MODELING AND CONTROL

A. Dean Forbes

Forbes AD. Modeling and control.
J Clin Monit 1990;6:227–235

I take the three tasks of the anchor to be (1) to situate—for noncognoscenti—the panelists' contributions in context, (2) to provide or elicit clarification of difficult points, and (3) to raise general questions to spark discussion. Consistent with these tasks, this discussion first comments on four aspects of physiologic modeling and control; it next raises general questions calculated to highlight further the differing perspectives of the panelists; finally, it summarizes each panelist's formal presentation and seeks expansion of material that may puzzle nonexperts.

FOUR ASPECTS OF MODELING AND CONTROL

Each panelist has elected to focus primarily on automatic control and has had relatively little to say about the problems of modeling. But, of course, control necessarily rests on a model of the response variable as a function of the controlling variable(s), be that model explicit or implicit in the controller. Two aspects of modeling and two aspects of control, each of which raises concerns among nonspecialists, merit brief comment.

1. Use of Simple Models of the Patient

To those acquainted with the elaborate models used in patient simulations, the models used by controller designers often seem highly simplified, so much so that one is tempted to wonder if such models can possibly suffice. This concern is not lessened by control practitioners who first introduce extremely general models and then simplify them drastically, giving assurances that all will be well.

The crucial point is that such simplification can be justified and, indeed, is essential. In estimating the parameters of a model, one is solving an inverse problem. Such problems exhibit unstable solutions, the more so when the number of parameters to be estimated is large. Thus, there is a strong impetus for a controller to use the simplest meaningful model. A strategy of simplification can succeed because the goal is reliable control, not determination of subtle patient characteristics. The challenge is to select from the vast array of distributed and lumped biophysical models or from the extensive library of phenomenologic models (autoregressive-moving average models and their extensions) a solvable model just rich enough to encompass the essentials of the situation *for control*.

From the Medical Department, Hewlett-Packard Laboratories, PO Box 10490, Palo Alto, CA 94303-0971.

2. Risk of Model Switching

The need for simple models does not imply that one can be cavalier about model specification. The fault tolerance of biologic systems is achieved, in part, through compensatory "backup mechanisms." When the organism believes itself to be *in extremis*, heroic mechanisms may come into play. These may render the selected simple model structure itself quite misleading. Control engineers seek to avoid the dire consequences to which an invalid model might lead by carefully choosing a benign clinical-environment so the patient is not likely to depart profoundly from the descriptive model, by including a "safety shell" to sense and avert impending catastrophe, or by engaging a vigilant human monitor to intervene should the controller become unable to cope reliably.

3. Controller Method Preference

Since sensitivities to controlling variables (such as drug infusion rates) can vary markedly across patients, model parameters must be individually tailored. When intrasubject variation is substantial, ongoing learning of model parameters—adaptation—becomes necessary. Just what is the best approach to adaptation is a focus of considerable debate. One should be wary of arguments such as this:

Controller method x fails (on patient group X in environment X).
Controller method y succeeds (on patient group Y in environment Y).
Therefore, controller method y should displace method x.

If patient groups as well as environments are sufficiently similar, the argument has force. But when there are substantial differences, the premises do not entail the conclusion.

4. Appropriate Medical Contexts for Automatic Control

The earliest studied arena for closed-loop control (and the first to market) is quite appropriate: postoperative control of blood pressure via sodium nitroprusside. In terms of anesthesia's time-honored flight analogy, postoperative blood pressure management is akin to "autotaxi." The flight (operation) has ended and instead of having a copilot taxi the plane to the gate, one switches to autotaxi.

There will also likely be a place eventually for nonexperts to oversee closed-loop blood pressure control in the operating room for simple surgery on medically simple subjects. To judge from the ample cautions supplied by Dr Prys-Roberts, that time has not yet arrived. But, on a "flight" that gives every promise of proceeding smoothly, there should be a place for "autopilot" control of blood pressure.

Leaps to closed-loop control of complex patients undergoing complex procedures do give pause. (Note that none of the panelists has advocated such a leap.) The flight analogy makes the point. Normally, a pilot is thoroughly familiar with the plane, autopilot, and their interactions; lengthy experience with that type of plane in actual flight and in wide-ranging simulations has seen to that. There is a substantial risk that a pilot—if unduly insulated from the flight characteristics of an unfamiliar plane by automatic control system(s)—might not gain an adequate feel for that craft's dynamics, hampering manual control should it become necessary.

Summary of Basic Points

1. For models underlying control, Einstein's dictum applies: "Everything should be made as simple as possible, but not simpler."
2. In the intensely nonlinear world of some patient crises, switching of model *structure* (rather than simple parameter drift) can occur and must be catered to.
3. Arguments for the superiority of a particular method of control often have a wishful tone. Most control methods—given sufficient Ptolemaic elaboration—can be made to work, to a point.
4. You can—perhaps—automatically control all of the patients some of the time, and you can—perhaps—automatically control some of the patients all of the time, but *no way* can you automatically control all of the patients all of the time. [Apologies to Mr Lincoln.]

GENERAL QUESTIONS

What are the panelists' responses to the four previous assertions and to the following general questions:

1. Are false mean arterial pressure readings still a problem for controllers or can bad data be reliably rejected?
2. Among the various categories of controllers, such as proportional-integral-derivative, rule-based, self-tuning, and hybrids, are any clearly superior to the others? How and when?
3. What is your assessment of work on multivariate

control, such as that of Voss et al [1] on the simultaneous control of cardiac output and mean arterial pressure using sodium nitroprusside and dobutamine?

4. How do you respond to the observations of Wittenmark and Åström [2] regarding self-tuning regulators:

   There is a large discrepancy between simulation or academic algorithms and practical algorithms. In the idealized environment of simulations it is easy to perform well. In practice the situation is quite opposite. The adaptive or self-tuning controller must be able to handle nonlinearities, unmodeled dynamics and unmodeled disturbances over a wide range of operating conditions.

5. How extensible are your methods? Will autonomous control in the operating room ever become feasible?
6. Where will the future focus of research lie: models, parameter estimation, robustness, multiple variables, new medical contexts?

REFERENCES

1. Voss GI, Katona PG, Chizeck HJ. Adaptive multivariable drug delivery: control of arterial pressure and cardiac output in anesthetized dogs. IEEE Trans Biomed Eng 1987;34:617–623
2. Wittenmark B, Åström KJ. Practical issues in the implementation of self-tuning control. Automatica 1984;20: 595–605

SUMMARY OF PAPER READ BY DR PRYS-ROBERTS

To minimize blood loss and improve surgical visibility, operative hypotension can be induced by decreasing cardiac output or systemic vascular resistance, or both, via increased airway pressure, postural changes, and specific drugs. Induction of hypotension must be slow enough that strong autoregulatory reflexes are not triggered and must be such that perfusion remains adequate. Although sodium nitroprusside decreases systemic vascular resistance without affecting cardiac performance unduly, individuals differ by thirty-to-one in their sensitivity to this drug. Hence, closed-loop control necessarily entails interpatient adaptation. If simple models are to be used, intrapatient adaptation is important. For each of the possible routes to adaptation, drug effect time lag must be taken into account. But sodium nitroprusside time lags vary from 20 to 55 seconds. Prys-Roberts and Millard assume a standard dead time of 30 seconds and use this as their data sampling interval. They then model mean arterial pressure as a function of sodium nitroprusside infusion by an equation wherein the present mean arterial pressure is given in terms of its value at the previous sampling instant, the value of the infusion rate two sample instants prior to the present, and an offset parameter. Control is achieved via the Clarke-Gawthrop generalized variance method, a method that produces smooth transitions and that has provision for manually decreasing the rate of response when model oscillation impends.

Via sodium nitroprusside infusion during ear-nose-throat surgery in 19 subjects, mean arterial pressure was automatically held within ±5 mm Hg "for a median value of 74% of the time" (as compared with 72% of the time in 10 similar subjects in whom mean arterial pressure was controlled manually). Via isoflurane during ear-nose-throat surgery in 8 other patients, mean arterial pressure was automatically held "within ±5 mm Hg of the set point 88% of the time" (as compared with 89% of the time in 12 subjects in whom it was controlled manually). In addition, the controller allowed rigorous assessment of the impact of oral premedication with clonidine, demonstrating its efficacy in reducing the intraoperative isoflurane required for automatic blood pressure control.

Questions for Dr Prys-Roberts

1. How great an impediment to reliable automatic blood pressure control do you believe the risk of triggering strong autoregulatory reflexes is?
2. You (and others) repeatedly warn that the proper choice of sampling interval relative to drug dead time is crucial. Why do you fix the sampling interval to a mean delay rather than tailoring it to the given patient?
3. For your results on automatic blood pressure control via sodium nitroprusside, could you expand the explanation that "the actual mean arterial pressure was maintained within ±5 mm Hg of the target level *for a median value of* [italics mine] 74% of the time"?
4. Since isoflurane can induce coronary steal in patients given to subendocardial ischemia, should control be avoided in such patients or should control incorporate information as to ischemic status?

Dr Prys-Roberts and Dr Roger Millard Reply

The Four Aspects

SIMPLE MODELS. The broad features of the response of a patient's mean arterial pressure to a constant infusion of sodium nitroprusside (SNP), considered in the context of cause and effect in isolation from other internal and external influences, are relatively simple. On time and

pressure scales with resolution adequate for the purposes of anesthesia, a coarse model of the response, such as an exponential level change after an initial delay, may often be an acceptable approximation for the goal of inducing and controlling hypotension with the drug. This model will usually be adequate and parsimonious with regard to the estimation of its defining parameters. However, one then has to ensure that these parameters of the model of the response to the drug are estimated unambiguously (for the purposes of control), and also that any other internal or external influences do not become incorrectly identified with the drug response parameter estimates.

RISK OF MODEL SWITCHING. We consider the controllers as aids to the anesthesiologist in performing his or her duties. The anesthesiologist must decide that he or she is using the drug in the manner for which it is intended. Other drugs, such as beta-adrenergic blockers, may be needed to bring the blood pressure into a range where its response to SNP infusion is of the normal, simple form. The controller should be insensitive to such interventions.

CONTROL METHOD. The controller design specifications should be adequate for the circumstances in which it is to be used, with regard to variations in patient sensitivities, response time scales, and environments.

MEDICAL CONTEXTS FOR CONTROL. The clinician decides the treatment appropriate for a particular patient. If an automaton is recruited in this treatment, then the operator must be sufficiently well trained in its use. If the controller is to be used in situations in which the state model of the patient's drug response may change abruptly, then its design must cater to these possibilities.

Answers to Specific Questions

1. We find no difficulty with self-tuning control of hypotension using isoflurane in the presence of the renin-angiotensin-aldosterone reflex [1]. If a baroreceptor reflex occurs, then a beta-adrenoceptor antagonist may be needed as an adjuvant, but this does not degrade the controller's performance [2]. When we used SNP for inducing hypotension, baroreceptor or epinephrine-norepinephrine reflexes initially gave us problems [3], but intervention with other adjuvants or with volatile anesthetics is the correct clinical remedy. In such a situation a switched-model design for the controller should be more robust [1,4,5].

2. Choice of sample time should be appropriate to the rate of response to the drug. This choice is also influenced by the versatility of the particular control method selected. Some control laws are more sensitive to variations in drug response time delay than others. In the long run it would be inconvenient to obtain knowledge of a patient's drug response lag time prior to choosing a sampling time at which to run the controller.

3. Alternatively, the pressure was maintained within ± 5 mm Hg of the target level for 67 $\pm$ 23% (mean $\pm$ SD) of the time by automatic control, and for 71 $\pm$ 18% by manual control, $p > 0.05$, not significant.

4. Both SNP and isoflurane have been associated with coronary steal [6,7]. Induced hypotension cannot be routinely recommended for patients with overt ischemic heart disease but may be indicated in some patients in whom the risk of ischemia may be outweighed by benefits of modest decreases of arterial pressure. In this context, the use of an automatic controller may be preferable to manual control.

Answers to General Questions

1. False pressures will cause a problem with any kind of controller if they are not rejected. Ways of reliably detecting them usually involve checking the status of the ancillary equipment [1,2]. For the parameter estimation technique, outliers on the normal linear model due to artifacts or abrupt pressure changes may be rejected by a statistical test [2,5].

2. Some controllers are better with processes showing a response dead time (e.g., self-tuners); others may be better when the regulation requirement is fairly slack (PID, rule-based). Some methods are robust if the process has no dead time and the range of possible process gains is small (e.g., fixed gain PID). Self-tuners give tighter control than robust controllers provided the PARAMETER ESTIMATION has been good. For some clinical applications a robust controller should often suffice. In this class there are, for example, passive-adaptive (gain-scheduled), H_∞, or multiple-model controllers [8]. The quality of control provided by clinical personnel and deemed acceptable should provide a guide as to required performance.

3. Reliable, continuous measurement of cardiac output is not yet feasible in humans, though noninvasive methods, such as pulsed Doppler, or thoracic impedance methods, have been described. Difficulties arise in the estimation of many parameters for arbitrary multivariable models of the cardiovascular system. This problem is confounded by the fact that the sign of the response to the said drugs is dependent on the state of the cardiovascular system (through their interaction on the measured variables) [9]. One possible alternative

would be to use separate single-input single-output regulators for each drug, with CROSS-COUPLED FEED FORWARD compensation for any interactions. Constraints on the signs of the response parameters (SWITCHED-MODEL), as appropriate, and on their ranges, would also be necessary. This approach requires fewer parameters to be estimated [4,5]. One must in any case use a method, such as Voss and colleagues imply, that is insensitive to uncertain process time delay and a considerable range in possible input sensitivity.

4. The problems mentioned by Wittenmark and Astrom in 1984 are now much better understood. (See [10,11] for a more up-to-date appraisal.)

5. The amount of supervisory software needed for reliable functioning of self-tuners can be very large. With the advent of the transputer, however, hardware with dramatically increased capabilities will become available. Software will remain a bottleneck [11].

6. All of these issues will be focused on in future research. Robustness and integrity are fundamental in any medical application. The clinical situation is multi-input multi-output, and any parameter estimation must be robust in this context. It will certainly be possible to control automatically the delivery rate of more than one drug. New medical contexts will be explored with current methods.

GLOSSARY

PID Proportional integral derivative controller. A type of control strategy in which the value of the control input delivered to the system is the sum of three functions of the current error. The first component is a constant, *Kp*, multiplied by the error (proportional). The second is a constant, *Ki*, multiplied by the time integral of the error (integral). The third is a constant, *Kd*, multiplied by the derivative of the error (derivative). Here, error is defined as the target value of the system's output variable minus its current actual value. The proportional term is used to achieve speed of response and accuracy. The integral term compensates for small persistent errors. The derivative term is used to damp out excessive transients.

PARAMETER ESTIMATION A procedure for statistically determining mathematical functions that very closely match experimental data. The procedure has two parts. First, one specifies a mathematical form (e.g., z = a* y + noise) that contains measured variables (z and y in this case), and parameters to be estimated (a in this case). Then one uses a statistical procedure to estimate the true value of the parameter a. In the context of this discussion, parameter estimation is very useful in on-line learning by the controller. Essentially, the parameter estimator uses past input/output information from the actual system being controlled to determine the difference equation exhibited by the system. The resulting difference equation is then used to design the appropriate controller. In adaptive controllers, the latter design is also done on-line.

CROSS COUPLED This term refers to interconnections among systems. Two systems are said to be cross coupled if the output of each system is affected by the input to the other.

FEED FORWARD A type of control in which control energy is inserted into the system without reference to an error signal computed from the current system output. Feed forward control is very helpful in compensating for measurable disturbances. It can be used to anticipate disturbances, such as any adjuvant drugs that may affect blood pressure. For comparison, most control strategies are of the feedback type. There, the current system output is fed back to its input to compute an error and specify some control action (e.g., with the PID strategy described above).

SWITCHED MODEL An expansion of the approach used in parameter estimation (see above). In some problems, rather than specify a single model form and estimate the best parameters, it is more convenient to specify a set of possible model forms and use hypothesis testing to specify the currently appropriate model, given the data. The result is a sequence of model switchings in response to system changes. The conventional parameter estimation approach would have resulted in more smoothly changing parameter estimates as the system changed.

REFERENCES

1. Millard RK. Studies in self-tuning control of hypotension during anaesthesia and surgery. PhD thesis, University of Bristol, January 1990
2. Millard RK, Monk CR, Prys-Roberts C. Self-tuning control of hypotension during ENT surgery using a volatile anaesthetic. IEE Proc D 1988;135:95–105
3. Millard RK, Hutton P, Pereira E, Prys-Roberts C. On using a self-tuning controller for blood pressure regulation during surgery in man. Comput Biol Med 1987;17:1–18
4. Millard RK. Experiences with self-tuning control of blood pressure. In: IFAC Symposium on modelling and control in biomedical systems, Venice. Oxford: Pergamon, 1988:134–139
5. Millard RK, Monk CR, Woodcock TE, et al. Some applications of self-tuning control to blood pressure regulation. IFAC Symposium on modelling and control in

biomedical systems, Venice. Oxford: Pergamon, 1988: 89–100
6. Monk CR. Isoflurane and induced hypotension. Curr Opin Anaesthesiol 1988;1:88–93
7. Monk CR, Millard RK, Hutton P, Prys-Roberts C. Automatic blood pressure regulation using isoflurane: a comparison against manual control. Br J Anaesth 1989;63:22–30
8. Astrom KJ, Neumann L, Gutman PO. A comparison between robust and adaptive control of uncertain systems. IFAC Workshop on adaptive systems in control and signal processing. Oxford: Pergamon, 1986:43–48
9. Roy RJ. Adaptive cardiovascular control using multiple drug infusions. IEEE Frontiers of Engineering in Health Care, 1982:459–464
10. Astrom KJ, Anton JJ, Arzen KE. Expert control. Automatica 1986;22:277–286
11. Astrom KJ. Towards intelligent control. IEEE Control Systems 1989;9(3):60–64

SUMMARY OF PAPER READ BY DR SEBALD

The usual difficulties of closed-loop drug infusion for blood pressure control are severely compounded during cardiac surgery. A patient's sensitivity to sodium nitroprusside may vary a hundredfold over the course of the operation. Noise levels often exceed 10 mm Hg peak-to-peak. External nonmeasured blood pressure disturbances such as bleeding and level of anesthesia can be substantially greater than the effects of sodium nitroprusside. This, coupled with the long and variable response delay and sensor noise, renders adequate patient modeling very problematic. In this hostile environment, proportional-integral-derivative control and simple adaptive control must be supplanted by sophisticated self-tuning or rule-based control. Robust control must regulate the controlled variable based on a model tailored to the given patient via estimated parameters. But regulation impoverishes the information in the drive signals used for estimation, making estimation more difficult. A controller should aggressively correct a real safety problem but be timid where a mistake might produce a catastrophe. Deciding which is which is tricky, as is the inculcation of other common sense behaviors. Self-tuning controllers as well as rule-based controllers involve very large sets of parameters. Their design currently proceeds via simulation. Sebald and colleagues use a minimax strategy for controller design, picking that controller whose worst performance across all patients is better than that of all other controllers. Two sorts of tools are essential for pursuing this strategy: (1) good system identification procedures (methods of model parameter determination) and (2) means of solving the minimax problem. The latter is accomplished by selecting a rich controller structure and a flexible patient model and then using simulated annealing, a computational method for finding global extrema, to determine the controller parameters that minimize the maximal loss across a large group of simulated patients.

Questions for Dr Sebald

1. You refer to the need to build common sense behavior into the controller. Is not the capture of common sense one of the outstanding unsolved problems of artificial intelligence? How do you propose to reach that goal?
2. Will not the single minimax–optimal controller be too much of a compromise too much of the time? Why not define a bank of controllers?
3. How many parameters do your controllers have? Your patient model(s)? In the face of the contaminating effects you discuss, will reliable real-world identification of patient model parameters be feasible?
4. The research scenario you paint seems complex. Are there no ways around the complexity?

Dr Sebald Replies

1. It is true that the capture of common sense is an unsolved topic in artificial intelligence. There are two important aspects here. First, some degree of common sense is mandatory to gain acceptance from clinicians. They become extremely nervous when the controller continues to infuse sodium nitroprusside even though the mean arterial pressure is below target and dropping. Second, there is a large range of degrees of common sense. We are striving for better and better performance in this area. It is not difficult to incorporate rudimentary common sense. One can easily establish rules of behavior (e.g., do not permit additional infusion if pressure drops more rapidly than some threshold) to provide at least basic capability. Such rules commonly exist in safety shells watching over controllers.

2. Minimax is an optimization criterion, not a control strategy. The two are totally independent. Any control strategy (e.g., proportional integral derivative, PID) has parameters that can be chosen by a minimax criterion. Multiple model controllers can also be designed with the minimax criterion. Banks of controllers can provide advantages. However, two problems arise. First, if the decision strategy is poor (e.g., the standard MMAC strategy) these controllers exhibit tremendously aggressive tendencies and much transient activity [1]. Second, banks of controllers have relatively large numbers of design parameters. One must choose the

number of controllers, their set points, and the parameters of the decision mechanism. In fact, our work on minimax design tools was originally motivated by a desire to properly design multiple model controllers.

3. Typically, our controllers have 30 to 50 parameters. Our patient models typically have about 10 parameters. Although system identification in these problems is very challenging, excellent tools do exist. There is a great deal of art in proper system identification, but—and perhaps more importantly—experts exist. Collaboration with experts like Michael Parti, PhD, an econometrician, has convinced me that estimation will not prove to be the limiting factor.

4. I would claim that the proposed approach leads to conceptually simpler solutions. Good computer-aided design schemes provide more exhaustive testing and greatly facilitate proper parameter choices. This frees the controller designer to put more effort into more capable and perhaps simpler architectures. The tendency today is to build hybrid systems with safety shells, estimators, controllers, artifact detectors, and so on designed in more or less piecemeal fashion. My feeling is that better design tools will lead to simpler, more integrated designs. Furthermore, the minimax approach proposed here permits the controller designer to view the problem at the appropriate degree of complexity, that of choosing a control strategy that must work across a variety of patients. Given the computational ability to carry out the design details, such a world view actually helps simplify the design effort by helping to focus on the real problem in one picture.

Answers to General Questions

1. Progress is being made in this area on a variety of fronts. We have had success with both rule-based and neural network systems. However, artifact detection will probably always be a problem. For example, as we move to multiple drug infusions, and multiple variable sensors, new kinds of artifacts will probably arise. In contrast, the use of multiple variables will aid in the detection of artifact. For example, a doubling of the blood pressure without a change in any other variable almost certainly represents artifact.

2. Essentially, one should use the simplest architecture that can perform the desired task. Unfortunately, one cannot separate the performance impacts of the architecture type from those due to the explicit choices made in its implementation. One can usually solve a control problem with any of a variety of architectures, provided they all have the requisite intelligence and provided the idiosyncratic details have been properly worked out. Most practical control algorithms are hybrids, including safety shells and possibly even fuzzy control or neural net components. The real problem, it seems to me, is the relative lack of useful tools for computer-aided design and testing.

3. Multivariable control seems simply to present more complex design environments. There are more artifacts to worry about. There are more variable interactions to be explicitly estimated. The same design considerations apply. One still wants robustness to interpatient and intrapatient variation. One still wants safety and a degree of common sense. One still wants exhaustive testing results to preclude unpleasant surprises. The difference is that all of these things become more complicated in the multivariable environment. My conclusion is that computer-aided design tools will become even more critical in the multivariable case, but there is no reason to believe that multivariable control will prove to be an insurmountable problem.

4. Wittenmark and Astrom are absolutely correct. The real design environment cannot rely exclusively on theoretical results, since most do not apply to the environments described by them. The only way to make these things work reliably is through the intensive use of computer-aided design and testing.

5. I believe that autonomous control in the operating room will certainly be possible in the reasonably near term. There are no fundamental technical obstacles to be overcome. It will take time, however, to do the needed research. I am far less clear on whether autonomous control is a desirable goal from all important points of view, including medical, legal, and the patient's. It is a really nice technical problem, which will be solved if such a solution seems advantageous from the larger perspective.

6. Closed-loop control in these environments is a multifaceted problem. All of the components—estimation, control, computer-aided design/testing, models, and so on—will receive attention. I believe that much basic work needs to be done to get this field on a solid design footing, and I would emphasize tool development before expanding to a variety of new medical contexts. There are, however, many possible application areas for a reasonably mature technology of closed-loop biomedical control. I especially think that systems capable of use by ambulatory patients will prove very beneficial.

REFERENCE

1. Udomkesmalee S. Intelligent parallel adaptive controllers for uncertain systems. PhD thesis, University of California, San Diego, 1985

SUMMARY OF PAPER READ BY DR WESTENSKOW

Into an anesthesia work station have been integrated patient monitoring, alarms, and closed-loop regulation of total breathing circuit volume, inspired oxygen concentration, and end-tidal anesthetic concentration. The required flow rates of oxygen and anesthetic agents then yield the associated uptakes. Fixed total breathing circuit volume is maintained by controlling the ventilator bellows position at end-expiration. Oxygen control is achieved via a classical proportional-integral-derivative controller. Anesthetic agent delivery is maintained by controlling the end-tidal concentration of gases.

In both a dog study and a clinical study, the system maintained the inspired oxygen concentration within ±0.2 vol%. Enforcing a step change took one minute. Steady-state oxygen flow was found to measure a patient's oxygen consumption to within ±6%. The anesthesia controller was tested with an oil lung model, "inducing" twenty-four times with an initial end-tidal concentration of 1 vol% halothane reached with little overshoot, on average, in 4.4 minutes.

Overall, the controller performed statistically better than experienced anesthetists in six of fourteen categories.

Questions for Dr Westenskow

1. Would you sketch how you established the ±6% accuracy in using oxygen flow rate to yield oxygen consumption (number of subjects, gold standard, and so forth)?
2. Is end-tidal anesthetic concentration a (the) sufficient measure for control, or are there plans to augment it?
3. As regards the overall performance of the controller, would you give all fourteen evaluative criteria used and compare automatic control and anesthetist performance for each?

Dr Westenskow Replies

1. The 6% accuracy was established by measuring the $\dot{V}O_2$ of 14 volunteers, first by the oxygen flow rate method, and second by a standard open circuit method using gas collection and mass spectrometer analysis [1]. Each volunteer was asked to breathe from the anesthesia circuit through a tightly fitted mouthpiece. After allowing 10 minutes for the volunteer to relax, his $\dot{V}O_2$ was measured by the oxygen flow rate technique. The volunteer's $\dot{V}O_2$ was next measured by an open circuit method using a 120-l Tissot spirometer for collection of exhaled air. Oxygen consumption rates ranged from 221 to 388 ml/min (standard temperature and pressure dry). Within this range the correlation coefficient was found to be 0.94 and the sample standard deviation from regression 17.7 ml/min. At the mean $\dot{V}O_2$ value of 275 ml/min this represents an error of ±6%.

2. Volatile anesthetic delivery can be controlled using sensors that measure a direct physiologic effect of the agent or, more simply, a sensor that measures the agent's end-tidal concentration. Direct physiologic monitoring of the electroencephalogram (EEG) has not been particularly useful, because the EEG changes only with very deep levels of anesthesia [2]. Blood pressure and heart rate have also been used to control volatile delivery, but rate and pressure are influenced by many physiologic factors other than anesthetic depth [3–5]. The end-tidal concentration of the volatile agent is easily measured and predicts the patient's reaction to surgical stimulation, the motor response, baroadrenergic reaction, cerebral (EEG) depression, and readiness for intubation [6–11]. Thus a desired physiologic effect can be achieved by controlling the volatile agent concentration.

3. Automatic control and manual performance have been compared during anesthesia by Smith and colleagues [12]. They found superior performance of the automatic controller in six of fourteen categories. We compared performance during induction of hypotension and found the computer brought the mean arterial pressure to the desired value in an average of 4.1 minutes compared with 6.3 minutes for the anesthesiologists. There was no statistically significant difference between the computer and the anesthesiologists for 25 performance criteria. The anesthesiologists performed better during two imposed hypertensive perturbations by turning off the sodium nitroprusside sooner when the pressure rapidly decreased. In general, the computer controller performed as well as experienced anesthesiologists who devoted full attention to the control of blood pressure [13].

Answers to General Questions

1. Many common occurrences external to and beyond the control of a controller may effect the validity of the mean arterial pressure feedback signal [14]. These factors include blood sampling, clotting of a catheter, improper seating of a pressure dome, air bubbles trapped in the pressure measurement line, and transducer or amplifier failure. The arterial pressure signal can be disqualified if the pulse pressure is too large or too small, if the waveform is abnormal, or if the heart rate is too fast or too slow. If the pressure signal becomes invalid during regulator control of mean arterial pressure, the regulator can be placed in a standby mode and the infusion rate held constant.

We tested a blood pressure controller during challenges imposed by dysrhythmias, respiratory therapy, hypovolemia, and blood transfusion [14]. Following the challenges, the controller returned the pressure to the desired value in less than 10 minutes. The controller properly rejected invalid pressure signals from a damaged catheter or pressure transducer as well as those caused by blood sampling and arterial line flushing. The controller rejected all invalid pressure signals during our testing.

2. PID controllers have been used very successfully to regulate oxygen concentration, circuit volume, and CO_2 in breathing systems [15]. Adaptive PID or self-tuning type controllers have been used for the control of the volatile anesthetic agent. Adaptive control provides an automatic means whereby the controller gains are adjusted, depending on the application. Controller learning or self-tuning can take place during anesthesia induction [5,16]. A disadvantage of this approach is that a learning period is required, during which the measurements of end-tidal anesthetic agent must be relatively noise free. In the operating room, it is difficult to find this noise-free period, particularly during induction. Model-based controllers, which do not need a learning period, need information, some of which is unknown, about the patient (oxygen consumption, breathing circuit size, volume, etc) [17]. It appears better to use stable fixed gain controllers, and to use a blower in the breathing system to cause uniform mixing and short transport delays, and to reduce patient-to-patient variability [18–22].

REFERENCES

1. Instrumentation for measuring continuous oxygen consumption of surgical patients. IEEE Transactions on Biomedical Engineering BME 1977;24:331–337
2. Coles R, Brown WA, Lampard DG. Computer control of respiration and anesthesia. Med Biol Eng 1973;11:262
3. Lampard DG, Coles JR, Brown WA. Electronic digital computer control of ventilation and anaesthesia. Anaesth Intensive Care 1973;1:382–392
4. Suppan P. Feed-back monitoring in anaesthesia II: pulse rate control of halothane administration. Br J Anaesth 1972;44:1263
5. Fukui Y, Smith NT, Fleming RA. Digital and sampled-data control of arterial blood pressure during halothane anesthesia. Anesth Analg 1982;61:1010–1015
6. Zbinden AM, Frei F, Westenskow DR, et al. Control of end-tidal halothane concentration. Part B: Verification in dogs. Br J Anaesth 1986;58:563–571
7. Smith NT, Schwede HO. The response of arterial pressure to halothane: a systems analysis. Med Biol Eng 1972;10:207
8. Saidman LJ, Brandstater B. Minimum alveolar anesthetic concentration: a standard of anesthetic potency. Anesthesiology 1965;26:756
9. De Jong RH, Eger EI II. MAC expanded: AD_{50} and AD_{95} values of common inhalation anesthetics in man. Anesthesiology 1975;42:620
10. Tinker JH, Sharbrough FW, Michenfelder JD. Anterior shift of the dominant EEG rhythm during anesthesia in the Java monkey. Anesthesiology 1977;46:252
11. Yakaitis RW, Blitt CD, Angiulo JP. End-tidal halothane concentration for endotracheal intubation. Anesthesiology 1977;47:386
12. Smith NT, Quinn ML, Flick J, et al. Automatic control in anesthesia: a comparison in performance between the anesthetist and the machine. Anesth Analg 1984;63:715–722
13. Westenskow DR, Meline L, Pace NL. Controlled hypotension with sodium nitroprusside: anesthesiologist versus computer. J Clin Monit 1987;3:80–86
14. Meline LJ, Westenskow DR, Somerville A, et al. Evaluation of two adaptive sodium nitroprusside control algorithms. J Clin Monit 1986;2:79–86
15. Westenskow DR. Automating patient care with closed-loop control. MD Computing 1986;3:14–20
16. Tatnall ML, Morris P, West PG. Controlled anaesthesia: an approach using patient characteristics identified during uptake. Br J Anaesth 1981;53:1019–1026
17. Vishnoi R, Roy RJ. Adaptive control of closed circuit anesthesia. IEEE Eng Med Biol Soc 10th Intl Conference, 1988
18. Richie RG, Ernst EA, Pate BL, et al. Closed-loop control of an anesthesia delivery system: development and animal testing. IEEE Trans Biomed Eng 1987;34:437–443
19. Westenskow DR, Zbinden AM, Thomson D, et al. Control of end-tidal halothane concentration. Part A: Anaesthesia breathing system and feedback control of gas delivery. Br J Anaesth 1986;58:555–571
20. Furst SR. A computer-controlled anesthesia delivery system. Master's thesis, University of Utah, 1985
21. Hayes JK, Westenskow DR, East TD, et al. Computer-controlled anesthesia delivery system. Med Inst 1984; 18:226–231
22. Schepp R. Lung volume measurement during closed circuit anesthesia. Acta Anaesthesiol Belg 1984;35:295–305