

Computer Control of Mechanical Ventilation

Robert L Chatburn RRT FAARC

Introduction

Computerized Output Control

Open-Loop Control

Closed-Loop Control

The Hierarchy of Ventilator Control Systems

Summary of Control Systems

Future Possibilities

Computerized Input Control

Summary

Computer control of mechanical ventilators includes the operator-ventilator interface and the ventilator-patient interface. New ventilation modes represent the evolution of engineering control schemes. The various ventilation control strategies behind the modes have an underlying organization, and understanding that organization improves the clinician's appreciation of the capabilities of various ventilation modes and gives an idea of what we can and should expect for the future. The operator-ventilator interface has received little attention in the literature, despite the fact that there is a whole science of human-computer interaction. This report suggests a methodology for the study of ventilator interfaces.

Key words: computers, ventilators, ventilation, user-computer interface, fuzzy logic, neural networks, artificial intelligence. [Respir Care 2004;49(5):507–515. © 2004 Daedalus Enterprises]

Introduction

Computers have played a major role in the evolution of the mechanical ventilator. Microprocessor technology has revolutionized ventilator-patient interaction by allowing advanced ventilation modes. It has also made the graphical user interface a standard feature of ventilator-operator interaction. This report discusses both the ventilator-patient interface and ventilator-operator interface and offers some predictions.

Robert L Chatburn RRT FAARC is affiliated with the Respiratory Care Department, University Hospitals of Cleveland, and with the Department of Pediatrics, Case Western Reserve University, Cleveland, Ohio.

Robert L Chatburn RRT FAARC presented a version of this report at the 33rd RESPIRATORY CARE Journal Conference, Computers in Respiratory Care, held October 3–5, 2003, in Banff, Alberta, Canada.

Correspondence: Robert L Chatburn RRT FAARC, Respiratory Care Department, University Hospitals of Cleveland, 11100 Euclid Avenue, Cleveland OH 44106. E-mail: robert.chatburn@uhhs.com.

The 2002 RESPIRATORY CARE Journal Conference on mechanical ventilation had a great deal to say about ventilator-patient interaction. In particular, the new ventilation modes were described in detail.¹ I don't intend to rehash that information. Rather, I intend to put it into perspective by introducing a new conceptual framework. Ventilation modes are often described as unique and unrelated developments, but they actually represent the evolution of engineering control schemes. The various control strategies behind the modes have an underlying organization. Understanding the hierarchy of control types will improve our appreciation for the modes' capabilities and give us some idea of what we could and should expect for the future.

Computerized Output Control

A ventilator must have a control circuit to manipulate pressure, volume, and flow. The control circuit measures and directs the ventilator's output. Early ventilators used simple mechanical components for control systems. Today

microprocessors allow much more accurate and flexible control of breathing variables, which has led to a wide array of ventilation modes.

Open-Loop Control

To lay the groundwork for discussing how ventilators operate I will review some basic control theory. The simplest type of control is called open-loop. Its advantage is low cost. Its weakness is that it is unable to cope with disturbances in the system. About the only example of open-loop control in mechanical ventilation was the early jet ventilator. Though the operator could set the input driving pressure, the pressure and flow delivered to the patient were highly variable and depended on the changing respiratory system impedance. Figure 1A is a diagram of an open-loop system. Shown are 3 subsystems connected in series: a controller, an effector, and a plant. The plant is the subsystem being controlled. Its output is the controlled or output variable. The effector is the mechanism that drives the plant to respond in a given way. Its output is the variable that is manipulated to control the behavior of the controlled system. We will refer to it as the manipulated variable (it is called the control variable in relation to mechanical ventilators). The effector, typically, is the prime mover or the device that drives or powers the controlled system. In a ventilator, the effector might be a piston pump or an electronic flow control valve.

The control circuit or controller contains the logic used to interpret or translate the input signal into a signal to which the effector responds. Figure 1B shows the controller and effector together within the ventilator. One of the few examples of open-loop control of mechanical ventilation is the type of jet ventilator used experimentally in the early 1980s.^{2,3} The operator could set a driving pressure, and the controller would turn a valve on and off at a set frequency and inspiratory-expiratory ratio. Gas was metered to the patient, but the actual pressure and volume

delivered were dependent on the moment-to-moment changes in the patient's respiratory system impedance.

Thus, an open-loop control system cannot correct for disturbances in the conditions affecting the controlled plant. It "goes on its merry way," oblivious of its surroundings. Disturbances are influences on the system that make the output unpredictable. During mechanical ventilation the major disturbances are changes in the patient's ventilatory drive, respiratory system mechanics, and leaks.

Closed-Loop Control

All modern ventilators use closed-loop control to maintain consistent pressure and flow waveforms in the face of changing environmental conditions. Closed-loop control is accomplished by using the output as a feedback signal that is compared to the operator-set input. The difference between the two is used to drive the system toward the desired output. For example, pressure-controlled modes use airway pressure as the feedback signal to control gas flow from the ventilator. Manufacturers typically do not use flow at the airway opening as a feedback signal, because they do not trust the flow sensors available for that purpose. Instead, they measure flow inside the ventilator, near the main flow-control valve.

Closed-loop control (also called feedback control) uses a sensor to measure the output of the effector. This signal is passed to a comparator (represented by the circles in Fig. 2) that essentially applies a simple equation: $\text{error} = \text{input} - \text{output}$. If the error in the effector output is large enough, an error signal is sent to the controller. The controller then adjusts the effector so its output is closer to the desired input (ie, makes the error smaller). The advantage of closed-loop control is that the output is continuously and automatically adjusted so that disturbances are not a problem. The greater complexity of that system makes it more expensive to build and maintain.

A feedback signal can be electrical (eg, from an electronic pressure transducer) or mechanical (eg, pressure regulators and continuous positive airway pressure valves). In mechanical devices a spring provides the input setting, and the position of the diaphragm (a measure of the gas pressure) is the feedback signal. When the force caused by the pressure exceeds the spring load, the diaphragm deflects and vents gas to the atmosphere to relieve the pressure.

The Hierarchy of Ventilator Control Systems

The basic concept of closed-loop control has evolved into at least 7 different ventilator control systems (set-point, auto-set-point, servo, adaptive, optimal, knowledge-based, and neural network control). These control types are the foundation that makes possible several dozen seem-

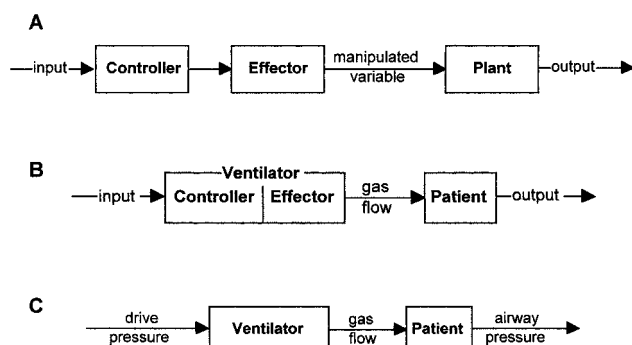


Fig. 1. Schematic diagrams of open-loop control of a mechanical ventilator. A: Basic control circuit. B: Open-loop control circuit for a ventilator. C: Example of open-loop control of a jet ventilator.

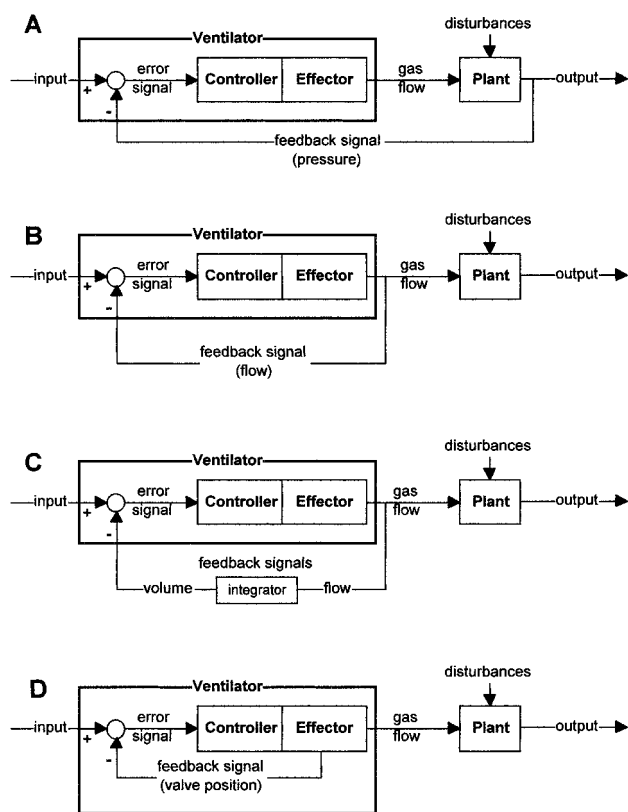


Fig. 2. Schematic diagrams of closed-loop control of a mechanical ventilator. A: Pressure control. B: Flow control. C: The flow signal is integrated to provide a signal for volume control. D: Flow/volume control using a calibrated gas-control valve instead of a flow sensor.

ingly different ventilation modes, but once we understand how these control types work, many of the seeming differences are seen to be similarities. With that understanding we can transcend the confusion that can arise out of ventilator marketing hyperbole and we begin to appreciate the true clinical capabilities of various ventilators.

Set-Point Control. All ventilators use at least set-point control (Fig. 3). In set-point control the output is constrained to match a constant input (ie, a set maximum pressure or flow value). This makes possible the standard volume- (actually flow) or pressure-controlled breaths. The operator sets either a fixed pressure or a flow limit, and the ventilator then maintains a consistent pressure or flow wave-




Fig. 3. Set-point control.


form output. (Recall that the term “limit” means that a control variable reaches a preset maximum value before inspiration ends.⁴) This type of control is similar to the “cruise control” function of an automobile.

Auto-Set-Point Control. Auto-set-point control is a more advanced version of set-point control. It gives the ventilator the decision of whether the breath will be flow-controlled or pressure-controlled, according to the operator-set priorities. The breath can start out as pressure-controlled and automatically switch to flow-controlled (eg, the Bird VAPS [volume-assured pressure support] mode⁵) or vice versa (eg, the Dräger Pmax mode).

Servo Control. Whereas set-point control attempts to maintain a constant output to match a constant input, servo control tracks a moving input, much like power steering in an automobile. Servo control was developed during World War II to aim ship’s guns and radar equipment. Servo control makes possible the proportional assist mode,⁶ in which the ventilator’s output follows and amplifies the patient’s own flow pattern. The ventilator can thus support the abnormal load imposed by disease while the patient’s muscles handle the normal load of the respiratory system’s normal resistance and compliance (Fig. 4).

Adaptive Control. Adaptive control means automatic adjustment of one set-point to maintain a different operator-selected set-point (Fig. 5). One of the first examples of a mode that used adaptive control was pressure-regulated volume control on the Siemens Servo ventilator. Adaptive control is an evolutionary step because it permits the ventilator to determine a set-point level independent of the operator. Set-point control operates *within* breaths, whereas adaptive control introduces another feedback loop that operates *between* breaths. It is that second feedback loop that prompted the term “dual control.” The feedback of exhaled tidal volume allows the ventilator to adapt to changes in the patient’s lung mechanics. Despite having various names for the specific modes it allows, adaptive control to date has been implemented as a way for the ventilator to automatically adjust the pressure limit of a breath to meet an operator-set volume target over several breaths. Notice that the operator’s influence has subtly moved away, in a sense, from direct control of the breath.

Optimal Control. Optimal control takes adaptive control a step further by allowing the ventilator to set both volume and pressure set-points (Fig. 6). Optimum control takes its name from the fact that a mathematical model is used to find the best (in this case, minimum) value of some performance function. The Hamilton Galileo is the only ventilator with this feature. It allows the ventilator to make all subsequent adjustments after the operator sets the target



$$P_{mus} = Load_{normal} + Load_{disease}$$

$$P_{vent} = K_1 \times V + K_2 \times \dot{V}$$

$$P_{mus} + P_{vent} = Load_{normal} + Load_{disease}$$

Fig. 4. Servo control is the basis for the proportional assist mode. In this mode the operator sets targets for elastic and resistive unloading. The ventilator then delivers airway pressure in proportion to the patient's own inspiratory volume and flow. When the patient's muscles have to contend with an abnormal load caused by disease, proportional assist allows the operator to set amplification factors (K_1 and K_2) on the feedback volume and flow signals. By amplifying volume and flow the ventilator generates a pressure that supports the abnormal load, freeing the respiratory muscles to support only the normal load that is due to the normal elastance and resistance of the respiratory system. P_{mus} = pressure generated by the muscles. $Load_{normal}$ = respiratory system's load during normal resistance and compliance (no disease). $Load_{disease}$ = extra load caused by disease. P_{vent} = pressure generated by the ventilator. V = volume. \dot{V} = flow.

minute volume. Again the operator is moved another step away from direct control of the breath.

An experimental form of optimal control gives the ventilator even more authority.⁷ The exhaled carbon dioxide signal allows the ventilator to estimate the patient's minute volume needs (Fig. 7). With this control system the oper-

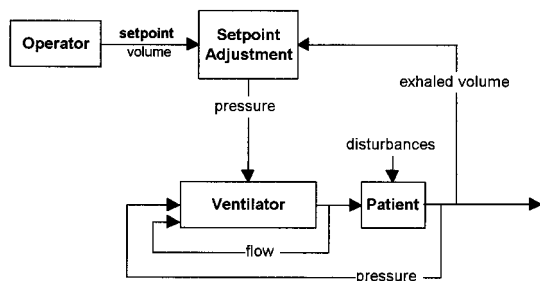


Fig. 5. Adaptive control. Notice that the operator has stepped back from direct control of the within-breath parameters of pressure and flow. Examples of adaptive control are pressure-regulated volume control (PRVC) on the Siemens ventilator and Auto-flow on the Dräger Evita 4 ventilator.

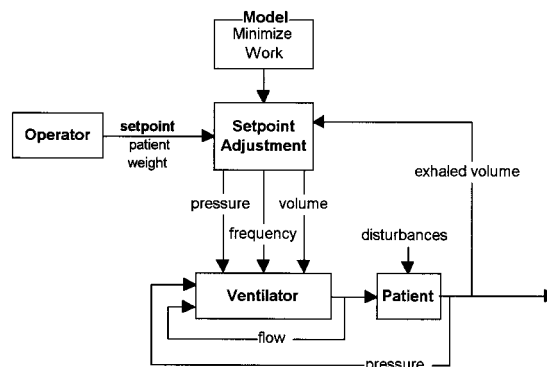


Fig. 6. Optimal control. A static mathematical model optimizes some performance parameter, such as work of breathing. The only commercially available form of optimal control is the adaptive support ventilation (ASV) mode on the Hamilton Galileo ventilator.

ator has stepped completely out of the picture, though, of course, we are only talking about eliminating the operator in the sense of establishing the ventilatory pattern. Operator input is still required for all the other variables, such as positive end-expiratory pressure, fraction of inspired oxygen (F_{IO_2}), and alarm settings.

Knowledge-Based Control. Knowledge-based control is yet another evolutionary step; it gives the ventilator more knowledge than what can be contained in a simple, static, mathematical model. In fact, knowledge-based control attempts to capture the experience of human experts and thus expand the scope of control to potentially all of the ventilation mode variables. An experimental application of this type of control has been described for automatic adjustment of pressure support.⁸ An even more sophisticated approach coupled a knowledge base with fuzzy logic (Fig. 8).⁹ In that control system the ventilator used

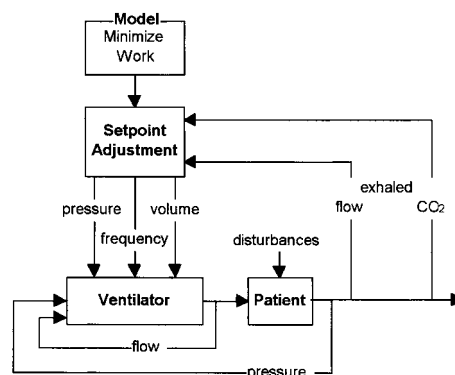


Fig. 7. An experimental form of optimal control allows the ventilator to estimate the patient's minute ventilation needs based partly on the exhaled carbon dioxide signal. This eliminates the need for the operator to set any of the major breath parameters. Set-points for fraction of inspired oxygen and positive end-expiratory pressure are still required, however.

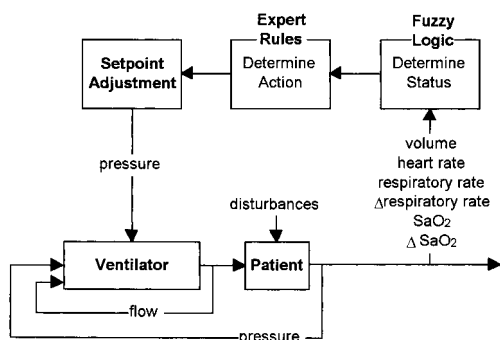


Fig. 8. A knowledge-based control system for automatically adjusting pressure support levels. ΔS_{aO_2} = change in arterial oxygen saturation.

both instantaneous measurements of physiologic values (eg, respiratory rate and oxygen saturation) and their rates of change. Fuzzy logic¹⁰ was used to integrate the measurements with predefined ranges of values representing the patient status. Once the patient's status was determined, appropriate expert rules were selected from a lookup table and used to adjust the ventilator. Though this was a limited application, it proved the concept.

No doubt the most convincing proof of concept was presented by East et al.¹¹ They used a rule-based expert system for ventilator management in a large, multicenter, prospective, randomized trial. Though patient survival and duration of stay were not different between human-controlled and computer-controlled ventilator management, computer control was associated with significantly less multi-organ dysfunction and lower incidence and severity of lung overdistention injury. The most important finding, however, was that expert knowledge can be encoded and successfully shared with institutions that had no input into the model. Note that the expert system did not directly control the ventilator but rather made suggestions for the human operator. In theory, of course, the operator could be eliminated.

Artificial Neural Network Control. The ultimate in ventilator control to date is the artificial neural network (Fig. 9).¹² Again, this experimental system does not di-




Fig. 9. Artificial neural network control. PEEP = positive end-expiratory pressure. F_{IO_2} = fraction of inspired oxygen.

rectly control the ventilator but acted as a decision-support system. In the Snowden et al report,¹² the comment I found most interesting was that the neural network was capable of learning, which offers substantial advantages over static rule-based systems.

Neural networks are essentially data modeling tools used to capture and represent complex input-output relationships. A neural network learns by experience, similar to the way a human brain does, by storing knowledge in the strengths of inter-node connections. As data modeling tools neural networks have been used in many business applications and in medical applications for both diagnosis and forecasting.¹³ A neural network, like a brain, is made up of individual "neurons." Signals (action-potentials) appear at the unit's inputs (synapses). The effect each signal has is approximated by multiplying the signal by a number (a weight) that indicates the "strength" of the synapse. The weighted signals are then summed to produce an overall unit activation. If this activation exceeds a certain threshold, the unit produces an output response. Large numbers of "neurons" can be linked together in layers (Fig. 10). The nodes in Figure 10 represent the summation and transfer processes. Each node contains information from all "neurons." As the network learns, the weights change and thus the values at the nodes change, affecting the final output.

There are 2 basic kinds of neural networks: supervised and unsupervised. A supervised network must be trained

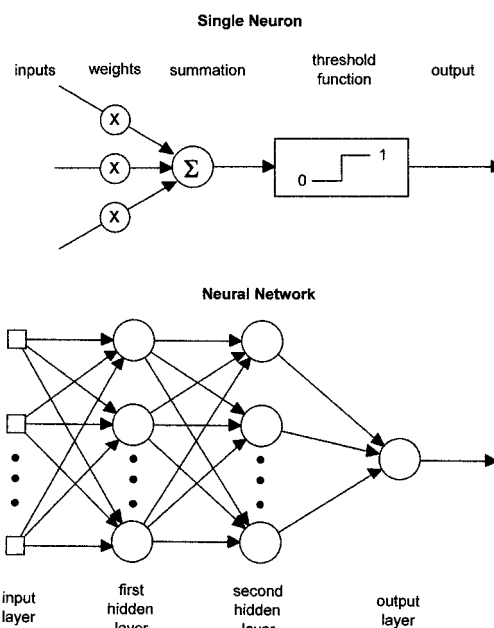


Fig. 10. Neural network structure. A single "neuron" accepts inputs of any value and weights them to indicate the strength of the synapse. The weighted signals are summed to produce an overall unit activation. If this activation exceeds a certain threshold the unit produces an output response. A network is made up of layers of individual "neurons."




Fig. 11. Neural network training.

by providing an input and comparing the output to the desired output. This feedback process is repeated until the output error is acceptable (Fig. 11). Unsupervised networks can train themselves.

Summary of Control Systems

In summary, ventilator control schemes have a hierarchy of evolutionary complexity. At the most basic level, control is focused on what happens within a breath. We can call this *tactical control*, with which there is a very direct need for operator input of static set-points. The next level up we can call *strategic control*, in which set-points are dynamic in that they can be automatically adjusted over time by the ventilator, according to some model of desired performance. The operator is somewhat removed in that, instead of individual breath control, the inputs are entered according to the model of desired performance and they take effect over several breaths. The highest control level so far is what might be considered *intelligent control*, in which the operator can be eliminated altogether. Not only dynamic set-points but dynamic models of desired performance are permitted. There is the possibility of the system learning from experience so that the control spans between patients (Fig. 12).

Future Possibilities

Of course, the real challenge in closed-loop control of ventilation is defining and measuring the appropriate feedback signals. If we stop to consider all the variables a human operator assesses, the problem looks insurmountable. Not only does a human consider a wide range of individual physiologic variables, there are also more abstract evaluations of such things as metabolic, cardiovascular, and psychological states. Add to that the various environmental factors that may affect operator judgment and we get a truly complex control problem (Fig. 13). I would like to speculate now about a response to that challenge.

The ideal control strategy would have to start out with basic tactical control of the individual breath. Next we

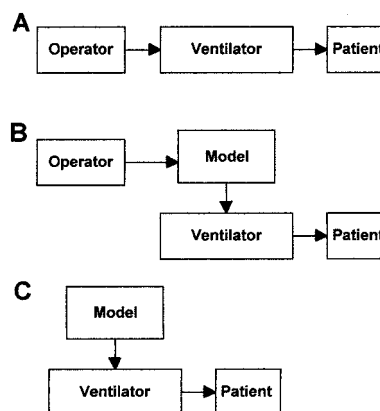


Fig. 12. Summary of control system hierarchy. A: Tactical control (within breaths) involves static set-points (set-point, auto-set-point, and servo). B: Strategic control (between breaths) involves dynamic set-points (static models, adaptive, and optimal). C: Intelligent control (between patients) involves dynamic set-points (dynamic models, ability to learn from experience, knowledge-based, artificial neural networks).

would add longer-term strategic control to adapt to changing load characteristics. Mathematical models could provide the mode's basic parameters, and expert rules would set limits to ensure lung protection. Next we would sample various physiologic variables and use fuzzy logic to establish the patient's immediate condition. That information would be passed on to a neural network that would

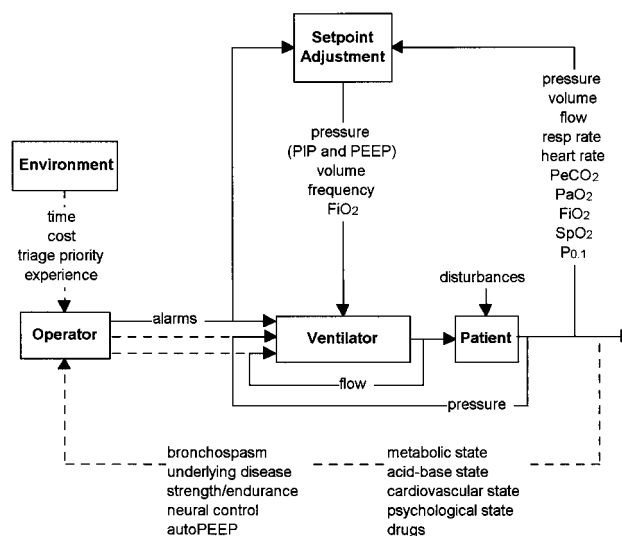


Fig. 13. The challenge of total computer control of mechanical ventilation. Solid arrows depict signals that have been used at least experimentally. Dotted arrows represent potential feedback signals. PIP = peak inspiratory pressure. PEEP = positive end-expiratory pressure. $F_{I_{O_2}}$ = fraction of inspired oxygen. P_{eCO_2} = partial pressure of carbon dioxide in exhaled gas. S_{pO_2} = arterial oxygen saturation measured via pulse oximetry. $P_{0.1}$ = airway occlusion pressure 0.1 s after the onset of inspiratory effort. autoPEEP = intrinsic positive end-expiratory pressure.

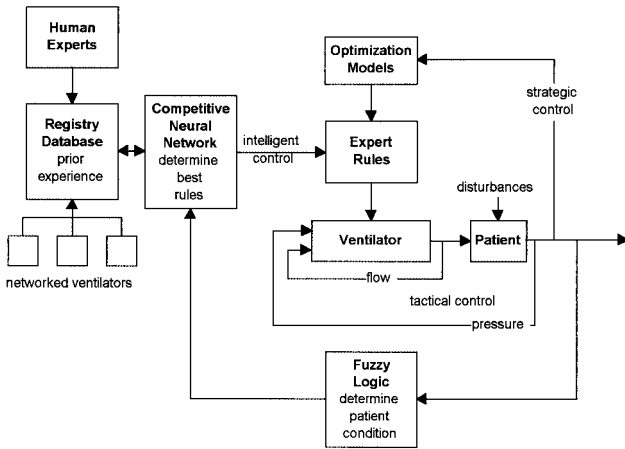


Fig. 14. A potential approach to the challenge of fully automated control of mechanical ventilation.

then select the best response to the patient’s condition. The neural network would have access to a huge database of expert rules and actual patient responses to various ventilation strategies. This arrangement would allow the ventilator to learn not only from its interaction with the current patient but also to learn from and contribute to the database. Finally, the database and the ventilator could be networked with other intelligent ventilators to multiply the learning capacity (Fig. 14).

Computerized Input Control

Let’s switch gears now and talk about the opposite end of the ventilator—the operator-ventilator interface. Computers have brought a high level of sophistication to ventilator input control. We have come a long way from using a crank to adjust the stroke of a ventilator’s piston to set tidal volume. The operator-ventilator interface must provide for 3 basic functions: to allow input of control and alarm variables; to monitor the ventilator’s status; and to monitor the ventilator-patient interaction status. Computers have automated the once manual procedures of producing waveform graphics and calculated patient status indicators (Table 1)

Ventilator manufacturers have a spectrum of choices when designing the operator-ventilator interface. At the low end, some ventilators still use hard-wired knobs and buttons in conjunction with digital light-emitting diode (LED) displays. Replacing the LEDs with a computer monitor costs more but adds a great deal of flexibility. Replacing the hard-wired knobs reduces the cost of future upgrades. The ideal situation in terms of flexibility and upgradeability is to use a touch-screen to create a “virtual instrument.” The virtual instrument concept has been quite popular among researchers and electrical engineers for many years. A software package called LabView is commonly used to replace a wide variety of stand-alone mon-

Table 1. Functions of the Operator-Ventilator Interface

Manipulate input control parameters
Mode parameters
Alarm thresholds
Monitor operational status
Mode status
Alarms status
Monitor ventilator-patient interaction
Control variable graphics (eg, pressure, volume, flow waveforms)
Calculated variables (eg, respiratory system mechanics)
Integrate adjunctive monitors (eg, oxygen saturation and exhaled carbon dioxide)

itors and displays. For example, the Ingmar Medical ASL5000 lung simulator’s interface is a virtual instrument based entirely on LabView. (By the way, ASL stands for “active servo lung” and it represents a control strategy that is essentially proportional assist ventilation in reverse.)

It amazes me that with all the research on mechanical ventilation there seems to be nothing written about the usability of the input screens. We, as users, just seem to take for granted the operator-ventilator interfaces the manufacturers have offered. This is particularly curious given that there is a huge industry built around the psychology of human-computer interaction.¹⁴ As I see it, this is a rich subject for future research and continued improvement. In my opinion, there are several problems with current ventilator screens. Some are too small, most are too cluttered, and all could better organize the information they present.

To facilitate future research I would suggest that we learn from studies of human-computer interaction, a field known as usability engineering. We can define usability as ease of use plus usefulness. An interface that has high usability is easy to learn, easy to remember (even after a period of absence), allows efficient use of the operator’s time, causes few errors, and is aesthetically pleasing.

Perhaps it is not a coincidence that there have been no studies of operator-ventilator interfaces. According to experts in the field, not only are manufacturers generally ignorant of usability design tools, but also the available methods for evaluating finished designs leave something to be desired. However, one prominent expert has proposed a testing method that might be appropriate for ventilator evaluation.¹⁵ It is composed of 3 fairly informal methods to identify problems and quantify their effects (Table 2).

Simplified thinking aloud is an easy method of getting user feedback. The basic idea is to observe users (one at a time) interacting with the ventilator as they perform typical operational procedures (eg, setting the mode). The user thinks aloud as he or she operates the ventilator. In formal settings videotapes are used, but good results can also be obtained by simply taking notes.

Table 2. 3 Tools for the Evaluation of Operator-Ventilator Interface Design

Simplified thinking aloud
Observe operators doing standard tasks
Heuristic evaluation
Evaluate product using rules of thumb for good design
Severity rating
Rank importance of design flaws

The key method is “heuristic evaluation”—the use of rules of thumb for good design to identify difficulties that users encounter. Research indicates that individual evaluators have substantial difficulty finding even major problems, so it is recommended that at least 3 evaluators be used.

There are hundreds of usability heuristics. Usability expert Jakob Nielson¹⁶ proposed a simplified list of 10 heuristics for general use. I have pared that list down to 5 rules that could be applied to ventilators (Table 3). Most importantly, the interface should clearly indicate the status of the ventilator’s operation, focusing on the mode parameters. The interface should use logical connections among ideas and provide definitions where applicable. It may be too much to expect that there will ever be a standard nomenclature among manufacturers beyond that imposed by regulatory agencies. Nevertheless, all the information the user needs to make the clinical decision should be available without having to resort to memory. As one design expert said, “. . . in the world of sales, if a company were to make the perfect product, any other company would have to change it—which would make it worse—in order to promote its own innovation, to show that it was different. How can natural design work under these circumstances? It can’t.”¹⁷

Once evaluators have identified potential problems, using heuristics, the problems can be collated on a list. The

Table 3. Heuristics for the Design and Evaluation of Ventilator Interfaces

Visibility of system status
Always keep the user informed of what is happening
Consistency and standards
Provide definitions and logical connections
Recognition and recall
User should not have to remember information from one screen to the next
Help and documentation
Make help easy to find, context-sensitive, and complete (include entire operator’s manual as an electronic document)
Minimalist aesthetic
Present only the information that is relevant

Table 4. Ordinal Scale for Ranking Severity of Problems*

0 = not a problem
1 = cosmetic problem
2 = minor problem
3 = major problem

* Use mean of at least 3 evaluators.

evaluators then review the list and rank the severity of the problems on a simple ordinal scale (Table 4). Outcome scores can then be easily calculated to identify needed design changes or to compare ventilators.

Summary

There is no doubt that the evolution of microprocessors has stimulated a parallel evolution in ventilator design. Computers have made the most tangible impact on the way new control systems are designed, leading to many new ventilation modes. But the promise of computer control lies in the power of the learning machine. Our slowness in improving outcomes may be caused by our lack of an organized way to learn from the experience of clinicians around the world. Powerful relational databases linked to learning ventilators that use fuzzy logic to categorize experience may be one solution. Such databases may provide enough information to create “virtual studies” by using resampling techniques¹⁸ instead of costly prospective, randomized, controlled trials. Or, perhaps more simply, we may never know exactly the reasoning that produces the output we desire but only which neural networks seem to be most successful.

As for the operator-ventilator interface, there is much room for improvement in usability. Just take a look at the interfaces of some of the most successful software tools, or even computer games, to see what might be possible. With the reduction in staff development resources we have experienced in this economy, we need to take full advantage of the teaching ability of the computers that are built into ventilators. I think we can expect to see great and exciting developments in the near future.

REFERENCES

1. Branson RD, Johannigman JA, Campbell RS, Davis K Jr. Closed-loop mechanical ventilation. *Respir Care* 2002;47(4):427–451; discussion 451–453.
2. Carlo WA, Chatburn RL, Martin RJ, Lough MD, Shivpuri CR, Anderson JV, Fanaroff AA. Decrease in airway pressure during high-frequency jet ventilation in infants with respiratory distress syndrome. *J Pediatr* 1984;104(1):101–107.
3. Branson RD, Hurst JM, DeHaven CB. Use of high frequency jet ventilation during mechanical hyperventilation for control of elevated intracranial pressure: a case report. *Respir Care* 1984;29(12):1221–1224.

4. Chatburn RL. Fundamentals of mechanical ventilation. Cleveland Heights Ohio: Mandu Press Ltd; 2003:31.
5. Amato MB, Barbas CS, Bonassa J, Saldiva PH, Zin WA, de Carvalho CR. Volume-assured pressure support ventilation (VAPSV): a new approach for reducing muscle workload during acute respiratory failure. *Chest* 1992;102(4):1225–1234.
6. Younes M. Proportional assist ventilation, a new approach to ventilator support. *Theory. Am Rev Respir Dis* 1992;145(1):114–120.
7. Laubscher TP, Frutiger A, Fanconi S, Brunner JX. The automatic selection of ventilation parameters during the initial phase of mechanical ventilation. *Intensive Care Med* 1996;22(3):199–207.
8. Dojat M, Harf A, Touchard D, Lemaire F, Brochard L. Clinical evaluation of a computer-controlled pressure support mode. *Am J Respir Crit Care Med* 2000;161(4 Pt 1):1161–1166.
9. Nemoto T, Hatzakis GE, Thorpe CW, Olivenstein R, Dial S, Bates JH. Automatic control of pressure support mechanical ventilation using fuzzy logic. *Am J Respir Crit Care Med* 1999;160(2):550–556.
10. Bates JHT, Hatzakis GE, Olivenstein R. Fuzzy logic and mechanical ventilation. *Respir Care Clin N Am* 2001;7(3):363–377.
11. East TD, Heermann LK, Bradshaw RL, Lugo A, Sailors RM, Ershler L, et al. Efficacy of computerized decision support for mechanical ventilation: results of a prospective multi-center randomized trial. *Proc AMIA Symp* 1999;251–255.
12. Snowden S, Brownlee KG, Smye SW, Dear PR. An advisory system for artificial ventilation of the newborn utilizing a neural network. *Med Inform (Lond)* 1993;18(4):367–376.
13. Gottschalk A, Hyzer MC, Greet RT. A comparison of human and machine-based predictions of successful weaning from mechanical ventilation. *Med Decis Making* 2000;20(2):160–169.
14. Clemmensen T. What happened to the psychology of human-computer interaction? A review of 17 years of research presented in one journal. In: Khalid HM, Lim TY, Lee NK, editors. *Proceedings of SEAMEC 2000, Kuching*. Available at <http://www.cbs.dk/staff/clemmensen/SEAMEC%202003%20Clemmensen.pdf>. Accessed February 27, 2004.
15. Nielsen J. Usability engineering. San Diego: Morgan Kaufmann, an imprint of Academic Press; 1993:17–20.
16. Nielson J. Ten usability heuristics. http://www.useit.com/papers/heuristic/heuristic_list.html. Accessed February 26, 2004.
17. Norman DA. *The design of everyday things*. New York: Doubleday; 1988:145.
18. Simon JL, Bruce P. Resampling: a tool for everyday statistical work. *Chance* 1991;4(1):22–23.

Discussion

MacIntyre: I want to discuss the control mechanisms of ventilators. It seems to me that the Achilles heel of this whole concept is the “expert rules” and where they come from. To be honest, I don’t think we as clinicians know what we want. I don’t think we know what the desirable output is.

In an exercise I did with a consensus group a few years ago,¹ we came up with the idea that there are only 4 things we’re worried about in mechanical ventilation: P_{O_2} , pH, plateau pressure, and F_{IO_2} . That’s it! We’re trying to maximize P_{O_2} and pH and, on the other hand, protect against unnecessarily high F_{IO_2} and plateau pressures, but we found trouble when we went around the room at that meeting and tried to rank “equi-toxic” levels of hypoxia, acidosis, supplemental oxygen, and plateau pressure. What was equi-toxic to an F_{IO_2} of 100%? Was it a plateau pressure of 35 cm H_2O ? Was it a P_{O_2} of 55 mm Hg? The thing that struck me in that exercise was that everybody had a different idea of what the proper balances were.

The problem with computer-controlled systems based on expert-rules is that one clinician might have one

set of expert rules that he or she thinks is appropriate, while another physician or therapist team might have a substantially different approach to balancing P_{O_2} , pH, plateau pressure, and F_{IO_2} . I think neural networks can get hopelessly confused, depending on who is setting up the desirable output variables.

The problem is that it’s not the *individual patient* that we learn the most from. We can learn a few things from an individual patient, but the *important* learning process in trying to figure out how to balance those 4 variables comes from large outcomes studies with large populations. I’m concerned that the only way to do that would be to network multiple institutions and see how various approaches work on a much larger scale, and to use important clinical outcomes such as long-term lung damage and mortality.

REFERENCE

1. Slutsky AS. Mechanical ventilation. American College of Chest Physicians’ Consensus Conference. *Chest* 1993;104(6):1833–1859; *Respir Care* 1993; 38(12):1389–1422.

Chatburn: As usual, you’ve hit the nail right on the head. But I would argue that we really don’t need to know a priori what the optimum scheme is. Take an analogous example: there was a great deal of interest in using artificial intelligence to create chess programs to play against humans. People had different strategies on how to do that. Ten years ago nobody believed that a computer could beat the top-ranking human chess player. Now it’s obvious that that can be done. The top-ranking chess player has stated it won’t be long before no human can beat a computer. Computer programmers didn’t arrive at that kind of evolution by sitting down and saying, “This is the single best way for a computer to play chess against a human.” A bunch of different people tried what they thought would be best and then competed.

So we should have several different groups of experts decide how they think it should be done and then let them “go at it” and see what the outcomes are. If we insist on huge multicenter trials, I don’t think it’s ever going to happen, because that would be expensive and take too long. But what if we just did a virtual experiment? Suppose we had a huge data

base of patients' specific reactions to specific ventilation control strategies. Then we could use a statistical procedure called resampling, in which, essentially, you take your database and sample with replacement, such that you have a virtually infinite sample size. Then you can at least approach that way of learning from your experience without having to know beforehand which is the best way to go.

MacIntyre: I would agree. You don't have to do this in a formal, prospective trial. But I think you and I would agree that the way to do this is to get a *huge* sample with meaningful outcomes such as death and lung injury, which you could do retrospectively and then design your expert systems.

Chatburn: It's fascinating and I just can't wait to see what the future holds because we're going *somewhere*. That's for sure.

MacIntyre: I'm also interested in your thoughts on alarm systems and feedback control. There are so many false alarms that some of the real alarms are ignored because they're lost in the noise of false alarms.

Chatburn: I didn't have time to go into that and I know you've written a lot on the subject. The problem relates partly to our passive approach to manufacturers: we just take what they give us; we don't demand intelligent controls and intelligent alarms. There's no reason that alarms can't learn in the same way that control theories have. I don't know of much research on developing intelligent alarm systems based on currently available engineering control theory, at least not in our industry.

Gardner: Regarding alarms, at my institution we use an electronic medical information bus to gather alarm data (such as a ventilator disconnect) from the ventilators.¹ Some alarms

sound alike so we're designing something that will flash every screen in the unit. The issue of alarm setting on bedside monitors is really problematic. Are ventilator monitors any better?

REFERENCE

1. Gardner RM. Computerized clinical decision-support in respiratory care. *Respir Care* 2004;49(4):378-386; discussion 386-388.

Chatburn: You mean stand-alone monitors?

Gardner: Yes. What is the false alarm rate on ventilator monitors?

Chatburn: It's high. I think the design problem is that there is so much noise in the alarm signal that you can't tell if you even have a signal or if it's a signal that you should be paying attention to, let alone how to process it and what the signal means. I think that's the key to the issue. To my knowledge there's no good way of addressing that problem at this point.

Gardner: I love your idea for evaluating the devices. I encourage you to go home and get those 4 experts and get 5 ventilators and do that, and publish the results, which I think would impact the manufacturers. It's a wonderful idea. In recent meetings I've been at people have said, "How did Microsoft Word get so widely accepted? I think it has a terrible user interface but almost everyone uses it."

Belda: I have a comment about the designing of the operator-ventilator interface. You mentioned Jakob Nielson (<http://www.useit.com>), who is a "usability guru," and I agree that he does fantastic work on user interface design and defining requirements for applications. One of the approaches he advocates for designing user interfaces and incorporating user-feedback is an approach towards ultimate simplicity. He urges designers to draw out on paper the interface's various elements,

whether those elements are knobs or buttons. Then he suggests having users push the mock-up buttons and follow through a branching logic scenario of, "Turn this button and this is what happens." Do you think that approach could work for getting manufacturers to reconsider designs based on users' needs and requests?

Chatburn: That brings to mind stories I've read about Jeff Hawkins, the creator of the Palm Pilot. He built a model out of cardboard, drew little buttons on it, and carried it around in his pocket all day. He'd whip it out and use it as if it were real, so he got a real user-interface evaluation by first-hand experience.

Obviously, that was a highly successful design approach, so perhaps what we can do in terms of scientific research is just clear the slate and say, "All right, forget what we know about what's available on the market today. What would we want in an ideal ventilator?" We could ask experienced clinicians, "Show me what you would use at the bedside. What would you like to have? What would be ideal?" and have them brainstorm. You could probably even do it on a spreadsheet on a computer or a slide show. You could really simulate to a great degree what the user would like to have. I think that would be a fascinating study. But if we don't communicate what we want to the manufacturers, they'll just keep making what they think they can sell.

Pierson*: In your discussion about ways to control mechanical ventilation and to have the machine measure and monitor things and adjust itself, and in all of the models you showed us, those were all mechanical and physiologic measurements. At 2 previous journal conferences John Han-

* David J Pierson, MD FAARC, Editor in Chief, *RESPIRATORY CARE* Journal, Seattle, Washington.

sen-Flaschen made an impassioned appeal for what he calls “patient-centered mechanical ventilation,”^{1,2} which he believes should be built into our routine at the bedside. Every time the ventilator is checked or any kind of bedside assessment is made, the patient should be asked, “Are you short of breath?” The patient’s experience of being ventilated, although it isn’t necessarily measurable in physiologic variables, may be pretty important. I would think that incorporating the patient’s subjective responses to ventilator adjustments would be really important, but also pretty difficult to do.

REFERENCES

1. Hansen-Flaschen JH. Palliative home care for advanced lung disease. *Respir Care* 2000;45(12):1478–1486; discussion 1486–1489.
2. Hansen-Flaschen J. Chronic obstructive pulmonary disease: the last year of life. *Respir Care* 2004;49(1):90–98.

Chatburn: I would say it’s probably not possible yet. But I would ask you whether he said how frequently that needs to be done? Hourly? Every 2 hours? Every 4 hours? Every time you check the ventilator? I think that would be the key. We’re never going to get rid of the operator entirely.

Pierson: Dr Hansen-Flaschen thinks it should be done every time *anything* is done that interacts with the ventilator, because sometimes patients are in substantial distress and he thinks that even another minute of that is undesirable. So, ideally, I wish that *this* monitor or some other neurophysiologic device could measure something that correlated fairly well with patient distress and that that could be incorporated into this as well. Because we know from the ARDS [acute respiratory distress syndrome] Network and other data that better physiologic measurements don’t necessarily correlate with better outcomes. Certainly nitric oxide and prone positioning and other things that improve P_{O_2} don’t seem to improve survival. In the ARDS Network study the patients in the control group, who received larger tidal volume, had better P_{O_2} and higher P_{aO_2}/F_{IO_2} ratios and therefore would have been judged as doing better during the initial days of their management, but they had a 20% lower survival rate.¹ It’s the problem of what you measure and whether it is the *right* thing to measure.

REFERENCE

1. The Acute Respiratory Distress Syndrome Network. Ventilation with lower tidal vol-

umes as compared with traditional tidal volumes for acute lung injury and the acute respiratory distress syndrome. *N Engl J Med* 2000;342(18):1301–1308.

Chatburn: It’s also a problem with paradigm-shift, too, because if we decide to check the ventilator every 2 hours because we have to manually make these simple within-breath adjustments, it becomes natural to say, “Well, while I’m doing that, I should also assess how the patient feels.” But if the ventilator is making minor adjustments by itself, then maybe we need to step back and say we need to do patient assessment uncoupled from what we do with the ventilator. Set the standard for how often to do patient assessments based on the needs of the patient and set the standard for how often to check the ventilator based on its capabilities.

Pierson: Patient distress is a complex thing that can be produced not only by how the ventilator is set, but also by whether you have the sedation right and various other things. So incorporation of some assessment of patient distress would have to be integrated with systems beyond just the ventilator.

Chatburn: Exactly.