Man-Machine Systems, The Hague, The Netherlands, 1992

HUMAN-COMPUTER INTERACTION IN CONTEXT: PHYSICIAN INTERACTION WITH AUTOMATED INTRAVENOUS CONTROLLERS IN THE HEART ROOM

E. Moll van Charante *,, R.I. Cook *,**, D.D. Woods*,**, L. Yue* and M.B. Howie****

**Cognitive Systems Engineering Laboratory, Department of Industrial and Systems Engineering,*
The Ohio State University, Columbus, Ohio, USA
***Department of Anesthesiology, The Ohio State University, Columbus, Ohio, USA*

Abstract. One result of recent research on human error and disaster is that the design of the human-machine system, defined broadly, modulates the potential for erroneous action. Clumsy space use of technological powers can create additional mental burdens or other constraints on human performance that can increase the chances of erroneous actions by people especially in high workload, high tempo operations. This paper describes studies of a computer based automated device that combine critical incident analysis, bench test evaluations of the device, and field observation in order to understand physician-device interaction in the context of heart surgery. The results link, for the same device, user group and context, three findings. First, the device exhibits classic human-computer interaction flaws such as lack of feedback on device state and behavior. Second, these HCI flaws actually do increase the potential for erroneous actions and increase the potential for erroneous assessments of device state and behavior. The potential for erroneous state assessment is especially troublesome because it impairs the user's ability to detect and recover from misassemblies, misoperations and device failures. Third, these data plus critical incident studies directly implicate the increased potential for erroneous setup and the decreased ability to detect errors as one kind of important contributor to actual incidents. The increased potential for error that emanates from poor human-computer interaction is one type of latent failure that can be activated and progress towards disaster given the presence of other potentiating factors in Reason's model of the anatomy of disasters.

Keywords. Human Error; Computer Interfaces; Medical information processing; Automation; Man-machine Systems; Human Factors.

INTRODUCTION

Improvements in our ability to monitor, understand, and influence complex processes have encouraged the development and introduction of more automated systems. For example, advances in the understanding of blood pressure mechanics and heart rate regulation, improved ability to monitor these variables and newly developed drugs with specific influences on these variables, have permitted the design of sophisticated intravenous infusion controllers capable of more precise intravenous drug delivery. Similar developments are underway in commercial aviation (new automation in the flight management systems of glass cockpits; cf., Wiener, 1989; Sarter and Woods, 1991), air traffic control (plans for closer, more "efficient" aircraft spacing), manned spaceflight (Woods et al., 1991), and electrical power generation and distribution. The objective of developing these new more autonomous systems is to augment operators by shifting activity from the operator to the automated device. It is an article of faith among device developers that their introduction will reduce demands and workload on the operator by taking over some part of the operator's task.

We have been studying the impact of technology on practitioner performance in four domains – commercial aviation, spaceflight, electrical power generation and anesthetic management during surgery. These studies have shown empirically that the assumption that new automation automatically reduces practitioner workload and improves overall performance is not justified. New automated devices create new cognitive and ergonomic "costs" at the both the individual operator and team level. The costs may involve new or changed tasks such as device setup and initialization, configuration control (e.g., software version control), or operating sequences. Cognitive demands change as well, including tracking automated device state and performance, new data management tasks, new attentional demands, new communication tasks and new knowledge requirements. These costs represent new levels and types of operator workload. Frequently, new automated devices produce a condition called clumsy automation by Earl Wiener (1989). Clumsy automation or the clumsy use of technology is a form of poor coordination between the human and machine. Here the automation benefits accrue during workload troughs while automation costs occur during high criticality or high tempo operations. Overall, the clumsy use of technology creates new operational complexities that appear to produce increments in workload and decrements in practitioners' ability to track changes in their environment, especially during high tempo, high criticality periods. Significantly, these deficits can create new opportunities for human error and produce new classes of system failures.

In this paper we report on our studies of a new automated intravenous infusion controller used to achieve more precise administration of vasoactive intravenous drugs in order to influence patient hemodynamics during cardiac surgery. Our investigations combined critical incident analysis, bench test evaluations of the device, and field observation in order to understand physician-device interaction in context. The results show that (a) the device exhibits a variety of classic human-computer interaction flaws (Norman, 1988), (b) these flaws increase the potential for various types of erroneous actions and impair the physician's ability to detect and recover from errors, and (c) these "latent failures" (Reason, 1990), when combined with other factors, contribute to critical incidents.

This paper focuses on the cognitive analysis of device use in context based on field observation of device use coupled with bench test evaluations. The results of a critical incidents analysis is reported for the anesthesiologist community in Cook, Woods and Howie (1992). A re-design of the device that illustrates how to correct error increasing HCI deficiencies has been completed as well (Yue, Woods and Cook, 1992).

COGNITIVE ANALYSIS OF DEVICE USE IN CONTEXT

Critical incidents reports obtained in a large teaching hospital (Cook, Woods and Howie, 1992) identified a group cases associated with an infusion controller (Rateminder, Critikon Corporation) used to control the flow of blood pressure and heart rate medications to patients during heart surgery. Each incident involved delivery of drug to the patient when the device was supposed to be off or halted. Detailed debriefing of participants suggested that, under certain circumstances, a device would deliver drug (sometimes at a very high rate) with little or no evidence to the user that the infusion was occurring.

Bench tests and field observation were used to identify (1) characteristics of the device which make its operation error prone and difficult to observe and (2) characteristics of the context of cardiac anesthesiology which interact with the device characteristics to provide opportunities for unplanned delivery of drug. The bench tests of device behavior and operating room observations of device use were mutually supportive. Observation suggested additional investigations about how the device behaved in particular circumstances. The human-computer interface related deficiencies identified in bench testing helped to direct operating room observations to test whether these problematic aspects did in fact lead to greater potential for erroneous actions and erroneous assessments of device state.

Infusion Control and Cardiac Surgery. Generally, in cardiac surgery, the anesthesiologist monitors the patient's physiological status (e.g., blood pressure, heart rate) and administers vasoactive drugs to control these parameters to desired levels based on patient baselines, disease type and stage of cardiac surgery. The vasoactive drugs are administered as continuous infusion drips mixed with intravenous (IV) fluids.

The device in question here is one type of automatic infusion controller which regulates the flow of fluid by adjusting a mechanical occluding clamp around standard intravenous (IV) tubing. The rate of flow (drops per minute) is determined by counting the drops that form in a drip chamber, comparing the measured rate to a user entered value, and adjusting the occluding clamp to achieve the target rate (Figure 1). If the device is unable to regulate flow or detects one of several different device conditions, it is programmed to cease operation (close the occluding clamp to stop flow) and emit an audible alarm and warning message. Three front panel controls are available: the target rate setting control, an ON/OFF button and a START button; however, the latter two perform multiple functions. A front panel display indicates the user entered target rate, a brief alarm message, and a "running" indicator (a large icon which flashes when the device is running but does not indicate actual flow or flow rate). The front panel also includes an unlabeled green light emitting diode (LED) that blinks when a falling drop of fluid is detected in the drip chamber. In clinical use up to six devices may be used simultaneously to control six different drugs (Figure 2); the individual drug flows are mixed together in a manifold and delivered to the patient in a single IV line (Figure 1).

Methods. Bench tests were used to identify characteristics of the device which make its operation error prone and difficult to observe. These tests required instrumenting the device to provide recording of its operation and then performing experiments to demonstrate device characteristics over a wide range of normal and abnormal conditions. Abnormal conditions included possible misoperations and device failures. In particular, we determined the device's ability to regulate flow with different rate settings, its ability to function in the presence of back pressure (a clinically relevant condition), and the speed with which it adjusted flow rates and detected deliberately induced faults. The device possesses multiple alarm states, many of which relate to detectable faults in setup. These alarm states and normal operation were mapped using simplified finite state diagrams (Figure 3 shows the finite state diagram for just three of the possible alarm messages; note that this is only one of many finite state diagrams that were needed to capture the complete space of device states/modes and operating sequences). The performance of the device was also determined when it was in motion; the devices are used during transport of patients from the cardiac surgical suite to the intensive care unit. The testing covered a period of three months and included over one hundred different experimental combinations.

To inform and complement the bench tests of the device, several physicians were observed setting up and using the device under actual conditions. Two authors observed cardiac surgical cases, at first together and later independently, and recorded user interactions with the devices. Users gave informed consent regarding the data collection and the study was approved by the institutional committee regulating research with human subjects. Each discrete user interaction with the device was recorded, along with details regarding the part of the surgical procedure and other user activities. Particular attention was given to the occurrence of alarms and changes in drip rates. The early combined records suggested high concordance between independent observers, and prior research on the conduct of cardiac surgical cases (Cook, Woods, Howie, 1990), supported the validity of the data gathering method.

Over 25 device setups and 10 complete surgical cases were recorded. User setup and operating sequences were transcribed and user and case identifying characteristics removed. The cases were heterogeneous; in some the infusion devices were not used at all while in others several infusions were started and stopped throughout the procedure. Data from this study were analyzed using process tracing techniques (Woods, 1992). In particular, physician-device interaction was mapped using diagrams which compare actual device operation, the image the device presents to the user, and user actions and intentions in context (Suchman, 1987).

Bench Test Results on Physician-Device Interaction. External indicators of the device operation make it difficult for users to assess or track device state and operation: (1) there is no indication of the actual rate of fluid delivery (the large digits on the front of the display refer to target and not actual rate), (2) a large moving element on the display suggests the actual rate of delivery but is misleading (it's flash rate seems to indicate the target rather than actual rate). Actual delivery of drug is indicated by the small unlabeled. The field study results confirmed that users confused indications of demanded drop rate with measurements of actual rate.

During operation, the controller initially overshoots the desired rate of flow, providing a small bolus of drug during startup (Figure 4). While this amount is small (in most cases amounting to 12 to 20 drops of fluid), the effect may be clinically significant when the target rate is low (e.g., 5 drops per minute) and the drug potent. Moreover, as the field observations established, in nearly all cases physicians will start infusions at low levels and titrate to a desired effect. In addition, the device requires variable lengths of time to and reach steady state at the target value. For example, at slow rates of flow the adjustment may take up to sixty seconds to reach steady state. Because the period during which the device searches for control is relatively long and not indicated to the user, users sometimes made changes in drip rates faster than the controller can accommodate them. The observation data showed that the indication to the user of target rather than actual rate strongly suggests to users that the device reaches its target rate very quickly.

Any motion of the device during operation can cause drops to be missed and leads the controller to lose its setpoint and begin searching to regain control. This inevitably results in either restarting with an initial overdelivery of drug (as a safety feature the device will close the occluding clamp, reducing flow to zero, and then begin to search for the target that has been set) or shutdown and an alarm. As a result, during transport from the operating room to the intensive care unit drug delivery will be quite erratic with frequent stoppages and overshoots of target rate. The field data showed that, given the lack of feedback, physicians were unaware of the seek mode of device behavior and the erratic control during patient transport.

Overall, the lack of visible feedback and the presence of hidden modes combined to hide device state and behaviors from the physician, i.e., they were unaware of various controller behavioral characteristics such as overshoot at slow target rates, seek behavior, erratic control while moving.

The bench testing revealed that there is a problem with ambiguous alarm messages in this device. Some messages are nuisance alarms. Several different alarm messages can be displayed for the same underlying problem; the different messages depend on operating modes of the device which are not indicated to the user (see Figure 3 for one example).

It was relatively easy to fool the device droplet sensor with a continuous stream of fluid. The sensor depends on the lens effect of a droplet of fluid but at high flow rates the droplets coalesce into a continuous column of fluid which does not register as a drop. This condition of free flow is detected as the absence of drops and indicated with either NO FLOW or OCCLUDED alarm messages. Because the droplet sensor surrounds the droplet chamber which contains the only indication of fluid flowing in the system, the use of the controller obstructs the users view. Moreover, the sensor is located some distance from the display panel and controls, making visual checks of the device operation even more difficult. The combination of the opaque interface and the inability of the sensor mechanism to discriminate between no/low flow and free flow proved to be a major factor in one of the critical incidents.

The front panel controls are complex and perform multiple functions. For example, the ON/OFF button performs different functions depending entirely on device states which are only partially indicated to the user. The START button provides both an initiation function and the ability to smoothly adjust the target rate while an infusion in progress (this unlabeled feature requires pushing two buttons simultaneously).

The bench tests also confirmed a specific engineering design flaw (as opposed to design flaws from a human-computer interaction point of view). This flaw, which appeared to be the result of a software error, prevents the device from completely occluding the IV tubing under specific circumstances and can permit flow of drug even when the device is apparently powered off. Details of this fault, which were discovered following a clinical incident, were communicated to the manufacturer and were corrected in an upgrade supplied to this institution. The manufacturer insisted that the fault would only be significant if users failed to follow certain setup procedures contained in the user manual.

The bench tests also revealed a variety of mapping problems. The mapping between fluid bags, tubing sets, the different parts of the infusion controller, and the connectors to the manifold are inconsistent and ambiguous (Figure 2). This type of deficiency can be expected to lead to misassemblies where part of one assembly is connected to a portion of a different infusion controller set up.

Overall, the bench test results identified several classic human-computer interaction deficiencies in the infusion controller (Norman, 1988; Cook et al., 1991):

~ Lack of feedback on device state and behavior.

~ Complex and arbitrary sequences of operation.

~ Ambiguous alarms.

~ Poor mappings between state and actions and between the parts of a single assembly.

Field Study Results on Physician-Device Interaction. Users encountered virtually every difficulty identified during the bench tests while using the devices in the field. Many of these problems were related to the initial setup of the device. The data from the field study also was examined to understand how users tailored their behavior to compensate for device deficiencies or weaknesses (Woods and Cook, 1991).

Setup of the devices is complex, involving up to 20 separate steps for each controller (up to six controllers may be used for a single surgical case). Because setup is complicated and because one or more of the controllers may need to be used unexpectedly and quickly to control patient hemodynamics, the infusion controllers are setup well in advance of the surgical procedure. This means that errors in device assembly and configuration may not become apparent until the devices are used, in some cases hours after setup.

Interestingly, users seemed quite aware of the potential for error and difficulties associated with device setup. Two main strategies for setup were observed in which similar components of each device were setup for all of the controllers before proceeding to the next component, and space, in which a complete device was setup before moving to the next. The setup of two different device sets by different users is shown in Figure 5. A critical step (close occluding clip wheel) which produces vulnerability to the software flaw described above was frequently missed by both users (an omission of an isolated action).

A variety of misassemblies were observed including loading the wrong occluding clip assembly, i.e., the assembly for one infusion controller setup is loaded into another infusion controller, and mounting the wrong drop detector, i.e., the detector of one infusion controller is mounted on the drop chamber for a different infusion controller assembly. Both misassemblies produce a dissociation between regulation and sensing of flow rate across two drug/IV sets and the potential for inadvertent drug delivery. A similar misassembly did lead to unintended drug delivery in one of the critical incidents.

Alarms were remarkably common during device operation. In one sequence of about five minutes duration there were at least a dozen alarms from a single device. It is significant that these alarms were not simply repeats of the same message but a variety of different messages; this sequence provoked a variety of user responses. It is important to note that, given the lack of feedback, when alarms recur, it is very difficult for the physician to determine whether or not the device has delivered any drug in the intervening period.

Users were sensitive to the fact that the devices could, under what appeared to them as ill-defined circumstances, deliver drug unexpectedly and sought to protect the system from these states by keeping various stopcocks and clamps in the IV lines closed until the devices were actually needed. This strategy itself was seen to provoke alarms since users sometimes started the devices without first clearing the IV line of obstructions; this invariably resulted in an alarm from the device indicating that flow was obstructed.

It was also observed that users were seldom able to devote undivided attention to the operation of the infusion controllers. The operating room is a busy place and the circumstances under which the drugs were required were also ones in which many other physician actions were required and where physician attention was focused elsewhere. Because the pace of this activity was faster than the time required for the controller to detect and announce faults, alarms always constituted an interruption of physician attention to other tasks. Users would start (or adjust) an infusion and then proceed to another task, and later they would be alerted by an alarm that the previous action had been unsuccessful (alarms generally occurred about 2 minutes after the triggering condition began). This alarm necessitated attentional scheduling, switching attention to the tree of infusion devices, locating the appropriate device, recall of the prior actions, and troubleshooting of the device. The most common troubleshooting approach was simply recycling the device (effectively restarting the controller's search for the target rate). This approach seemed to be efficacious, probably because it incidentally silenced alarms for some time and because intermittent occlusion (temporary occlusion or backpressure) was common. The complex cycle of adjusting the device, turning to other tasks, device alarm, and refocusing attention to the device group led in at least one instance to a user misidentifying the faulty device.

Particular factors in the environment constrain operators to use the devices in certain ways (e.g. set up precedes use by as much as two hours). The device design does not support these constraints. For example, the START ME alarm appears about two minutes after the device is setup and is a strong prompt to turn the device off since setup and use are always separated by much much more than two minutes. In this environment, START ME is in fact an indication that the device needs to be turned off, not that it needs to be started. Moreover, since it is delayed by two minutes (a long time in this world) the operator will be doing other things by the time it appears. The message was never observed to indicate correctly what the operator should actually do.

The device design, when considered in the context of use, encourages the user to turn the controller off following set up until that device is actually needed. However, misassemblies can occur that exist as latent failures in the system waiting for a trigger to become an active failure, i.e., an inadvertent drug delivery. Since the device is off, no alarms will be generated. This dynamic played out in one of the critical incidents investigated in Cook et al. (1992).

The most intense periods of device use also were those time periods of highest cognitive load and task criticality for the physicians, i.e. the time period of coming off cardio-pulmonary bypass. It is precisely during these periods of high workload that the automated devices are supposed to provide assistance (less user workload through more precise flows, smoother switching between drip rates, etc.). However, this is also the period where the largest number of alarms occur and where device troubleshooting is most onerous. As we have found in other operating room devices (Cook et al., 1990) and in other complex domains (Wiener, 1989; Sarter and Woods, in press; Woods et al., 1991), close study of device use in context revealed that this automated device exhibits the classic properties of the clumsy use of technological possibilities.

The field observations showed that the HCI deficiencies identified in the bench tests actually do increase the potential for erroneous setup and operation of the device. The field observations also showed that device HCI deficiencies increase the potential for erroneous assessments of device state and behavior. The potential for erroneous state assessment is especially troublesome because it impairs the physician's ability to detect and recover from erroneous actions and setups.

DISCUSSION

One assumption of recent research on human error and disaster (e.g., Reason, 1990) is that the design of the human-machine system, defined broadly, modulates the potential for erroneous action. The results of this series of studies directly link, for the same device and context, HCI design deficiencies to increased potential for erroneous actions and impaired ability to detect and recover from errors. The results, when combined with a study of critical incidents for the same device and context, directly link the increased potential for erroneous setup and the decreased ability to detect errors as important contributors to critical incidents. The increased potential for error arising from poor human-computer interaction can be seen as one type of latent failure that can reside within a complex human-machine system. Activating this type of latent failure in the presence of other potentiating factors leads incidents nearer to disaster. The data in this particular case provide direct support for Reason's (1990) pathogen model of the anatomy of disasters (cf. also, Woods, 1990).

As new automated devices are developed and deployed the apparent simplicity of the computer interface tends to hide significant underlying complexity (Cook et al., 1991). Furthermore, one can see in this, as in other cases that we have investigated (Woods and Cook, 1991), how new technology that glosses over substantive issues in human interaction can increase the operational complexity of the overall system. This technology induced complexity increases the demands on practitioners despite superficial putative benefits claimed for the new device or technology.

The field study portion of this research also reveals that practitioners are not passive recipients of new technology; rather, they actively tailor the device itself and their work to accommodate the new device. Because practitioners are responsible agents in the domain, they work to insulate the larger system from device deficiencies and peculiarities of the technology. This occurs, in part, because practitioners inevitably are held accountable for failure to correctly operate equipment, diagnose faults or respond to anomalies even if the device setup, operation, and performance are ill-suited to the demands of the environment. In this study users tailored their behavior to avoid problems and to defend against device idiosyncrasies. However, the results also show how these adaptations may be only partly successful. The adaptations, while useful in narrow contexts, were often brittle and unable to cover the wide range of circumstances actually found in the domain.

One may ask how the present study was able to reveal so many design deficiencies and problems associated with device use. The answer seems to us to lie in the methodology of study which focuses on understanding device use in context. This type of context bound approach depends on a hybrid of techniques involving critical incident investigations, evaluating device function and behavior under controlled conditions, and field studies of device use. Understanding the demands placed on practitioners by the environment is key both to the design and the evaluation of infusion controllers and, by extension, of other computer-based devices as well.

There is a tendency to view new automated devices as operator aids which make fewer demand on operator cognition than the manual tasks they replace. While this may be true in a narrow sense, the new device, when considered in the full context of the user's environment, also creates new cognitive and physical demands on practitioners. The device described here is particularly difficult for operators to setup, monitor and supervise effectively. Many of its design flaws are readily apparent in retrospect. Yet devices such as this one are commonplace and widely used. Interestingly, flaws in the device are not so apparent that it is rejected by users at first glance. Moreover, because the environment is complex and fast paced, when problems do occur their source is frequently difficult to identify:

~ faults may be missed – it was the presence of medically knowledgeable cognitive engineers working in close cooperation with anesthesiologists on another project that enabled recognition and follow up of the critical incidents that occurred,

~ faults may be ascribed to other agents (often the users!) – it was only after a series of incidents had occurred that there was general recognition that the problem lay in device design given the context of the heart room and not with individual users,

~ faults may be rationalized as the "learning curve" for introducing any new device, or

~ faults may be seen simply as "unusual" events with no larger significance.

Especially in the operating room, where poor quality human-computer interface is the rule rather than the exception, the poor performance of a single device may seem to be routine, i.e., just another obstacle that physicians are supposed to work around to provide safe and effective patient care.

Discovering the characteristics of the device presented here required substantial time (the present research represents over two person-years of technical effort) which is far beyond the resources of most users who are, after all, asked to work with many different devices. The apparent simplicity of the device given the interface characteristics was a formidable barrier to understanding how the device actually works both in general and in a particular context. For example, extensive bench tests were required to penetrate the opaque barrier in order to map the state transition space for the device.

More importantly, the interactions between device characteristics and user environments are necessarily heterogeneous. As the complexity of the domain increases, device roles must necessarily become more specialized as is required by Ashby's Law of Requisite Variety. The use of general purpose devices therefore forces the practitioner to undertake the task of shaping the device to conform to the requirements of the task. The heterogeneity of user environments must be matched by the heterogeneity of devices; as the tasks become more diverse, the user aids must diversify in step. Paradoxically, this is a time in which many devices are becoming generic: the controller described here is used in many different parts of the hospital, for many different purposes and it may function well in some of them.

The purpose of this paper is not to indict a device or even a class of devices, nor is it to call into question the value of automated controllers. Rather it seeks to focus attention on the difficult problem of providing useful assistance to operators of complex, high risk processes (Norman, 1990). It is clear, from this study, that it is possible to produce controllers which have an apparent quality of engineering "reasonableness" but which actually combine with the operating environment to produce the potential for new kinds of faults and failures. The challenge to the engineering and human factors communities is to find ways to characterize the practitioner's context with sufficient precision and detail to permit the design of user aids which assist rather obstruct high workload, high tempo task performance.

ACKNOWLEDGEMENTS

The opportunity to investigate and follow up the critical incidents with the research reported here arose because of another research project to build a Corpus of Cases on Human Performance in Anesthesia sponsored by a grant from the Anesthesia Patient Safety Foundation to the third author. Additional support for this study was provided by The Ohio State University Department of Anesthesia. We thank the Universiteit van Utrecht, Geneeskunde Faculteit, Utrecht, The Netherlands and Professor Smalhout for making it possible for Eric Molle Van Charante to do his research practicum in the Cognitive Systems Engineering Laboratory and the Department of Anesthesia at The Ohio State University. We are very grateful for the assistance of the many physicians who gave their time and shared their experiences with us in the course of this research. Address correspondence to D. D. Woods, Department of Industrial and Systems Engineering, The Ohio State University, 1971 Neil Avenue, Columbus, OH 43210, U.S.A. (email: woods@csel.eng.ohio-state.edu).

REFERENCES

R.I. Cook, S. Potter, D.D. Woods, and J.S. McDonald. Evaluating the human engineering of microprocessor controlled operating room devices. Journal of Clinical Monitoring, 7:217-226, 1991.

R.I. Cook, D.D. Woods, and M.B. Howie. The natural history of introducing new information technology into a dynamic high-risk environment. In Proceedings of the Human Factors Society, 34th Annual Meeting, 1990.

R.I. Cook, D.D. Woods and M.B. Howie. Unintentional delivery of vasoactive drugs with an electromechanical infusion device. Journal of Cardiothoracic and Vascular Anesthesia, 6:1-7, 1992.

D.A. Norman. The Psychology of Everyday Things. Basic Books, New York, 1988.

D.A. Norman. The 'Problem' of Automation: Inappropriate Feedback and Interaction, Not 'Over-Automation', Philosophical Transactions of the Royal Society of London, B 327, 1990.

J. Reason. Human Error. Cambridge University Press, England, 1990.

N. Sarter and D.D. Woods. Pilot Interaction with Cockpit Automation: Operational Experiences with the Flight Management System. International Journal of Aviation Psychology, in press.

L.A. Suchman. Plans and Situated Actions: The Problem of Human-Machine Communication. Cambridge University Press, Cambridge, England, 1987.

E.L. Wiener. Human Factors of Advanced Technology Glass Cockpit Transport Aircraft. Technical Report~117528, NASA, 1989.

D.D. Woods. Risk and Human Performance: Measuring the Potential for Disaster. Reliability Engineering and System Safety, 29:387–405, 1990.

D.D. Woods. Process Tracing Methods for the Study of Cognition Outside of the Experimental Psychology Laboratory. In G. A. Klein, J. Orasanu and R. Calderwood, editors, Decision Making in Action: Models and Methods, Ablex, New Jersey, 1992.

D.D. Woods and R.I. Cook. Technology-induced complexity and human performance. Proceedings of IEEE International Conference on Systems, Man, and Cybernetics, IEEE, 1991.

D.D. Woods, S.S. Potter, L. Johannesen and M. Holloway. Human Interaction with Intelligent Systems: Trends, Problems, New Directions. Cognitive Systems Engineering Laboratory Report 91-TR-02, February 1991.

L. Yue, D.D. Woods and R.I. Cook. Reducing the Potential for Error Through Device Design: Infusion Controllers in Cardiac Surgery. Cognitive Systems Engineering Laboratory Report 92-TR-01, The Ohio State University, Columbus OH, January 1992.

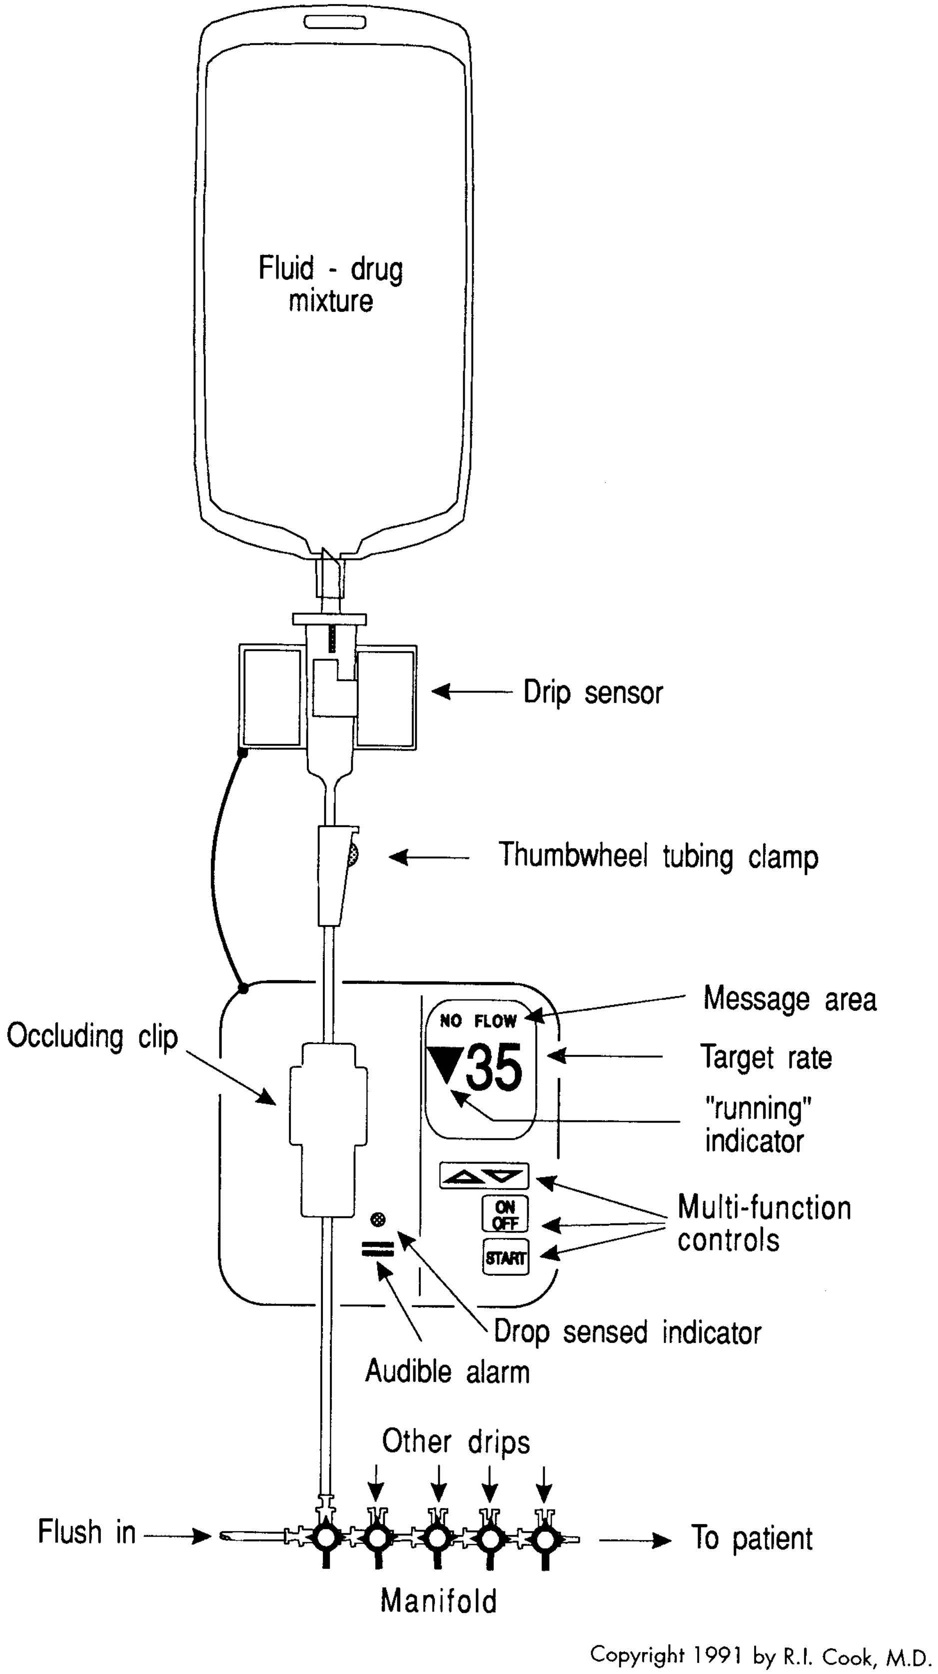

Figure 1. The assembled drug administration set up including the infusion controller. The latter includes three components: the drip sensor, the occluding clip and the controller unit with multi-function controls and LCD display area.

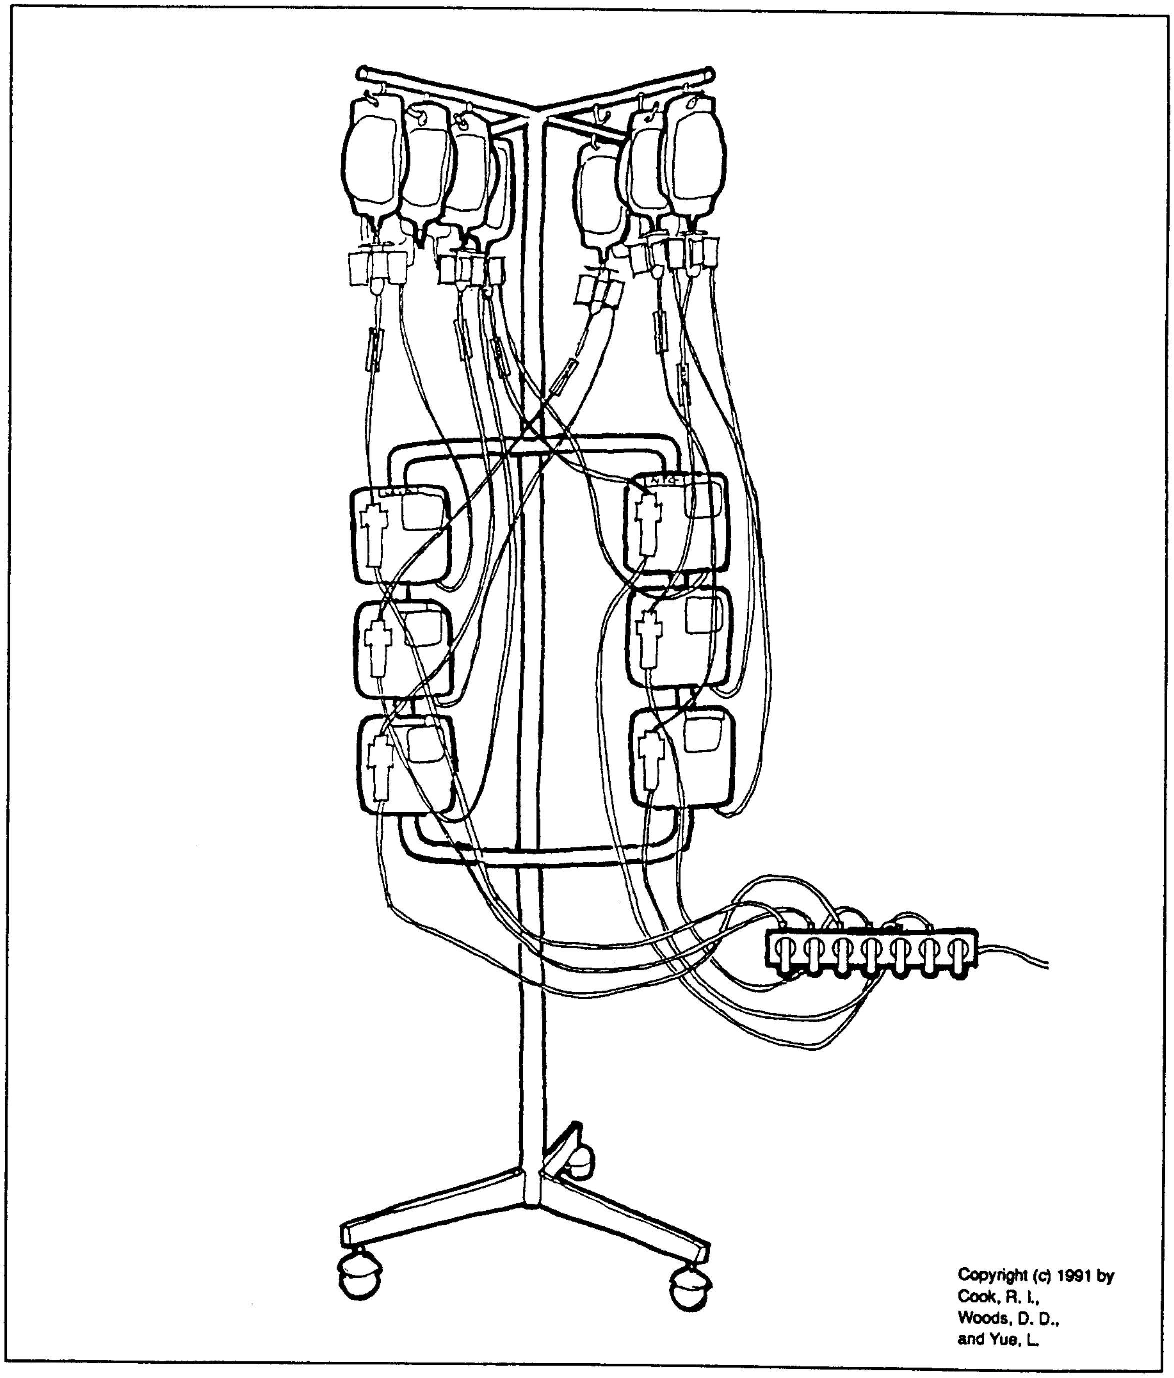

Figure 2. A drawing of the tree of six infusion controllers as would look in preparation for use in cardiac surgery.

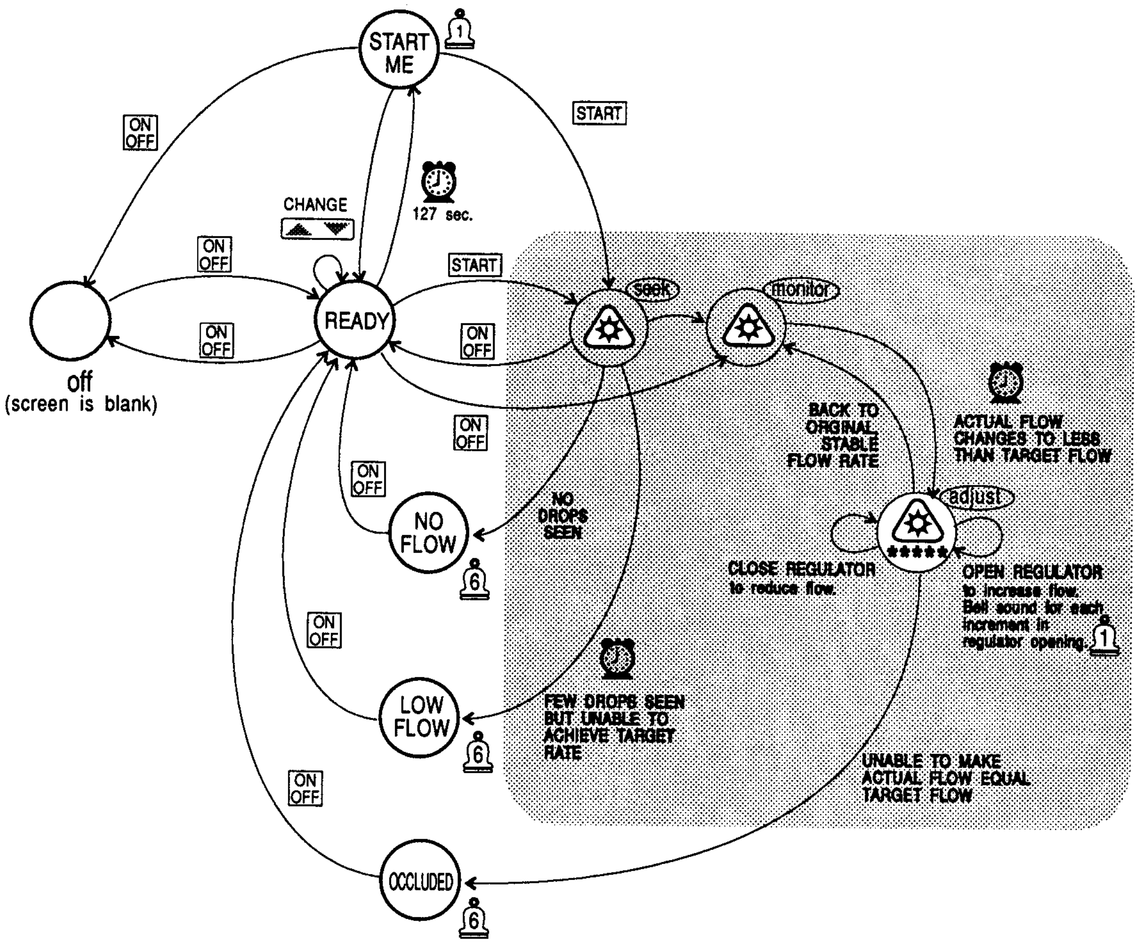

Obstruction to flow alarm states

Only messages in circles are visible to the user. Note that the triangular icon is used to indicate that the device is operating and gives a rough indication of the target but not actual drip rate. Note the multiple functions of the ON/OFF button. For simplicity, numerous states and state transitions not shown.

LEGEND

(message)	Distinct state in which "message" appears on display
⚠	Display icon indicating running state
⏰	An interval of elapsed time
🔔6	Alarm bell rings six times
(state)	Internal state (not visible to user)
[name]	Press of the named control button

Figure 3. State transition diagram for alarms related to obstructions to flow.

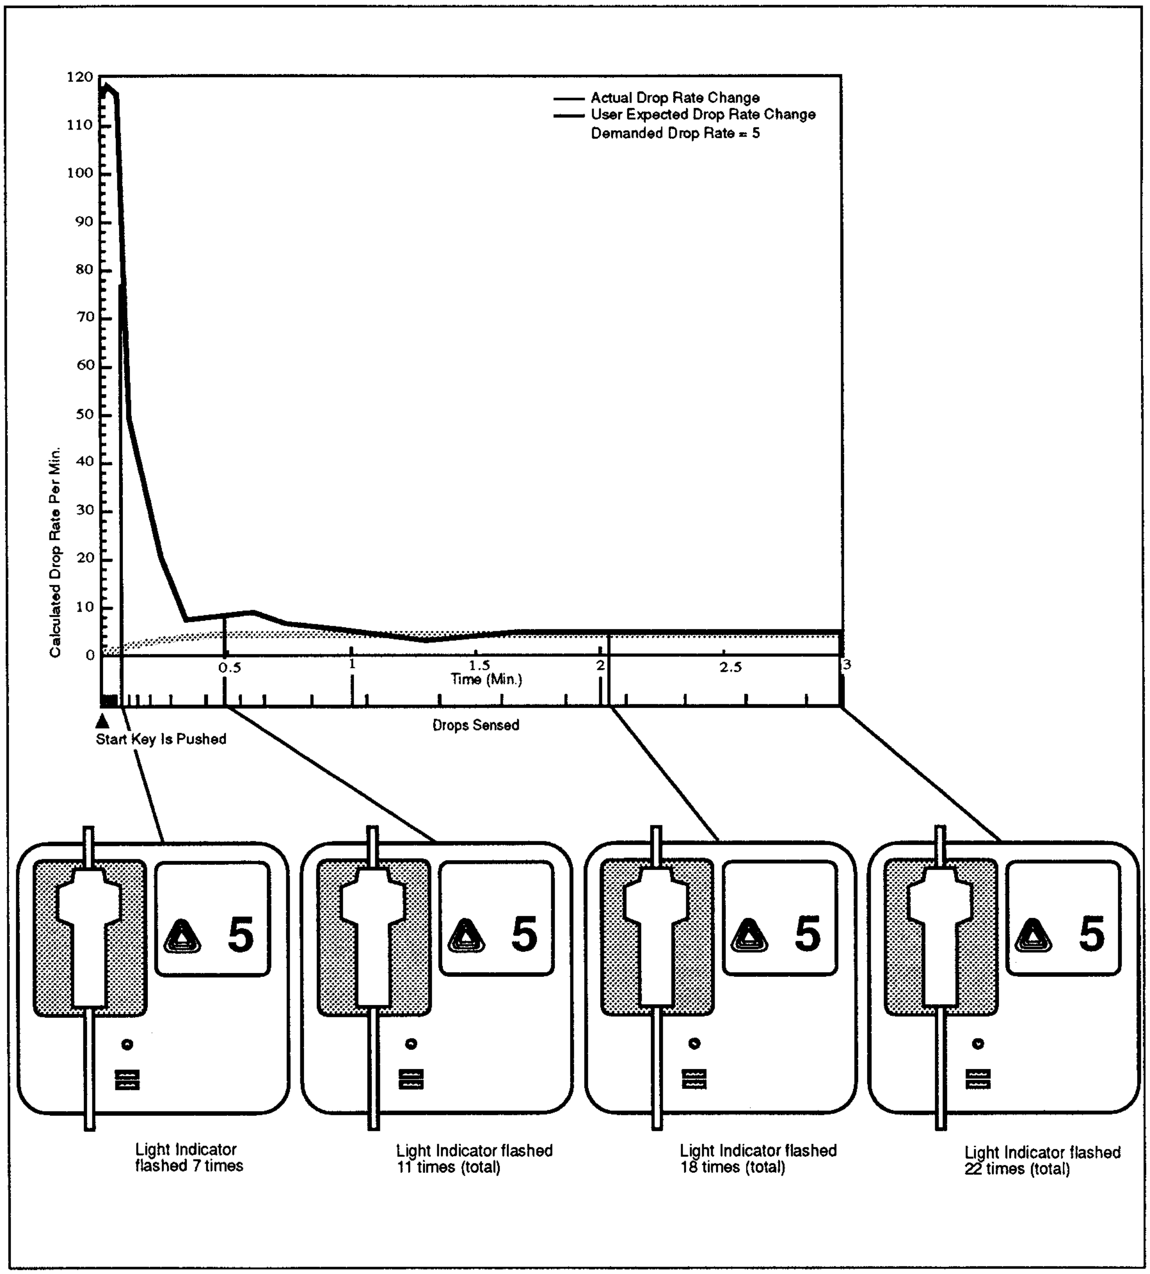

Figure 4. Compares data about the initial overshoot behavior (blackcurve) of the infusion controller at low rates and the anesthesiologist's conceptual model of what the device does (graycurve). The triangle indicates when the device was started; the tickmarks along the time axis indicate when drops were actually sensed.The data plotted is from one typical case run in the bench tests. The drawings of the device represent four snapshots during the overshoot behavior. One can see how the interface hides controller behavior.

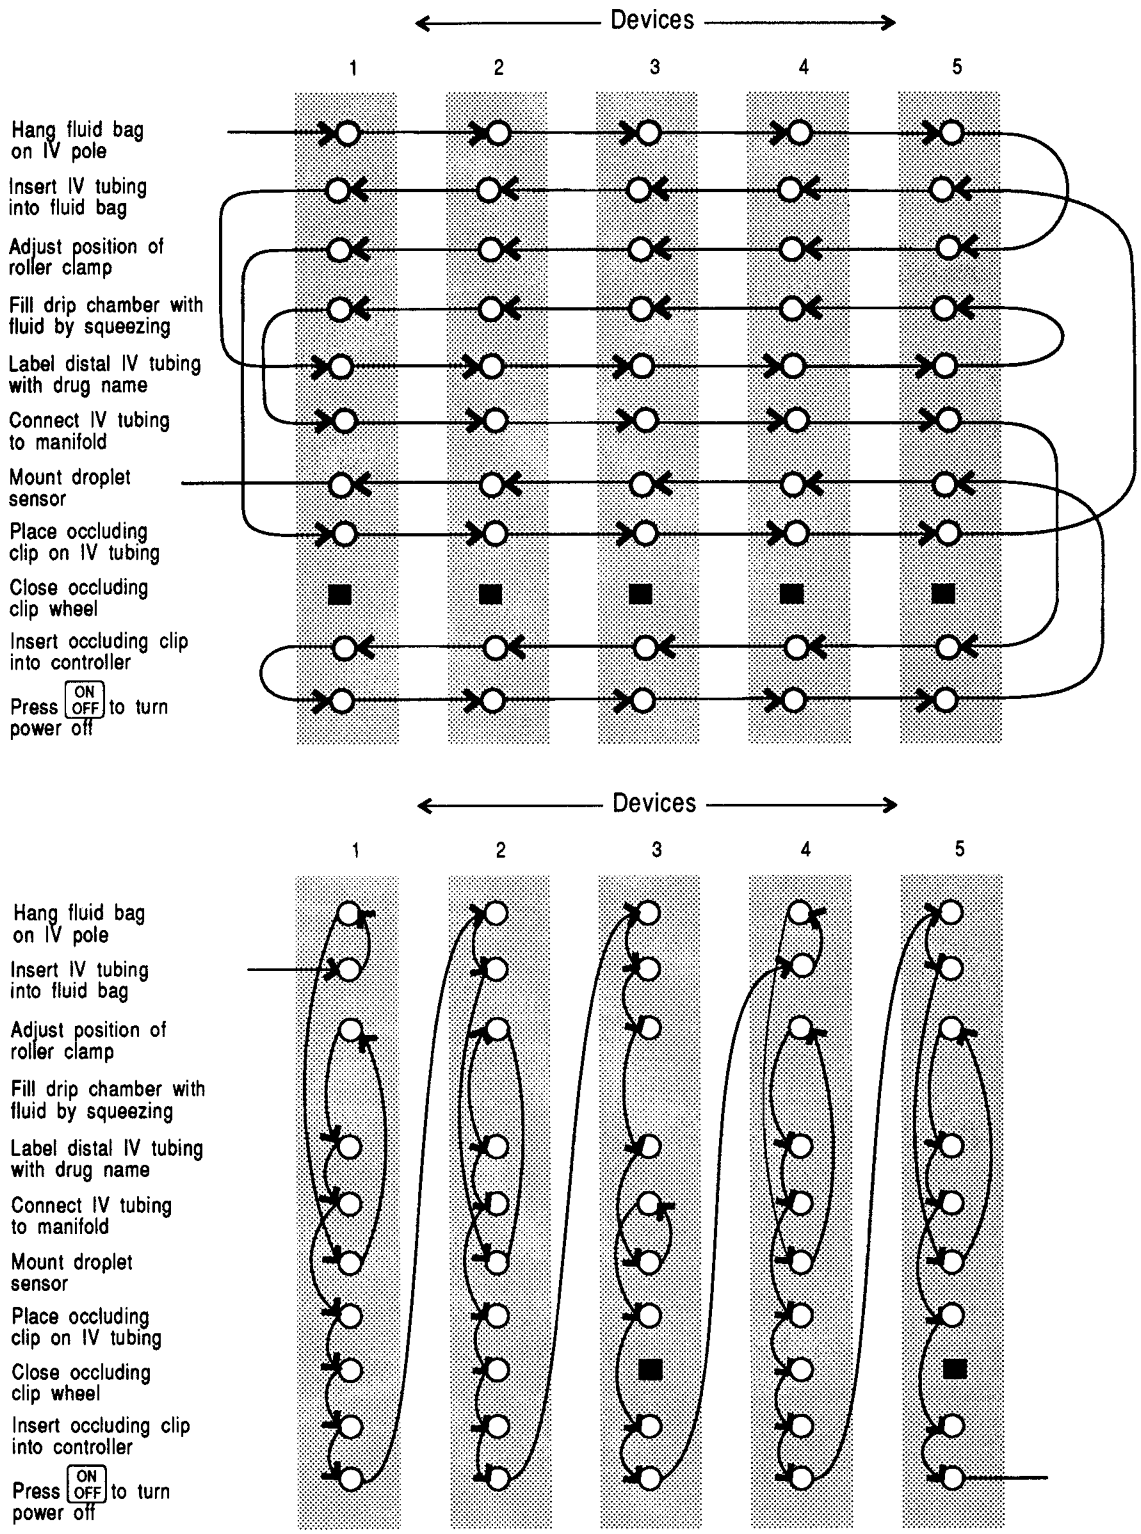

Figure 5. Charts two different device setup strategies used by physicians. The data is from field observations of actual device set up. Note that one step in the "procedure" is frequently omitted in both strategies.