CHAPTER 5

Sociotechnical Systems

In Chapter 1, we argued that technical artefacts such as aeroplanes, electric drills, computers and ballpoint pens differ from both physical objects and social objects in that they embrace something of both. Technical artefacts are tangible objects with physical properties, but they are also objects with a function, which they have in virtue of their embeddedness in use plans aimed at the achievement of human purposes. In Chapter 2, we examined the way in which technical artefacts come into being, by way of a design process, and in Chapter 4, we considered the knowledge that this requires. In this chapter, we show that to view technology as merely a 'collection' of technical artefacts would be an immense oversimplification. In so doing, we would completely fail to acknowledge the layeredness that is such an important feature of modern technology: the technical artefacts discussed so far are building blocks in wholes of a far greater complexity. Although one cannot build 'loose' technical artefacts that are as big as the earth itself – where would one be able to assemble them? – the wholes or *systems* central to the present chapter do, in fact, span the entire globe. We shall see that as a result of the character of these sorts of systems, the principles traditionally adhered to by engineers when designing technical artefacts, and the kind of knowledge upon which they rely in the process, cannot remain unaltered when it comes to the matter of the designing and implementing such complex things. As has already been noted in Chapter 3, this also has consequences for the ethical dimension of technology.

5.1 HYBRID SYSTEMS

Imagine that you are walking around a major airport because you are about to travel to a distant destination. You will make use of numerous artefacts. Some of these, like the baggage trolleys and the benches in the cafés and waiting areas, will be simple artefacts, probably designed by just one or a few persons. Others will be incredibly complicated, designed by whole teams of engineers who will probably have worked on them for years, for example, the computers used for the checking-in system where the details of all the flights are stored and, of course, like the showpiece of engineering ingenuity, the actual aeroplane itself in which you will ultimately fly. However, alongside all these separate things you will also be making use of something much more encompassing, just as much created and maintained by human hands as the prototypical artefacts just mentioned, but at the same time, it is much more impalpable and harder to fathom. That 'thing' is the *world civil aviation system*. All the artefacts just mentioned constitute a part of it, but there is incredibly much more to it than that. Some of the components are concrete, like the buildings in which the passengers are subjected to a whole range of routine procedures before entering the plane or after stepping out of it, and the runways used for take-off and landing. Additionally all kinds of equally tangible people are involved: the cabin crew, of course, but also the personnel who work at the check-in desks, operate the X-ray equipment, check the passports and load the baggage onto and off the plane. But then, there are also, in very divergent ways, numerous other components of a more abstract nature: the air corridors within which the aeroplanes fly, the airline companies responsible for the flights, the regulations that pilots and airline companies have to observe, the organisations that draw up and enforce these regulations, the treaties agreed to between countries that make it possible for planes to fly from one airspace zone to another, the companies that insure the system for the different ways in which it can fail, and so forth, and so forth. Each of these things contributes to the functioning of the world civil aviation system. Without this system, it would be impossible for you to travel by plane to your chosen destination and afterwards to fly back home again. Without this system, you would undoubtedly not find at your final destination the places justifying the trip in the first place, like a holiday hotel, a conference resort or a business centre. Without each of the listed components, the system would not be able to function in the way it does at present, certainly not on the scale and with the level of efficiency that most of us completely take for granted.

The world civil aviation system is an example of a *sociotechnical system*. The fact that we refer here to a system will not come as a surprise to most. A system is understood to be an entity that can be separated into parts, which are all simultaneously linked to each other in a specific way. An aeroplane is itself a system. All of its components, however, are 'hard' things, the behaviour of which is governed by various natural laws. A thorough knowledge of these laws is required in order to comprehend how an aeroplane operates from the way in which its components work together and to design and connect the components in such a way that the resulting plane does precisely what is expected of it. Admittedly, in present-day aeroplanes, there are all kinds of control systems that are computer-steered, running on software systems that were designed in isolation of the computer and, in that respect, definitely not tangible. However, once downloaded, this software ensures that the onboard computers are set up in a specific physical state, which is essentially no different from the way in which a thermostat is adjusted to a specific setting for it to maintain a certain temperature.

The presence of software thus changes nothing whatsoever about the character of the aeroplane as a physical entity. What makes the aviation system a special sort of system is the fact that it includes all kinds of components – the organisations and institutions, the conditions and rules – that are not tangible things. For all of these components, no thorough knowledge of the natural sciences will do to understand how they work and to fit them into the system in an effective way. The relevant employees mentioned, in their capacities as human beings made of flesh and blood, are tangible, but that is hardly relevant to the position that they occupy within the system. One does not need to possess knowledge of biology to fit the staff working at the check-in desks into the system as a whole. The most basic everyday knowledge is sufficient to do justice to their character as physical-biological organisms that cannot walk on air or pass through walls, so that they need accessories like doors, seats and floors to be able to carry out their work. Which is not to deny that occasionally, specialised knowledge about the biological side of humans is relevant, for instance, when designing cockpits where a pilot's attention has to be divided between a wide range of instruments, or – to briefly step outside the framework of civil aviation – when designing fighter planes where the pilots are exposed to extreme acceleration speeds.

In the case of people employed within the civil aviation system, matters such as their height, weight, gender or stamina are not the issue at stake; what counts is the fact that they are *persons*, capable of understanding and carrying out instructions and also of understanding the purposes served by these instructions. The regulations and the organisations and institutions that contribute to the aviation system presuppose the status of human beings as persons. Rules and regulations are drawn up by people and can be observed or flouted by people. Organisations and institutions are created and maintained by people. In order to assess the way in which people function within the civil aviation system, we largely rely on a general picture of the way in which people go about doing the things they do in everyday life, supplemented with knowledge obtained from the fields of sociology and psychology. None of the natural sciences, however, has anything useful to say on the matter. We cannot find either rules or organisations anywhere in the mineral, plant or animal world.

This, thus, brings us to the heart of the matter of what makes the world civil aviation system such a special sort of system: it is a *hybrid* system. It consists of components which, as far as their scientific description goes, belong in very many different 'worlds'. This is what makes them essentially different from even the most complex of technical systems, like for instance civil aircraft. Even though the engineers who were involved in the designing and the manufacturing of the Airbus A380 had very different backgrounds – mechanical engineering, materials science, aerodynamics, electronic engineering and computer engineering – all these disciplines share a form of describing the world rooted in natural science. However, the aviation system into which such an Airbus A380 operates involves numerous other things – people, institutions, rules – about which the natural-scientific way of describing the world has little to say and for which a social-scientific way of formulating matters is therefore required.[26] Hybrid systems, in which certain components, are described and researched using the natural sciences and other components, are described by drawing on the social sciences are called *sociotechnical systems*.

This hybrid character of sociotechnical systems needs to be distinguished from the dual nature of the technical artefacts that was central to the discussion in Chapter 1. This dual nature even applies to the most simple of technical artefacts, such as a screwdriver or a nutcracker: they are all objects that, apart from their physical properties, also have their particular function and their embeddedness in a context of human use plans as non-physical features. Sociotechnical systems, being technical artefacts with an extremely high degree of complexity, are, in principle, also subjected to this dual nature. One can view a sociotechnical system as a particular 'thing' having certain causal properties, and one can examine the function that it has in a context of human actions. By contrast, the hybrid nature of a sociotechnical system has to do with the *composition* of the 'thing' it is, as a result of

---

[26] Unfortunately the realm of human and social sciences is much less well-organised than the realm of the natural sciences. We use here the term social sciences as the most general term for a set of disciplines containing the humanities as well as the social, economic and cultural sciences and also parts of psychology such as cognitive, social and organisational psychology. We shall not deal here with the various interrelations between these specified disciplines.

which it can no longer be unambiguously seen as a single tangible thing to be picked up, in a manner of speaking, and held up to the light for further inspection. This hybrid character makes the examination of its causal aspects a much more problematic matter than in the case of traditional technical artefacts, with all the consequences that this entails for the designers of such systems. It also gives the dual nature of technical artefacts a much more complex character, however. On the one hand, the social aspect, the context of human action, is manifested in every system aspect of the total system within which people are involved. On the other hand, it is for sociotechnical systems, in particular when they reach country or continent scale, hard to identify the function of the system as a whole, as it is simultaneously embedded, at any one moment, in the context of action of numerous different individuals.

Even though it is eminently this hybrid character – the presence of components requiring a physical description and components requiring a social description – that characterises sociotechnical systems, the designing, implementing and maintaining of these systems remains predominantly in the hands of engineers, who have been educated in pronouncedly natural-scientific ways. That is why these systems constitute a major challenge for the engineering sciences. All kinds of traditional notions about what constitutes the designing of a technical artefact, how the design process should be structured, what kind of knowledge is required and how one should assess the functioning of a designed artefact, become very problematic whenever they are literally transplanted to the context of designing sociotechnical systems. The reasons for this are presented in the following section.

5.2 SYSTEM ROLES FOR PEOPLE: USER AND OPERATOR

The special character of sociotechnical systems is not grounded simply in the interaction between man and machine. Virtually all technical artefacts have an on/off knob for the purpose of starting or stopping their functioning, and on top of that all kinds of knobs, switches and handles allowing the user to adapt their functionality. Technical artefacts are, after all, manufactured to be used and this presupposes that there is a *user* who makes use of the artefact by manipulating it. That does not mean that using an artefact necessarily requires continuous manipulation. Sometimes the manipulative aspect only extends to the installing of the artefact, inserting it into some network of causal connections, like when a memory chip is installed, a smoke detector is connected or a communication satellite is launched. In the artefact's use plan discussed in Chapters 1, 2 and 4, the kind of manipulation suitable for the artefact in question is specified. But without some form of physical interaction, we can have no use. What makes sociotechnical systems special is, first of all, that they have many users at any one moment and, secondly, that they involve people in two different ways, namely, not only in the role of user of the system but also in the role of operator. The word 'role' is apt because one and the same person can simultaneously be a user of and an operator in a system. If a pilot plans to spend some time at his or her destination for a holiday, after having flown the plane there, he or she is, on the one hand, performing the role of operator by flying the plane and, on the other hand, he or she is performing the role of user by using the aviation system to get to his or her holiday destination.

The role of operator may well seem to appear out of the blue. Does the pilot, in his or her role as flyer of the aeroplane, really do something different from what the owner of a coffee-maker does when turning the machine on to make a cup of coffee? There is, indeed, an important difference: the user of the coffee-maker operates the machine to realise a goal of his or her own that would be much harder or even impossible to realise without the device. The pilot of an aeroplane that is on its way to Singapore does, we presume, have instructions to arrive there within a reasonable period of time, but he or she is not doing so in order to realise his or her own goal of arriving in Singapore within a certain amount of time. It is the passengers who have this as a goal and by taking that plane they realise that goal. To do so they not only use the aeroplane but also the entire aviation system, which they 'operate' by purchasing a ticket. The pilot of the plane is a *component* of that system, a component that is necessary for the whole system to be able to fulfil its function, just as much as the plane that the pilot flies is. After all only few of us have at their disposal their own private air transport to take them all the way to Singapore, and the operation of civil aviation aircraft is too complex to be left to the travellers themselves. It is also (as yet) too complicated and too expensive to manufacture an aeroplane that is capable of operating completely automatically, taking off without any human intervention after the last passenger has boarded and closed the door behind him- or herself, and then landing all the passengers safely in Singapore some ten hours later.

Pilots may be the most obvious indispensable human link in the civil aviation system, but they are certainly not the only ones. Just as indispensable are the air traffic controllers who sit in the control rooms at airports and supervise the taking off, landing, cruising and taxiing movements. The great speed and limited manoeuvrability of civil aircraft, combined with the limited vision of pilots would make all air traffic involving more than one plane in the air at any given time virtually impossible if there were no air traffic controllers to maintain radio contact with the pilots and to monitor, by means of their radar equipment, the position of all the aeroplanes. In much the same way, all the staff members mentioned above – the staff working at the check-in desks, those who operate the X-ray equipment, the passport controllers and the ones who load the baggage in and out of the holds – are components of the aviation system, each contributing a specific function to the operation of the entire system.

All sociotechnical systems have operators who fulfil such roles, because it is too difficult or even impossible to build a system consisting merely of interconnected technical artefacts – that is, machines or 'hardware' devices – and guarantee its adequate functioning. Every large chemical plant or power station has control rooms that are permanently occupied by one or more operators or controllers. Even in road transport systems – which, even though forming another sort of transport system, are very different from aviation systems, as we shall see – operators are coming to play an ever bigger part, with the increasing volume of traffic on roads, by monitoring the flow of traffic and attempting to control it, through the imposition of speed restrictions, the opening or closing of traffic lanes and the provision of specific information.

Human operators started to receive serious design interest with the rise of *systems engineering* during the first two decades after the Second World War. At the instigation of many military organisations engineers were increasingly becoming involved in the development of complex systems consisting of numerous components of a divergent nature in which knowledge drawn from very different disciplines was processed. The main challenge facing those responsible for the overall design process was to coordinate the behaviour of the separate components. Since the development of the computer was very much in its early phases, the only 'mechanisms' available that were capable of realising such complex coordinating activities were human beings. Thanks, though, to the rapid pace at which the computer has developed since, human beings are becoming increasingly redundant. The control rooms of factories and plants have thus become gradually emptier and within the foreseeable future the fully automated flying of aircraft will be a reality.[27]

This should not, however, be taken to mean that large-scale systems involving human operators will be merely a temporary phenomenon. The complexity of sociotechnical systems, particularly large infrastructural systems, consists just as much in the fact that, unlike a typical technical artefact such as a coffee-maker, they have very many different users. The functioning of the system as a whole, as it appears to each of its users, not only requires coordination between the technical or *hardware* aspects of the system and the behaviour of the users – like the driving of the vehicles by their chauffeurs – but also, and especially, the mutual coordination of the behaviour of the many users. To achieve this coordination, it will not to do to implement one or another causal mechanism that is attuned to the physical characteristics of the users. Successful coordination comes about through agreements, rules, laws, habits, in short, precisely the sort of things that are studied in the social sciences and not in the natural sciences.

5.3 RULES AND COORDINATION MECHANISMS

It is in the notions underlying agreements, rules, laws, and so on, that one becomes most sharply aware of how it is that sociotechnical systems differ from traditional technological systems. If you want to 'direct' people, it is common to do this through *rules* or *instructions* and not through causal stimuli and signals.[28] A *rule* is a directive or norm that has the underlying intention of bringing about a behavioural pattern, irrespective of whether that pattern actually occurs. A rule can be observed or ignored, just adhered to from time to time or abided by depending on the circumstances. If a rule is not followed in a particular case, this indicates that it is apparently not in everyone's best interests to behave in the manner that the rule dictates.[29] Such deviations will not necessarily undermine the intended behavioural pattern, but they are likely to do so, especially if many people are tempted to breach a rule because acting according to the rule requires an effort or is costly in one form or another. What is especially characteristic of a rule is, therefore, the existence of *sanctions* relating to the breaching of the rule. As a result, rules also presuppose the existence of a social group in

---

[27] See, for instance, Weyer, J. [2006].

[28] To be sure, direction by way of causal stimuli and signals is occasionally applied, and successfully, but it requires precisely engineered, stable circumstances to be effective.

[29] We distinguish, therefore, between a rule and a *convention*, where a convention is an actually occurring behavioural pattern, which has come about through mutual coordination of behaviour among different people and which continues to exist, even without being prescribed, for the simple reason that it is in everyone's interests to abide by that convention.

which the rules are considered to apply and through which sanctions can be enforced. After all, an individual cannot be expected to enforce sanctions upon himself if a rule is deviated from.

In order to allow a sociotechnical system to function properly, therefore, rules have to be thought up or drawn up and imposed as a coordination mechanism. Consequently, these systems take on a kind of complexity that is lacking with technical systems. The role of an operator within a sociotechnical system is defined by a set of rules or *instructions* that determine in what situation the operator must take what particular measures. These instructions constitute, as it were, a – more or less forcefully prescribed – use plan for those components of the system that the operator has under his or her control. In addition to this, there are use plans for the users of a sociotechnical system that tell them how they may and must use the system and how they have to behave when using the system if it is to retain its functionality for others. Just how the system as a whole will function depends crucially on the way in which these rules are formulated. On the one hand, the actions that an operator or user carries out, in accordance with the rules, must indeed give the results corresponding to the functionality of the system, given the actions that these same rules impose upon other users. On the other hand, the rules must be of such a nature that it may reasonably be expected that operators and users will actually carry out the actions that have been stipulated in them. In other words, once laid down rules must also be followed. This latter aspect is irrelevant in the case of purely hardware systems.

The physical behaviours of the *hardware* components of a system follow law-like regularities that can be researched and established by means of experiments and tests. As long as there is uncertainty concerning the physical behaviour of a component no engineer will be prepared to include it in the design. To be sure, it is inevitable a matter of judgment whether or not there is sufficient knowledge available about the behaviour of a particular component or material, and unexpected things leading to the failure of the system can never be completely excluded. Still, when designing technical systems, an engineer can rely on the existence of natural laws at various levels. Engineering design is completely based upon verifiable reliability. This extends to technical systems whose functionality is supported by software: once loaded with a certain program, a computer system will go through a sequence of states in an entirely predictable way, and provided that the software is free of errors, these states will be precisely the states that a component must have, according to the design, to make its contribution to the realisation of the functionality of the entire system.

At first sight, the situation that applies to sociotechnical systems does not appear to be significantly different. You could 'calculate' what behavioural dispositions must be given to the people within the system if you want the system, including its human components, to function as intended. These behavioural dispositions could subsequently be inculcated in the relevant people, either through training or purely by instruction. This could be seen as a process of 'loading' a certain 'program' into a person. If we could assume that an individual will behave exactly in conformity with a set of behavioural rules, once 'loaded', then the design problem will be restricted to the 'designing' of these dispositions or instructions, which is not very different from the process of designing software for the components of a purely technical system. The problem, however, is that this situation, in which people will unfailingly execute specific instructions and completely and meticulously follow instructions, seldom or never occurs, and can hardly be expected to occur, in view of the constitution of the people required to execute or follow the directives. Every individual is, as it were, a computer on which a great many programs are running all the time, programs about which the designers of the sociotechnical system in question have only a very general and limited idea. Rather than loading operators with a single program that secures the precise execution of the operator's role as the one and only thing that the operator cares about, designers can do no more than add to the vast amount of already installed software a mere handful of subroutines.

5.4 SYSTEM DESIGNS AND SYSTEM BOUNDARIES

With this understanding, we come to the problems that confront those who are responsible for designing, implementing and maintaining sociotechnical systems, and we come to the challenges they pose for the way in which engineers are traditionally taught to tackle the designing of artefacts and systems. These problems will be dealt with from two angles. First, there is the problem of how to draw the system boundaries and, accordingly, to establish the extent of the design task. In the second place, there is the problem of the predictability of the system's behaviour and the extent to which this can be controlled.

The most important thing that can be said about the boundaries of a system is that these are not given beforehand but have to be decided upon on the basis of various considerations. Or, to put it more precisely, it is only once a boundary has been drawn that it is clear exactly which system is the subject of research and design. A system is, after all, a collection of interrelated components, each of which can be broken down into smaller components, which makes these components, in their turn, also systems. Furthermore, every system that we define by drawing a boundary is connected beyond that boundary with other things so that we can view the system under investigation and these other things as components of a still wider system. In the case of natural systems, this series of embeddings spans the entire spectrum from the most elementary particles at one end to the entire universe at the other end. For hybrid systems of the type discussed here, however, the spectrum is obviously much more confined. The smallest component of a hybrid system is one single person or a single technical component while the largest system that could be described as an entire sociotechnical system must remain, for the time being, confined to the earth.[30]

As far as the designing of systems is concerned, the important question is not what boundaries can be drawn in an existing reality, but what place the new system is to occupy in the existing reality? As long as this is unclear, the design task with respect to the system is still undetermined. Imagine that, as an engineer, you are requested to design a new type of aeroplane engine. In order to be able to do that, you not only need to know of which technical system this engine is to be a component – for what type of plane it is required – but also of which (socio)technical system the plane will be a

---

[30] The position of a sociotechnical system in terms of time and space is not a foregone conclusion. There are, for instance, all kinds of spaceships dotted around the solar system, spaceships that are linked to control centres on earth and which, in most cases, can also to an extent be controlled by those centres. Does this mean that our largest sociotechnical system extends to our entire solar system? It is an intriguing question, but we shall not endeavour to address it.

component. If the aircraft is merely destined to fly in the private airspace above the land of a large estate owner (let us presume that something like that is possible), then the functional requirements can be limited to the exclusively technical requirements: power, thrust, weight, and so forth. If, on the other hand, the aeroplane in question is destined to fly as a component of the existing world aviation system, then an accordance with all kinds of legally stipulated norms and standards will also form part of the functional requirements for the engine. An aeroplane that does not meet these requirements cannot be integrated into the system, just as little as a bolt will fit into a nut with an incompatible screw thread. These norms and standards are preconditions to which you, as a designer, must conform in much the same way that designers, in general, have to accept the laws of nature, which they cannot change.

Matters become slightly more complicated when the issue is the design of an entire aeroplane. If what is intended here is a private plane for the same large estate owner, which will exclusively be flown in his private airspace, then again the designing is merely a technical artefact. Naturally you will provide the owner with a manual with instructions for use, or you might even offer him a training course, but, in essence, it will be no different from the instructions for use you would provide for a coffee-maker. After delivery, the owner is then free to vary matters as he wishes, to develop alternative operating methods or to discover by himself how the aircraft can be used for all kinds of stunts. In the case of an aeroplane that is destined to become a component in the world civil aviation system, matters are again different: it now even becomes necessary to decide whether the design task will be confined to just the technical artefact 'aeroplane' or whether it should be extended to the sociotechnical system 'aeroplane plus flight crew'. In the first approach, the responsibility for the drawing up of exhaustive and adequate operational instructions and for the training of the pilots lies with a different 'designer' and, as aeroplane constructor, you are only concerned with external (safety) norms concerning the layout of the cockpit, the instruments and the operation panels. In the second approach the drawing up ('designing') of the rules that define the pilot's role and the coordination between the way the aeroplane itself is designed and the way in which it is operated also form part of your design task. Evidently, the second design approach could produce improvements with respect to safety, which will not be recognised, or not so easily, in the first approach. Which of the two approaches is adopted is in no way predetermined; it remains a question of choice. It is a choice that does not generally lie with the designer but rather with the owners or managers of the (smallest) system within which the aeroplane under design will function as a component.

These examples hardly touch on the complexity of the civil aviation system, which contains a large number of subsystems, each with its own operators. For all those subsystems, separate sets of rules have to be designed. These rules then have to be attuned to each other – especially those that define the role of the pilot and the rules that define the role of air traffic controller – but what must be perpetually borne in mind is the fact that the civil aviation system is embedded in at least one larger system – that of the network of sovereign states. The rules that bind people in their various roles within the aviation system must therefore also accord with the rules that national and international legislation imposes upon individuals. It is a kind of complexity that notably comes to the fore at moments when something goes wrong, which is why we shall now examine more closely a tragic aviation accident that occurred on July 1st, 2002, in the airspace above southern Germany. At an altitude of ten kilometres, a Tupolev 154 from Bashkirian Airlines crashed with a Boeing 757 from the freight carrier DHL. Such an accident always has a whole chain of causes and is, in that respect, the outcome of an unfortunate sequence of events, but from the design perspective, this particular accident demonstrates that also, at the highest levels of complexity, we have to bear in mind the system character of the technology we jointly create.

5.5 THE MID-AIR COLLISION ABOVE ÜBERLINGEN[31]

In 2002, in the wake of previous mid-air collision incidents, aircraft making use of European airspace were all equipped with a TCAS, *traffic collision avoidance system*. In the nose of the cockpit, there is an instrument that sends out a signal but can, at the same time, pick up signals sent out by other aircraft. When the received signal, in combination with the plane's own position and cruising speed, indicates that unless one or both aeroplanes change course, they will collide, the TCAS equipment transmits coordinated instructions to the pilots: one of the crews is given a spoken instruction to descend and the other is instructed to ascend. The TCAS is intended as a last resort in an emergency: it is the task of the relevant air traffic controllers to notice, at a much earlier stage, that two aeroplanes are flying at the same altitude, on courses that will lead to disaster and to rectify the situation by directing one of the two to a different altitude. The TCAS was introduced for situations where air traffic controllers fail to do so. This was indeed the case in the Überlingen incident, but there is more to the story. The air traffic controller on duty in the area where the aeroplanes were flying had indeed failed to notice the impending accident in time and had thus not intervened when he should have. Eventually, however, he did notice the problem and intercepted by instructing the Russian aeroplane to reduce its altitude. But the air traffic controller's instructions came so late that by then, the TCAS on board both aeroplanes had been activated: on the basis of the signals that had been exchanged, the software had generated instructions to the effect that the captain of the American plane should descend while the Russian aeroplane had been instructed to ascend. There was just one second's difference between the Russian captain receiving the message generated by the TCAS to ascend and being instructed by the controller on the ground to descend. Of course, this led to great confusion and debate among the Russian pilots, but there was very little time available for finding out what to do, and after air traffic controller and TCAS had repeated their conflicting instructions, the captain of the Tupolev 154 decided to follow the instructions from the ground and not the message generated by the TCAS. As the Boeing 757 had only been instructed to descend by its own TCAS, this aircraft also started to drop altitude, with the result that shortly afterwards, the two planes collided, causing the death of everyone on board both aeroplanes.

---

[31] For this section use has been made of the official investigation report published by the German Bundesstelle für Flugunfalluntersuchung [2004] and Weyer's book on the subject [Weyer, J., 2006]. Another publication that discusses the role of the TCAS in this particular case, and in general, but which does not view matters from a sociotechnical perspective, is that of Ladkin, P. [2004].

What this accident makes painfully clear is that the TCAS is designed for a closed system involving just two aeroplanes and their crew, but not for a wider system that also includes air traffic controllers. It was designed for situations in which, for one reason or another, air traffic control has dropped out and is no longer involved, but it is implemented nevertheless as a component in a system where air traffic controllers are also present as system components. In the instructions defining the roles of flight crews and air traffic controllers, allowances should have been made for the fact that a crew could receive instructions both from the TCAS and from air traffic control, a possibility that is evident if matters are considered from the perspective of the aviation system as a whole. The existing aviation system regulations, however, provided no answer to the burning question with which the Russian captain had briefly grappled, namely that of which of the two instructions to follow. The whole question ultimately tied up with the embedding of the civil aviation system in the social system of national and international legislation and regulation because during the investigation into the disaster, it emerged that a pilot could refer, in such cases, to no less than five different documents, the exact status of which remained vague and which were not, moreover, in unison with each other. Since this tragic accident, the instructions within the aviation system have been amended by emphatically stipulating that whenever a flight crew receives contradictory instructions from the on-board TCAS and the traffic controllers on the ground, they must ignore the instructions from the ground and follow those issued by the TCAS. But this ruling still appears to ignore the embedding of the world aviation system in the overall social system. In situations where, for instance, three planes are flying in close proximity in the same airspace the TCAS might well fail, either because one of the three planes has no TCAS due to its being, for example, a private or a military plane, or because the software algorithm of the TCAS is not correct, for it has been proved correct only for two-plane situations. Aided by their radar equipment, the air traffic controllers would have the power to correctly direct matters in such situations, but the new ruling obliges captains to abide by the TCAS, even in cases where the captain has good reason to doubt the instructions generated by the TCAS on the basis of his or her own observations. This new ruling therefore, it is asserted, contravenes international legislation, which lays down that the pilot is at all times fully responsible for the safety of the passengers and staff in the aeroplane that he or she flies.

Apparently, then, it is extremely difficult to guarantee coordination between system components up to the highest levels of complexity. What severely aggravates this problem is that such large-scale sociotechnical systems, in fact, have no designers. Rather they evolve through historical processes of spontaneous and directed linkages of subsystems that are integrally designed by engineers or, more to the point, by teams of engineers. If such subsystems become connected, the links between them are adjusted or possibly even replaced by new links, but there is no single organisation which, from a design aspect, contemplates the entire system and investigates whether the functionality of the system is guaranteed by the way in which the components engage. Even a technical system that is developed on a world-wide scale like, for instance, the American GPS system and its intended European counterpart Galileo, is quickly integrated with other existing and newly created systems once it has been implemented, so that yet other systems with new functionalities can emerge for purposes of, in this particular case, telecommunication, navigation and cadastral registration. But also the embedding in the relevant national and international umbrella system is not appraised from a total all-round perspective – simply because there is no authority that represents such a perspective – but rather from the angle of subsystems each with their associated vested interests.

5.6 SYSTEM DESIGN AND CONTROLLABILITY

This brings us to the second problem that sociotechnical systems pose for the traditional engineering approach to design, namely a loss of predictability and control. When developing technical artefacts, the external circumstances within which the system has to fulfil its function are explicitly included in the requirements. Traditionally, it is the task of the designer to produce an artefact that functions as long as the circumstances obtain as specified. Whether the circumstances within which the system is used or implemented, in fact, meet this specification is not so much the designer's problem but rather that of the customer or user, even though, as a designer, one should make certain that the functional requirements taken to define the design task are a correct translation of what the client or commissioner has in mind for the system. With sociotechnical systems, this is completely unattainable. The overarching social system within which every sociotechnical system functions as a component is in a perpetual state of flux. Even if we imagine that a particular large-scale sociotechnical system – say, the world civil aviation system – is developed all in one go, it would still be impossible to precisely specify the institutional context within which that system has to function. This problem that can hardly be resolved by broadening the definition of the design problem and expanding the system borders by involving the institutional context. There is, ultimately, an overarching system, which is the system of sovereign states, but this is a social system, not a sociotechnical system, and at that level, one cannot speak of designing in the sense of engineering designing. The institutional context of national and international legislation and regulation has come into being and functions in a completely different way from technical artefacts. This is even unavoidable because society as a whole cannot be abandoned in favour of a new design; designing is an activity that occurs within the context of society. In that respect, society is like a ship that has to be repaired while sailing on the high seas but kept afloat in the meantime, to use a metaphor introduced by the philosopher Otto Neurath.

The source of the uncertainty with regard to the functioning of a sociotechnical system like the one being considered here is shared by all such systems, though not always to the same degree. The problem was felt acutely during the past two decades with respect to the energy infrastructures, when deregulation turned the institutional context of the existing infrastructure upside down whilst the public expected of the managers of these infrastructures that the functionality of the system be preserved. This situation created major challenges for the operators because it forced them to alter their view of the nature of the system that they are managing and to create new conceptual frameworks and models. There is, however, another source of uncertainty, which hardly comes into the picture in the aviation sector but is very evident in, for instance, the road transport system. In the air transport system, the end-users or passengers are hardly capable of influencing the behaviour of the system. For the behaviour of the aeroplane and for the way in which the pilot and the air traffic controllers go about their work, it does not really make much difference if their aeroplane is loaded with passengers or with freight. It is a system which, in its day-to-day operations, lies close to the engineer's ideal of a completely controllable and predictable system, into the running of which the human operators, through their punctual execution of seamlessly coordinated instructions, smoothly fit.

The same definitely cannot be said of the road transport system. The individual drivers of all the different cars and lorries and the motorbike riders each fulfil a double role: they are users of the system, but they also operate a small part of it. By using the system, a car driver also simultaneously modifies it by introducing a temporary technical component in the form of a car and herself as an operator of that same car. These operators are tied to much less stringent rules than airline pilots and have much greater freedom. The rules that define the driver's role consist, to a large degree, of instructions to react in certain ways to certain signals such as traffic lights, road signs, road markings and the lay of the road. By automatically (e.g., traffic lights) or occasionally (e.g., incidental speed limits and lane control on motorways) altering these signals, the global system controllers attempt to adjust the system in such a way that the constant changes in the configuration of the system are compensated and its functionality maintained. To what extent their efforts are actually successful is a matter of dispute. The fact remains that because of the very nature of the system, the global system controllers must always be uncertain about the degree to which their endeavours will succeed. Sociotechnical systems have unavoidable *emergent properties,* that is to say, properties that admittedly emanate from the properties of the components and from the way the system is structured but which are not predictable, for the simple reason that in order to predict them, one would need to have access to knowledge which, at least in practical terms, is unavailable or, if in principle available, cannot be accessed in the available time.

There are many different schools of thought on what is the best way to design and manage these kinds of dynamic sociotechnical systems. On the one hand, there is the strong tendency to try to force the system in the direction of the aviation system by restricting the role of the individual user through various sorts of automatic vehicle control systems. In that way, the system would become more controllable in line with traditional engineering norms. On the other hand, there are also small-scale experiments where the task of successfully coordinating all vehicle movements is laid entirely with the road users by deliberately eradicating all the instruments that are customarily used to direct their behaviour, such as road signs, give-way road-marking, traffic lights and all the other paraphernalia. In cases where the global control possibilities of a sociotechnical system are fundamentally limited because for principled or for practical reasons the users of that system are granted a high degree of freedom to involve themselves in the system, it may well be advisable for authorities to resist the temptation to make maximum use of the possibilities for top-down control. Global control does not, of necessity, lead to better results than locally coordinated actions between individual users do.

5.7 CONCLUSION

With the examples given in this chapter, we have shown that traditional engineering opinions about the designing of technical artefacts and about the knowledge that such designing requires is no longer adequate when the artefacts attain a form of complexity that leads us to introduce the notion of sociotechnical systems. The designers and operators of such systems are confronted with numerous aspects that are not easily or not at all describable within the traditional engineering approach, which is overwhelmingly oriented toward the natural sciences. This traditional approach and the accompanying conceptual frameworks, models and theories therefore need to be enriched with knowledge that has been and is being developed within the domain of the social sciences.

As we have seen, one of the features of sociotechnical systems is that they are less predictable than traditional technical artefacts; sociotechnical systems can display unexpected behaviour even if the end-users set out to use the system in a 'neat' or 'tidy' way. Although the notion of a use plan is as problematic for a sociotechnical system as the idea that such systems are designed on the drawing board as a whole, it remains intuitively clear that there are intended and unintended or proper and improper ways of making use of a sociotechnical system. In the next chapter, it will become clear that the presence of emergent properties plays a part in what we can say about the way in which technology develops. Sociotechnical systems also force us to think anew about the way in which we attribute responsibility to designers and users in the ethical questions that surround technology, as will also be borne out in the final chapter.

5.8 A FEW MORE ISSUES

In this chapter, we have argued that the crucial quality of sociotechnical systems is their *hybrid* character. Such systems include both 'hard' technical components and people, but the way in which people fulfil the roles that have been designed for them is fundamentally different from the way a technical component does what it has been designed to do in accordance with the laws of nature. Even though this hybrid character is most apparent in the large-scale infrastructural systems in our society, being large-scale and complex are not prerequisites for having such a hybrid character. Try to imagine the smallest possible system that qualifies as being hybrid. What is minimally required? Can you give an example that meets these requirements? How does it differ most from a purely technical system?

One extremely difficult problem is to describe precisely what a typical sociotechnical system looks like. What are its components? Just 'people' and 'machines', or possibly also other sorts of things? Can, for example, organisations and institutions, which in our society, in their role as 'legal persons', are sometimes treated as being equivalent to persons of flesh and blood, be components of a sociotechnical system? Are the rules and instructions that define the roles of operators perhaps to be treated as components of such systems? But how can abstract things, like rules, and to some degree also institutions, together with tangible things, like people and machines, be components of one and the same thing? But then, is a sociotechnical system itself something tangible? Currently, we lack really satisfactory answers to these questions. Even in the social sciences, where one might expect to find some help, there is, in fact, little agreement on the answers to these kinds of questions.