



ELSEVIER

International Journal of Medical Informatics 00 (2002) 1–18

International Journal of
**Medical
Informatics**

www.elsevier.com/locate/ijmedinf

Interaction design theory

Enrico Coiera

Centre for Health Informatics, University of New South Wales (UNSW), Sydney, NSW 2055, Australia

Abstract

Objective: This paper presents a framework for the design of interactions between human and computational agents working in organisations, mediation by technological systems. *Design:* The design of interactions within an organisation is viewed from the point of view, not of the technology mediating the new interaction, but of the human and computational agents who interact with each other. *Results:* Understanding the limits to individual agent resources permits an analysis of the impact that a new interaction will have in a given setting. When we look beyond simple interaction settings, we can use the notion of interaction equilibria to predict the impact of new information and communication technologies within an organisation. Economic supply and demand curves, for example, may allow us to make both qualitative and quantitative predictions about technological adoption of communication systems. *Conclusion:* Rather than focusing solely on characteristics of individual technologies, or psychological and social issues, these can be combined to explain the overall decisions that individuals make when using technologies. Without necessarily understanding all the local decision criteria used by any individual, we can make robust predictions about how a group as a whole will interact.

© 2002 Published by Elsevier Science Ireland Ltd.

Keywords: Human–computer interaction; Communication space; Cognitive load; Economic agents; Emergent behaviour

1. Introduction

“In the next 50 years, the increasing importance of designing spaces for human communication and interaction will lead to expansion in those aspects of computing that are focused on people, rather than machinery. . . . The work will be rooted in disciplines that focus on

people and communication, such as psychology, communications, graphic design, and linguistics, as well as in the disciplines that support computing and communications technology. . . . Successful interaction design requires a shift from seeing the machinery to seeing the lives of the people using it”. Winograd (1997) [1].

Traditionally information systems are designed around an idealised model of the task that needs to be accomplished, and failure in system performance is explained away by

E-mail address: ewc@pobox.com (E. Coiera).

blaming human social and cultural ‘barriers’ to technology adoption. In this world, the newly engineered system is always impeded by the barrier of human frailty. The design of the computational machinery is the scientific high ground, and understanding the mess of implementation in the real world is left to ‘soft’ social science and happenstance.

But people are part of the system. The web of interactions needed to make anything work in a complex organisation always involves humans solving problems with limited resources and working around imperfect processes. Designing the technological tools that humans will use, independent of the way that the tools will impact the organisation, only optimises local task-specific solutions, and ignores global realities. The biggest information repository in most organisation sits within the heads of those who work there, and the largest communication network is the web of conversations that binds them. Together, people, tools, and conversations—that is the ‘system’. Consequently, the design of information and communication systems must also include the people who will use them. We must, therefore, design interactions that reflect the machinery of human thought and communication, sometimes mediated by communication channels, sometimes in partnership with computational agents.

Interaction design is a newly coined discipline, and it focuses on constructing the ways people interact with objects and systems, and the product of interaction design is almost entirely the quality of the user’s experience [2]. In this world, the effectiveness of a piece of software is not an internal attribute of the software, but emerges from the way the software is interpreted by users, and that interpretation is dependent upon the user’s specific context, culture, knowledge and resources.

It has been stated that at the level of an individual user, interaction design will always

be an ambiguous or subjective process for the designer [2]. In this paper I will argue that, from an organisational or large population perspective, the process of interaction design can be a much more principled one. Based upon models of the way populations of agents interact with each other, mediated by technology, but bounded by scarce cognitive and physical resources, we are able to model the effect of introducing new technologies, or indeed design technologies in the anticipation of their effect on populations of users.

This paper is the third in a series, which sketch out a general theory of mediated agent interaction. In the first paper, the communication space was introduced as the major component of the overall interaction space for humans working in complex organisations, and the gap between the current focus on interactions outside of the communication space and the overall interaction needs of humans was highlighted. A framework for mediated interaction theory was also developed, based upon the notion of bounded mutual knowledge, or common ground, that is shared between communicating agents [3]. This was developed to help us decide whether an interaction was best served by communication or information technology. In the second paper, an examination of the costs of the grounding process lead to the view that grounding is a kind of cost minimisation process, based upon maximisation of agent utility, and resulted in cost-minimising equilibria developing between agents [4]. This is important because it suggests that, although there may be a wide variation in the nature of individual interactions, we can over time predict the characteristics of interactions at a population level, and design systems to effectively support the majority of interaction at the organisational level. In this third paper I will turn to the task of designing systems to support interactions, whether in the commu-

nication space or not. Specifically, I will look at how, through a formal theory of interaction design, we can design and implement technological systems to support individual interactions by modelling the wider interaction space within which individuals operate.

2. Designing the interaction space

“In creating tools we are designing new conversations and connections”. Winograd and Flores (1986).

Typically information system design occurs with a single task assumption that the user is going to be wholly focussed on interacting with the system that is being designed. Such ‘in vitro’ laboratory assumptions do not translate ‘in vivo’ when individuals use systems in working environments. Since individual agents in an organisation are working in a complex environment, they may at any one time be carrying out a variety of tasks and interacting with other agents to help execute those tasks. With the exception of environments where there is rigorous workflow control, this means that we cannot predict what interactions will actually be occurring at the same time as any interaction we specifically design.

While we cannot predict every specific interaction that will occur in the real work environment, the design process can model the typical interaction space within which any new system will be introduced. The interaction space can be modelled to include the most important interactions that will be concurrent with the new designed interaction. Within the overall interaction space, the communication space forms a subset containing interactions

that are informal person-to-person exchanges. We can also conceive of an information space, which by corollary contains those formally structured interactions that typically rely on information models and computational systems. I have argued elsewhere that the notions of the communication space and information space are probably driven by a dichotomy in communication and information technologies, and that at a more abstract level, they actually are not separate spaces but form different ends of a continuum of explicit task modelling [3].

To construct the interaction space, we start with a general description of an interaction between two agents, which may either be human, or computational. An agent typically has a number of tasks that need to be carried out, and a pool of resources available to accomplish those tasks. An interaction occurs between two agents when one agent creates and then communicates a message to another, to accomplish a particular task. A mediated interaction occurs when a communication channel intermediates between agents by bearing the messages between them. For example, e-mail can be used to mediate an interaction between two individuals, just as can an electronic medical record, which is as much a service to communicate messages between clinical staff, as it is an archival information repository.

The first step in modelling the interaction space surrounding a new interaction we wish to design is to note which other agents will be local to the new interaction, and then to examine the likely effects any existing interactions might have on the new interaction. By doing so we enhance the chance that new interactions will succeed when they are eventually introduced into the interaction space. In general terms the impact of one interaction on another may be to:

- *Compete with another interaction as a direct substitute.* For example, a user could use an on-line database to seek information, or instead, use the telephone to call a colleague to ask for the same information. The human-human interaction mediated by the telephone competes with the database to meet the user's information needs.
- *Compete with another interaction for the resources of an agent.* An agent has limited resources and if they are expended on one interaction they may not be available to another. For example, an information system may be well received in 'in vitro' laboratory tests, but when placed in a work environment we may find the users have insufficient time to use the system, because of competing tasks. The new system is then 'rejected' not because it does not do what was intended, but because the impact of the real-world interaction space on its intended users was not modelled. Concurrent interactions can also subvert the execution of a designed interaction. For example, a user may be interrupted in the workplace and take up a new task, and not log off the information system they were using, causing a potential breach of security.
- *Create new information transfer pathways, through a combination of interactions.* Each interaction connects agents, and each new interaction enables novel conversations between agents. If these combinations are not factored into system design, then the introduction of a system may produce unexpected results. For example, consider the interaction between a human agent and an electronic record system or EMR. Computational agents that might co-exist with the EMR could include other applications like e-mail or a word-processor. If the design process fails to include these additional computational agents, then unintended interactions made possible via these other agents may subvert the original EMR design. For example, it may be possible for a user to copy a section of text from the medical record to a word-processor, where it can be edited, and then re-inserted into the EMR. However, since this interaction with the word-processor is not part of the original design scope, it may introduce problems. A user could inadvertently paste the text into the record of a different patient, and no formal mechanism would be in place to prevent this context-switch error. Similarly, text might be copied from an EMR that has been designed with powerful security features to prevent unauthorised access, and then copied into an e-mail message, which is insecure. In both cases, the co-existence of an unmodelled additional computational agents introduce interactions beyond the scope of the original system design, and permit behaviours which would be prohibited within the designed interaction, but which are permitted in the interaction space.
- *Support the new interaction by providing resources that are critical to its execution.* The designer of a medical record system usually focuses on sculpting the interaction between a single clinical user and the record. However, other human agents also populate the EMR interaction space. The EMR user is often not the sole author of the content that is captured in the record, but is recording the result of a set of discussions with other clinical colleagues. If the goal of designing an EMR is to ensure the highest quality data is entered into the information system, then it may be even more important to support the collaborative discussion between clinicians, than it is to engineer the act of record transcription into the system. Failing to model the wider EMR interaction space means we

may over-engineer some interactions with diminishing returns, when we could be supporting other interactions that may deliver substantial additional benefit to our original design goals.

3. Interaction dynamics and the competition for agent resources

Since it is not possible to predict every specific interaction within the space in which a designed system will operate, it may appear that there is little that can be done to account for all possible concurrent interactions. However, we can be quite specific about the typical impact that concurrent interactions will have on agents. As discussed above, one of the ways interactions impact each other is to compete for the resources of individual agents. Based upon cognitive psychological models, we should be able to say something about the cognitive resources available to human agents, the cognitive loads they will typically be under in a given interaction space, and the types of errors that may arise because of these loads. Consequently it should be possible to craft information or communication systems that are tolerant of the typical interaction load of users will be under.

We begin by noting that accomplishing a task normally consumes resources, and this also holds during an interaction. An agent has internal resources, which in the case of a human are cognitive resources like memory capacity and knowledge. These resources are limited, and in some cases will deplete with use. One consequence of operating with finite resources is that agents will need to minimise costs on resources and maximise the benefits of any action they contemplate.

An interaction between agents is thus dependent upon:

- The task at hand, which has a resource cost and an implied benefit upon completion, and which must be evaluated against other tasks the agent might have to determine whether it is worth doing.
- The internal resources available to the agent to attempt tasks, which include finite computational or cognitive resources.
- The external resources available to the agent within their interaction space, which includes communication channels and other agents. Such resources each have a cost, a benefit, and limitations on their capacity.
- The impact of concurrent interactions.

The cognitive resources of a human agent include prior knowledge, which is stored in long-term memory (LTM), and attentional resources, which include working memory (WM). The cognitive capacity of individuals to successfully carry out a task are limited by the resources of attention, which can store and process only a limited amount of data at any one time [6]. As attentional resources are loaded with new tasks, the capacity to enact further tasks is diminished, and current task execution may be compromised [7].

As a consequence, the design of any single interaction needs to take account of an agent's inherent capacity to carry out the task. However, in complex organisations, where we wish to introduce information and communication systems to support work, the design task is more complex. Specifically, each new tool introduces a new dialogue, a new interaction [5].

We thus create a setting in which there are multiple possible interactions that could potentially occur at any one time, driven by the task at hand for individual agents. As a consequence, the agents involved in carrying out tasks both need to initiate interactions as well as receive interruptions requesting inter-

actions. The actual sequence of which interactions occur will be dictated by the circumstances local to each agent at that time, and multiple interactions may be maintained over time (Figs. 1 and 2).

This multitasking setting has a number of consequences. Since a human agent's cognitive resources are limited, the amount of attentional resources available to any single interaction will be reduced on average with each new interaction we introduce. It should thus come as no surprise that introducing a computer onto a Doctor's desk will result in less attention being devoted to the patient. In one study, the presence of a computer during doctor–patient consultations had detectable effects on the focus of the doctor's attention [8]. While they were at the computer, doctors confined themselves to short responses to

patient questions, delayed responding, or glanced at the screen in preference to the patient. The response delays probably were an effect of the costs involved in switching between interactions. We know that switching between tasks requires a suspension of the processing of the current task, and alterations to memory. Consequently, interruptions have the potential to disrupt memory processes, and lead to tasks being mis-executed, forgotten or repeated [7].

Interruptions, which are requests for a human agent to stop their current task and commence a new interaction, are common in clinical work places. In studies of hospital-based clinicians, interruption rates can on average be as high as 30% of the total time devoted to communication [9]. Consequently, the introduction of any new interaction into such organisations needs to recognise the existing interaction patterns of agents, and design the new interaction in such a way that it neither has a substantial deleterious effect on the existing interactions, as well as ensuring that the new interaction will have sufficient attention from agents to ensure that it will be well executed.

When one task dominates attention, then other concurrent interactions may fail to get the attention they need. For example, studies have shown that when a doctor is interacting with a desk-based computer, they may not hear what is being said to them by the patient sitting opposite the desk [10]. Sometimes, a single interaction so consumes a human's attention, that all other interactions are unable to get their attention. In a study of the task-execution behaviour of anaesthetists, the study subjects were found to completely miss significant events displayed on the screen of their monitoring devices. It was hypothesised that this was caused by the anaesthetist's attention being swamped by other tasks [11].

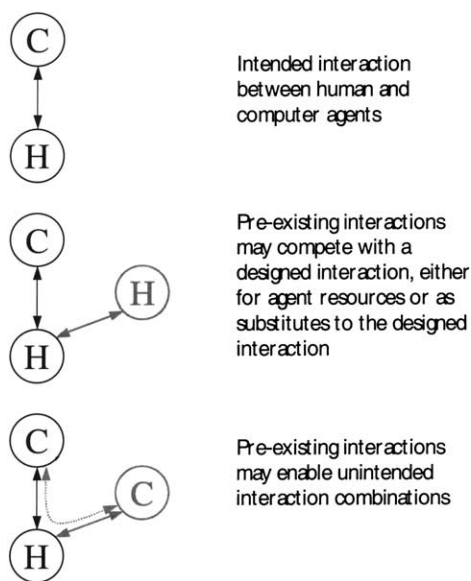


Fig. 1. When a new interaction is designed it does not exist in isolation but is placed within a pre-existing interaction space. Other interactions that exist in the interaction space impact the new interaction in a variety of ways, and if they are ignored in the design process may have unexpected consequences when the new interaction is implemented.

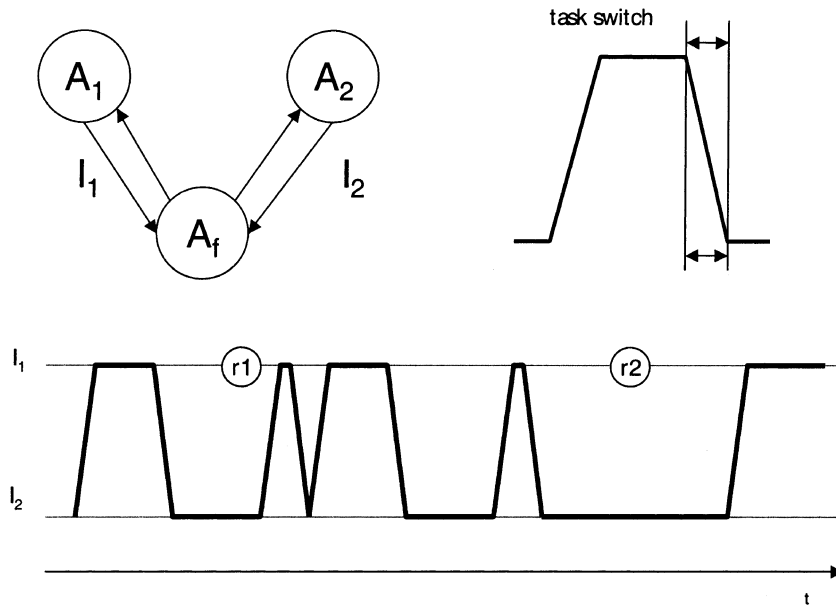


Fig. 2. Interaction history for agent A_f with two other agents A_1 and A_2 . The two interactions I_1 and I_2 compete for A_f 's attention, and A_f switches between them over time. Each switch is not instantaneous but requires a refocusing of attention and carries its own cognitive switching costs. Agent A_1 makes two interrupting requests to A_f during this period. The first request r_1 succeeds in getting A_f 's attention and causes a switch to interaction I_1 . The second request is not successful, because A_f 's attention is fully loaded by the current interaction I_2 and the request is missed.

3.1. Requisite attentional resources

This leads directly on to the issue of how one crafts an individual interaction to ensure it minimise the load it imposes on the attentional resources of agents, as well as maximising the likelihood that it will be well executed. This reduces down to the following questions:

- Which agent resources will be required to accomplish the new interaction?
- Given the existing work environment, what is the base-line cognitive load going to be on the agents who will be expected to participate in a new interaction? As a result, what agent resources are actually going to be available for the new interaction, given other likely concurrent interactions?

- When resources are scarce then either the new interaction or existing interactions will suffer. Consequently, given multiple possible interactions available to an agent, and the agent's local priorities, where will be the agent's focus of attention be?
- What will be the overall impact of introducing the new interaction into the existing interaction setting?

I have argued elsewhere that we can model many of the resource requirements of a given interaction by considering the size of message that needs to be sent between agents, the channel that will be used to convey the messages, and the relationship between agents. I have also argued that we can best model relationship between agents as the

degree to which they share models of the world, or have common ground [3,4,12]. This is because the greater the common ground between agents, the more succinct and accurate the communication is between them.

How does common ground help us model attentional resources? We know from cognitive psychology that although working memory has a small and finite storage capacity to manage immediate tasks, long-term memory can substantially enhance its performance. An item in working memory may actually point to a complex schema in long-term memory, and with growing expertise individuals can substantially enhance task performance by building such schemas [13]. We can intuitively think of this as working memory using LTM schemas as ‘virtual memory’ to supplement its own very restricted resources.

We now can see how common ground enhances task execution. For a given task, an agent will need to send a message to another agent. The more the agents already share, stored as schemas in LTM, the less a specific message needs to contain, and the more that can be assumed. By simplifying the task of constructing and sending a message, common ground minimises the use of attentional resources.

We can formalise this discussion by noting that for interacting agents a and b , a message of length m transmitted over a time t with a grounding efficiency of G_E , the requisite channel capacity C is [4]

$$C = \frac{m}{t} G_E \quad (1)$$

Grounding efficiency is a measure of the resource requirements for an interaction between two agents, all else being equal, and can be thought of as the average length of messages sent between them divided by the ‘true’ message length. When a pair of agents shares much common ground they can ex-

change terse messages since much can be assumed, and as a consequence their interactions are shorter than the ‘true’ message, which needs to include everything needed for it to be understood.

Consequently, for interactions which by definition require that agents exchange messages to accomplish a task, Equation 1 should provide us a basis for modelling the requisite resources to accomplish an interaction. It provides us with observable measures of internally unmeasured resources.

We next note that the likelihood that an agent will be able to handle an interaction if given it, or even recognise the request for an interaction, will be determined by the resources available to the agent. If the agent’s cognitive load is high on a current task then new tasks may be missed, poorly understood, or poorly executed. The threshold for a successful interaction is thus when the available agent capacity exceeds the requisite channel capacity i.e.:

$$C_{\text{avail}} \geq C_{\text{req}}$$

With Eq. (1), we can connect together the triad of channel characteristics, context-specific message requirements and the nature of the relationship between communicating agents. It allows us to make three different types of interaction design inference:

- *Channel-centric inference*: For a given pair of agents and a message, we can specify channel capacity requirements.
- *Relationship-centric inference*: For a given agent, channel and time resource, we can say something about the common grounding that would help select an appropriate second agent to communicate with.
- *Task-centric inference*: Time can be considered a key task-limiting resource. For a pair of agents and a given channel, we can say something about the complexity of mes-

sages that can be exchanged between them over a given period of time.

The last of these suggests that we can generalise the previous equation further to incorporate other resources. Resource-specific versions of the equation can be crafted, substituting bandwidth for other bounded resources like time, monetary cost, utility of knowledge received etc.

4. Global interaction design

Fundamentally, every agent makes some assessment of the costs $C(x)$ and benefits $B(x)$ of a potential interaction x , and when the benefits exceed costs, elects to interact [14] i.e.: if $B(x) > C(x)$ then do x else do not

If a number of interactions are available to the agent, it will probably choose the one that it perceives to deliver the maximum cost-benefit trade-off. What emerges then is a picture of a series of cost-benefit trade-offs being made by individual agents across different possible interactions. The goal for the agent is to accomplish its tasks through frugal utilisation of its resources, maximising the utility of any given interaction, through a maximisation of benefit, and minimisation of cost.

As an interaction designer looking at an organisation, it will not be possible to model all possible interactions nor determine what the outcome of individual agent interaction choices will be. What we would like, however, is to determine the overall state of interactions within an organisation, and make design choices that reflect the current interaction load. This will allow us to design new interactions that suit the current context, as well as predict how likely agents in the organisation will choose the new interactions.

Consequently we would like to be able to make statements about the likelihood of an interaction occurring, given the organisational context. Specifically, the probability that an interaction will occur is driven by the difference between cost and benefit. For most agents, this decision is based upon evidence of past interaction choices. We can say that, for an agent which has a history of n previous interactions of type x , that the probability of it carrying out x again next time should be related to the average overall past benefit B_{av} in the past i.e.

$$B_{av} = \frac{\sum_{i=1}^{i=n} (B(x_i) - C(x_i))}{n}$$

However, an agents available capacity C_{avail} to execute a new interaction also determines whether a new interaction will be taken up or succeed. Consequently we can say that the likelihood that an interaction will occur is a function of the agent's estimate of its benefit as well as the agents current capacity to handle it i.e.

$$P(\text{interaction}) = f(C_{avail}, B_{av})$$

We now set up a situation where we are measuring populations of interactions across groups of individual agents in an organisation, to determine likely future choices. This would seem to be an impossible task, given the complexity of the systems we are trying to model. However, it is likely that agents will behave in fairly stable ways over time, because each interaction is not an independent event, but depends upon past and future interactions. For example, an agent may have invested in developing a relationship with another in the past, or expended resource learning to use a piece of software, which benefits a present interaction.

Similarly, the anticipation of a future interaction may alter an agent's behaviour in the present. For example, if the agent is going to have a series of similar interactions, then it may be cheaper in the long run to establish a high level of shared common ground, and minimise grounding during individual interactions. One of the consequences of such grounding behaviour is that, over time, agents should choose a degree of shared common ground with other agents that will minimise future interaction costs, but will still be sufficient to accomplish the goal associated with interactions.

The two agents in an interaction either give or receive ground, and the decision to participate in the interaction occurs when both agree that the 'price' or benefit of doing so is mutually acceptable, much like the way two agents would agree to exchange other goods of value. This leads to a hypothesis called the law of the mediated centre [4] which states that communicating agents will be driven to an intermediate level of model sharing with other agents over a series of interactions, where the total costs of maintaining shared models and communicating models at interaction time are minimised. The mediated centre represents an interaction equilibrium point.

5. The economic interpretation of interaction

The grounding equilibrium is one example of agent's adapting their interaction behaviour over time, and they should also make similar adaptations to other aspects of interactions, such as choice of interaction task, other agents and channels. If agents do display such equilibrium behaviours over time, then we have a powerful set of analytical methods to assist in the modelling of organisational interaction to assist design. Specifically, we

can look to other disciplines that have had the same task of inferring the likely outcome of the multiple, often conflicting individual choices of agents, yet which collectively coalesce over time into definable population behavioural equilibria.

Specifically, we can look to disciplines like economics that offers an array of analytical techniques to predict decision equilibria. A number of different economic methods seem applicable to the task of predicting interaction outcomes including supply–demand analyses, barter modelling, game theory and computational simulation methods. To illustrate the nature of this type of analysis, the remainder of this paper will present a supply–demand analysis of interaction, and illustrate its potential to support interaction design with in organisations.

Since we are dealing with the choices to initiate or receive an interaction by free-willed and independent agents, under conditions of scarcity of resources such as time, we can think of interacting agents as behaving according to economic principle of supply and demand. Sending agents are 'buyers' with the need to interact, and receiving agents 'supply' them with their time and information by making themselves available for the interaction. There is thus at least an implicit, and sometimes explicit, negotiation by agents over their willingness to interact with each other, based on their personal assessments of cost and benefit from past interactions and current goals.

We can thus structure the supply and demand for interactions by graphing the quantity of product 'sold'—in this case the number of interactions made—against the price of interacting for the sender and receiver (Fig. 3). A key idea of economic analysis is that these curves represent the emergent behaviours of populations of individual agents. Without necessarily understanding all the




Fig. 3. Supply and demand curves for interactions. Demand is generated by a population of potential sending agents, and supply provided by a population of agents who are potential receivers of those interactions.

local decision criteria adopted by any individual agent, we can still make robust predictions about how a group as a whole will behave.

The curves indicate that when interactions are cheap, then senders are willing to make many of them. Conversely, when the cost of an interaction is high, senders only make a few. For receivers, the cost increases with the number of interactions, as interrupted tasks and newly acquired tasks compound the demands on their resources like time. Sender and receiver curves may not always have this simple shape, but so long as the curves are non-identical and intersect, the following analysis should generally hold.

The point at which supply and demand curves intersect represents the point of market equilibrium, where most buyers and sellers have maximised their outcomes. Any transactions occurring at prices away from the equilibrium result in unhappy buyers or sellers, who then drive the market back to equilibrium as they try and minimise their dissatisfaction with the outcome [14].

This is clearly a simplification of what occurs in reality, which we will retain for the present discussion. However, just as some

simple economic analyses assume agents have perfect information about the market place to decide what to buy and sell, we are making a similar assumption that our interacting agents have perfect information about which interactions to choose to make or receive. Yet we have already seen that under conditions of high cognitive load, some agents may be 'blinded' to the available interactions around them. Economic methods for handling situations in which information is 'asymmetric' between agents should also allow us to incorporate such details in interaction equilibrium models.

5.1. Getting specific about interaction costs

We can assume that agents are willing to collapse their cost and benefit assessments into a 'price' at which they are willing to interact. In traditional economics, price is usually expressed in monetary units, but in this case the price of an interaction is a complex amalgam of personal costs and benefits. Amongst these, time is an easily quantifiable and identifiable component of the overall costs and benefits. As we saw earlier, message length with respect to a specific task is a proxy measure for cognitive effort, mediated by the available common ground.

We can now be more precise about the notion of interaction costs, based upon the general model of agent interaction presented in the previous section. Specifically, the channel across which an interaction occurs, and the state of grounding between agents are the two major external costs that shape the interaction curves for a given task.

Thus, it is cheaper to converse with an agent that understands the issues being discussed, compared with one that needs to be filled in with the necessary background examples. However, developing that back-

ground knowledge is not a free process, and in many situations an agent may feel that the cost of doing the work in building prior common ground is not worthwhile. We can summarise by noting that grounding costs relate to:

- The grounding efficiency of the interaction, which is reflected in the cost of constructing a message given the relative ground, and the extra ground that needs to be transmitted within an interaction to facilitate message interpretation.
- The cost of building and maintaining common ground prior to the interaction.
- These two costs are traded off over time to arrive at the equilibrium point of the mediated centre.

Similarly, there are costs associated with the channel chosen to mediate an interaction. Some channels may allow rapid and high-quality transmission of a message, but be very expensive. Others may be slow but cheap, and in many circumstances sufficient for the task at hand. Channel costs are numerous and include [4]:

- the actual price charged for using the channel;
- the effect of noise on the transmitted message;
- the effort involved in locating channel access;
- the number of connection attempts required before another agent replies;
- the opportunity costs involved in message delay due to the channels transmission characteristics like bandwidth and latency;
- and similarly, the time it takes to transmit a message.

The interaction equilibrium point in Fig. 3 should represent the point in an organisation at which most senders and receivers are

making mutually acceptable use of the communication system. Thus, for a given organisation and its communications infrastructure and practices, the interaction equilibrium point represents the optimum communications outcome. However, while an organisation over time should reach equilibrium, and that equilibrium represents the optimal use of the existing communication system, the wider impact of that equilibrium may not be ideal. For example, the level of calls in a hospital may not be ideal if the traffic at equilibrium is adversely affecting patient care.

6. Predicting the impact of a new interaction class

Changing the communication infrastructure or practices in an organisation will alter caller and receiver costs and benefits when making calls, and thus alter the shape of the supply and demand curves, resulting in a new equilibrium. Thus the communication goals for any organisation are twofold. Firstly, they should free up the ‘communications market’, so that individuals can make optimal personal use of the existing system. Secondly, given other organisational goals and constraints, the ideal level of interaction should be estimated. Then, the underlying structure of the communication infrastructure and the way it is used should be altered to shift the equilibrium towards the ideal.

The notion of freeing up the communication market deserves some further comment. In one clinical user study, people were characterised as acting selfishly, reasoning only locally about their personal costs and benefits when interacting with others [15]. This was portrayed in the study as being in some sense a negative characteristic of the study subjects’ behaviour. However, in the present analytic framework, such locally motivated selfishness

does not preclude a system moving to the optimal allocation of resources. In fact, it ensures it.

Thus, if we believe in this ‘free market’ model, an organisation shouldn’t use the tactic of fixing the costs of calls by edict e.g. limiting the time or number of calls. Individuals will simply behave in ways that circumvent these rules [14]. What we can say is that if we want resource allocation that is close to ideal, we should not focus on the behaviours of individuals, who will always seek to optimise their individual circumstances, but we should seek to improve the underlying communication market’s structure. For an organisation, that means changing the available tools, channels or allocated tasks structures of agents.

An individual’s assessment of the costs and benefits of initiating and receiving interactions over the long run determine how they use a communication system. The only avenues for changing the resulting equilibrium level are to either change the decision-making process of individuals or the costs and benefits that drive those decisions. New technologies introduced into an organisation could plausibly do either of these two things. By making it easier to make a call, for example, a new technology will reduce costs and increase the number of instances in which people judge it worthwhile to call. Alternatively, by making hidden global costs and benefits explicit, a visualisation of the call traffic levels within an organisation may influence individuals to alter their local decision criteria.

Even though the supply and demand model of call behaviour presented so far is very simple, it is enough to permit us to model the effects of technologies, and predict the communication patterns that will emerge in an organisation once a particular new technology is introduced.

6.1. Example 1—introducing mobile telephones into the workplace

What would the effect be of giving hospital workers mobile telephones? Using the current framework, introducing mobile phones can only either alter the supply of calls being answered or the demand to make calls. In fact, simple analysis suggests that it is the supply that is mainly altered. When someone carrying a mobile phone is called, there is an increased likelihood that they will answer, compared with someone whose phone is in an office or responds to a pager by finding a shared phone.

So, we can say that the supply of calls received increases, and this has the effect of shifting the supply curve to the right (Fig. 4). Consequently, the quantity of calls made will also increase, as the equilibrium shifts. This is because callers now find it cheaper to make a call in terms of the time taken for a call to be answered, and the number of call retries, increasing the overall level of interruption in an organisation.

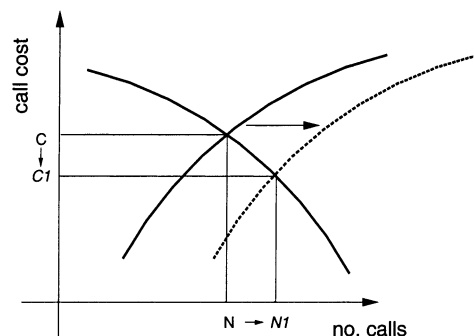


Fig. 4. Introducing mobile phones increases supply of calls, decreasing cost of calling and increasing the number of calls made.

6.2. Example 2—*asynchronous message services*

So far the analysis has looked at the effects of technological interventions directly affecting equilibrium by changing characteristics of call supply or demand. In contrast, asynchronous services like voicemail or e-mail do not alter the supply of calls nor the underlying demand for calls to be made. Rather we can consider asynchronous services as direct product substitutes, competing directly with synchronous calls to satisfy some underlying communication need.

In this case, individuals make local cost and benefit assessments about whether to use a synchronous or asynchronous service to satisfy their communication need. When it is perceived that one type of service is cheaper and as effective as another, individuals will over time adopt the most economic modality of communication.

It is likely that there may be some inherent biases in this decision making, for example a habitual preference for synchronous systems simply because they were introduced first into the communication service ‘marketplace’. Certainly, there has been an observed bias to synchronous systems in the example data that suggests that individuals considered synchronous methods to be better than the existing asynchronous ones [15].

If an asynchronous system were introduced into an otherwise synchronous dependent environment, and the new system was perceived to often be a cheaper alternative, then it would steal some of the existing synchronous market share, and have the effect of reducing the overall call demand. The net effect would be a decrease in the level of interruption and as well as call cost (Fig. 5).

In practice, while asynchronous systems might be cheaper to use, for example not having any line busy or phone unanswered




Fig. 5. Asynchronous message services reduce call demand, resulting in a decrease in calls made and call cost.

costs, they may have other hidden costs that make them unattractive. For example, the lack of feedback about whether a message has been received, read or acted upon may mean that even if there is no immediate need to complete a task, a caller cannot readily consider that sending an asynchronous message completes the communication task. The costs here include checking back to see if the task has been completed, as well as the cognitive costs of carrying an open task.

6.3. Example 3—*call filtering services*

A system that permits a call receiver to automatically filter an incoming message can be built by detecting caller attributes such as past call history, telephone number, or receiver attributes such as current task, location and so on. Messages are filtered if they fail to meet specified attribute values for allowed calls. Such filters are commonly available today to help manage the growing amount of email many individuals have to manage as part of their organisational duties.

To a first approximation we can treat a call filter as a barrier, letting in selected messages and excluding the remainder. So, without changing the inherent demand for making calls, it does restrict supply. Call receivers thus

limit the number of calls they take, at the increased cost of having to specify the behaviour of the call filter on their communication system, and perhaps also at the cost of the filter failing in some circumstances and blocking important messages. Equally, callers have to work harder to get a call through a filter on average as they build up a model of its behaviour and try and ensure their call meets the criteria for call success.

Consequently filtering services decrease the number of calls made, but at an increased cost per call for callers and receivers over time (Fig. 6).

6.4. The effect of introducing more than one technology

The current qualitative analysis could now be transformed into a quantitative one, for example explicitly measuring costs to those making calls, and those receiving them. However, as we have seen, the qualitative argument is already sufficiently powerful to draw some robust conclusions. The analysis becomes more difficult however, when more than one new technology is introduced.

In particular, when opposing qualitative directions of change in supply or demand occur, we cannot say what the overall trend




Fig. 6. Filtering services restrict supply, reducing number of calls made and increasing the cost per individual call.

will be. Thus, introducing mobile phones and asynchronous messaging will decrease overall interaction costs, but we are unclear what the effect will be on the total number of interactions at the new equilibrium (Table 1).

To resolve the issue, the opposing effects need to be relatively ordered in magnitude [16]. So, if a call filtering service was coupled with asynchronous messaging, we could have opposing effects on call cost, but reinforcing effects reducing the number of calls made. One would need to measure the cost of using the filtering service, and the changed demand on calls with competing asynchronous systems, to resolve the issue. Even simple order of magnitude measurements may be sufficient.

7. Supporting the economic analogy for interaction

The central idea presented in the last section is that we can profitably draw an analogy between the communicative interactions between agents and their economic interactions. More specifically, we can adopt some of the analytic framework of neoclassical economics to help in the analysis of general agent interactions. Interactions have been presented here as a set of ongoing adaptations by individual agents to optimise their overall resource utilisation and benefit. Cliff and Bruten have made a similar argument in economics, where they argue trading behaviours can also be considered to be an example of long-run adaptive behaviour [17].

Table 1
Qualitative changes in call costs and overall call traffic

	Number of calls	Call cost
Mobile telephony	Up	Down
Call filtering	Down	Up
Asynchronous messaging	Down	Down

However, the conception of communicative interaction between agents as trading may appear problematic for a number of reasons:

- *There is no competitive marketplace within which agents can trade.* For example, an agent may only want to interact with another specific agents. How can that be competitive, when no other agents get the opportunity to interact? This criticism is answered by noting that there is no necessity for inter-agent competition implied in the interaction equilibrium model. Rather, the competition occurs between the various channels, ground and messages the agent can bundle together during interaction. One could even envisage an interaction equilibrium emerging with only two agents over time. Certainly, experiments in auction trading have shown that very few agents are needed to replicate normal market equilibrium [18,19] Having said this, the opportunities for interaction are enlarged when there is inter-agent competition, so it is a sufficient rather than necessary condition for the emergence of interaction equilibria.
- *Money is not exchanged during an interaction, so there is no external measure of utility, nor compensation available for engaging in the interaction.* A simple response to this is to simply note that in many cases a true price can indeed be attached to an interaction, and money exchanged. For example, providers of information services charge for answering specific questions. In most cases, however, agents certainly do not charge each other for interaction. They do however make judgements about the cost and benefits of engaging in the interaction. In such circumstances, when costs and benefits are not readily reflected in monetary terms, economists use what is called a ‘reservation price’ for the activity [14]. This is the amount of money at which

an agent would no longer be indifferent between doing and not doing something. So, people can be characterised as making decisions ‘as if’ everything has a price attached. In fact, we often consider the time spent in telling someone something as a loosely transferable thing. For example, after explaining something to a colleague, you can consider that they now ‘owe you one’.

- *Agents may not be able to make free choices to interact, nor benefit from those interactions.* For example, working within an organisation, an agent may be compelled to answer telephone calls, thus having neither direct choice in the response, nor direct benefit. In response, one could simply observe that many buyers in the economic marketplace are themselves constrained intermediaries. For example, a buyer for a clothing store may gain no specific benefit from purchasing stock for the store, nor have the choice to elect not to purchase stock. Economists thus need to invoke several models of self-interest [14]. In one, the self-interest standard of rationality, agents only consider the costs and benefits that directly accrue to themselves. In the present-aim standard of rationality by contrast, they simply act efficiently in satisfying their current goal e.g. working for an organisation, they will maximise results for the employer. Both are suitable foundations for developing economic models. One can also note again that even when the choice to interact and select a channel is removed, an agent still has choice about message constructing and grounding. One can choose to answer a question to the best of our abilities, or do the minimum possible.
- *There are limited choices of channels, agents, and messages in many organisations that prevent optimal interaction equilibria*

developing. For example, an organisation may only supply telephones, or limit the time available to answer a telephone call. It is indeed hard to argue against the existence of such structural impediments in some organisations. In practice we approach communication in more or less optimal ways, depending upon organisational communication processes, and mechanisms for making communications decisions. However, just as a command and control economy doesn't preclude the existence of a free market in other countries, a restricted interaction environment in some organisations doesn't preclude their 'market' being freed up or other organisations adopting a different model. And the more freely determined agents are in their interaction choices, the more 'efficient' the 'market' should become.

- *In a free market, no individual intentionally enters into a loss making deal.* Agents enter into a trade with the expectation of some benefit. However, when an agent answers a telephone, for example, the receiver has restricted choice in answering, and knows nothing about the 'deal' being offered which could be costly. This criticism is actually directed at the phone as a 'trading' mechanism, rather than the notion of interaction as trading. When agents answer a telephone, they indeed do not have information about the proposed interaction. But we could construct scenarios in which recipients do have the choice to answer, and information about the 'deal' being offered. For example, an answering machine can be set to screen calls. Based upon the 'bids' of people talking to the machine, an agent can choose to answer or not. There is even greater freedom in choosing to interact with other communication services like e-mail. However, it is only recently that we have the means to do

so. This makes for interesting predictions about the optimising benefits of communication services that permit freer choice (e.g. call screening or e-mail) compared with those that make choice opaque such as the telephone

8. Conclusion

In this paper I have shown how we can approach the design of interactions within an organisation by viewing them from the point of view, not of the technology mediating the new interaction, but of the agents who are asked to use the new technology. Understanding the limits to individual agent resources should allow an analysis of the impact that a new interaction will have in a given setting. When we look beyond simple interaction settings, we can use the notion of interaction equilibria to predict the impact a new interaction class will have on the interactions within an organisation. Economic supply and demand curves, for example, may allow us to make both qualitative and quantitative predictions about technological adoption of communication systems.

Rather than focusing solely on characteristics of individual technologies, or psychological and social issues, these are combined to explain the overall decisions of individuals using technologies. In particular, we have seen that the supply and demand curves represent the emergent behaviours of populations of individual agents. Without necessarily understanding all the local decision criteria used by any individual, we can still make robust predictions about how a group as a whole will behave. Future work will extend this analysis to a broader technological framework, and show how a deeper analysis of local communication behaviours can help us generate economically inspired population-level

explanations of information and communication technology use.

References

- [1] T. Winograd, The design of interaction, in: P.J. Denning, R.M. Metcalfe (Eds.), *Beyond Calculation—The Next 50 Years of Computing*, Springer, New York, 1997.
- [2] G. Crampton Smith, P. Tabor, The role of the artist-designer, in: T. Winograd (Ed.), *Bringing Design to Software*, ACM Press, 1996.
- [3] E. Coiera, When communication is better than computation, *Journal of American Medical Informatics Association* 7 (2000) 277–286.
- [4] E. Coiera, Mediated Agent Interaction, in Quaglini, Barahona, Adreassen (Eds), *AIME 2001—Proceedings of the Eighth Conference on Artificial Intelligence in Medicine Europe*, Springer Lecture Notes in Artificial Intelligence No. 2101, Berlin, 2001.
- [5] T. Winograd, F. Flores, *Understanding Computers and Cognition*, Addison-Wesley, Boston, 1986.
- [6] A. Miyake, P. Shah, *Models of Working Memory*, Cambridge University Press, 1999.
- [7] J. Parker, E. Coiera, Improving clinical communication: a view from psychology, *Journal of American Medical Informatics Association* 7 (2000) 453–461.
- [8] D. Greatbatch, P. Luff, C. Heath, P. Champion, Interpersonal communication and human–computer interaction: an examination of the use of computers in medical consultations, *Interacting with Computers* 5 (2) (1993) 193–216.
- [9] E. Coiera, R. Jayasuriya, J. Hardy, A. Bannan, M. Thorpe, Communication loads on clinicians in the emergency department, *Medical Journal of Australia* 176 (2002) 415–418.
- [10] N. Booth, P. Robinson, J. Kohannejad, Identifying successful communication skills in computer use in the consultation—The Information in the Consulting Room project (iiCR), PHCSG Annual Conference Proceedings, Cambridge, UK, 2001, pp. 90–97.
- [11] E. Coiera, V. Tombs, T. Clutton-Brock, Attentional overload as a fundamental cause of human error in monitoring, Hewlett Packard Laboratories Technical Report, 1996 (<http://www.coiera.com/publica.htm>—accessed 1 October 2002).
- [12] H. Clark, *Using Language*, Cambridge University Press, Cambridge, 1996.
- [13] J. Sweller, Cognitive load theory, learning difficulty, and instructional design, *Learning and Instruction* 4 (1994) 295–312.
- [14] R.H. Frank, *Microeconomics and Behaviour*, third ed., McGraw Hill, New York, 1997.
- [15] E. Coiera, V. Tombs, Communication behaviours in a hospital setting—an observational study, *British Medical Journal* 316 (1998) 673–677.
- [16] E. Coiera, Qualitative superposition, *Artificial Intelligence* 56 (1992) 171–196.
- [17] D. Cliff, J. Bruten, Shopping ‘til you drop 1—market trading interactions as adaptive behaviour, Hewlett–Packard Laboratories Technical Report HPL-98-58, 1998.
- [18] V.L. Smith, An experimental study of competitive market behaviour, *Journal of Political Economy* 70 (1962) 111–137.
- [19] D. Cliff, J. Bruten, Minimal-intelligence agents for bargaining behaviours in market-based environments, Hewlett–Packard Laboratories Technical Report HPL-97-91, 1997.