

Towards a Rational Analysis of Work Synchronisation

Position Statement for CHI 2004 workshop on "Work Design"

Andrew Howes (Cardiff University)

Alonso Vera (NASA Ames Research Center)

Richard L. Lewis (University of Michigan)

contact HowesA@Cardiff.ac.uk

A typical morning for one of the authors. Emails PhD student and asks her to phone him when she arrives in order to arrange a meeting that morning. Emails Head of Department's secretary and requests meeting time on Mon or Tue with HoD. Starts writing research proposal. Interrupted by knock on door (project student requesting help). Explains too busy, consults diary, could she come back at 3pm? Interrupted by phone ringing, colleague requesting help with computer problem. 3.30pm? Agreed. Returns to conversation with student. Returns to research proposal.

In a recent survey of causes of workplace stress amongst senior managers, interruptions came top (BBC News, February 2000, survey published conducted by PPP Healthcare and the Institute of Management). Another survey showed that most office workers were interrupted every 10 minutes by telephones, faxes or emails (Pitney Bowes report cited in BBC web site article; see also O'Connell and Frohlich, 1995). Interruption is perceived amongst knowledge workers as a major cause of inefficiency. In other fields, e.g. flight deck safety, interruption is recognised as a source of critical failure. Findings in the scientific literature indicate an effect of latency on the degree of interruption, and that time is required to re-establish the primary task after an interruption has terminated. Others have argued that the disruption caused by an interruption is mediated by similarity of content to the primary task. In the extreme, interruption can lead to behaviour in which so much time is spent managing switches of task that little effective work can be done on meaningful tasks.

In modern organisations, the transition to synchronised work can be achieved by face-to-face communication or using one or more of a number of systems, including phones, answer machines, voice-mail, secretaries, and email. In addition, there are more innovative technologies to manage the transition. Some of these have been directed at making a potential initiator of a synchronisation aware of their recipient's situation, either at the present time or in the future. For example, RAVE (Gaver et al, 1992) provided video links between workers and to shared spaces and allowed them to 'glance' into offices and work spaces in order to assess whether it was a suitable time to interrupt. Variations on this theme, in attempts to ameliorate privacy concerns, provide blurred video links, or audio-only links or allow some user control over the quality of information available. Other systems allow the user to set an explicit level of 'interruptability' which for example in Teamwave (Greenberg & Roseman, 1998) is indicated to potential initiators with a door-icon that may be in a state between shut and fully open. Also, on-line calendars can be used in order to help people schedule meetings. These have seen some success but require workers to keep them up to date so that others can determine when in the present or the future they are likely to be available (Grudin & Palen, 1995).

Other technologies have been directed at giving the recipient of a synchronisation more information about the initiator and the topic of the desired synchronisation. For example, cell phones typically give information about the caller identity (assuming it is in the phone's database) while the phone is ringing. Also, while the primary purpose of an answer phone is to take calls while a recipient is absent, they can be used to screen calls: The recipient is present but chooses not to answer until having considered some part of the message content.

Research on technologies that support work synchronisation has raised interesting questions. One is, why do people choose a particular technology for synchronisation, and what constrains how they choose to use it? Importantly, it is clear that users are flexible and creative in generating strategies for integrating new technologies into work practice (Charman and Howes, 2003). The adaptability of users reveals an acute sensitivity to the costs and benefits of different patterns of use. For example, users of RAVE were observed to place the video cameras on the floor so that they provided awareness information but at a reduced cost to privacy (Gaver et al., 1992). Users of Instant Messaging (IM) used the technology in order to establish whether it was a convenient time for a phone conversation (Nardi, Whittaker and Bradner, 2000), suggesting that

either the initiators of synchronisation were sensitive to the costs of phone interruptions for the recipient or that they sought to avoid the costs of phone calls that do not connect (Nardi et al. claim that about 60% of office phone calls fail in this manner). One of the authors has observed users at Price Waterhouse who did not leave voice-mail for each other when calls did not connect, rather they would rely on the recipient to see a display of their number and call back at a convenient time. Users discovered that the fact that the technology made the callers number available meant that recipients could avoid the cost of leaving messages.

In fact, the costs and benefits of a synchronisation strategy are not wholly dependent on the technology but also on the protocols that determine how the technology is used. McFarlane (2002) describes four classes of protocol for establishing synchronised work (NB. his interest was in the effect of interruption and not in synchronisation per se): Immediate; negotiated; mediated; and scheduled. To illustrate these he used an example in which a person is indirectly driving a car by supervising a robot while conversing with a passenger. Consider cases where the robot is the initiator and the person the recipient. Immediate coordination would involve the person responding as soon as the robot requested supervision. Negotiated coordination would involve the robot announcing that it needed supervision and then the person deciding when to respond. Mediated coordination would involve the robot communicating the need for coordination to a third-party (perhaps a PDA). The third-party then decides when and how to initiate coordination. Lastly, scheduled coordination involves the robot limiting its requests for supervision to specific time slots, e.g. every 15 minutes, or on the hour. For this particular task, McFarlane found that people using 'negotiated' synchronisation were able to achieve much lower disruption costs to the ongoing task than participants using the other protocols. Scheduled synchronisation produced the greatest disruption of all. While his task and protocols lacked ecological validity, McFarlane's results point to the potential importance of synchronisation protocols on worker performance.

The disruptive effects of interruption then derive from the fact that one activity (the synchronised communication) intervenes in the contiguous processing of another (the pre-existing activity) and from the fact that the synchronisation protocol and the technology used to support it, has in and of itself costs.

There is then a recognised applied problem and a research community dedicated to providing technologies to ameliorate the problem (McFarlane, 2002). However, there is little work aimed at providing a systematic theoretical account of the constraints that people consider when choosing protocols and technologies for initiating synchronous work. Without such an account the field is locked into a high-cost cycle of technological development and refinement that is driven by individual hunches rather than by informed and systematic engineering theory. To work towards such an account, the aim of this proposal is to investigate the technological, cognitive and organisational constraints that determine how people make choices about synchronisation in the work place. Work synchronisation supports complex sociotechnical tasks that constrain human behaviour in many ways. Our interest is in how people balance the competing sets of costs and benefits for themselves, for the people with whom they wish to synchronise (or who wish to synchronise with them), and for the organisation. Through this work we aim to inform (a) models of how people generate and choose between synchronisation protocols; and (b) the development of new technologies and protocols for work synchronisation.

Theoretical Framework

We adopt, and seek to extend, three existing theoretical contributions towards understanding human-human and human-computer interaction. The two contributions form strands of our framework: (1) A rational analysis of the constraints on human decision making in protocol and technology selection; (2) A view of the structure of discourse, and (3) a framework for formal modelling of constraints on behavior. In this abstract we focus on strand 1 and 3.

How do people choose protocols and technologies? A protocol for synchronisation is a strategy for achieving a goal. When new communications technologies are introduced into an organisation, or as workers gain experience with work and with their colleagues, they may adapt synchronisation strategies to fit. The importance of strategy change has been recognised by a number of researchers into human-computer interaction. The research focus has primarily been on the effect of perceived performance time on selection between known strategies and between

a known strategy and planning/learning. Given adequate knowledge one might expect people to minimise the time taken to achieve a task. However, Olson and Nilsen (1988), for example, have observed that computer users do not always choose the faster of two known methods, and reasoned that this is because of additional mental time costs imposed by switching between methods. O'Hara and Payne (1999) found that a small increase in the time taken for a computer to complete a copy operation resulted in more participants generating an alternative, lower cost, solution to a word processing task. The increased cost of some operators led to the participants spending more time thinking about how to achieve the task. More recently, Gray and Fu (2001) studied people learning to program a simulated VCR. For one group of participants, the perceptual-motor cost of consulting instructions was increased, rather minimally, by making the instructions invisible until mouse-clicked. Gray and Fu found that as the experiment progressed, participants in the high cost of consultation condition often preferred not to consult instructions at all, instead relying on their imperfect memory, and thus making more errors than the other help condition, who could (and did) consult the help window simply by looking at it. It is clear then that people are sensitive to small changes in the time costs of the methods that they use to achieve goals. As Gray and Boehm-Davis (2000) stated, "milliseconds matter", even very small changes to time cost can lead people to adopt radically different behaviours.

An assumption in the strategy change literature is that people approximate rational behaviour in order to improve utility and that they select between options according to the equation $E=PG-C$, where E = expected utility, P = probability of gain, G = gain, and C = cost (Anderson, 1990). Payne, Howes and Reader (2001) have argued that this basic decision making framework is sufficient to explain the choices that individual computer users make about which strategies to use, even in highly interactive tasks (see also Dix, Howes & Payne, in press). Time taken may be one of the costs that people take into account when attempting to make a choice but it is not clear whether findings generalise from experiments where participants were engaged in standalone tasks (e.g. programming a VCR) to the kinds of collaborative task that we are interested in here. While probable, it is an empirical question as to whether time costs are an important factor for work synchronisation. All other things being equal, milliseconds matter but in complex organizational settings, a simple comparison of the time cost of say sending an email versus making a phone call will have to be off-set against other factors. We can divide these factors between the costs and benefits to the initiator, the recipient, and the organisation.

Σ *Expected utility for the initiator.* Potential gains G for the initiator include the achievement of the desired goal, and the reduction in mental effort derived by the removal of the need to remember the goal. Gains are offset by probability of failure ($1-P$) (e.g. phone calls that fail to connect). The costs C to the initiator may include time costs (total time taken for the synchronisation) and workload costs (e.g. planning a voicemail message). Further, values of P and C will be more favourable for situations where a technology is matched with an appropriate protocol.

Σ *Expected utility for the recipient.* Potential gains G include the achievement of the recipient's goals that are dependent on synchronisation initiated by others, offset by the probability that the content will not be relevant (O'Connell & Frohlich, 1995). There is also the possibility of reciprocal gain deriving from local altruistic acceptance of synchronisations with otherwise low or zero gain for the recipient. The costs C to the recipient include the time cost of the synchronisation activity and the overhead of the interruption to the ongoing activity (e.g. time to re-establish mental context).

Σ *Expected utility for the organisation.* The organisation gains if workers achieve synchronisation by safe, productive, and healthy means (see Vicente, 1999). However there may be mismatches between what counts as a gainful synchronisation for the organisation and what is gainful for either the initiator or the recipient.

Much of the CSCW literature on awareness technologies assumes that people do consider the costs to others. While there is some evidence from empirical studies of social decision making that people apply a sense of fairness in making decisions about the allocation of rewards (e.g. the social utility model), there is little evidence as to how people account for distributed time costs under work conditions.

Even if people do wish to take into account the costs to others little is known about the extent to which they are able to assess the extent to which others are interruptible. A basic

assumption of many CSCW systems aimed at increasing awareness of initiators or recipients is that people are able to make these assessments on the basis of video or audio information. The alternative of interrupting someone to find out if they are interruptible has obvious drawbacks.

Modeling synchronisation with constraint logic programming

One of the reasons that it is difficult to understand how to design systems to support work synchronisation is that the constraints on behavior derive from a multiplicity of sources. One possible response to this problem is formal modelling (e.g. Duke and Duce, 1999). Here we argue that formal modeling offers a particularly useful technique for inferring the implications of a diverse set of interacting constraints.

While substantial advances have been made in cognitive modeling over the past ten years, there are critical problems that can be addressed by adopting an approach based on the mathematical specification of psychological theories of the relationships between technology design and strategy selection. We have developed a prototype tool that will be reported in a CHI'2004 paper (Vera, Howes, McCurdy, Lewis, 2004). The tool demonstrates how constraints on the temporal and resource properties of cognition, task, and device can be encoded and used to derive predictions of optimal skilled behavior.

The tool, called CORE, uses a Sicstus Prolog implementation of Constraint Logic Programming for Finite Domains (CLP FD) to calculate the time course and sequencing of skilled human performance. It calculates a prediction of the asymptotic bound on optimal performance time of human performance on a task given a declarative statement of task and psychological constraints. Example input/output from CORE is given on the project web site (<http://www.Cardiff.ac.uk/psych/howesa/ccm/>).

The task constraints can be described to CORE in terms of a hierarchical analysis of the interaction between a person and the task environment. The psychological constraints are described in terms of the temporal and resource properties of a distributed set of processors each with its own processing capabilities. In this respect, CORE is influenced by the Model Human Processor (Card, K., Moran, P., & Newell, 1983). Each processor is defined in terms of a set of parameters. For example, a processor might typically execute an operator with a cycle time of 50 ms. In a CORE model each processor is a resource for a defined set of processes (e.g. the gaze processor is a resource for shift-gaze-direction and fixate processes). Each process has parameterized limits (min, max) on its duration. The duration of motor movements can be automatically determined by a calculation of Fitts' Law. The duration of cognitive and perceptual processes may be directly determined from the empirical literature (e.g. estimates of the time required to switch attention), or by functions that, for example, model hypothesized decay rates.

We have used CORE to model a number of scenarios in human-computer interaction and cognitive science. These include a call centre operator who interacts with both a device and a customer, user's of a simulated automated teller machine, and various dual task scenarios. We are currently seeking to extend this range of scenarios to tasks, such as work synchronisation, that are constrained by factors other than those that are imposed by the immediate cognitive context. A crucial step in achieving this goal will be extending the flexibility of the payoff function against which CORE determines optimality. At present we have mostly focused on minimizing the total time required to achieve a task. While we have also included some aspects of memory usage in the cost function, this needs elaborating before we can start to model methods for work synchronisation.

Acknowledgments

The first part of this position statement was taken from unfunded research proposals that were written by the author in collaboration with Kenton O'Hara and with Simon Banbury.