Pace and interaction

Alan Dix

at time of publication: University of York

currently: School of Computing, Staffordshire University, UK.
A.J.Dix@soc.staffs.ac.uk
http://www.soc.staffs.ac.uk/∼cmtajd/

this paper appeared as: A.J. Dix, (1992). Pace and interaction,
in *Proceedings of HCI'92: People and Computers VII*,
A. Monk, D. Diaper and M. Harrison (eds),
Cambridge University Press, pp. 193-207.

January 30, 1998

Abstract

Channels of communication are presented as an emergent property of cooperative work. During actual interaction channels of communication are typically used in an intermittent fashion. Thus bandwidth is not an appropriate measure. Instead *pace*, the measure of the rate at which individual communications occur through a channel, is proposed as a primary property. We can relate this to the pace of interaction between participants, and to the pace of the common task. Any mismatch of pace will result either in the participants being forced to adopt coping strategies or in the worst case a complete breakdown in the cooperative work.

Keywords communication, CSCW, group work, formal analysis, channels, information theory

1. Background

The general theme of much of our work at York has been the formal analysis and conceptual modelling in aspects of human-computer interaction. Initially this was the traditional single human – single computer systems, however for the past few years we have also been looking at issues arising in the areas of computer conferencing and CSCW. This has involved the production of various conferencing and group editing systems (McCarthy *et al.* 1990, 1991a, b, c) which have then been used in experimental studies of various collaborative tasks. The models we used to analyse single party interaction concentrated on modelling the computer system from the user's viewpoint and expressing usability properties of the system. The fluidity of interaction in multi-party systems adds many new dimensions, some of the existing models still shed insight, but new concepts are required.

This paper concentrates on the properties of the communication channels between the participants and how these interact with the task. In particular we will look at the *pace* of the channel and of the interaction, a measure of the rate at which interaction proceeds. This complements other analyses such as examining the transcripts of experiments using conversational analysis techniques (McCarthy *et al.* 1991c), and using the concepts of shared data and messaging to analyse different modes of communication (Dix 1991, Ch.$\hat{1}$0).

The predictions made concerning breakdowns when there is a poor mismatch between the pace of communication and task has particular relevance to the analysis of cooperative tasks where there is a strong time-critical element such as process control.

2. Channels

Channels physical and virtual

The analysis in this paper concerns channels of interaction. However, in both real world communication and computer based communication it may not be at all obvious what the channels are. At a purely physical level we can look at the human input and output channels: eyes, ears, touch, speech, gesture; at the computer: screen, keyboard, stylus, mouse, sound synthesis; and between people and computers: telephones, networks, postal systems. However, humans are capable of selective attention and systems have layers of abstraction built upon them. So, the telephone system has the appearance of a totally connected network. It appears that we have a direct line to everyone we know, limited purely by the availability of the callee, but really this is implemented by a small number of lines between major cities and countries, each line itself carrying several calls either mixed at different carrier frequencies or interlaced with one another.

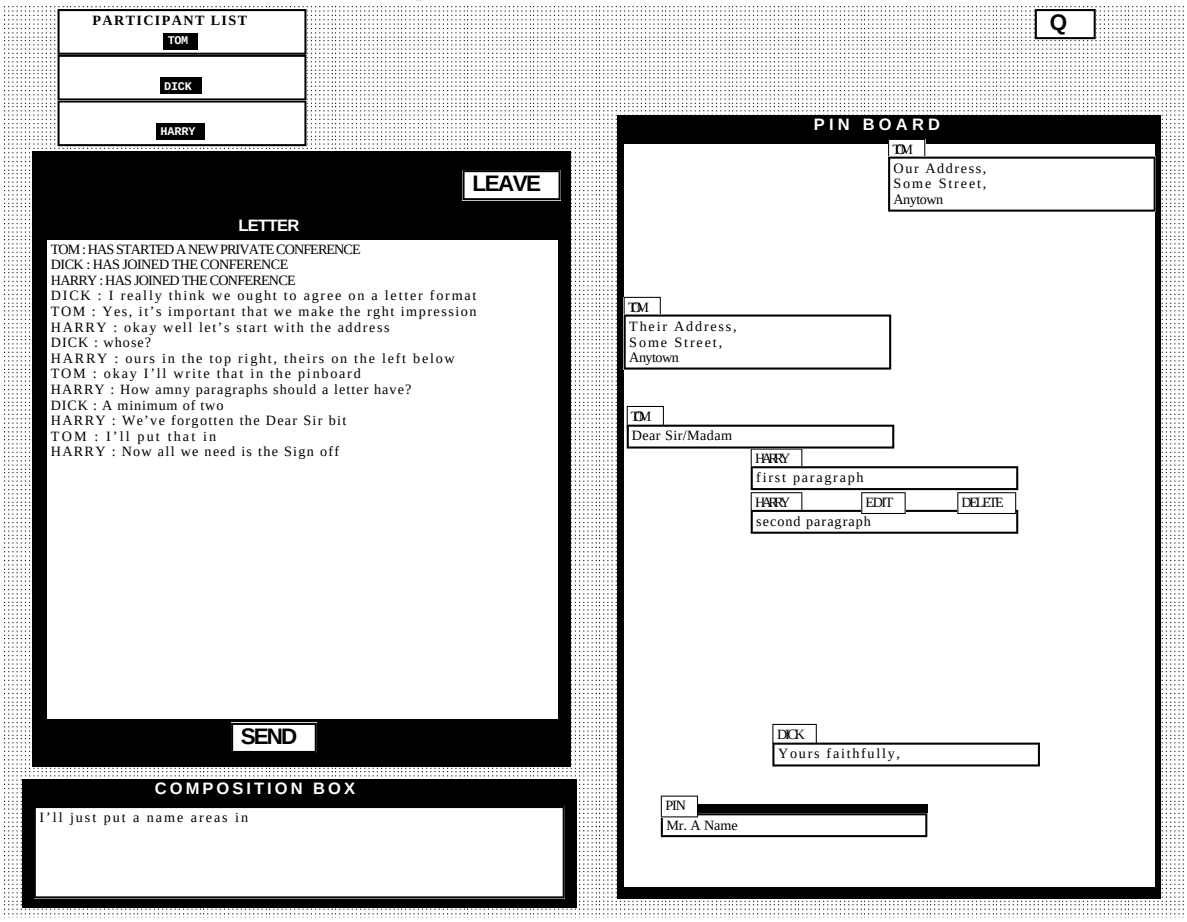

Figure 1: Conferencer screenshot: showing text transcript and pin-board

In the experimental conferencing systems we have been developing at York, the fundamental communication channel between workstations is an ethernet and the major channel between workstation and participant is screen, keyboard and mouse. However, that is of course not the system the participants perceive. In one system, the Conferencer (described

in detail in McCarthy *et al.* 1991), the participants had two major areas for communication: one was a linear conversation area where textual contributions are added at the end, the other was a simulated pin-board, where contributions are added like postcards to a pin-board (see fig. 1). In addition there were various status flags and private areas. We called these simulated channels of communication *virtual* channels (McCarthy *et al.* 1990). However, the story does not stop there. Just as in a letter one may carry out several simultaneous conversations, so in the system a single virtual channel as supplied by the system may be used for several conversations, and can thus be seen as being several virtual channels for the user. In addition, the different channels may be used in concert, to produce one virtual channel for the user. Again this happens in the real world, at the physical level binaural hearing or parallel cables, or at a higher level one may have a single thread of communication using whiteboard and speech.

Thus virtual channels are themselves an *emergent* property of interaction. This is somewhat a problem for *a priori* analysis as we may not be able to accurately predict how participants will use the channels we provide as designers. However, it has important implications for design in that we should consider how capable the communication environment is at supporting users' own virtual channels. Spatial arrangement, as in the pin-board, is especially flexible in this respect allowing users to set up areas which act as virtual channels. For example, Whittaker *et al.* (1991) describe how in "XScrawl", a shared electronic whiteboard, the participants defined private areas of the screen and set up shared areas for specific purposes. The linear transcript alone, on the other hand, is particularly poor in this respect. Combinations of linear and spatial channels, as in the Conferencer, allow the emergence of even richer virtual channels. Conversation in the linear transcript rarely included explicit references to the pinboard, but was only comprehensible using it as context.

So determining the channels of communication is not a trivial process. We cannot look at a system and say precisely what the channels will be. Only when the system is used will the actual channels become apparent. However, from now on we shall assume that we know what they are for a particular system and concentrate on their properties.

Bandwidth and pace

The classical measure for a communication channel is bandwidth, measuring the rate at which information can be passed through it. For human systems this must obviously be tempered by perception, for instance the tens of MHz of a video screen are reduced substantially when we consider it must be viewed by the human visual system. However, with the application of some basic psychology we can obtain adequate figures

A more fundamental problem with bandwidth as a measure is the assumption of continuous transmission. This is useful for a telephone company calculating peak transmission rates, but does not reflect real interaction. When a channel is used for interaction it is rarely used continuously, instead there are bursts of transmission interspersed with periods of inactivity. This can be taken into account by an average bandwidth figure, but this does not really reflect the feel and nature of the channel.

Take for instance conference papers. The rate of writing a paper may be of the order of say a thousand words per day. However, if a project presented say three papers of 3650 words each per year the average bandwidth would be only 30 words per day (is this a measure of the average information production of a research project!). If we compared this to say a weekly 200 word report, we get a similar bandwidth, however the channels have vastly different characteristics.

As another example, imagine you have a transatlantic cousin. Compare mailing an annual 60 minute video cassette (with greetings from the family, silly songs etc.), as against having a continuous (24hour a day!) phone call throughout the year..

Video:

one video per year
information per video = 1 hour * 25 frame/sec * 5Mbit/frame

bandwidth = information / (365 days * 24 hours)
= 14.3 Kbaud

Phone:

bandwidth approx 10 Kbaud

The video cassette and the telephone conversation have roughly similar bandwidths, but this bears little relation to their usefulness as media of interaction. Which is preferable, of course, depends on how you get on with the cousin.

So a typical channel is used in chunks. The rate at which chunks are produced we shall call the *pace* of the channel. Thus the pace of conference papers is yearly, the video was weekly and the continuous telephone is perhaps seconds (depending on the participants). We can give a very precise measure of pace for a channel by counting the number of chunks within a given period. However, for most purposes, although the quantitative nature of pace is very important, an order of scale is all that is required.

Pace of reception

Pace was defined above as the rate at which chunks were *produced*. As one of the reviewers pointed out, this rate may often differ from the rate at which chunks are *received*. For instance, after a few days absence, one may find several email messages or letters waiting, all from the same person. These are then read together – the pace of reception is slower than the rate of production. The example at the end of the next section demonstrates the same effect. On the other hand, one may read a long letter in several sessions and so the pace of reception could be argued to be faster in this case.

These mismatches between production and reception may occur for several reasons. The chunks may correspond to several topics – this *multiplexing* is discussed later and could be seen as several virtual channels. The mismatch may correspond to a breakdown. Once the Inland Revenue were sending letters to me at the wrong address. Eventually I received several months worth of tax demands and reminders at the same time! In all cases, it will usually be appropriate to take the *slower* of the two pace figures as defining the pace of the channel – however, one would imagine that, in most cases, the two would be of roughly similar magnitude.

Other properties

Although this paper concentrates on pace, it is, of course, not the only property of interest. Other interesting properties of individual channels include the size of chunks (granularity), the nature of the medium (visual, auditory etc.), and structural properties (linearity, persistence, etc.). Clark *et al.* (1989, 1991) discuss the ways such properties help or hinder the parties in their achieving mutual understanding.

Pace is not even the only temporal property of interest: the time taken to compose a message, the lag between transmission and reception and the time to perceive (read or hear) the message are all important. To some extent these temporal properties are subsumed in the pace of a channel as the tendency is for the participants to modify their use of a channel to take account of the other temporal properties. For instance, if I know that posts take several days to get to America, I am unlikely to issue letters more than once every week or so. However, although these factors may be subsumed by the *emergent* pace of the channels, knowledge of these factors may help one to predict this emergent pace.

In addition to temporal properties of individual channels, there are properties of multiple channels. For example, some channels can be attended to concurrently or even transmitted on concurrently. Others, such as two text based electronic channels, cannot be used simultaneously (try typing in two places at once!). Even these judgements are based on an understanding of pace. For example, if I am preparing two e-mail messages (where the pace of the channel is slow) I may regard myself as preparing them simultaneously even though I don't actually type in both at once.

Despite the importance of all these other factors, the paper restricts its focus to pace: the other properties could have, and in some cases have had, a paper written about each! However, before moving on we will look at an example, where the simple physical pace of the channels was *not* the dominant factor.

In a small pilot study groups played a game requiring not split second, but still real-time response. Different groups were given versions of the conferencer using either chunked or character by character text messages. That is, their messages were either sent explicitly with a 'SEND' button, or every few characters as they typed. We thought that the difference in pace of the textual channels might have significant effects on the behaviour. However, there appeared to be no difference at all. Why didn't the pace of the channels affect the interaction?

The fact that the participants had to give attention to the task and also to composing their own contributions meant that they could not continually attend to the incoming messages. So they only consulted the text 'conversation' area periodically and thus hardly ever saw the difference between the two types of channel. One way of looking at this experiment was that the emergent pace was the same in both channels, it is then not surprising that there is no behavioural difference. However, it does emphasise that the physical pace is *not* dominant and in this case it is concurrency constraints between various elements of the interface which determine the style of communication.

3. Interaction and task

Pace of interaction

Although channels may be one way transmitters of information, interaction requires a two way traffic. Consider two users with (at least) one channel in each direction between them. Their interaction is typically composed of turn-taking. This involves one party transmitting information to the other and then getting feedback. The rate of this turn-taking is, by extension, the pace of the interaction. It can be precisely measured in a similar fashion to the pace of a channel. In fact, the measured pace of interaction is likely to be a more robust measure than that of individual channels. Problems such as the differing rates of production and reception are naturally composed into a single rate as soon as one focuses on conversation rather than communication.

Even allowing the to and fro channels to operate concurrently the pace of interaction is clearly limited by the slower of the two channels. Another way of looking at this, more sympathetic to the view of pace as an emergent property, is that the pace of the individual channels is governed by the pace of the desired interaction. Thus if I am conversing by letter I would not usually send and receive a letter everyday, even though the channel would support this rate of interaction. Rather, I would be likely to write monthly (or less) to my friends. The measurable pace of the postal channel between me and my friend would then be monthly even though its theoretical maximum rate is so much faster.

In different circumstances the intrinsic properties of the channel may dominate the interaction or the interaction may dominate the channel. For an example of when the channel dominates the interaction, we could return to conference papers. In most conferences, it is realised that the pace of interaction determined by an annual presentation of papers is too slow (the channel is dominating the interaction) and thus after the presentation of papers a period is allocated for faster pace interaction. Here the inadequacy of the primary channel has been compensated for by the addition of a higher pace channel. We shall discuss later the various coping strategies which participants adopt when this is not possible.

In the Conferencer system both effects were observed. The ethernet has a bandwidth of tens of MHz and low levels of software can communicate at a (rather unpredictable) pace of several exchanges per second. However, at the interface the pace of the textual channels is limited by the time it takes to type a contribution, click on the send button and then for the displays to be updated. From this, assuming some lowish message length and typical

typing speed, we can calculate a theoretical maximum pace for the channel. Sometimes this pace was too slow, user's could not type fast enough and various breakdowns occurred. During other experiments the participants spent more time working independently, and the pace of interaction was far slower than this theoretical maximum. This is reflected in a slower measurable pace for the channel.

Pace and task – the control loop

To say that the same channel has too slow a pace in some circumstances, whereas the pace is governed by interaction in another is not all that useful. Why the difference? The crucial element is task. When we talk of cooperation there is implicitly or explicitly some job of work to be done. This work has some associated pace and we should expect that this pace and the pace of interaction are well matched.

Let's look at a simple situation to demonstrate this (fig. 2). We have two users. One (U1) is the manager who understands the work domain and the other (U2) who is actually in contact with the artifacts of work (W). We assume that U2 has sufficient control of the artifact, W.

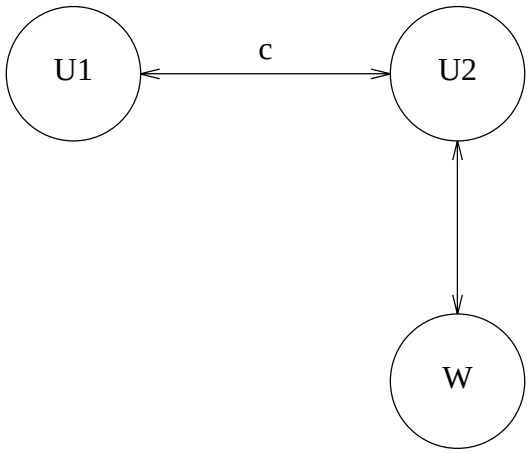

Figure 2: Novice/instructor interaction as a control loop

U1 and U2 communicate via a bidirectional channel (c). In the case that the pace of interaction of the channel c is at least as fast as the pace of the interaction, we have, in control theoretic terms, a closed loop control situation. U1 can give instructions to U2 who can observe the results, relay these as feedback to U1 who can then take appropriate responses. However, if the pace of the task is too great then there is insufficient time for this control cycle to occur. That is we revert to an open-loop control system. U1 must give commands without having feedback as to the effect of each command before giving the next. If the task artifacts are at all unpredictable then this is clearly unsatisfactory. It is a fairly obvious result, that open loop control is only possible if the controlled process is very predictable. If there is any element of judgement or reactiveness required, as in most creative tasks or in open systems, such as process control, then there will inevitably be some breakdown.

Consider a learner driver. The instructor can give instructions such as: "drive a little closer to the kerb Mr Dix". However, she cannot give detailed instructions for the individual movements of the wheel, there are many such movements each second and the pace of spoken commands is just too slow. In such cases the instructor must allow some control to rest with the learner, or (in extreme circumstances) grab the wheel herself.

Communication through the artifact

This novice/instructor example is of course not typical of cooperative working, but the general lessons about control apply to the more general situation. Here we would expect that the participants would each have control of different (or shared) aspects of the artifacts of work. They would receive feedback concerning their own actions and possibly about the effects of others actions (fig. 3).

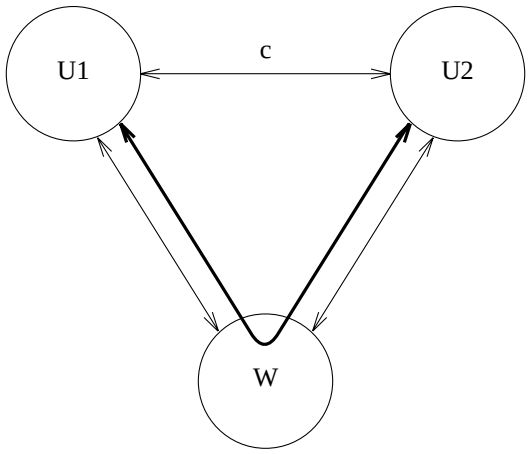

Figure 3: Communication through the artifact

This of course opens up the possibility for communication between the participants *through* the artifact. This form of communication is especially obvious when manhandling objects, such as piano moving. Another interesting example is to return to the motor-car, and drivers' behaviour at junctions. Some communication is explicit: indicators, flashing headlights, facial expressions and gestures; some is implicit in the road positioning and speed of the cars. Whether drivers' behaviour at junctions is a good example of cooperative work is of course a point of debate.

The importance of communication through the artifact to this paper is that it often forms a communication channel whose pace is intimately connected to the task at hand. In the case of piano moving the feedback about other's actions is almost as fast as that about one's own. The pace of communication is thus almost that of an individual's motor control. However, the pace need not always be sufficient if the artifact is less solid. We can compare moving a bookcase to moving a mattress. Apart from the difficulties of moving floppy things in general, the mattress is difficult because it does not supply the instant feedback about the other persons actions.

This mismatch in feedback is not confined to physical objects. It occurs whenever the pace of each individual's interaction with the artifact is greater then their feedback about each other's actions. A case in point is when you try to pass someone in the street and each repeatedly chooses the same direction to move. The sensing of each others movements is of a much slower pace than the framing of intentions and effecting of one's own movement. From a control theoretic standpoint, we have two adaptive controllers, both adapting at the same rate. This usually results in chaos (Jacobs 1974). Resolution of the dilemma is often achieved when both parties recognise the breakdown and eye contact is made. The resulting channel is of sufficient pace to enable a cooperative decision to be made.

Interactivity

In the last two sections we have seen how pace of interaction and pace of task are interrelated and in particular how cooperation can only succeed if the pace of interaction is sufficient for the task. We could go further and define interactivity in this way. We usually regard a computer system as being interactive when we get feedback within a few seconds or faster. Batch systems have a much slower feedback rate, perhaps in the realm of hours or days. However, this distinction disregards the task entirely. Instead, we could say a system is interactive *for a task* when the pace of the interaction is greater than the pace of the task.

This has important implications. If we want to cooperatively sketch a picture, an e-mail connection would be rather inadequate. However, if we consider the process of writing a book, e-mail contact becomes highly interactive (although it lacks other crucial features such as deixis). This definition of interactivity also gives one a handle on single-user applications examining the tasks involved and determining which sub-tasks are best handled using (traditional) interactive techniques such as menus and direct manipulation, and which are better suited to more batch oriented work.

Thus in both single-user and cooperative systems obtaining a correct level of interactivity can become a design goal. However, we may be faced with an existing system, or a novel design may be limited by the available technology. In such cases we may not be able to match the pace of the channels to that of the task.

4. Coping strategies

How do people cope if the channels they are provided with do not match the pace of the task they wish to perform? There are subtle problems if the channels support too great a pace, but here we shall concentrate on the more obvious problem when the pace is insufficient. People are very adaptable and only rarely does the mismatch result in a breakdown. More commonly the participants are able to use various coping strategies to ameliorate the effect of the inadequate channel. Basically, the strategies reduce the total number of interactions required to perform the task. They do this either by changing the cooperative nature of the task, or by increasing the size of each chunk of communication.

Reducing the pace of the task: delegation and roles

An obvious coping strategy is to simply perform the task slower. Almost always this changes the nature of the task somewhat. For instance, writing and debugging programs using an interactive debugger is not simply faster, but different in kind from the original submit-compile-printout cycle. In fact, the nature of the task has changed so much that it is not at all certain that the highly interactive version performs its purpose any faster than the original.

At a social level we may choose to reduce the necessary pace of interaction by delegating aspects of the tasks to different participants. This assignment of roles reduces the need for feedback in two ways. Firstly, it reduces the cooperative nature of the task, as the actions of the individuals no longer impinge so much upon each other. Secondly, it increases determinism. We saw that an insufficient pace of interaction leads to an open loop control situation, which can only succeed if the task is sufficiently predictable. The participants may not be able to make the raw task more predictable, but by taking on roles they can make their own actions more predictable to each other. Each has expectations of the others and can therefore work on the assumption that these will be met, thus reducing the need for feedback. These roles may be worked out beforehand or may emerge as part of the interaction, even within the period of an experimental situation (Miles *et al.* 1991, Gibbs 1989). Thinking back to the driving instructor, she must obviously trust her pupil's ability to drive in a relatively straight line before attempting to drive on the open road. A special case arises when the participants know one another very well. In this case they may

be able to predict one another's actions well enough without any explicit establishment of roles.

Laziness and eagerness

When the pace of interaction is high participants can minimise the amount of information they transfer by being *lazy*.[1] That is they communicate ambiguous or minimal amounts of information on the assumption that misunderstanding will become apparent during later interaction, or that they will seek clarification from one another where necessary. For instance, if this were a journal article I might consider quoting a rather obscure reference to Thropleshaw and Hannigan's seminal work on communication within cormorant colonies, on the assumption that it could be dropped if the reviewer thought it irrelevant. However, with a conference paper there is no such interaction and the reference must be omitted. Laziness also reduces the amount of planning required during interaction. A conversation could take many courses looking at different topics. By being lazy in one's communications one need not look very far ahead in this tree of possible conversations to make each turn.

If the pace of interaction is too slow the participants can respond by being *eager*. They can foresee the possible course of the interaction and frame communications which encompass many of the possibilities: e.g., "If you don't pay within 7 days we will take you to court". One of the most extreme cases of eagerness is the framing of laws. In most countries the interaction is via the judiciary and there is little possibility of interaction with the legislature once the laws have been passed.

If we compare spoken with written communication we find that letters are far more eager than speech. When writing a letter one takes more care that the points are stated clearly, and one may even consider alternative responses of the recipient and state your position on each. For instance, a letter may say "if you marry me I will be happy for ever, but if not life will lie like bitter herbs upon my tongue" (lovers have a tradition of being over eager). In extremis one frames a communication which describes one's reaction in all possible situations. In this case one is essentially programming.

One can find similar incidents of eagerness in electronic conversations, for example, the messages (fig. 4) from Severinson Eklundh's corpus (1986) quoted by Bowers and Churcher (1988) both exhibit eagerness, the contingent part of each message being introduced by the key phrase "In that case".

A:Subject: Report C 123660
The above mentioned report is out of stock. The remaining ones are C 12366 + C 123660. What to do? Reprint? In that case, do you have any changes to suggest?
... ...
A:Subject: SIGSIM meeting
Are you going to Linkoping tomorrow?
In that case when are you leaving?
Does SIGSIM pay for the trip or what?

Figure 4: Excerpts from Severinson Eklundh's corpus (1986)

Multiplexing

Another common coping strategy is to carry on several conversations in parallel. This does not speed up the individual conversations but does allow them to complete faster than if

[1]The use of the terms laziness and eagerness comes from the realm of pure functional programming. See Thimbleby (1990) for their application to interfaces where bandwidth is the limiting factor.

they proceeded one after another. At a macroscopic level parallel conversations using the same channel can be seen as a form of virtual channel, but at the smaller scale of topics within a conversation it is best to regard it as a separate phenomena. As an illustrative example, assume we have a task which requires conversation about three topics a,b and c, and that each topic would normally require three turns from each participant to reach agreement. We will call the participants Jill and Tom. So a lazy conversation about topic a would proceed as:

Jill "a1", **Tom** "a2", **Jill** "a3", **Tom** "a4", **Jill** "a5", **Tom** "a6"

If the conversations were eager, Jill would try to guess Tom's replies (and possibly give contingent responses):

Jill "a1 a2? a3 a4? a5", **Tom** "a6"

Finally, we can imagine them discussing a, b and c in parallel:

Jill "a1 b1 c1", **Tom** "a2 b2 c2",
Jill "a3 b3 c3", **Tom** "a4 b4 c4",
Jill "a5 b5 c5", **Tom** "a6 b6 c6"

With either coping strategy the length of the total conversation is reduced from 18 turns to 6. (If they were combined it would be only 2 – letter based correspondence is often this terse.) In reality, the transcripts are not so disciplined: some turns involve several topics some only one. Some current analysis of transcripts from experiments using our electronic conferencing system is focusing on the structure of topic shifts and how they correlate with levels of mutual agreement. In the next section is an example from similar work elsewhere.

Text based communication

To observe most of these strategies, one need go no further than the typical written letter or e-mail message. It is interesting that synchronous textual communication, as used by several investigators including ourselves at York, sits somewhere between speech and standard written correspondence. The transcripts have some passages which are highly interactive, but others where letter-like coping mechanisms are employed. Many partial breakdowns occur, not because the coping mechanisms are ineffectual, but because the participants are confused as to which conversational strategies are operative.

For an example, consider a transcript (fig. 5) from the TMPI project (Hewitt et al. 1990) taken during a computer mediated game of diplomacy.

	text channel	audio channel
E:	I don't like the other three being in cahouts - shall I form an alliance with one to stitch them up? I'll move to holland first	
F:	it may be best to move your navy in london to the english channel then you'll have more controlthe sea	E: John, have you seen the reply to my last question F: ((Laughs))
F:	ok russia will be a good move as they are probably the biggest threat!	

Figure 5: Excerpt of transcript from Hewitt *et al.* (1990)

The turn by England (E) introduces two topics, that is it is an example of multiplexing, but France (F), which is played by John, only replies to the second topic.. This causes temporary confusion on England's part who then resorts to the audio channel to sort it out. France subsequently replies to the first topic in the text channel (for reasons of secrecy the audio channel could not be used for all communication).

Similar breakdowns occur in many of the transcripts from our own experiments. Usually, these breakdowns are quickly sorted out, emphasising the adaptability of the participants, but they obviously add an extra load to the conversation and thus to the task.

One of the most common breakdowns is overlap, when two contributions are composed simultaneously, thus breaking the linearity of the transcript. This is a failure to adjust the pace of the conversation to match that of the channel. It is hardly surprising that this is a problem, possibly more remarkable is that it is not usually so for letter writing (although I could tell a story ...). Possibly the reason for the discrepency is that for letter writing there is an objective timescale which allows the participants to predict each others turn-taking. For synchronous electronic text communication, the timescales are beyond the 'immediate' response time of a few seconds (Shneiderman 1984) but within the user's subjective timescales, which always feel long.

Confounding social factors

If we observe a situation where the pace of interaction is insufficient for the task we should expect to see some coping strategy being adopted. Such strategies interact with social factors. Eagerness can be used to force a conversation along a preferred path by one of the participants as it is hard in a linear conversation to return to by-passed parts of the conversation. This may be used to force a decision by not discussing alternatives and is thus is a controlling ploy in conversation. It may also be used to hide weak arguments: "you might think A is a problem, but in that case I can do B" The fact that the point A has been dealt with may hide the many other objections which could not be, and yet which are difficult to reintroduce. On the other hand laziness can be a means of delaying commitment or of enhancing divergence. Many brainstorming tools deliberately seek to maximise the ambiguity of participants' contributions and minimise their length.

Channels which encourage or discourage eagerness will thus have profound effects on group dynamics. If this is seen as a problem one may use other aspects of the channels to ameliorate the situation. For instance, if one is involved in remote text based communication, the pace of interaction may force an eager conversational style. However, if the conversation transcript is persistent, and especially if it is non-linear (as in shared hyper-text), it becomes much easier to refer back to old topics and thus investigate more of the conversation tree.

Predictions

The above discussion makes some precise predictions about the interaction of channels and task. If the two are mismatched, we expect to either see coping strategies or breakdowns. In previous studies (Chapanis 1988) there has been relatively little effect on performance of changing the medium of interaction (e.g. face to face, speech only, typed messages). This is to be expected as in most cases the participants can use coping strategies to alter the necessary pace and thus to perform reasonably well even on degraded channels. One approach to investigating this is to perform detailed analyses of the transcripts from such experiments. Problems with interpretation will inevitably arise because of the confounding social factors above, and because of the range of variation between groups. Despite these problems, qualitative analysis of the structure of mediated conversation has been used with considerable success, for instance, the study of the TMPI project diplomacy transcripts (Jirotka *et al.* 1991, Hewitt *et al.* 1990) and also our own analyses at York. Knowing which coping strategies are adopted are important if we are to assess the burden this places on the participants, and to predict when the coping strategies will fail.

Although people are very adaptable, we do expect to find situations where coping strategies fail and thus breakdowns occur. This requires situations where the task demands close interaction, ruling out delegation, and where the nature of the task makes eager communication difficult. This may occur where a task has many equi-probable branches: the number of possibilities making eagerness impractical. Another difficult area is time-critical tasks where the pace of the task cannot be altered and where external environmental influences preclude eagerness. This is the very nature of cooperative process control tasks. Such tasks are often of a safety critical nature and is obviously important to know whether peoples coping strategies will be sufficient for emergency situations. The quantitative measures of pace allow one to make some predictions about when such breakdowns will occur. However, further work is needed to determine how people react to such breakdowns – possibly changing one coping strategy for another.

5. Summary

We have seen how the virtual channels by which people communicate are influenced by the physical channels they are supplied with, but are also an emergent feature of interaction. Pace is itself an emergent property of channels and yet is quantifiable and measurable. This is important as we can compare the possible and emergent pace of channels with the pace of tasks which people perform. Thus we can analyse existing situations to predict breakdowns or where coping strategies may be needed. Furthermore, an understanding of emergent channels and pace of interaction can guide our design of systems so that users can create virtual channels which match the tasks they perform.

Acknowledgements

This work was funded by SERC Advanced Fellowship B/89/ITA/220 and SERC grant GR/F/01895. The author also wishes to thank his colleagues at York and the anonymous referees of this paper, for their helpful comments.

References