K. Kuutti, E. H. Karsten, G. Fitzpatrick, P. Dourish, and K. Schmidt (eds.). *Proceedings of the Eighth European Conference on Computer-Supported Cooperative Work, 14-18 September 2003, Helsinki, Finland*

Dependable Red Hot Action

Karen Clarke, John Hughes, Dave Martin, Mark Rouncefield, Ian Sommerville
Departments of Computing and Sociology, University of Lancaster
{k.m.clarke|j.hughes|d.b.martin|m.rouncefield|i.sommerville}@lancs.ac.uk

Corin Gurr, Mark Hartswood, Rob Procter, Roger Slack, Alex Voss
School of Informatics, University of Edinburgh
{corin|mjh|rnp|rslack|av}@inf.ed.ac.uk

Abstract. We present a brief observational, 'ethnographic', study of the Roughing Mill in a steel plant and use material from recorded activities to provide 'illustrative vignettes' of some aspects of the accomplishment and problems of everyday work. The account provides a 'bottom up' method for developing a more sophisticated and situated view of the problems of dependability. The paper documents the social organisation of work in the Roughing Mill, the interaction between the computer scheduler and the skill of the mill operator in accomplishing 'dependable' production of steel plates from slabs.

Introduction: dependability and socio-technical systems

> "Dependability is defined as that property of a computer system such that reliance can justifiably be placed on the service it delivers." (Randell, 2000)

As computer-based systems – embracing humans, computers and engineered systems- become more complex and organisationally embedded, so the challenges of dependability – of building systems involving complex interactions amongst computers and humans – increase. In these systems, failure, lack of dependability, can result in financial or human loss and, consequently, improved means of specifying, designing, assessing, deploying and maintaining complex computer-based systems would seem of crucial importance. Much of the work on dependability has necessarily, and naturally, focused on massive, extraordinary, public failures such as the London Ambulance Service failure of 1992, the space shuttle catastrophe of 1986, or the Ladbroke Grove train disaster of 1999. This paper is, however, concerned with rather more ordinary, everyday instances of dependability and failure. Instances of undependability in many settings are not normally catastrophic, but are rather mundane events that occasion situated practical (as opposed to legal) inquiry and repair. Dependability can then be seen as being the outcome of people's *everyday, coordinated, practical actions*. Workers draw on more or less dependable artefacts and structures as a resource for their work of achieving overall dependable results in the work they are doing (Voß et al., 2002; Clarke et al., 2002).

In this paper, we wish to explicate how overall dependability is practically achieved in the operations of a steel rolling mill, a rather different setting than most studies of dependability in IT systems have looked at. Here, the research is not situated in bright, clean offices of the services industries or the technologically advanced and safety critical sectors of the nuclear industry or aircraft safety, but in the noisy, dirty and dangerous world of steel manufacture. Although the focus of activity is transforming a slab of red hot steel into steel plate rather than, for example, the provision and control of information, similar dependability issues of timeliness, responsibility, security etc can arise, and need to be resolved, in the interaction between computer systems and human skills. Our research consists of a brief observational, 'ethnographic', study (Hughes et al., 1992; 1994) of the Roughing Mill in a steel plant. Although 'quick and dirty', the fieldwork covered all three daily working shifts and a number of roughing mill operators of varying levels of skill and experience. In the paper we offer 'illustrative vignettes' of aspects of this particular work in the Roughing Mill as an example of a more 'bottom up' method for developing a richer situated view of the practical problems of dependability (Suchman, 1995).

The paper provides us with an opportunity to respecify the problem of dependability, and hence the lessons for IT systems design, by documenting 'real world, real time' practices whereby dependability is rooted within the practical ongoing social organisation of work. Our argument is hardly radical in emphasising the point that any abstract 'rules for dependability' – such as procedures, models, proscriptions, prescriptions, etc. – have to be applied within the context of some socially organised work setting in which those who have to apply such rules have to deal with all the contingencies and other demands on their attention and effort. What is perhaps more radical is the intention to treat this observation seriously as both a research endeavour and as an input to system design. Our interest in the social organisation of work is in how the work activities (which are often glossed and idealised) are actually carried out and accomplished as day-to-day activities with whatever resources, including technological, are to hand and facing up to whatever contingencies arise. As far as system design is concerned – and as we have said elsewhere – such an interest seeks to understand the work as a first priority so that any innovation in system design better resonates with the work as actually done. This point of view is based on two suppositions: first, that most systems fail because they do not resonate with the work as it is actually done as a 'real world, real time' phenomenon; second, that even when the intention is to change the work (to make it more efficient, reliable, etc.), it is always best to have a good idea of what may be lost in doing so to put against any putative gains.

Image 1: Problems – getting 'turn-up'.

The research reported upon in this paper was motivated by several observed 'problems' in the Roughing Mill, some of which may be viewed as relevant to issues of dependability. The dependability issues were manifested in the complex interrelationship of skill, teamworking and awareness that could result, for example, in a range of 'troubles'. These included:

- 'Cobbles' or 'turn-up' of the part rolled slab that makes it difficult, and sometimes impossible, to manipulate the slab through the Mill (see Images 1 and 2).
- Badly shaped slabs coming into the Mill that produce, for example, 'fishtails' or other defects in the finished slab.

- Slab defects produced by the furnace, for example, ‘thermic shock’ requiring the Mill operator to make adjustments in how the slab is rolled and that may mean the final rolled plate will not yield all the ordered plates.
- Various kinds of marking etc. on the slab produced by difficulties in rolling that may influence the quality of the final plate.
- A variety of computer problems related to the identification, measurement and sequencing of the slabs.

Image 2: A ‘cobble’ being lifted from the line.

As in any tightly structured sequence of interdependent activities, such ‘troubles’, even though they are often regular and routine, are ‘troubles’ which detract from the dependability of the system by producing waste, slowing production, creating frustration and increasing overall costs. However, and again as with most systems of high interdependency, achieving ‘smooth’ operation day in and day out is extremely difficult and requires a great deal of experienced skill on the part of the operators of the technology.

The analysis that follows uses a framework of ‘sensitising’ concepts (Blumer, 1954; Blythin et al., 1997) that have been developed over the years in doing ethnography as a contribution to system design. It provides a means of bringing out the grounding of dependability on the social organised skill and competences of those involved in the work setting. As Popitz et al. remark in their much earlier study:

> “It is not sufficient to remark that the individual work activities are embedded within a larger work context. One must be more concrete and with each individual work activity demonstrate how and to what extent cooperation with other work activities is a requirement.” (Popitz et al. 1957, in Schmidt, 1994)

There are, of course, some important differences from the steel rolling mill studied by Popitz et al.; ours is predominantly computer controlled; operators have far greater overview of the whole process and facilities for communicating and coordinating work. It certainly is no longer the case that “operators simply take their cues from the state of the field of work and infer the actions and intentions of their colleagues from that." While it remains the case that "a strip is just a strip, and whatever is done to it is done with the single purpose of transforming it the proper way" (Schmidt 1994: 27), we wish to provide some important insights into just how this is done. The concepts we employ are distributed coordination, the situated orientation to plans and procedures, and the achievement of various forms of awareness of work. To begin, we describe the work of the Rolling Mill and the rolling process in a rather idealised fashion. In this way we can begin to bring out the situated actual activities done by the operators as routines of their daily work.

The roughing process

Like many unfamiliar work activities, the process of rolling a slab of steel appears complicated beyond belief. Ideally, however, it is simple enough. Slabs, or blocks of steel, are rolled into steel plates of varying thickness. The process begins with the available steel slabs being assembled in the slabyard and moved to furnaces. Usually more than one steel plate will be made from each rolled slab so 'as rolled' plated will be cut on one of the shear lines. 'Build rules' are used to determine how many plates of required sizes can be made from standard slab sizes with minimum wastage.

There are two furnaces each with two 'strands' of slabs passing through them. The temperature for each slab is calculated and passed to the furnace controller. The slabs are heated to around 1250°C. Each slab's temperature is repeatedly calculated as it is heated and when it has reached the required value is flagged as ready to roll. A another slab is then pushed into the furnace so moving the strand one along with the slab ready to be rolled dropping out of the discharge end. The mill – really the Roughing and the Finishing Mill – use reversing mills or rollers. The incoming slab will have already been specified, and displayed to the operator, as to be rolled in one of two orientations: length to width or length to length. However, this requirement is not imperative and is sometimes overridden, if, for example, there are flaws in the slab that make it difficult to follow the requirement.

The slab is reduced in thickness by a series of 'passes' back and forth through the mill until the desired thickness is reached. The Roughing Mill 'stand' (where our observations were concentrated) is a large structure that supports two steel rollers turned by two large electric motors. The distance between the rolls – the 'roll gap' – is adjusted by the 'screws'. Slabs are transported on roller tables that are controlled in sections to give more delicate control over the movement of the slab. The slab can be turned on the 'turning table' which consists of alternate rolles, thinned on alternate sides, and which can rotate in opposite directions. Moving side guides are used to square up and centralise the slab for passage through the rollers. The thickness of the slab can at any stage be inferred from the screw position the last time the slab passed through the roll gap. This process is typically fraught with problems since the whole mill is significantly elastic under the forces generated by rolling along with the fact that the rolls expand as they heat and wear as more steel is rolled. The process of 'zeroing' the mill – adjusting it so that the unloaded roll gap is actually zero when the indicator says it is – is difficult to carry out and often poorly understood. The length and width of the slan can be measured by an optical gauge known as the 'Kelk' or 'Accuplan' but only when it is held still on the turning table. At this stage it is important to achieve the final plate width and the keep the slab – by now almost a plate – as close to the ideal rectangle as possible.

The computer calculates the sequence of screw settings and turns. The screws are reset automatically after each pass. The computer requires a width reading when necessary and corrects to achieve an acceptable width and keeps track of what has been rolled. The operator is responsible for manipulating the slab to turn it and enter it through the rollers centrally and squarely after the screws have been automatically set. This involves hand and foot controls. If necessary the operator can also take control of the screws. At each 'pass' through the mill the steel is reduced in thickness by the 'draft'. As the volume remains the same, the other dimensions must increase. Most of this increase appears as extra length in the rolling direction and is quite easy to predict. (There is also some 'spread' outwards as it passes through the rollers. This may be large in terms of product tolerances but is always small in relation to the elongation of the slab. However, it is difficult to predict.) The first target is to elongate one of the slab's dimensions until it reaches the width of the final product. It may be rolled in both orientations until this is achieved. It is then turned through 900 and rolled in the same orientation from then on. The drafts from then on will, ideally, be the maximum possible in order to reduce rolling time and minimise heat loss. Different limits apply at different parts of the process.

Although there are variations according to the composition and quality of the slab, the general procedure in the Roughing Mill is as follows. The slab is pushed from the furnace, through the wash boxes to remove scale and then aligned and centred on the rollers. Information on the monitor in the 'pulpit' – the control room where the operator works – tells the operator the slab quality, its present width and length, the width and length required, the orientation, the 'turning point' (the measured point at which the operator should turn the slab to roll for final length), and the 'finish point' (the point at which the operator should send the slab through to the Finishing Mill).

The rolling operation itself begins with 'pre-broadside passes' through the Mill and the slab is sprayed to remove scale. The operator then 'goes for width' by rolling the slab to produce the desired width up to the 'turning point'. Measurement of the slab through the Accuplan is displayed in the 'pulpit'.

Image 3: Aligning the Slab ─ the mill lights are green.

One red light indicates that measuring is taking place, two that the slab has achieved width. Green lights are displayed for the operator to turn the slab to roll for length. As the operator puts the slab through the mill he turns and aligns it (see Image 3). The scheduler reduces the gauge at each pass – displayed on the overhead monitor and the 'clock') until finish point is achieved. The final pass is a reverse pass. The rollers are then lifted and the plate sprayed on its way to the Finishing Mill.

Of course, in actuality the process rarely goes as smoothly as this. 'Troubles' of various kinds are a regular feature. One prominent trouble is when the part rolled slab 'turns up' to form a U or W shape that makes it impossible to manipulate. There are a number of techniques, all involving heavy manual labour, to recover from such events, but they cause delays and do not always succeed. Although the process is not fully understood, the cure is straightforward. The screw settings should ensure heavy drafting at critical points in the process but this requires considerable experience on the part of the controller. Indeed, many of the more experienced operators will go into manual mode for the last few passes. A related problem is when the plan view of the plate is not the ideal rectangle. If the problem is severe the final rolled plate will not yield all the plates required. The operator does have a degree of control over this but the automatic controller gives no help. Sometimes the length or the width does not turn out as planned and further action is necessary.

In the pulpit, the operator has various monitors and controls. On his left is the furnace monitor (in this case misaligned after a mishap with a crane), load measures, mill light, screw inject (this can also be done through the central control pad) and levers for screwing the rollers up and down. To the right of the monitor is the main control pad, the main monitor, rack lever, amp meter, the monitor for sprays, and temperature gauges (see Image 4). To the front of the operator are a number of foot pedals for turning the lab on the rollers and sending it through to the Finishing Mill. There is also a head level display that provides the reference points for the slab currently being rolled. Outside, on the Mill itself, the mill 'clock' and measuring lights provide further information. Indeed, the mill 'clock' – the way in which the 'hands' alter to reflect the changing of the screws – was a clear and persistent focus as the operator worked.

Image 4: Pulpit Controls: Right Side

Some concepts for analysis

We now want to move onto some analysis of the fieldwork materials using the presentation framework developed out of previous ethnographic studies in a variety of domains. It is important to stress that the framework of concepts is in no way a theory of work, organisation or whatever. The rationale for this insistence on the framework not being a theory would take us too far afield, but, briefly, is connected to Garfinkel's critique of constructivist sociology. It is in an important part a heuristic and practical device for bringing out the generic everyday features of socially organised work settings and, at the same time, presenting these in a form useful for designers. In this particular instance we are interested in whether the heuristic of the framework facilitates the identification of dependability issues. The features we want to illustrate in this paper are as follows: 'distributed coordination', 'plans and procedures' and 'awareness of work'.

Distributed coordination

This points to how work tasks are performed as coordinated activities, that is, as activities that have interdependence with activities done by others who may not be co-located. It is clearly a notion closely tied to the idea of a division of labour but goes further in emphasising the ubiquity of the day-to-day need to achieve coordination within a division of labour. Depending on the work setting and the activities concerned, distributed coordination can take many forms, involve varied technologies and operate at different periodicities. As with the other concepts, 'distributed coordination' is a methodological injunction to treat work settings, persons and activities as embedded in an organised ensemble. The activities and the persons who perform them are interconnected as part of some organisation of activities and persons which has to be coordinated in order to 'get the work done'.

Plans and procedures

'Plans and procedures' refers to one of the more obvious means by which distributed coordination is achieved and supported. Project plans and schedules, manuals of instructions, procedures, workflow diagrams are all ways of enabling persons to use as resources for coordinating work activities. There is no implication here that any particular set of plans, etc., is successful at coordination, or conforms to some ideal standard. The explicit point of plans is to coordinate the work of different persons so that separate work activities, either in parallel or serially, have a coherence and, typically, through this meet other goals such as efficiency, meeting time constraints, beating the enemy, growing the company, and so on (Schmidt, 1997). Although 'plans and procedures' are, of course, about coordination – and often an important resource in its achievement – 'plans' are abstract construction that require implementation within the specifics of the circumstances in which it is to be followed (Suchman, 1987; Dant and Francis, 1998). The accomplishment of a 'plan' is dependent upon the practical understanding of what the plan specifies in *these* circumstances, using *these* resources, and facing up to *these* contingencies. In many cases of 'real time, real world work', accomplishing the plan often involves using local knowledge, 'cutting corners', 'bending the rules', even revising the plan in order to meet its overall objective.

Awareness of work

'Awareness of Work' refers to the way in which work tasks are made available to others and constitutes a major aspect of the means through which co-ordination of work tasks is achieved as a practical matter. As Popitz argues:

> "An operator only operates the system rationally and effectively if each operation is carried out with a view to the necessary cooperation with others … he has to take into account the preceding, concurrent and immediately ensuing operations." (Schmidt 1994: 26)

The various ways in which 'awareness' is developed, in which work is made public are of interest as essential ingredients in 'doing the work' as part of a socially distributed division of labour. It is intended to encompass a range of informal as well as formal phenomena for indicating the 'state of the work thus far'. It is also, as with plans and procedures, intimately connected to distributed coordination in that making others aware of the state of the work is an important resource for making distinct aspects of a division of labour work together. 'Awareness of work' does not point to some psychological property but, rather, to those visible features of the work and its setting by which those involved can make judgements about the 'state of the work'.

Before filling these concepts out using 'vignettes' from the fieldwork on the Rolling Mill, it is important to stress that these concepts direct attention to different features of the work process, features which are not necessarily discrete in the sense that one might see 'distributed coordination' and then 'awareness of work' and then 'plans and procedures'. The same work activity can, that is, often be seen from the perspective of any of these concepts. We want now to move on to illustrate the above concepts using fieldwork material from the Rolling Mill.

Everyday work in the Roughing Mill

Distributed coordination

It can be no surprise that the rolling process involves coordination. It is a process designed from the outset (no doubt based on less technologically sophisticated predecessors) to coordinate the various activities of men and machines involved in turning a slab of steel into a plate. The Roughing Mill is part of a series of work activities beginning with the Furnace, the Roughing Mill itself, the Finishing Mill and the Shear Lines where the 'mother plate' is cut into the various 'daughter plates' ordered by customers. (This series could, of course, be linked to other aspects of the organisation – producing the slab itself, loading, invoicing the customers, etc., etc. Our concern, however, is with the Roughing Mill.) Further, the process of reducing a slab of hot metal to a plate has to be done in 'real time'. The slab cannot be parked for any time until the Roughing or the Finishing Mill are ready for it otherwise it loses heat and becomes unworkable. Moreover, it is a dangerous place, noisy, full of steam and pieces of molten metal. Accordingly, communication and coordinating between the various stages of the process has to be dependable, reliable and unequivocal. Further, processing the slabs into plate is subject to various and not always controllable inconsistencies in the quality of the material being worked with: slabs are not always the right size or shape, not always at a workable temperature, cannot always be 'roughed' satisfactorily, and so on. Yet contingencies such as these need to be dealt with by the operators and in a way which keeps the process running as smoothly as it can be under the circumstances. Accordingly, communication and coordination between the processes needs to be as simple and as quick as possible.

For the Roughing Mill operator, coordination with the Finishing Mill (the next stage in the production process) is important and was achieved in a number of ways. For example, the lights on the Finishing Mill enabled the Roughing Mill operator to see when the Finishing Mill had done with the current plate and therefore ready for the next, as this short extract from the fieldwork illustrates:

> "(points to Finishing Mill) two lights … a red one and a white one … the white one means its finished … so that's a guide to us … if we can see through the steam."

However, as the comment suggests, it is not always possible to see the lights so 'at a glance' information about the state of the Finishing Mill was not always available. Hence, recourse was often made to the RT link (a microphone) in the Pulpit that was also used to alert the Finishing Mill to any problems with the roughed plate. This two-way link also functioned to alert the Roughing Mill operator to anything – such as particularly long slabs or the imminence of 'turn-up' – that might affect his work.

> "he was letting me know that the front end was up … so he was bringing it back just to knock it down … That's another thing we look for … this (slab) finished length is 12 metres long … I notice that (pointing at one in Finishing Mill) was 24 metres (that's) why I'm waiting for him to finish."

The process is not always smooth and, accordingly, coordination in this case between two closely connected processes is essential to regulating the pace of the process in 'real time'.

Coordination with the Furnace was done mainly through the microphone, a monitor and the Mill 'light', the latter being used to control the supply of slabs to the Roughing Mill.

> "I've put the Mill light off … or they might push another one before I start that one."

> "I turned my light off because … if I'd had problems with it I'd have had another one standing there getting cold and I'd have the same problems again."

However, the monitor and RT were mainly used. The monitor is especially important in checking which furnace a slab came from in case its size was incorrect.

> "looking at monitor to see where slab is coming from (*which furnace*) … so if I get another wrong size again I know where its come from. (*Speaking into microphone to Furnace controller)* That's a bad shaped slab, Pete … cut short on one side … I've got to send that one back, mate … it's stuck under the washbox … it's got a dirty great black line through it."

There were innumerable instances in which communication using the RT was used between all sections of the process trying to make the process as smooth as possible. Of course, it sometimes this fell short of the ideal – slabs the incorrect size, taking longer over a process than expected, roughed plates not sufficiently rectangular, etc., – and which have to be communicated to others up and down the process. What we see in these exchanges – and they are but an exceedingly small set of examples that could have been used – is the everyday work that goes into achieving coordination dependably in the process of changes slabs into plates. In this sense coordination is part and parcel of the working division of labour itself, using the resources to hand.

Plans and procedures

As the idealised description of the process suggests, transforming a slab of hot metal into a plate is designed as a linear, step-by-step process moving from the Furnace to the Roughing Mill, to the Finishing Mill and, finally, to the Shear Mill where the plate is cut into sizes for the eventual customer. Of course, much more goes on within these stages, however, the point is that being a step-by-step process certain conditions have to be met before moving onto the next stage. That is, key to the whole process is scheduling and pacing.

As we have already noted, scheduling and pacing were not always straightforwardly achievable. As one, highly experienced, operator wryly remarked about the 'them' who had designed the system: '… for them to design scheduling … is a bit like me trying to design a plane because I've flown in one." (A common enough sentiment, in many settings, that we simply note.) A number of problems inevitably arose when the computer and automatic systems went awry. In one case, for example, computer problems in the Finishing Mill produced wrong readings for number of passes and wrong measures on every pass. In another instance, the computer lost its reference point and the operator had to take over manually. In another case an operator noticed that the computer was failing to update:

> "it's not been giving us first draft reference … it's brought up the plate draft but kept it at whatever we sent the last plate at ... it's not updating on the screen at all…for some reason it's not updating …so there's obviously a fault somewhere … that's why I'm in manual … I don't trust it now because I don't know what it's doing … and the computer hasn't pushed now (provided another slab) because it thinks I'm still at 230 (the initial draft of the plate - 230)."

In this case, the pacing of the mill has been compromised by the computer failure and the operator is waiting for a slab from the Furnace.

The successful accomplishment of a 'plan' or a 'procedure' is dependent on the practical understandings about what the plan specifies in *these* circumstances, using *these* resources, and so on, not least when things 'do not go according to plan'. In such cases it means adapting to the situation at hand as in the following example.

> "… if a slab comes down and its all got thermal cracking … then we'd roll it the other way … tell them … make a note … they'll say … why did you roll it that way".

That is, if a slab appears with a thermal crack on one side, rather than following the computer's instructions and rolling the slab so that the crack appears at the side of the plate and effectively ruins the quality of all the 'daughter' plates that are cut from it, the operator will override the computer and roll the slab so that the crack appears at the end of the plate and may be discarded in the waste. Indeed, there were a number of occasions where operators used their own judgements rather than the computer in order to realise the aims of the procedure. For example, the work of the Roughing Mill critically depends upon slab quality, that is, the metallic composition of the slab and the relative proneness to 'turn up'.

> "I shan't give this a lot of water as it's 269 quality and liable to turn up … with 269 quality a lot of drivers drive with barrel water off to keep the heat in the slab."

> "going for manganese … real hard stuff … we don't use any water … you just have to work real fast."

The operators were aware that different slab qualities are liable to various defects, such as 'fishtails' and 'tongues' (see Figure 1) and the work was adjusted accordingly.

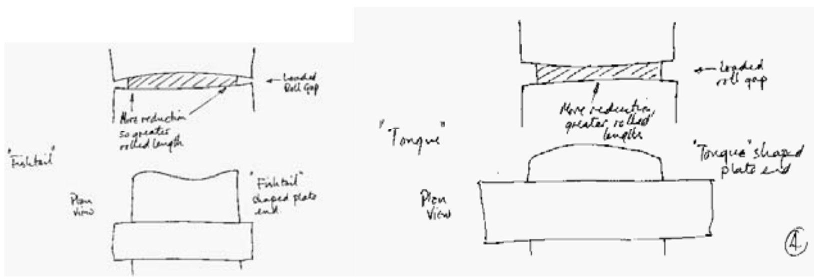

Figure 1: 'Fishtail' and 'Tongue' defects.

However, scheduling and pacing are not simply about doing one's job in isolation from how it might impact on others further down the process. The Roughing Mill operator concludes his task when the slab has been 'finished at measure', that is, rolled to a specified thickness, length and width. Occasionally operators were observed to ignore their 'finish at' measure – instigating an alarm – in order to send the plate through to the Finishing Mill in an adequate state.

> "If I send that at 49 … it's going to shoot up *(turn up in the Finishing Mill)*…it's 233 quality which is the worst one for turn-up … you need a minimum of 3 metres in length … because if you get less than that there's a good chance it could turn-up in the Finishing Mill."

> "instead of finishing at 35 I'll drive it down and put a bit more length on it … less chance of it turning-up then."

This clearly involves knowledge of the work of the Finishing Mill on the part of the Roughing Mill operator as well as knowledge of the Finishing Mill (that plates need to be over 3 metres to roll easily) and the relative capacities of the two mills:

> "… it wants to send it at 60 ... but it's a bit short ... so I take over manually and knock it down a bit (alarm) ... gone to manual ... it wanted to send it away at 50 ... it makes it difficult for the other (FM) drivers ... it'll take him 2-3 passes to get it down to that."

As these examples suggest, 'real time, real world' work often involves the utilisation of 'local knowledge' and 'local logics', commonly interpreted as 'cutting corners' or 'bending the rules', to support the overall objectives of the plan. In the Roughing Mill this was perhaps most obvious when things began to go wrong and operators were faced with slabs 'turning up'. As Rognin and Bannon (1997) suggest:

> "Despite our attempts to automate an ever larger set of control functions, and to build-in forms of automated reasoning and intelligence into these computerised control systems, there is still a crucial need for human agency to monitor and, if necessary, to over-ride computerised systems under special circumstances or unusual conditions." (Rognin and Bannon 1997)

Awareness of Work

> "... horrible plates these are ... they've been turning up all night."

One aspect of 'awareness' that seems of particular importance in the Roughing Mill, concerns having an awareness of what happens to a slab once it has been through the Mill – that is an awareness of what happens to a plate once the Roughing Mill operator has effectively 'done his job'. This involves some knowledge of how the end product is arrived at, what happens when the slab leaves the Roughing Mill and how what happens to a slab in the Roughing Mill may impact on the work done elsewhere in producing the finished product. That is, having such knowledge that not only can the Roughing Mill operator 'do his job' but also he accomplishes it in such a way as to enable others to do their work – by providing a plate with appropriate qualities. Observations revealed a number of successes and failures in this respect – the most obvious failure being lack of much immediate feedback on the quality of the plates produced.

Awareness: Professional Vision

> "… sometimes you can sit here and look at it and think, 'that one's going to be a bastard'."

One way of understanding and analysing awareness issues for operators in the Roughing Mill is in terms of Goodwin's (1994) notion of 'professional vision':

> "... a community of competent practitioners ... expect each other to be able to see and categorize the world in ways that are relevant to the work, tools and artifacts that constitute their profession."

There are a range of subtleties to this notion but in many ways it involves being able to see, at a glance, whether something is right or wrong, is going to be easy or difficult and having the knowledge to do something about it. Sometimes – as the phrase 'professional vision suggests – it does involve just looking:

> "Can you see around the edges of that plate ... how its cold (darker) ... that's when it starts turning up and you've got problems."

> "... turn the water off on this one ... it's not very bright ... turn the water off to keep the heat there."

> "... a lot of slab defects you see as you roll lines from the furnace … gauge variation ... sometimes we send back because of worry about overload."

This vision is based around an idea of achieving the appropriate 'ideal' shape, the 'dog bone', whilst recognising that this perfect shape is predicated on *"a 2600 slab and a nice set of rollers"*. When this is not the case, adjustments need to be made and professional vision deployed, for example:

> "It's Wednesday … I'm thinking of the state of the rollers (changed every Thursday) ... they'll be hollow in the middle now." Sometimes, for example, this involves altering the transfer references: "this one will want to turn at 120 ... I'll do it at 118 ... that will offset the roller."

The other important aspect of 'professional vision' is the coordination between what the operator sees happening in the mill and on the rollers, and the foot and hand various actions performed in the pulpit. This is where professional vision links with 'skill', where to ensure adequate throughput the operator needs to get a rhythm going that requires anticipating what is happening with the slab on the rollers.

Professional Vision: Interacting with the computer

> "We're stopping for a while ... Houston's got a problem."

A good part of the observed 'professional vision' in the Roughing Mill involved some kind of comparison between what the computer 'said' and the operator's experience and skill. This was most obvious where the operator went into manual or over rode the scheduler in some way. So, for example, it was common for the most experienced operators to go into manual for the last few passes – because their experience was that *"... because the computer at less than 45 pisses about ... does 4-5 passes ... that's what causes turn-up."* and "*because it says 45 the computer tries to do it in 3-4 passes when you can do it in 2 ... it's to do with the pacing of the mill ... we're rolling plates quicker than the computer thinks we are.."* In some ways the end product of this was – at least in the case of the most experienced operators though less so in the case of newer operators – a healthy suspicion of what the computer was 'telling them' or asking them to do. This was heightened by cases of the computer providing wrong slab sizes or instructions:

> (Concerned about wrong slab sizes being given by computer) "I've just had a look at that (the slab waiting) and it looks about right ... if the size here (monitor) said 30 metres I'd know straight away that was wrong." or (reference number gone) "it's not updating on the screen at all ... for some reason its not updating ... so there's obviously a fault somewhere ... that's why I'm in manual .. I don't trust it now because I don't know what its doing."

It also generally involves a reliance on the (inaccurate) Mill 'clock' rather than the head display for an understanding of what the Mill is doing:

> (watching the clock) "the clock is out but only by about 3mm ... we use the clock because its easier to read … we can anticipate the speed of the screw ... (compared with head display) ... if it's going down in a pattern … and it suddenly puts 15 on you know something's wrong."

Discussion: dependable red hot action

> "... how important it is to accept the reality of human fallibility and frailty, both in the design and the use of computer systems ... all too often, the latest information technology research and development ideas and plans are described in a style which would not seem out of place in an advertisement for hair restorer." (Randell, 2000)

When defined as "The ability to deliver service that can justifiably be trusted" – dependability has a number of attributes – many of which apply in this particular case. These include: availability (readiness for correct service); reliability (continuity of correct service); safety (absence of catastrophic consequences); integrity (absence of improper system state alterations); maintainability (ability to undergo repairs) and more. But as we consider broader, socio-technical, notions of "system", the ability to achieve a clear and documented understanding of the intended service of the system – and hence some view of dependability – becomes increasingly difficult. Once we start taking into account the actual practice of a socio-technical system rather than any idealisation of it, it seems increasingly difficult to determine with sufficient precision what is meant by the "service" the system offers. Thus it also becomes difficult to determine what is meant by a "failure" of that service, and thus what is meant by "dependability" in this broader context. In these circumstances, we may need to broaden our understanding of what dependability means beyond the simple "absence of failure", particularly if we consider 'quality of service'. In this case this would, for example, include the quality of the eventual 'daughter' plates that could be cut from the plate; the amount of waste; the timeliness with which the plate is presented to the next stage in the rolling process and so on. Moreover, in this instance many of the 'failures' observed are 'low consequence' that may not directly cause dependability-critical problems. However, they can contribute to a reduction in quality of service, and also give rise to a more fragile operating environment in which dependability-critical problems may become more likely.

To improve system dependability, we can reduce the number of human errors made, include system facilities that recognise and correct erroneous states, and so on. But, when we start considering people using a computer-based system, the notion of failure becomes rather more complex. In a situation where computer-based systems are used by groups or teams of people, usually in conjunction with other systems then recognising failure becomes even more difficult because different users may have (in this case did have) different models of how the system is supposed to behave. Unexpected behaviour to one user – such as the novice mill operator – is normal behaviour to another. Some users – the experienced mill operators – may have learned how to work-round problems in the system, others may not have.

One of the most obvious dependability issues to have emerged from this research concerns various forms of awareness and its impact on dependability – in particular a lack of awareness in several, perhaps crucial instances. So, for example, while the computer system is configured to ensure the manager knows the composition of the slabs in the furnace and the order in which they may appear, none of this information is conveyed to the Roughing Mill. Here operators simply respond to whatever slab appears in front of them. Such an awareness – of what's coming out of the furnace – may prove useful both for pacing and teamwork in the Roughing Mill. This applies both in the sense of developing a rhythm of work but also to ensure either that the more experienced operators roll the more difficult slabs or are available for assisting the less experienced operators. At the same time, there appears to be little in the way of form of 'reverse awareness' – from the Shear Lines to the Roughing Mill, for instance, in terms of information about the quality of finished plates. This might, for example, enable a mill operator to decide that a plate should be scrapped before it goes through the Finishing Mill because the defects in it – such as lines – make it worthless. At the moment there appears to be a touching faith that any such defects will be remedied somewhere else in the production process. At present there appears to be no real, useful feedback to ensure that poor quality plates are taken off early in the production process ensuring that production time and resources are not wasted.

Another aspect of improving dependable production relates to the setting of the controls and the information provided (Andersen, 1999). To some this may appear to be an essentially ergonomic problem in terms of the best positioning of the available controls. The issue of modifying the pulpit controls raised a number of interesting, though different, opinions. Some operators would prefer the measurement gauges to be on or nearer the monitor (so that they did not need to turn their head), others appeared to have incorporated the head turning seamlessly into their work. Others felt that the mill load gauges should be more easily visible to the operator. The head display was dismissed by many as 'going too fast' to be of much use in the work of rolling but a useful indicator for when the computer goes down. As for the monitors, apart from the main display monitor, the other displays appeared to be rarely used, especially by the less skilled drivers. This links with some of the points made above and highlights the issue of generating displays that are appropriate to the right people at the right time and in the right place. This may also be related to dependability issues of 'diversity' – of providing a range of measures by which operators can obtain relevant information. In practice this issue of diversity rarely arose as an everyday concern but it became important when things began to go wrong, when the computer started giving the wrong measure or the wrong slab or the wrong dimensions.

In terms of problems with the plates in the form of cobbles, or faults, or quality the observations revealed an interesting tension and trade-off in terms of dependable production between human skill and computer scheduling. The problem of cobbles was seen by the operators as a product of particular steel features – such as high manganese, no washes, poor sizing etc that were exacerbated by scheduler problems. This meant that the more experienced operators routinely and regularly over rode the scheduler and went into manual to drive down faster and prevent turn-up. At the same time problems with the pacing system seemed to result in reduced productivity with slabs not being pushed through fast enough; alongside poor combinations of steels and sizes and bad slab planning such as too much rolling in one direction. Finally, a number of teamwork and more general 'human factors' issues appear relevant to issues of dependability. Clearly the observations revealed remarkably different levels of operator skill. Of particular interest was the different working 'tactics' adopted by the differently skilled workers with the less skilled (and often younger) operators rigidly following the schedule – which may cause problems in the Finishing Mill in terms of the length and gauge of the plates. Skilled (and generally older) operators were far more likely to go into manual to prevent or reverse turn up, and to over ride the turning and finishing measures in order to reduce the problems elsewhere. In some ways the most problematic were the 'intermediate' skilled who were prepared to override the computer but occasionally lacked the requisite skills or experience.

Various incidents observed highlight the issue of dependability and trust that arose in everyday work – whether it was trusting the technology, trusting the process or trusting others in the working division of labour. While there are a number of different theoretical approaches to the study of trust (Axelrod, 1997; Kipnis, 1996; Luhmann, 1979; 1990), Luhmann points to the problem of conceptualising trust, suggesting that most approaches fail to pay attention to the *social process of trust production*, i.e., they leave unspecified "the social mechanisms which generate trust" (1990:95). Our study accommodates Luhmann's recommendation to look at trust accomplishment as a social process, and rather than emphasising theoretical accounts of trust, our investigation is concerned with how trust is achieved, how it can be seen in action. Our interest is in how trust is woven into the fabric of everyday organisational life – the workaday world – as part of the 'taken for granted' moral order (Garfinkel, 1967) and the impact this might have on 'dependable' production. As Fogg and Tseng (1999) suggest, "trust indicates a positive belief about the perceived reliability of, dependability of, and confidence in a person, object or process." (Fogg and Tseng 1999) Of particular interest is the inter-relationship between trust and technology.

Although based on a relatively small number of observations in a short period of time, a number of conclusions concerning 'dependability' and the everyday working of the Roughing Mill are readily apparent. Considered in terms of improving dependability, while there are clearly things that need to be maintained – such as the levels of skill, and aspects of awareness in the Mill team; others may need to be done differently or supported differently – such as pacing and the scheduler or more general awareness in the plant as a whole. However, any changes need to be carefully considered in terms of their interactional effect - for example changes in the scheduler may make greater demands on the operators skill and may thereby impact on the quality of the finished product.

Observations suggest that problems were an everyday, commonplace feature of work – mundane, generally low consequence failures. When we consider problems or 'failure' as an everyday fact of life, we shift our ideas of problems, failure and dependability. We are, perhaps, drawn away from thinking about creating *failsafe* systems for such complex environments. Indeed it may be, as Law (2000) suggests (though in a rather different context) that imperfection is *necessary* to effective system functioning. Instead, the problem becomes one of *dynamically responding in the best way* to problems as they arise; and providing support for this response. This suggests that achieving or accomplishing dependable production involves far more than simply reconfiguring the scheduler or pacing system – even if this can be done – but needs to attend to the complexities of collaborative working and the classic CSCW problem – what to automate and what to leave to human skill and ingenuity.

Acknowledgements

This work was funded by the ESRC/EPSRC Dependability Interdisciplinary Research Initiative. We particularly thank the Roughing Mill operators for their tolerance of the 'professional stranger' amongst them and Roy Nicholls for facilitating access. Thanks also to Kjeld Schmidt for some insightful comments and references.

References

Andersen, P.B. (1999): 'Elastic Interfaces: Maritime instrumentation as an example', in *proceedings of CSAPC'99*, Valenciennes, France, pp. 35-41.

Axelrod, R. (1997): *Complexity of Co-operation: agent based models of competition and collaboration*. Princeton. NJ. Princeton University Press.

Blumer, H (1954): 'What Is Wrong with Social Theory?', *American Sociological Review*, vol. 19, pp. 3-10 (reprinted on line at: http://www.calvin.edu/academic/crijus/courses/blumer.htm

Blythin, S., Hughes, J., O'Brien, J., Rodden, T. and Rouncefield.M. (1997): 'Designing with Ethnography: A presentation Framework for Design', in *Proceedings of Designing Interactive Systems '97*, ACM Press, Amsterdam.

Clarke, K., Hartswood, M., Procter, R., Rouncefield, M. and Slack, R. (2002): 'Minus nine beds: some practical problems of integrating and interpreting information technology in a hospital trust' in *Proceedings of the BCS Conference on Healthcare Computing*, Harrogate, March 18-20.

Dant T. and Francis, D. (1998): 'Planning in organisations: Rational control or contingent activity?', *Sociological Research Online*, vol. 3, no. 2.

Fogg, B. J. and Tseng, H. (1999): 'The elements of computer credibility', *Proceedings of CHI 99*, New York, NY: ACM, pp. 80-87.

Garfinkel, H. (1967): *Studies in Ethnomethodology*, Englewood Cliffs, New Jersey, Prentice Hall.

Goodwin, C. (1994): 'Professional Vision', *American Anthropologist*, vol. 96, no. 3, pp. 606-633

Hughes, J., King, V., Rodden, T. and Andersen, H. (1994): 'Moving out of the control room: ethnography in system design', *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94)*, Chapel Hill, ACM Press, pp. 429-438.

Hughes, J., Randall, D. and Shapiro, D. (1992): 'Faltering from Ethnography to Design', *Proceedings of CSCW'92*, North Carolina, ACM Press.

Hughes, J,. King, V., Rodden. T. and Andersen, H. (1994): 'Moving out of the control room: ethnography in systems design', *Proc. CSCW '94*, North Carolina, ACM Press, pp. 429-438.

Kipnis, D. (1996): 'Trust and Technology', in R. M. Kramer and T. R. Tyler (eds.): *Trust in Organizations: Frontiers of Theory and Research*, London: Sage, pp. 39-50.

Laprie, J-C. (1995): 'Dependable Computing, Concepts, Limits, Challenges', Invited paper to *FTCS-25 25th IEEE International Symposium on Fault-Tolerant Computing*, Pasedena, USA.

Law, J (2000): 'Ladbroke Grove, or How To Think about Failing Systems', Centre for Science Studies and the Department of Sociology, Lancaster University at http://www.comp.lancs.ac.uk/sociology/soc055jl.html

Luhmann, N (1990): 'Familiarity, Confidence, Trust: Problems and Alternatives', in Gambetta, D, (ed.): *Trust: Making and Breaking Cooperative Relations*, Oxford. Basil Blackwell. Available online at www.sociology.ox.ac.uk/trustbook.html

Luhmann, N. (1979): *Trust and Power*, Chichester Wiley.

Popitz, H., Bahrdt, H., Jures, E. and Kesting, H. (1957): 'Technic und Industriearbeit', *Soziologische Untersuchungen in der Huttenindistrie*, J.C.B. Mohr, Tubingen.

Randell, B. (2000): 'Facing Up to Faults', Turing Lecture, January 31st.

Rogers, W.F. (1986): *Report of the Presidential Commission on the Space Shuttle Challenger Accident*. http://science.ksc.nasa.gov/shuttle/missions/51-l/docs/rogers-commission/table-of-contents.html

Rognin, L. and Bannon, L. (1997): 'Sharing Information: The Role of Teams in contributing to Systems Dependability Constructing Shared Workspaces through Interpersonal Communication', *Proceedings of Allocation of Functions (ALLFN'97)*, October 1-3, Galway, Ireland.

Schmidt, K. (1994): 'Modes and Mechanisms of Interaction in Cooperative Work: Outline of a Conceptual Framework'. Riso National Laboratory.

Schmidt, K. (1997): 'Of maps and scripts: the status of formal constructs in cooperative work', *Proceedings of GROUP'97*, ACM Press.

Suchman, L. (1995): 'Making Work Visible', *CACM*, vol. 38, no. 9, pp. 56-64.

Suchman, L. A. (1987): *Plans and Situated Actions: The Problem of Human-Machine Communication*., Cambridge, Cambridge University Press.

The Ladbroke Grove Train Inquiry http://www.hse.gov.uk/railway/paddrail/lgri1.pdf

Voß, A., Slack, R., Procter, R., Williams, R., Hartswood, M. and Rouncefield, M. (2002): 'Dependability as Ordinary Action', *Computer Safety, Reliability and Security: Proceedings of the 21st International Conference, SAFECOMP 2002*, Catania, Italy, September. Reprinted in S. Anderson, S, Bologna, M. Felici (eds.): *Lecture Notes in Computer Science* vol. 2434, Springer Verlag, pp. 32-43.