17. LaMarca, A., Edwards, K., Dourish, P., Lamping, J., Smith, I. and Thornton, J. (1999). Taking the Work out of Workflow: Mechanisms for Document-Centered Collaboration. *Proc. European Conf. Computer-Supported Cooperative Work ECSCW'99* (Copenhagen, Denmark). Dordrecht: Kluwer.
18. Medina-Mora, R., Winograd, T., Flores, R., and Flores, F. (1992). The Action Workflow Approach to Workflow Management. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work CSCW'92* (Toronto, Ontario). New York: ACM.
19. Randall, D., Rouncefield, M., and Hughes, J. (1995). Chalk and Cheese: BPR and Ethnomethodologically-Informed Ethnography in CSCW. *Proceedings of the Fourth European Conference on Computer-Supported Cooperative Work ECSCW'95* (Stockholm, Sweden). Dordrecht: Kluwer.
20. Robertson, G., Card, S. and Mackinlay, J. (1993). Information Visualization using 3D Interactive Animation. *Communications of the ACM*, 36(4), 56-71.
21. Schmidt, K. (1997). Of Maps and Scripts: The Status of Formal Constructs in Cooperative Work. *Proc. ACM Conf. Supporting Group Work GROUP'97* (Phoenix, AZ). New York: ACM.
22. Star, S. L. (1989). The Structure of Ill-Structured Problems: Boundary Objects and Heterogeneous Problem Solving. In Gasser and Huhns (eds), *Distributed Artificial Intelligence, Volume 2*. London: Pitman.
23. Star, S. L. and Griesemer, J. (1989). Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. *Social Studies of Science*, 19, 387-420.
24. Swenson, K., Maxwell, R., Matsumoto, T., Saghari, B. and Irwin, K. (1994). A Business Process Environment Supporting Collaborative Planning. *Collaborative Computing*, 1(1), 15–34.
25. Suchman, L. (1985). Office Procedures as Practical Action: Models of Work and System Design. *ACM Transactions on Office Information Systems*, 1(4), 320-328.
26. Suchman, L. (1987). *Plans and Situated Actions*. Cambridge: Cambridge University Press.
27. Suchman, L. (1993). Technologies of Accountability: Of Lizards and Aeroplanes. In Button (ed), *Technology in Working Order: Studies of Work, Interaction and Technology*. London: Routledge.
28. Suchman, L. (1994). Do Categories Have Politics? The Language/Action Perspective Reconsidered. *Computer-Supported Cooperative Work*, 2(3). 177–190.
29. Suchman, L. (1995). Speech Acts and Voices: Response to Winograd et al. *Computer-Supported Cooperative Work*, 3, 85-95.
30. Winograd, T. (1986). A Language/Action Perspective on the Design of Cooperative Work. *Proceedings of the ACM Conference on Computer-Supported Cooperative Work CSCW'86* (Austin, TX). New York: ACM.
31. Winograd, T. (1994). Categories, Disciplines and Social Coordination. Computer Supported Cooperative Work, 2, 191-197.
32. Winograd, T. and Flores, F. (1986). *Understanding Computers and Cognition: A New Foundation for Design*. Norwood, NJ: Ablex.