Theoretical Issues in Ergonomics Science
Vol. 10, No. 4, July–August 2009, 315–334

Minkowski spaces as models of human–machine communication

Neville Moray[a] and Peter A. Hancock[b]*

[a]*Department of Psychology, University of Surrey, Guildford, UK;* [b]*Department of Psychology, University of Central Florida, Orlando, FL, USA*

(*Received 18 March 2008; final version received 27 September 2008*)

Acquisition and use of information about system state to initiate action is central to the safety and efficiency of human–machine operations. A model of limits is presented on the rate at which past information can affect the perception of a system's state and thus the rate at which decisions affect the future. The model represents timescales of physical processes, operator interventions and management of decisions and actions. It is founded on Minkowski space–time diagrams showing cones representative of accessible and inaccessible past and future. This approach is generalised to establish widescale application to all human–machine systems. Unlike physics, for which this description was developed, human scale operations are bounded by lower information transmission speeds. Thus, the interactions in a multi-dimensional space are complex and one element of the system may have access to regions inaccessible to others. A computational version of the model may predict features of human–machine interaction.

Keywords: Minkowski spaces; human–machine interaction; communication model

1. Introduction

Recently, research on human error and accidents has emphasised the problems of integration between maintenance personnel, operators and management (Reason 1997). At the core of safety-critical human–machine systems is the technological implementation of more or less hazardous industrial systems, such as chemical plants, power stations or transportation systems. Events in such systems, particularly when they are highly automated, evolve on the time course of physical causality from fractions of a second to many hours and even days. Processes are monitored and controlled by human operators, who are dependent on sensors, processors, computers and interfaces to identify the state of the system, and are limited in their speed of response by psychological and physiological timescales (as outlined, for example, in Card *et al.* 1983). Management decisions in turn depend on the transmission of information from the level of operations and maintenance and from the external world of legislation and societal expectations. The effects of management decisions thus often take a long time to affect other levels of the system.

The timescales of information acquisition and of outcome action are critically important in the management of complex systems. This temporal dimension of human–machine interaction has been emphasised in recent years by the francophone school

*Corresponding author. Email: phancock@pegasus.cc.ucf.edu

ISSN 1463–922X print/ISSN 1464–536X online

DOI: 10.1080/14639220802281299
http://www.informaworld.com

of field studies in human–machine interaction (see, for example, De Cortis 1988, Cellier *et al*. 1996, Hoc 1996). If the dynamics of the physical system are fast, and the time needed for humans to acquire information and propagate action is long, then the system will often be in danger of running out of control or will be subject to actions based on inadequate information about system state. High level managerial decisions can be so delayed that they become irrelevant to the later stages in the evolution of a real-time event. If critical events are slow, then it may take a long time for changes in the system to become apparent and the process may have undergone an irreversible change before the state can be identified and operators (or automatic controllers) initiate appropriate countermeasures. As is known from control theory, mismatched time constants frequently give rise to instability (see, for example, Sheridan and Ferrell 1974, Sheridan 1992).

It would therefore be most useful to have a unified framework in which to represent the temporal dynamics of participants in complex hierarchical human–machine systems. It would be particularly useful to represent the temporal relations between information acquisition and the causal effects of intervention so as to identify critical situations in which temporal relations between information acquisition and action may lead to a loss of coordination in event management. The premises for the framework offered here are therefore:

(1) Complex human–machine systems depend on communication among their component subsystems to achieve coordination and stability.
(2) Both information acquisition and control actions can be thought of as forms of communication. The former is communication from the system to the controller about the past values of state variables, and the latter is the communication of causal influences into the future from the controller to the system.
(3) The present cannot be known with exactitude. At best, the very recent past can be known to a somewhat greater degree.
(4) Communication, coordination and the propagation of causality can be represented by Minkowski diagrams, which can be generalised beyond their original purpose to reveal weaknesses and strengths of complex systems, which have different timescales and which are informationally and causally coupled to one another. In particular, problems of temporal coordination can be specifically illustrated in such diagrams.

2. Minkowski diagrams

The German physicist Herman Minkowski (1908/1923) introduced the space–time diagram within which the history of any particular entity, identified as one 'world line', can be located. The physical universe can be represented as a four-dimensional space in which the horizontal axis is time-like and the other three axes are space-like coordinates. For simplicity a two-dimensional version of the diagram is usually used to present the ideas, with the x-axis being time (T) and the y-axis one dimension of space (S). The generalisation to three (or more) spatial dimensions is straightforward, as is the generalisation of 'space' from Euclidean three-dimensional physical space to conceptual spaces of arbitrary dimensionality, using the Pythagorean relation to calculate distances. As shown in Figure 1, the diagram has four distinguishable regions of space–time and a unique cut on the time axis. This cut is made by a vertical line which cuts the T axis where the time is the present and the location in space is the location of the observer-controller. This point, $P(0, 0)$ is the observer's present, here-and-now (sometimes called the 'specious

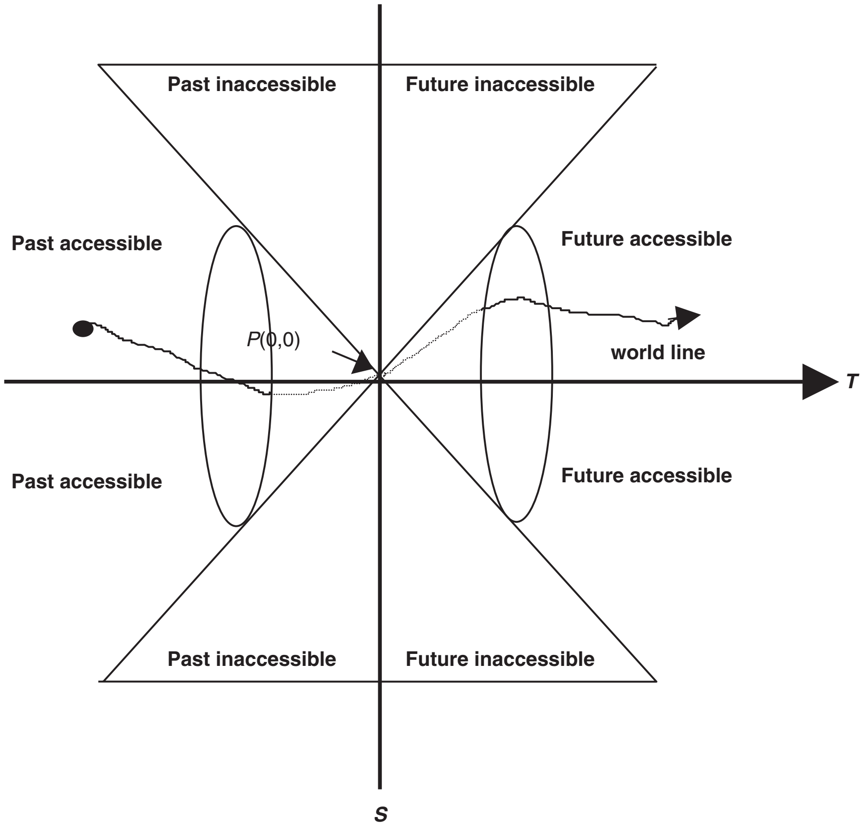

Figure 1. Prototypical Minkowski diagram for time and two-dimensional space. $P(0,0)$ is the Present 'here-and-now'.

present'). To the left (negative values of T) lies the past and to the right (positive values of T) the future. Two diagonal lines divide space–time into accessible and inaccessible regions, which lines intersect at $P(0, 0)$.

Space–time can be divided thus because there is a limit on the speed at which information can be transmitted through the universe, since according to the theory of relativity no information can propagate faster than the velocity of light. An event may be so remote from the observer (O), who is at $P(0, 0)$, that O cannot at that present know of the occurrence of the event because there will not have been time for information, even travelling at the limiting velocity of light (c), to reach O. Such events are located in the regions labelled 'past inaccessible'. Similarly, since the effects of action on the future state of the universe can be thought of as the transmission of causal information to some region of the universe, there will be locations at distances d such that they cannot be affected by any action of the observer until a certain time T^* has elapsed, because even if causality is propagated at a velocity c it will not reach those regions before time T^*, where $T^* = d/c$. These regions are labelled 'future inaccessible' in Figure 1. The history of a particular observer, O, or entity as he, she or it moves through space and time is represented as a 'world line', the locus of a point that at all times lies within the accessible regions of past and future and passes through $P(0, 0)$.

Given Wiener's (1961) definition of cybernetics as 'communication and control in the animal and the machine', communication and control can be unified as the transmission of information. Given the velocity of light as a limiting rate of this transmission, the geometry of the future accessible space determines what causal control can be exerted when and where by an observer who is an operator at $P(0,0)$. (This paper omits any consideration of analogies of quantum effects. Recent discussions of 'entanglement' in quantum theory (Bouwmeester *et al.* 1997) do not affect the arguments put forward here unless para-psychological phenomena exist, for which there is, the authors believe, no current convincing evidence.)

3. Generalised Minkowski space–time diagrams as models of complex human–machine systems

In the Minkowski diagram, the form of the cones, and in particular the angles that the boundaries between accessible and inaccessible past and future make with the S and T axes, are determined by the universal constant c, the velocity of light. The space represented by the diagram is neither space nor time, but 'space–time' in which velocity and distance are scaled by the velocity of light. As a result the angles that the boundaries of the cone make with the T and S axes are always 45°. This is typical of how physics views this formulation. However, one may consider a more general case of any (multi-dimensional) space in which there is a limit on the velocity of communication in that space, whether as communication or as control. The Minkowski space is then a generalised communication space and the representation of the physical universe treated in relativity theory becomes a special case of the general principle. This paper will describe the regions that are accessible as being within the observer's cones of accessibility (COA) and regions that are not accessible will be called inaccessible spaces (IAS). Now some general properties of such spaces are considered (see also Hancock and Chignell 1995).

First, consider the effect of different limits on the rate of information transmission. If the limiting velocity $\mathrm{d}S/\mathrm{d}T$ is very low, the accessible cones become much flatter. It requires a longer time to learn about the past and takes longer to influence the future. By contrast, if the limit on $\mathrm{d}S/\mathrm{d}T$ is high, more of the space is accessible both in past and future; see Figures 2 and 3. The rest of this paper will represent the limiting velocity for communication as $\boldsymbol{c}$, a symbol for a generalised limit of speed of communication.

A generalised space need not be symmetrical with regard to past and future. One can imagine spaces where the communication of information to the operator at $P(0,0)$ (reception of information) is limited by one velocity of communication, while the operator's causal communication with the future is limited by a different velocity, perhaps because the operator uses a different medium for communicating causality from that through which information is received. In such a case, it could happen that an operator receives information very rapidly about events in the past but can only very slowly affect events in the future. Such a case might be quite common if the perception of displayed information is very rapid but decision and response time is prolonged, perhaps because of the latency of system response (as in super-tanker control). In such cases the Minkowski space would resemble that shown in Figure 4. (The opposite could also occur, with slow assimilation of displayed information but fast control response.) In general, one can partition $\boldsymbol{c}$ into components governing the limiting rate of arrival of information from the past, $\boldsymbol{p}$, and the limiting rate of communication to the future by $\boldsymbol{f}$.

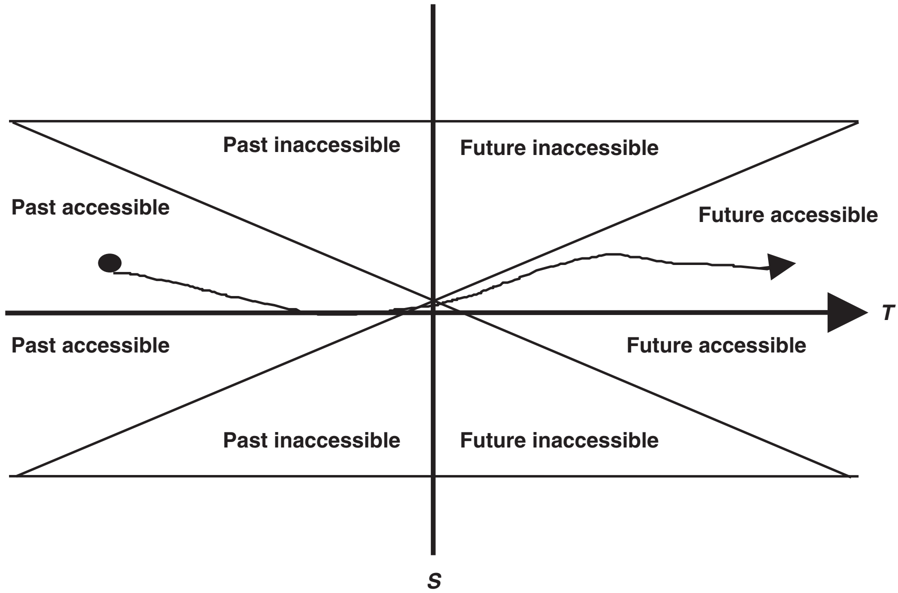

Figure 2. Minkowski space–time diagram: low dS/dT for a slow system.

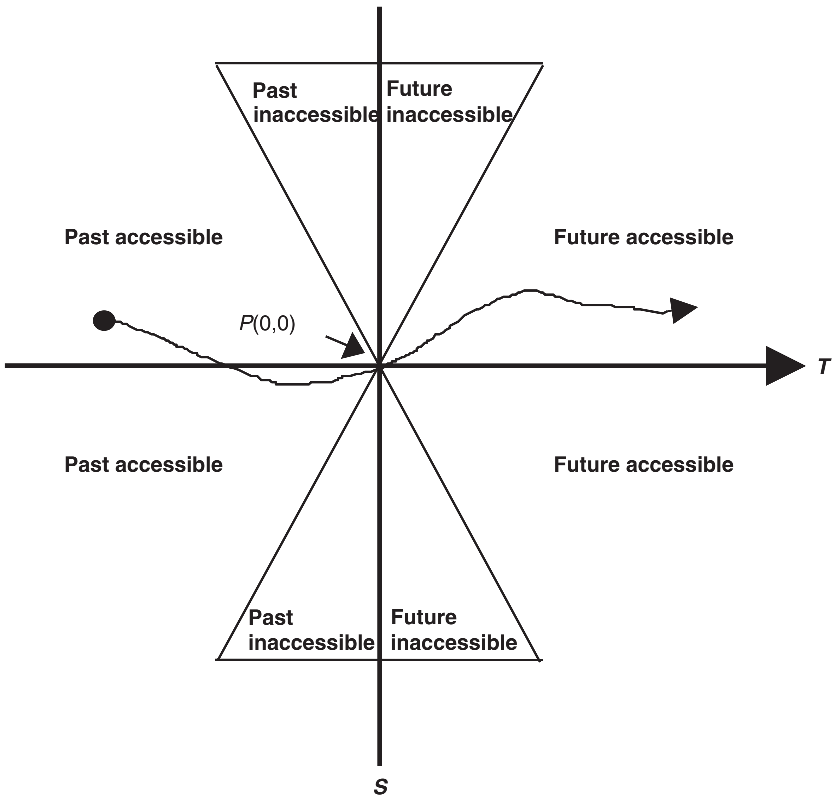

Figure 3. Minkowski space–time diagram: high dS/dT for a rapid system.

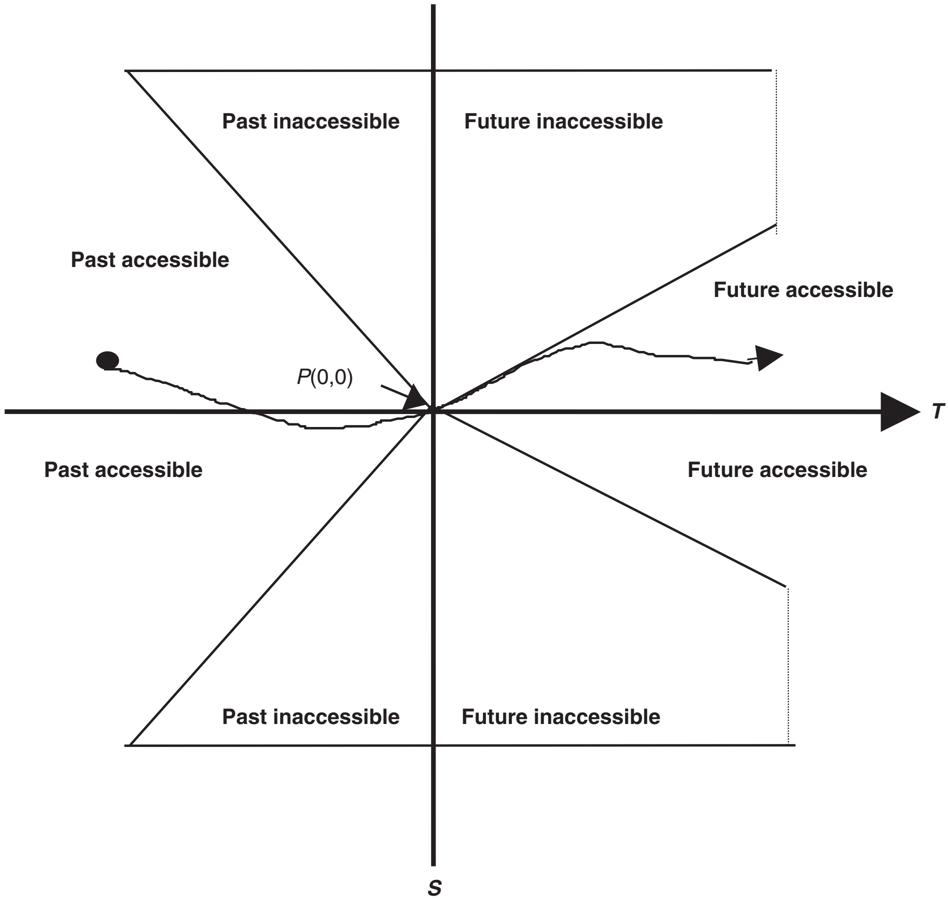

Figure 4. Minkowski space–time diagram: asymmetric limits to rate of acquisition of past information and future control.

Given the general description of operators and entities in human–machine systems, there will almost always be more than one entity (object, observer or operator) present at time P and they will necessarily and inevitably occupy different locations. The case of a two-entity space is shown in Figure 5. Here, for the observers O1 and O2, the present occurs at the same T_o for each, but since they are separated in space they are located respectively at (S_1, T_o) and (S_2, T_o). In this example, the rate at which information can be processed by O2 is less than that of O1 and hence O2's COA is shallower than that of O1.

An event (E) occurs at $E(S_E, T_E)$. By definition, O1 becomes aware of E at the moment information crosses the S axis of O1's COA, and O2 becomes aware of E when information crosses the S axis of O2's COA. Those moments can be identified by producing the dotted lines from $E(S_E, T_E)$ parallel to the slopes of the respective COA until they cross the respective T axes. In this example, E occurs much closer to O2 than to O1, that is ($|(S2)o - S_E| < |(S1)o - S_E|$); but because the slope of the O2 COA is much less than that of O1, E is processed by O2 much later than O1. To estimate the minimum delay before knowledge of E is experienced by O, a line is drawn from E parallel to the slope of the COA and where that line crosses the S_0 axis, the value of T defines the 'date' at which the event can be known to O and hence when the information becomes knowledge useful

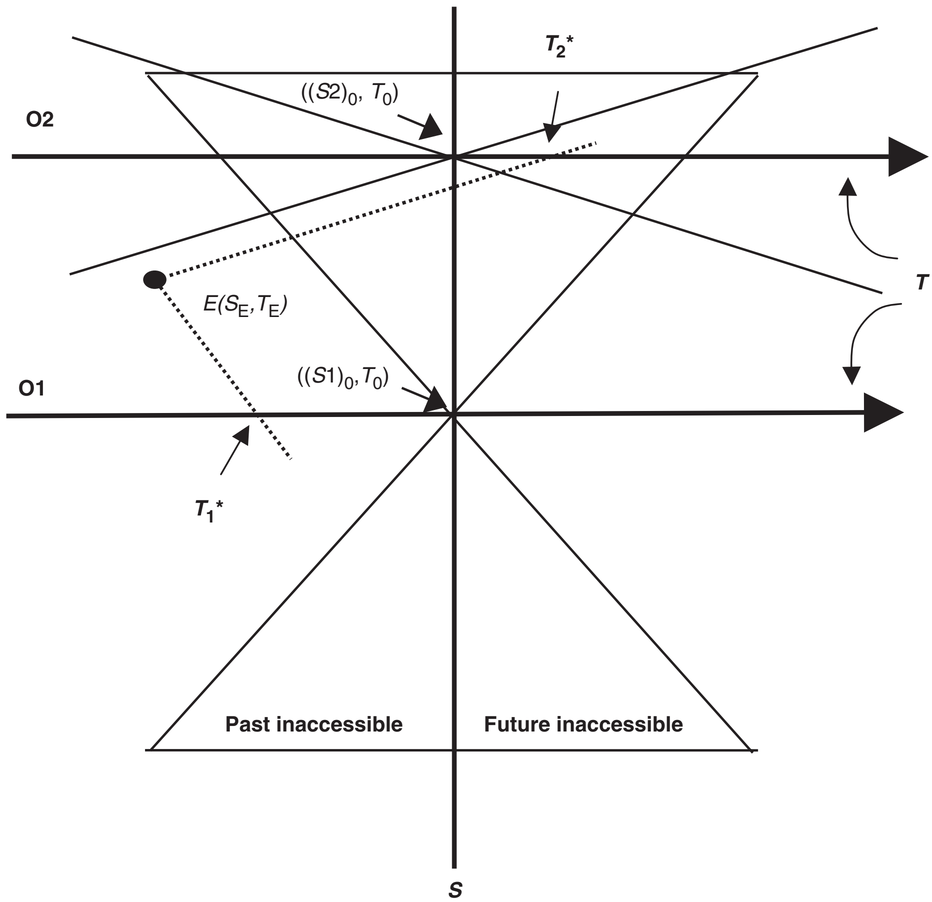

Figure 5. Minkowski space–time diagram with two observers O1 and O2 located at $P(S_1,T_{(\cdot)})$ and $P(S_2,T_{(\cdot)})$. E is an event to be observed. Although the T axis is drawn for two values of S, it is common to both observers. See text for further details.

for initiating action in the right-hand COA, in which the slope of the COA has a similar role for the propagation of action.

The present always occurs simultaneously for the two observers and is indicated by the vertical line at $S_{(\cdot)}, T_0$). The convention $X_{()}$ is used to indicate a variable that can take several values, to all of which the same comment applies. Hence, in this case, $S_{()}$ means that the comments apply to all S.

The time until which O1 or O2 can affect the future is similarly dependent on the slopes of their COA to the right of the present. The point at which an observer's COA meet is here-and-now, the present, T_0.

One is free to generalise the notion of 'space'. Space could be a multi-dimensional 'problem space' as used in discussing human decision making and problem solving, or in artificial intelligence, or it could in particular be the multi-dimensional space in which the history of a complex human–machine system unfolds. In this latter space, the entities that inhabit the space include the components (at all levels) of the physical plant, the state variables, the human operators, the managers and organisational entities; in fact, all the entities that reveal themselves in an abstraction hierarchy or a part-whole decomposition

(Rasmussen 1986, Rasmussen *et al.* 1995). In such a space, the rates of communication of different entities are limited in different ways. Instead of a single limiting velocity ***p*** or ***f***, there is a family of velocities, some physical, some physiological, some psychological, some sociological, which can be referred to as **p1**, **p2**, etc.

Where do these limits arise? Some come from the inertia involved in hardware-based actions. For example, suppose that an automated system monitors the system state by sequentially polling a large set of sensors. The minimum time at which an event can be communicated to any entity at (S_0, T_0) is then dependent on the cycle time of the polling sequence, plus sampling time, plus processing time in the process computer, plus the time to write to the display screen or activate an alarm. If a human operator is scanning displays in a control room by moving his or her eyes over the displays, then the acquisition-limiting rate is set by eye movement dynamics plus the time to read a display when visually fixated. If the control room involved is very large and sampling is performed by the operator moving round the room, then the limiting rate is further influenced by the time to traverse the room from one point to another. (Anyone who has watched a crew dealing with an unexpected incident in a large complex control room will have seen this limit affect behaviour as the operators move from one console to another across the room.) In a computerised control room, the limiting acquisition rate may be the time it takes to access the appropriate page of information in the computer's database and alarm display system. Furthermore, the discrete digital nature of a computer, whereby new information is available only each time the screen is refreshed, means that the limiting cones may be discrete rather than smooth functions of time (the iterative nature of an update rate is illustrated in Figure 6). The slope of the cone will depend on a kind of

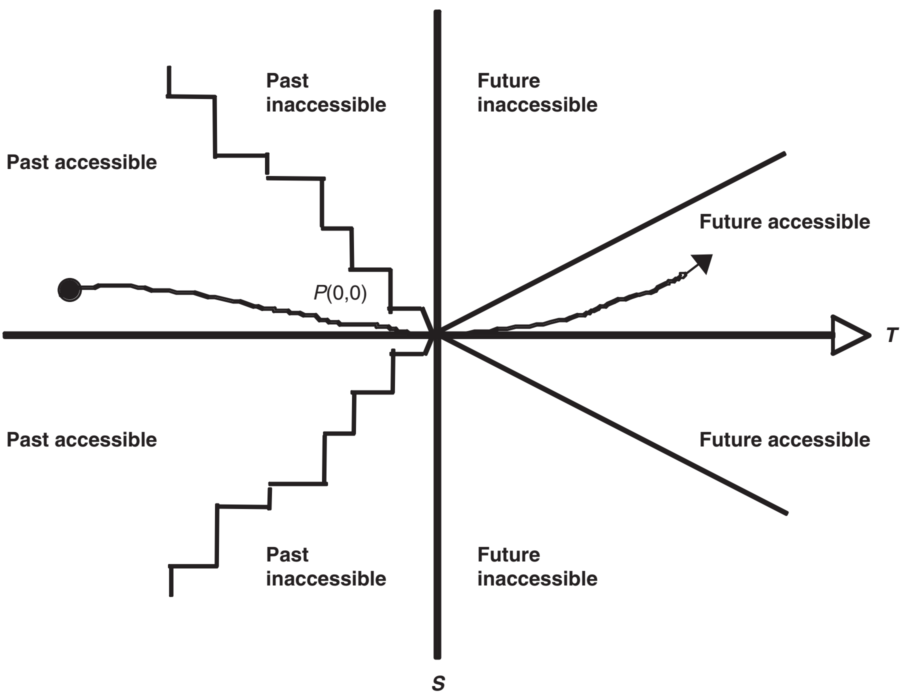

Figure 6. Minkowski space for a system in which the acquisition of information is digital, but control is continuous.

convolution of the computer hardware rate of operation and the movements used by the human to call up the appropriate page and use a trackball, mouse or keypad or other device (Hancock 1996). The time to influence the future (the slope of the limiting cone into the future) is similarly influenced, for example, by hardware limits, such as the time to operate a motorised valve or for control rods to be driven into the core of a reactor, sometimes by the speed of computational algorithms in software and sometimes by the time taken for the human operator to exercise his/her perceptual-motor skills.

A detailed analysis of the slope of the COA would show that it is constantly changing. It is flatter where the observer moves physically about the control room, steeper where it represents perceptual judgement and steeper still when it represents the time to paint information to a computer screen. Moreover, when modelling the human–machine system there should be a representation of the variability of the slope at all points, perhaps using standard deviation bars.

The absolute timescales of these differing constraints can vary considerably. Thus, update rates in fast computational systems can be in the order of micro- or milliseconds, while human locomotion around a large control room may, as has been noted, well take seconds or even minutes. While it may initially appear that only the slowest timescale provides the rate limitation, this is not necessarily so. This is because it also matters when any particular rate limitation occurs in the respective areas of accessibility, hence the need for the present over-arching descriptive framework.

Figure 7 shows an example of a Minkowski space whose communication limits are set by the rates of the physical processes in a large-scale system without any operators.

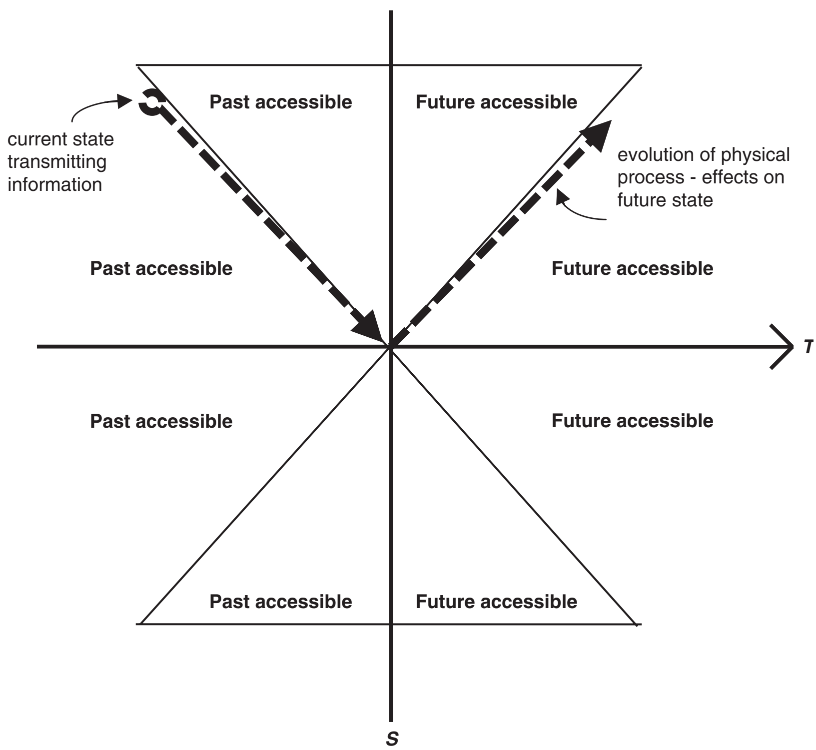

Figure 7. Minkowski diagram with limits set by rates of fast physical processes both for information acquisition and for control actions: a fully automatic system.

The space represents events in which one part of, for example, a large industrial plant can receive information about the states of other parts and in which one part can communicate causal effects to other parts. (For simplicity, in this and the following diagrams, symmetry between past and future will usually be assumed unless otherwise indicated, but this restriction is not a necessary one.)

4. Generalised Minkowski space–time diagrams as models of human–machine systems

This section considers the case of machine automation, where an operator monitors the plant and intervenes as needed. In industrial systems it will often (but not always) be the case that the reaction kinetics or thermodynamics of the process are faster than the ability of the operator to track the observable signals that define the position of the plant in the state space. This is due to certain inherent limitations, such as the relatively slow rate at which humans can switch attention, and the number of degrees of freedom in the system and the number of components, displays, etc. In fact, from here on this paper will concentrate on systems in which the rate limitation is always due to human operators. This case is shown in Figure 8, where the **p** and **f** lines indicated by long, heavy dashed lines show the limits on information processing and action due to the physics of the process, and the continuous lines with a lower slope the limits due to the human observer-controller.

In such a case, the operator O can eventually learn about every part of the system state space, but only after prolonged observation and analysis; and the operator can affect any part of the system, but again only after a delay in the case of the more 'distant' components. (The plant and the operator are shown at identical locations at $P(0, 0)$ for convenience. No loss of generality is implied.) The difference between the shapes of the limiting COA for the physical plant and operator implies problems for the operator in the form of constraints on understanding and controlling the plant in real time. While operators may be able to identify the plant state in principle, in general they will lag behind the evolution of events. In Figure 8, the physical event E affects the system at (S_0, T_0) but cannot be known by O until the information reaches them along the slope of O's personal COA at $P(S_0, T^*)$. Thus, there will be a delay in O's response to the event and the information from E will have begun to cause a physical effect by propagating causal information downstream before O has knowledge of the need to respond.

One can call the slopes of the physical plant hardware *f* function $\mathbf{p}^{\mathbf{plant}}$ and $\mathbf{f}^{\mathbf{plant}}$ and that of the operator $\mathbf{p}^{\mathbf{human}}$ and $\mathbf{f}^{\mathbf{human}}$. The region between the cone boundary of the physical plant dynamics and the cone boundary of the operator dynamics is a region of time that system designers must look to reduce. It is in this time region, between $\mathbf{f}^{\mathbf{plant}}$ and $\mathbf{f}^{\mathbf{human}}$, that operators lose control of the plant. One way to reduce this space is to use fast time computer models, which can be accessed by the operator, predictive displays or suitable operator mental models. There is, in fact, some evidence that the frontal cortex functions in this fashion, so helping humans to adapt to the vagaries of uncertain, and only partially predictable, environments (Hancock 2007, 2008). Figure 9 clearly shows the reason why fast time models and predictor displays are valuable: they increase the slope of the operator boundaries by providing an effective $\mathbf{f}^{\mathbf{human}}$ that is as fast as, or faster than, the $\mathbf{f}^{\mathbf{plant}}$.

Figure 10 introduces an even slower component. As is known from accident enquiries, such as that into the events at Three Mile Island or the sinking of the *Herald of Free Enterprise*, there is frequently a very long delay between the reporting of events by operators and their reception by management (Reason 1997). There can also be long

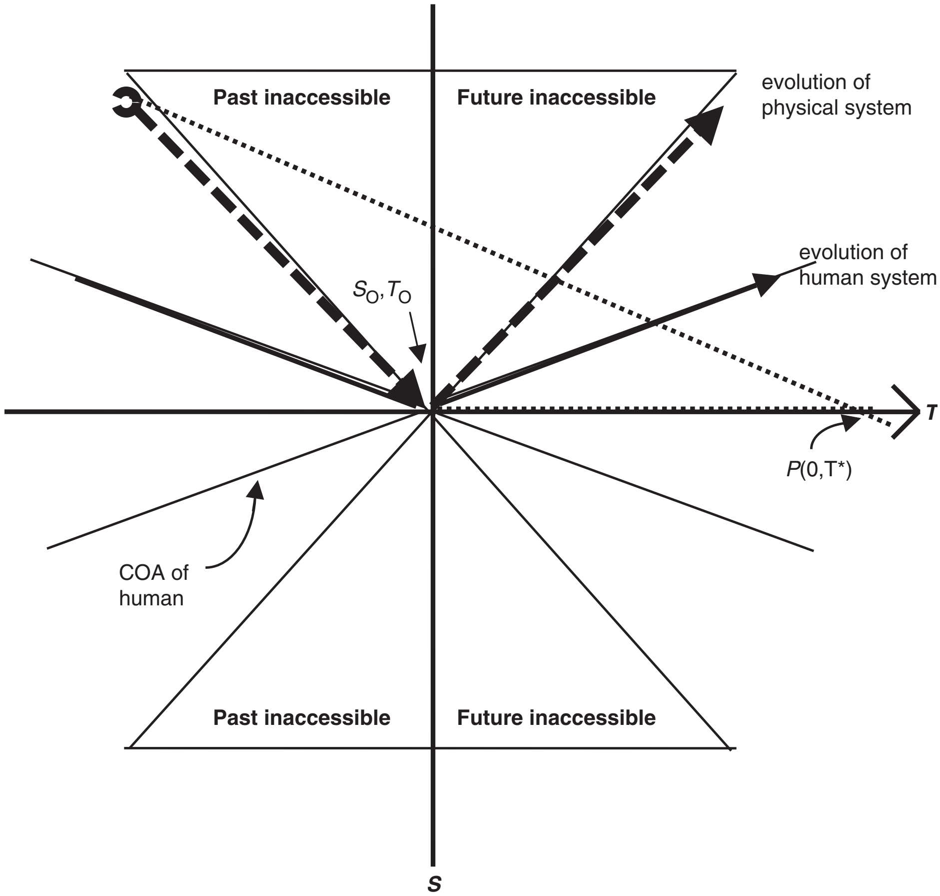

Figure 8. Minkowski space with limiting cones of fast physical processes and slower human information acquisition and action. The human at location (S_O,T_O) cannot know about the event that starts the physical process until the time T^* due to the slope of $\mathbf{p}^{\mathbf{human}}$. COA = cones of accessibility.

delays before managers act or before the instructions of management are implemented. In general, the physical plant is at an inner loop level, operators at an intermediate loop level and management at an outer loop of a series of nested closed-loops that comprise a hierarchy of control. The time constants of the outer loops are usually progressively longer than those of inner loop components and the associated COA are therefore flatter.

The implications of this are shown in Figure 10, which now includes the Minkowski space for management. The same comments can be made here in respect of the implications of the difference between the accessible regions of the respective spaces. Events that are seen as critical by operators cannot 'reach' management until a long time has passed, nor can decisions of management feed back as information to be used by operators.

The diagrams show clearly the importance of the design and philosophy of operating procedures. If operating procedures require rigid adherence to rules, and if operators may only deviate from prescribed rules with the permission of higher authority, the system will be unable efficiently to respond to fast transients of an unexpected nature, since the $\mathbf{f}^{\mathbf{plant}}$ of the physical system has a limiting rate that far exceeds that of the $\mathbf{p}^{\mathbf{mng}}$ and $\mathbf{f}^{\mathbf{mng}}$ cones

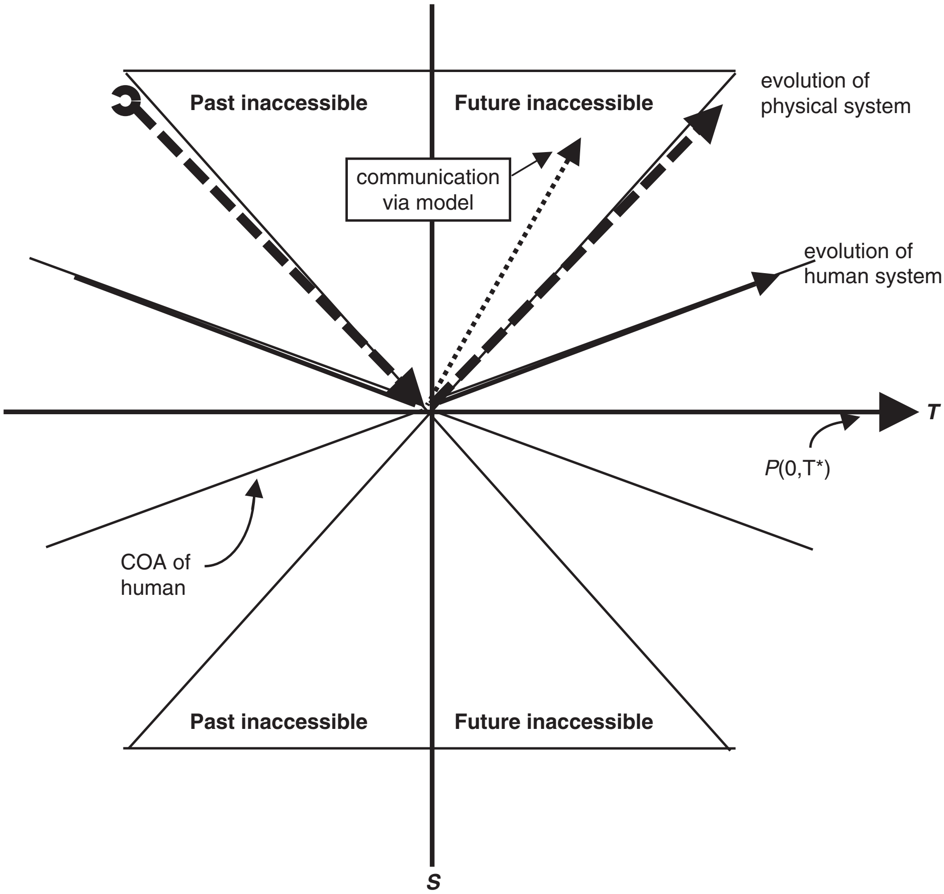

Figure 9. Effect of fast-time model. The rate of $\mathbf{f}^{\mathbf{human}}$ is increased to one greater than the rate of evolution of the physical process as the model is used to generate action in response to the information acquired. COA = cones of accessibility.

of management. Each level of authority adds a delay, so that the COA of the overall human–machine system becomes progressively flattened. Intervention by the operator will thus be progressively delayed and the danger that the overall system will 'run away', beyond the capacity for control, will increase accordingly. Accordingly, the problem of catching up with the evolution of the plant will tend to be less soluble and unless drastic action is taken to provide a short cut, control will be lost. Note also that if the operating rules insist on a period during which the operators may only observe the plant without intervening (as in the '20 min rule' common in nuclear power plants) the Minkowski diagram is further altered and a larger inaccessible region introduced (see Figure 11). Only if operating procedures permit some degree of flexible response without the appeal to a higher authority, and only if operators have some kind of access to fast time models, can unexpected fast events be dealt with in an effective manner with a spatio-temporal cushion of safety. Thus, although examples from the realm of process control have been featured, Minkowski diagrams can, in principle, be used to evaluate any complex human–machine system. The usefulness of this approach is not restricted to evaluating operating and management philosophies or to just analysing the human–machine system at the level of

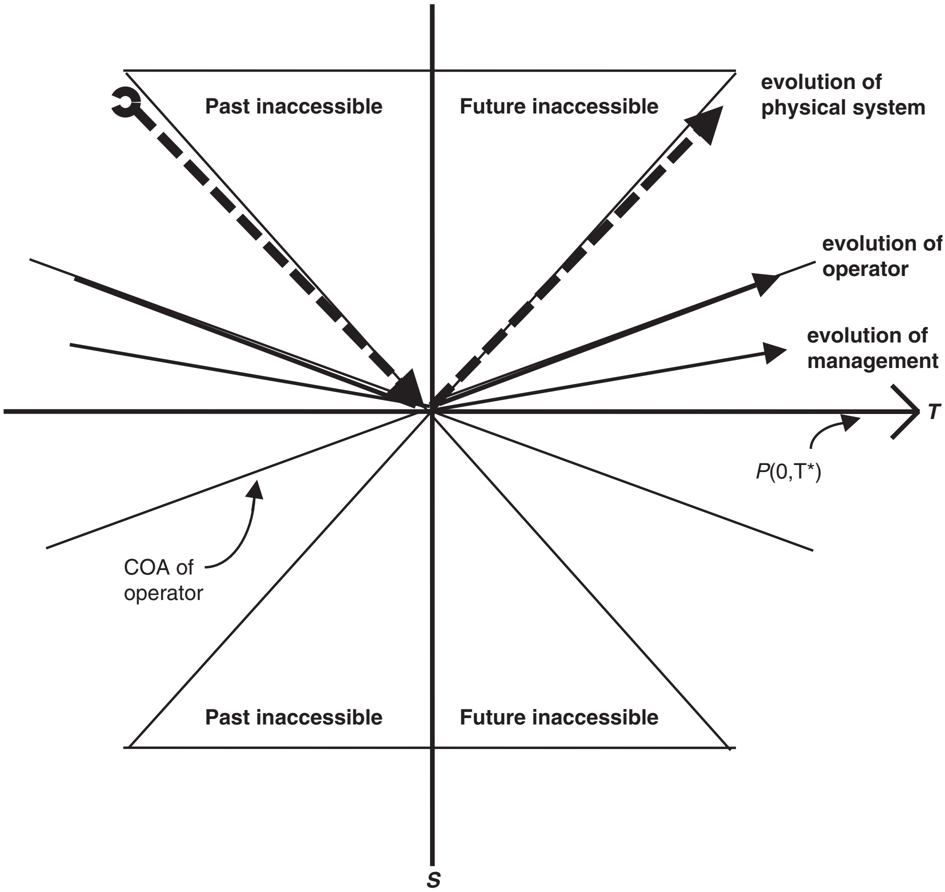

Figure 10. Owing to the time taken to communicate from inner to outer control loops in a hierarchical system, the response rates of management are even slower than those of operators. COA = cones of accessibility.

control room design. The authors believe that the framework that has been described is one of ubiquitous utility to all human–machine systems, their design and operation.

5. Towards a computational model

While these diagrams reveal qualitatively the origins of problems of communication and control in the different limiting rates of communication, it is desirable to turn this approach into a predictive model. The ultimate goal must be a computational model to allow one to locate the various entities of a complex human–machine system in multi-dimensional space and identify where problems may be expected to arise. This requires a quantitative approach to the construction and analysis of these diagrams. That in turn requires one to identify metrics for the velocity of communication and distance in the space; one needs metrics for the $\boldsymbol{p}$ and $\boldsymbol{f}$ functions. In addition, one needs to be able to deal with interactions among the various components. Can one represent, for example, the impact of a multi-person crew on the speed of response? Can one represent different strategies of management intervention? Can one represent closed-loop systems where

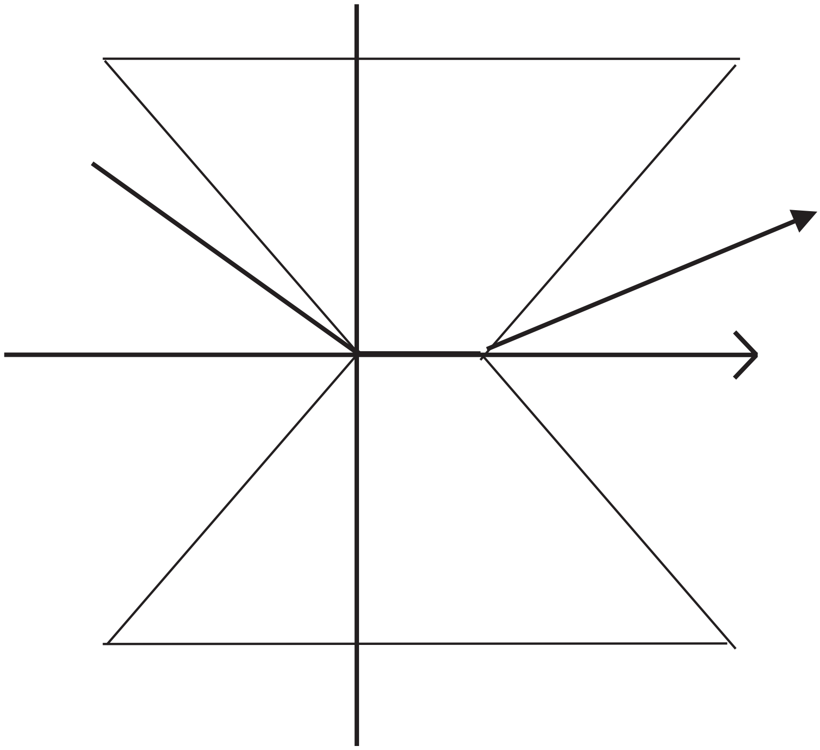

Figure 11. Enforced delay. Operating procedures impose a delay before action can be taken (see, for example, the '20 min rule' used in many nuclear power plants). The size of the inaccessible region for action is increased.

causality runs backwards as well as forwards if expressed in a conventional block diagram? Several approaches suggest themselves.

The *T* axis in these diagrams can be scaled in conventional time units (seconds, minutes, hours), but what units can be used for distances on the *S* axis? In Minkowski's original treatment, *S* would be measured in physical units of space such a kilometres, millimetres, etc., or in 'space-like' units in which these conventional units are rescaled by the velocity of light. As stated earlier, inherent physical constraint establishes the slope of the COA in relativity theory. However, this cannot be done in the application that is proposed here. Rather, the distance, $S_{(\cdot)}$, from the 'here' axis, is also time-like. It is measured in the time it takes to operate on information. Thus, an event requiring three images to be painted to a computer screen is 'further away' from O than an event needing only one painting. An event requiring O to walk across the control room for 10 seconds is 'further away' than one needing no movement and is given literally by the speed of movement of the operator. An event requiring O to make three eye-movements is 'further away' than one requiring only two eye-movements, and one needing two perceptual judgements is 'further away' than one requiring only one such judgement. A display with a low signal-to-noise ratio, for example, is 'further away' than the same display with a higher signal-to-noise ratio, because processing time is inversely related to the signal-to-noise ratio (Sheridan and Ferrell 1974).

Quantitative values for properties such as those just listed could be estimated, for example, from approaches such as the Model Human Processor (MHP; Card *et al.* 1983) or a similar database. In addition, the definition of distances and velocities requires a detailed task analysis for application to a specific system. The slope of a COA must be calculated step by step, as in the MHP, including, if necessary, changes in slope caused by a multi-person crew. This is easily represented on a diagram, as is the change resulting from having a crew rather than a single operator. Perhaps one may think of distance in terms of the strength of coupling between components. A weakly coupled component is more 'distant' from the component to which it is coupled than is a strongly coupled component (Perrow 1984). In that case there are various quantitative metrics that can be employed, such as Conant's information transmission analysis (Conant 1976).

The slopes of the limiting ***p*** and ***f*** lines describe the fastest possible responses to events. However, it should be noted that in designing any human–machine system one may also wish to know what will happen if the operators do not necessarily respond as rapidly as possible, either to the need to acquire information or to make a response. In those cases, the slopes of the actual information transmissions will be flatter than the limiting value itself. A slow acquisition will mean that information about a given will not be acquired at $P(S, 0)$ but at some later time, $T \neq 0$. The effects of action can similarly be delayed.

Classic Minkowski diagrams describe only open-loop systems, since they were originally proposed to describe one-way communication from the past to the future (and the history of the universe is presumably open loop!). However, by combining the world line histories of two components, one can also describe closed loop systems within this general framework. In Figure 12, the entity located at $P(0, 0)$ is an error sensor that receives information from the past about the deviation of the system from its desired state. It propagates a signal (command) to an effector with a delay T_1, and at location $p(x, t)$ at time $t = T_1$, the effector acts on the basis of that information. After a delay, T_2, the feedback signal arrives back at $p(0, t')$ and the cycle repeats.

As stated earlier, in principle one may construct the spaces using methods such as those derived from GOMS (Card *et al.* 1983) to calculate the expected durations of perception, movement, judgement, etc. Applied to a detailed task analysis, the COA of operations could be related to those of the physical processes for critical incidents of system operation. Team interaction could be modelled based on times for speech communication, properties of the channels of communication used, etc. (see Figure 13). Since GOMS and other similarly structured models are specifically designed to predict temporal characteristics of operator behaviour, they provide a natural starting point for developing metrics for human–machine interaction, but it can well be augmented from other quantitative models of human–system interaction (see, for example, Pew 2007).

If such metrics can be identified, a number of interesting analyses become feasible. For example, the change in the shape of mutually inaccessible spaces resulting from delays at the interface could be calculated and exploited in design and interactions between different levels of the human–machine system hierarchy could be represented. The shapes of the respective COA could be used as an important design heuristic. Further development of interpretations of Minkowski spaces would be needed, but could be made providing that the structure of the spaces was rigorously defined. Consider, for example, how one might represent interaction with management when the latter, in response to information that a change in plant status has occurred, either provides assistance or 'censors' the activity of operators by preventing certain actions from being taken. The question here is how to represent the interactions of the spaces inhabited by the respective entities. The properties of entities in human–machine system space have many characteristics that are quite unlike the linear nature of the original Minkowski space. For example, it is possible that following the reception of information by management, the latter might institute control

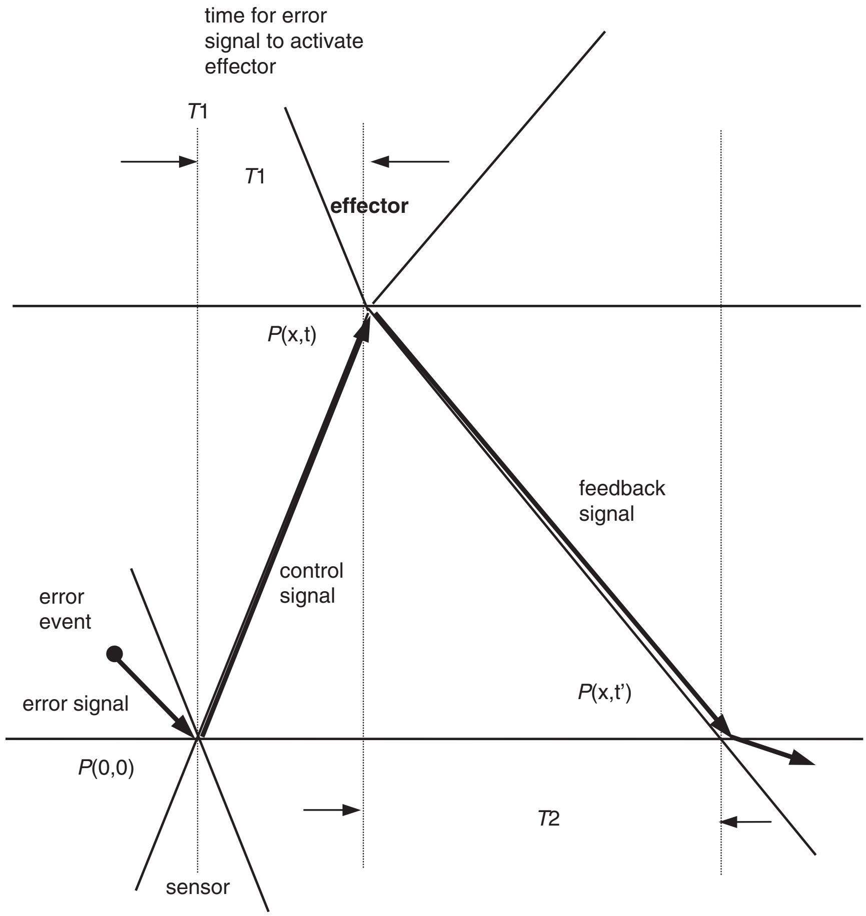

Figure 12. Representation of closed-loop feedback control in Minkowski space.

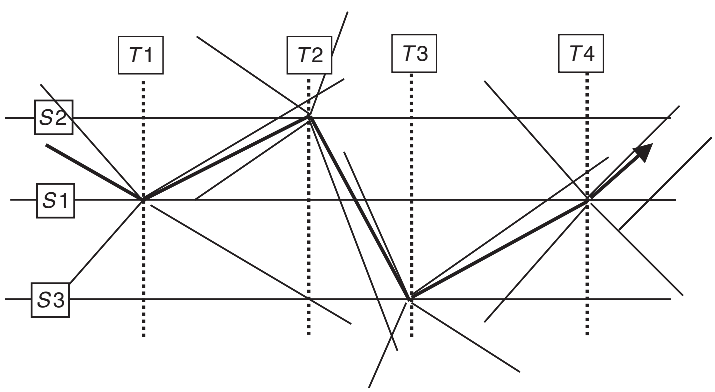

Figure 13. Intensive communication. Operator *S*1 receives information at time *T*1 and sends a message proposing action to operator *S*2, who receives it at *T*2 and in turn communicates with operator *S*3 at time *T*3, who finally confirms the choice of action by *S*1, who then responds at *T*4. Note that no operators act as rapidly as possible. Thin lines represent the ***p*** and ***f*** cones of operators. The thick line is the trajectory of information through the system space.

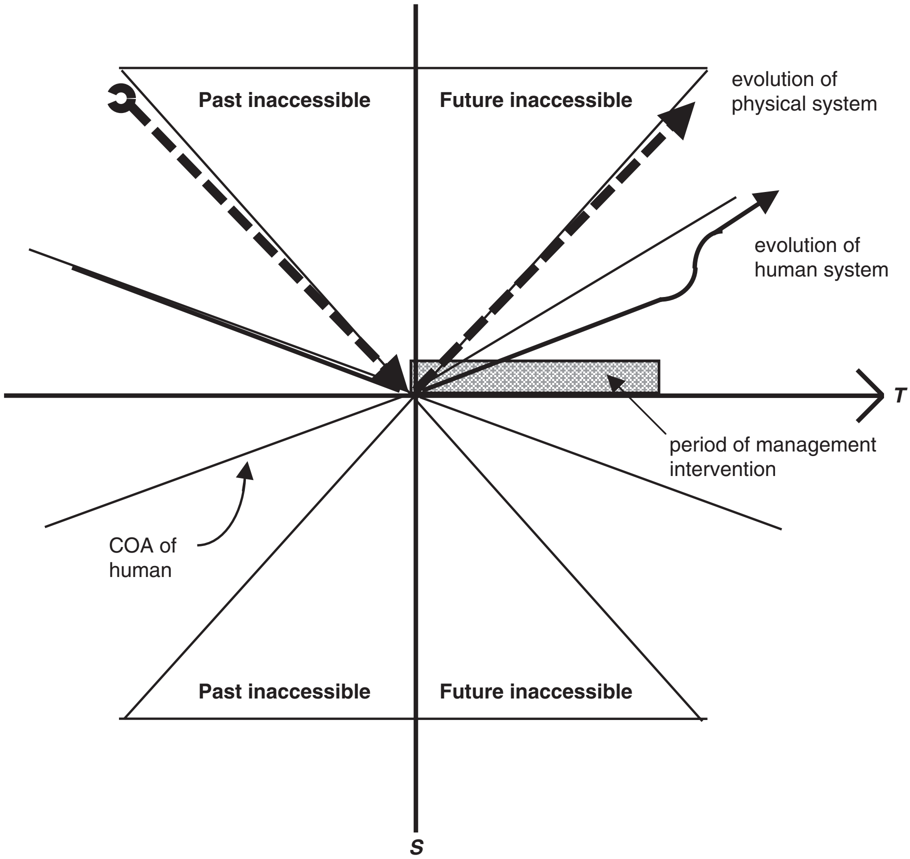

Figure 14. Effect of management inhibitory intervention. The $\mathbf{f}^{\mathbf{human}}$ is slowed, so that the rate of evolution is less than normal. After the management intervention ceases, the rate returns to a line parallel with the original uninhibited rate, but some overall delay remains. COA = cones of accessibility.

that is valid only for a limited time into the future, after which their intervention ceases to be relevant. They might, for example, divert personnel for a short period and then restore them. Such a case is represented in Figure 14.

How can this diagram be interpreted quantitatively? The effective cone of the human system as a whole is represented by the slowest component. Thus, the slope of the 'operator intervention' cone would be deformed by the intensity of the management interaction. The result of transient interference would be something like that shown in Figure 14. After the interference has ceased, the $\mathbf{p}^{\mathbf{human}}$ returns to its original value, but responses are always delayed compared with what they would have been had the management not interfered. If management intervention facilitates operator performance, the response COA can be steepened (see Figure 15).

6. Summary and conclusion

A central issue in the prevention of failure in all forms of human–machine interaction is the necessity for timely communication and control. Although there are several existing

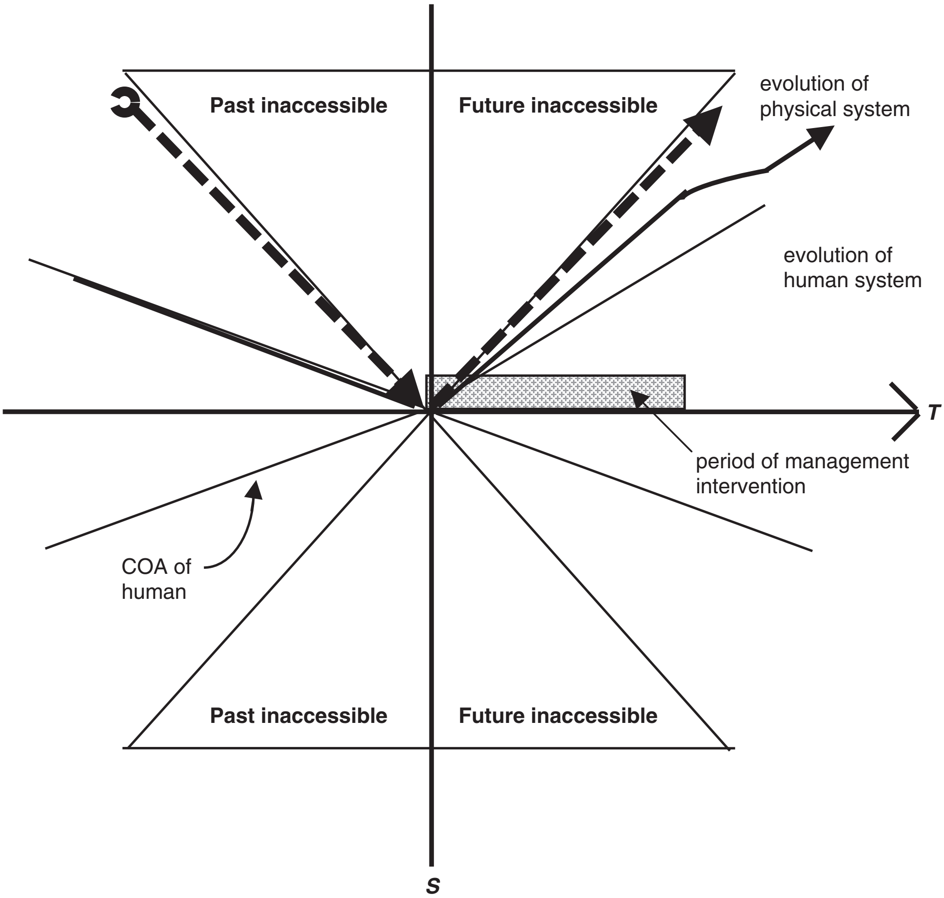

Figure 15. Effect of management facilitatory intervention. The $\mathbf{f}^{\mathbf{human}}$ is increased, so that the rate of evolution is faster than normal. After the management intervention ceases, the rate returns to a line parallel with the original unfacilitated rate, but some overall gain in response speed remains.

human–machine interaction models, they do not provide a full generalisable architecture of context in which to frame behaviour. Such a framework, which illuminates a number of traditional issues, is presented here. Hindsight makes it frequently surprising that an accident was not foreseen: given the state of the system, surely operators could see what was about to happen and surely they should have had time to intervene? It is known that in hierarchical systems with complex coupling between and within levels, communication and control are difficult. It is therefore desirable to have a general model of these processes and the phenomena associated with them when information is used for action in real time and in the face of system dynamics.

The use of Minkowski spaces offers a new approach to this problem of representation. Minkowski spaces allow one to see, in a general way, why certain information cannot reach the operator and why certain actions cannot succeed in time. They suggest models for interference and goals for the design of better interfaces and decision aids. A fundamental property is that the differences between the slopes of cones and the regions between the cones represent margins of safety or areas of instability and loss of control, depending on their relative slopes. More complex versions of the diagrams can therefore show the respective properties of feedback and feedback delays, team interaction and intervention by higher levels of management, regulators, etc. If metrics can be developed which reflect the spatio-temporal characteristics in detail of human–machine systems, they may have even more to offer, interpreting space as the multi-dimensional conceptual space that is required fully to describe the properties of humans, machines and their interaction.

References

Bouwmeester, D., *et al*., 1997. Experimental quantum teleportation. *Nature*, 390, 575–579.

Card, S.K., Moran, T.P., and Newell, A., 1983. *The psychology of human–computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Cellier, J.-M., DeKeyser, V., and Valot, C., 1996. *La gestion du temps dans les environments dynamiques*. Paris: Presses Universitaires de France.

Conant, R.C., 1976. Laws of information that govern systems. *IEEE Transactions on Systems, Science and Cybernetics*, SMC-6, 240–255.

De Cortis, F., 1988. Dimension temporelle de l'activité cognitive lors des démarrages de systèmes complexes. *Le Travail Human*, 51, 1215–1238.

Hancock, P.A., 1996. Effect of control order, augmented feedback, input device and practice on tracking performance and perceived workload. *Ergonomics*, 39, 1146–1162.

Hancock, P.A., 2007. On the nature of time in conceptual and computational nervous systems. *Kronoscope*, 7, 1–12.

Hancock, P.A., 2008. The battle for time in the brain, *In*: J.A. Parker, P.A. Harris and C. Steineck, eds. *Time, limits and constraints: The study of time XIII*. Brill, under review.

Hancock, P.A. and Chignell, M.H., 1995. On human factors, *In*: J. Flach, P.A. Hancock, J. Caird and K. Vicente, eds. *Global perspectives on the ecology of human-machine systems*. Vol. 1, Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Hoc, J-M., 1996. *Supervision et controle de processus: la cognition en situation dynamique*. Grenoble: Presses Universitaires de Grenoble.

Minkowski, H., 1908/1923. Space and time. *In*: H.A. Lorentz, A. Einstein, H. Minkowski and H. Weyl, eds. *The principles of relativity*. London: Dover.

Perrow, C.A., 1984. *Normal accidents: Living with high risk technologies*. New York: Basic Books.

Pew, R.W., 2007. Some history of human performance modeling. *In*: W. Gray, ed, *Integrated models of cognitive systems*. Oxford University Press.

Rasmussen, J., 1986. *Information processing and human-machine interaction: an approach to cognitive engineering*. Amsterdam: North-Holland.

Rasmussen, J., Pederesen, A.-M., and Goodstein, L., 1995. *Cognitive engineering: concepts and applications*. New York: Wiley.

Reason, J., 1997. *Managing risks of organizational accidents*. Ashgate: Gower Press.

Sheridan, T.B., 1992. *Telerobotics, automation, and human supervisory control*. Cambridge, MA: MIT Press.

Sheridan, T.B. and Ferrell, W.R., 1974. *Man-machine systems*. Cambridge: MIT Press.

Wiener, N., 1961. *Cybernetics*. 2nd ed. New York: Wiley.

About the authors

Neville Moray received his Bachelor and Doctoral degrees in psychology from the University of Oxford. He is a Fellow of the Human Factors and Ergonomics Society, a Fellow of the Ergonomics Society, a Fellow of the International Ergonomics Association, and a Certified Professional Ergonomist. He has received the President's Award of the International Ergonomics Association for services to teaching and research in ergonomics, and the Arnold M. Small President's Distinguished Service Award of the Human Factors and Ergonomics Society for lifelong contributions to human factors. He has taught in universities in the United Kingdom, Canada, the United States, France, and the People's Republic of China, and has held research posts in Italy and Japan. He served as the UK representative on the human factors committee of the NATO Science Committee, and has been a member of the US National Research Council's Committee on Human Factors. He has consulted in nuclear safety, chemical weapons destruction, ship design, control room enhancement, human error and safety, and surface transportation. He is currently Emeritus Professor of Psychology at the University of Surrey in the United Kingdom. He lives in the south of France, where he is a professional artist.

Peter Hancock is Provost Distinguished Research Professor in the Department of Psychology and the Institute for Simulation and Training at the University of Central Florida. He is the author of over 500 refereed scientific articles and publications as well as 15 books including the recent edited text *Performance under stress* with his colleague Jim Szalma, published in 2008, as well as the forthcoming text *Mind, machine and morality*, each published by Ashgate. He is a Fellow of HFES, IEA, APA, APS, and the Ergonomic Society of Great Britain, from which he received the 2008 Otto Edholm Award. His current experimental work concerns the evaluation of behavioural responses to high-stress conditions. His theoretical work concerns cultural influences on perception and action, investigations into the nature of time as well as human relations with technology and the possible futures of this latter symbiosis. He is a past President of the Human Factors and Ergonomics Society and collects and studies antique maps while being a committed Ricardian.

Moray, N.P., & Hancock, P.A. (2009). Minkowski spaces as models of human-machine communication. Theoretical Issues in Ergonomic Science, 10 (4), 315-334.