

EXTENDING CLOSED-LOOP CONTROL TO THE MANAGEMENT OF CHRONIC DISEASE

WILLIAM W. STEAD, M.D., and (*by invitation*), WILLIAM M. GREGG, M.D., M.P.H.,
and JIM N. JIRJIS, M.D.

NASHVILLE, TENNESSEE

ABSTRACT

A closed-loop control process assures that a system performs within control limits by direct feedback of the system's output to change the system's inputs. We developed methods for the closed-loop control of system-based practice, using ventilator management as a model or test bed. The control system has three components: 1) an explicit end-to-end plan; 2) a record of what is done as it is done; and 3) an instant display of the status of each patient against the plan for that patient. The status display provides process control by showing the clinical team where corrections are needed while the team still has the time needed to act prospectively. We are extending these methods to the management of chronic disease. Their extension requires engagement of the patient as a member of the team, a coordinated plan across the care continuum, informatics algorithms to stratify individual patients according to co-morbidities and their current level of control, and a means of detecting the presence or absence of a reaction to each action taken by the team.

A closed-loop control process assures that a system performs within control limits. In closed-loop control, the system's output is fed back directly to change the system's inputs. The way in which a thermostat works with a furnace to control room temperature is an example of closed-loop control. Closed-loop control starts with an explicit objective (e.g., the desired room temperature), a measure of the status of the system against that objective system (e.g., the difference between the actual and desired room temperatures), and a mechanism for adjusting the system's inputs to correct the difference and meet the objective (e.g., turning the furnace on or off). Medicine has used closed-loop

Correspondance and reprint requests: William W. Stead, M.D., Associate Vice Chancellor for Health Affairs, Chief Strategy and Information Officer, McKesson Foundation Professor of Biomedical Informatics, Professor of Medicine, Vanderbilt University Medical Center, Nashville, TN 37203, Tel: 615-936-1424, Fax: 615-343-2330, E-mail: bill.stead@vanderbilt.edu

Potential Conflicts of Interest: Dr. Stead is a co-inventor of two patient medical record products, one licensed to McKesson, Inc., and one licensed to Informatics Corporation of America, and receives royalties from these products through Vanderbilt University. He is a director of HealthStream, a public company, at which he is compensated by an annual option grant. Dr. Jirjis is a co-inventor of the patient medical record product licensed to Informatics Corporation of America, and receives royalties through Vanderbilt University.

control since Sheppard and Kouchoukos' pioneering work in the 1970s with post-operative fluid management (1). However, most such examples are at the subsystem level, where a computer can take complete control of the appropriate adjustments.

We are trying to develop closed-loop control for system-based practice through which teams of people, well-defined processes, and informatics tools work in concert to achieve a desired result, and each shortfall is turned into either timely course correction during the care of an individual patient or an iterative, adaptive development of the system of practice (2). The question is how to achieve closed-loop control when a human is in the loop applying compassion and judgment and when the system is evolving as it learns.

In 2007, we hypothesized that we could achieve the desired effect with a three-component control system consisting of: 1) an explicit end-to-end plan; 2) a record of what is done as it is done; and 3) an instant display of the status of each patient against the plan for that patient (3). The status display could provide process control by showing the clinical team where a correction is needed while the team has time to make the correction. Plans, displays, and performance could evolve together iteratively until the desired performance is achieved.

We tested this method in our work on ventilator management. The ventilator process-control dashboard has a row for each ventilated patient and a column for each element of the bundle of standard practices involved in ventilator control, colored to indicate the status of that element. Thus, green confirms adherence; yellow is an alert that action is needed to maintain adherence; and red highlights an element that is out of compliance. The human experts directing the ventilator process remain in charge, doing what they think best as long as the result is within the control limits established by the plan for the patient, or overriding the plan when the experts view deviation from the plan as being in the best interests of the patient. Control is achieved because every member of the team has the necessary information in time to initiate either corrective action or a process to override the plan. Patient-specific data on process performance and clinical outcomes are linked in the enterprise data warehouse to guide after-the-fact refinement of standard practices as the system learns. Immediate feedback is targeted at concurrent compliance with every element of the patient's plan, the Z100 score. We found that with this method of closed-loop control for ventilator management, the Z100 score improved from 23% to 87% and the rate of ventilator-acquired pneumonia decreased from 18.3/1,000 ventilator days to 11.4/1,000 ventilator days across Vanderbilt University's six adult intensive care units (ICUs) (4, 5).

EXTENDING THE METHOD FOR CHRONIC DISEASE

We are now trying to extend the closed-loop method to the management of chronic disease, with hypertension, congestive heart failure, and diabetes as the test beds for this. The differences between an inpatient population with a single clinical condition and a population with multiple chronic diseases and care sites, including the home, presents challenges to closed loop control in the form of system barriers and patient factors. System barriers include multiple specialists, malaligned reimbursement, fragmented plans and data, and limited direct observation. Patient factors include health literacy, financial resources, environments at home and work, and digital-technology resources.

Because the method based on the closed-loop control system provides process control through instant feedback to the system-management team, the first step in establishing the method is the formation of a coordinated team. Patients and their support communities must be engaged as full participants in goal setting, self-management, and monitoring, since most of the care delivered through the method takes place beyond the direct observation of health professionals. A care coordinator is an essential member of the team, working with patients and each of their providers to make sure that they have the information and communication needed to work as a team.

Unlike a quarterback model, in which a primary-care provider coordinates care among independent specialists, each with their own discipline-specific plan, the participants in the closed-loop control method work together to develop a common collaborative pathway. Patient-inclusion criteria are explicit. Thus, for example, a patient is evaluated or managed for diabetes if the patient's record includes the diagnosis of diabetes, an HbA1c above 7%, or a fasting blood glucose level above 125 mg/dL. Explicit stratification rules, based on both risk level and control status, guide the selection of diagnostic/therapeutic algorithms, the patient's self-management protocol, and the intensity of the care process. Continuing the example of diabetes, a patient is flagged as being at moderate risk if the patient has a complex insulin regimen, and as being at high risk if the patient has complications, co-morbidities, or is frail or elderly. A patient with diabetes is considered to have the disease in control if the patient's HbA1c is less than 7% or an individually assigned target value, and if the patient's plan is not being adjusted. Since hypertension does not have a measure like HbA1c that reflects physiologic control over a period of time, we consider hypertensive patients to be in control if 75% of their blood pressure home readings are below their individual target level. Explicit red flags are identified to guide patients about when to get help. For each process

step, the expected reaction is defined to support proactive monitoring and escalation of care. For example, if a prescription is written, the plan calls for a follow-up in 24 hours to see if it has been dispensed, rather than waiting until the patient’s next visit to verify this.

The collaborative-care pathway is adapted into an individualized plan for the patient as the patient engages with the care team. The entire team agrees on the goals of care and which goal should come first, for example, a patient had chronic pain and moderate hypertension, and the team decided to take up to 2 weeks to control the pain and to then reassess the patient’s hypertension. The plan specifies who is responsible for each of a patient’s medical conditions or for a set of actions, but responsibility is managed as a baton that can be passed among the team members. Thus, if a patient is seeing a specialist because of refractory hypertension, the patient’s preventive health maintenance can be completed during that visit. Similarly, once specialists make changes in therapy, they can depend on the care coordinator and ambulatory intensivist to manage the follow-up of the patient, knowing that they will be re-engaged if the patient’s disease-control targets are not achieved.

As highlighted in Figure 1, shared expectations, explicit accountabilities, and explicit alert and escalation criteria are features required for the plan component of the closed-loop control system. As in the ventilator work, the other two components of the control system, after the development of an end-to-end plan, are a record of what is done as it is done, and instant visualization of the patient’s status as compared with the corresponding criteria of the plan.

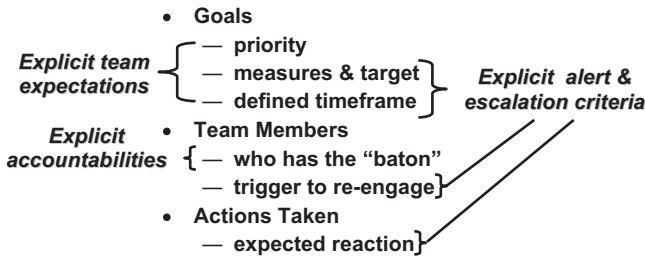


FIG. 1. Common plan as adapted to the patient with explicit team expectations, accountabilities and alert/escalation criteria to support closed loop control

improvement in the percentage of patients whose conditions are under physiologic control from a baseline of less than 50% control to more than 80% control. The Vanderbilt University Medical Center has agreed to fund the additional direct expenses of the coordinated team-based care plan and to make up any loss of volume-based compensation by participating physicians for the duration of the the test period for the proof of concept of the closed-loop control process.

Between April and July of 2010, we engaged an initial group of 56 patients. We started with a small number of patients with refractory conditions. For example, of the 26 patients who had hypertension, the disease was in physiologic control in only 34% at baseline, a figure below our average of 48% in physiologic control. We agreed that we would not add more patients until more than 80% of the initial group of 56 patients had achieved physiologic control. In the beginning, we paused briefly after engaging each additional patient to refine team roles, pathways, and the closed-loop control process to reflect our experience. Eighty percent of the hypertensive patients in the initial group reached the first target value for hypertension control in less than 5 weeks, and were in stable control at 8 to 9 weeks. Control of congestive heart failure and diabetes followed the same trajectory.

In September 2010 we stopped engaging patients for a month to make changes in the informatics tools so as to reduce the effort needed to gain the level of physiologic control described above. In October 2010, we began to engage a second group of 70 patients. Based on the experience with this latter group, we will decide whether we are ready to engage the full cohort of 568 patients identified for the proof of concept of the closed-loop control process or whether additional cycles of testing and refinement are needed to achieve a combination of roles, process, and tools adequate to handle that patient volume.

Throughout these cycles of system design and refinement, we are capturing workflow and resource utilization data to see if we can afford to scale the closed-loop method to cover our entire chronic-disease population. The answer will depend upon whether the increased expense during the initial phase of a patient's engagement and rapid cycle management to control limits is offset by savings due to efficiency or effectiveness. Early offsets will include an increase in size of the panel of patients that a physician can manage through the use of a coordinated team and technology, and a reduction in downstream management costs once a patient's condition is under control. Intermediate offsets will include reductions in the cost of managing complications.

CLOSED-LOOP CONTROL FOR A LEARNING SYSTEM

Closed-loop control for system-based practice requires both a means of instant control to ensure that the system performs as expected, and a feedback of downstream outcomes so that the system learns from the results of its performance. Our method uses a three-part control system consisting of an explicit plan, a record of what is done as it is done, and an instant display of the status of each patient against the plan for that patient. This display shows the care team where corrective action is needed while the team members have time to act.

During development of the system of plans, roles, processes, and technology for the closed-loop control of patients' care, we present this display to the team as a prototype as they make rounds to inpatients or see outpatients in the clinic and obtain feedback to guide rapid refinement of this system. If the team decides that the system, on its face, has validity, we begin to use the display to guide process performance. As process performance improves, residual process variation increasingly represents intentional decisions by the team to override points in the plan. Next, intermediate clinical outcomes are fed back to show whether the plan achieves the desired result, and whether decisions to deviate from points in the plan yield better results. In time, summative clinical outcomes allow comparison among different system-based approaches.

In the case of ventilator management, we learned that the Z100 score for measuring concurrent compliance with every element of a patient's care plan was more important as a process outcome than was compliance with the elements of the plan taken individually. The subsequent reduction in ventilator-associated pneumonia per one-thousand ventilator days provided an intermediate validation of process effectiveness, followed even later by summative benchmark data showing that our adult hospital was the best of its class within our peer group with regard to mortality, length of stay, and cost per discharge for patients on a ventilator more than 96 hours.

Table 1 compares the measures we are using to guide the proof of concept of closed-loop control for the management of chronic disease with the measures that we used with it in ventilator management.

We are currently using similar instant feedback to the team, but with additional dimensions to accommodate multiple providers and conditions. We are managing concurrent execution of each element of the plan, but with an additional focus on alerts to impending shortfalls in the reaction expected after an action to reduce lost cycle time, e.g. time from identification to engagement, time from engagement to target, and time from shortfall to change in plan. We are using physi-

TABLE 1

Comparison of measures used in guiding ventilator management with those used for chronic disease management

Outcome Measure	Ventilator Management	Chronic Disease
Summative outcomes	<ul style="list-style-type: none"> ■ Mortality ■ Length of stay ■ Cost of inpatient stay 	<ul style="list-style-type: none"> ■ Mortality ■ Condition-related morbidities ■ Cost per year at risk
Intermediate clinical outcomes	Incidence of complications per 1,000 ventilator days	% of patients whose condition is under physiologic control
Process outcomes	Z100 Score % of time that all 7 elements of the bundle are done for every patient	Cycle Time <ul style="list-style-type: none"> ■ Time from identification to engagement ■ Time from engagement to target ■ Time from failure to change of plan
Instant feedback	Status of patients on a ventilator	<ul style="list-style-type: none"> ■ Status of panels under management ■ Status of patient across providers

ologic control as the intermediate clinical outcome, and monitoring mortality, morbidity, and cost per year at risk as draft summative measures.

ACKNOWLEDGEMENTS

Our ideas for extending closed-loop control to the management of chronic disease evolved through the Vanderbilt University Medical Center's innovation and quality improvement initiatives. C. Wright Pinson, David Posch, and Robert Dittus have provided executive oversight of our project. Sharon Mullins has provided project management, and Julie Scott, Rebecca Bumm, and Joan Peterson have led the coordination of care. Betty Akers and Racy Peters have led the operations component of the project; Pete Powell coordinated the compensation aspect; and Collin Mothupi performed the data analyses. Dario Giuse and Lianhong Tang have provided informatics support.

REFERENCES

1. Sheppard LC, Kouchoukos NT. Automation of measurements and interventions in the systematic care of postoperative cardiac surgical patients. *Med Instrum* 1977; 11:296–301.
2. Stead WW, Starmer JM. Beyond expert-based practice. In: McClellan MB, McGinnis JM, Nabel EG, Olsen LM, eds. *Evidence-Based Medicine and the Changing Nature of Health Care*. 2007 IOM Annual Meeting Summary. Washington, DC: National Academies Press, 2008;94–105.

3. Stead WW, Patel N, Starmer JM. Closing the loop in practice to assure the desired performance. *Trans Am Clin Climatol Assoc* 2008;119:185–95.
4. Zaydfudim V, Dosssett LA, Starmer JM, et al. Implementation of a real-time compliance dashboard to help reduce SICU ventilator-associated pneumonia with the ventilator bundle. *Arch Surg* 2009;144(7):656–62.
5. Talbot TR, Carr DS, Parmley CL, et al. Reduction of ventilator-associated pneumonia using real-time course correction via a ventilator bundle compliance dashboard. Oral presentation at the Fifth Decennial International Conference on Healthcare-Associated Infections, Atlanta, GA, March 21, 2010 (in press).

DISCUSSION

Stokes, Iowa City: That was a very interesting presentation. I have a question that revolves around an observation. Computerized ventilator assistance and outpatient management are of course very different, and the expected outcomes are going to be very different. In an outpatient setting with multiple disorders, you are going to expect changes in multiple outcomes. I noticed that in your ventilator outcomes, the primary outcome seemed to be ventilator-associated pneumonia, and I also noted that with time, although your Z100 score improved dramatically, the ventilator-associated pneumonia outcome didn't seem to improve after you began to initiate the series of interventions you described. Maybe you could comment on that and how it might relate to the outpatient experience.

Stead, Nashville: Well, what we have tried to do is work with a tier of measures. So our process measure, which is the Z100 score, is instantly available. We have it in real time. The ventilator-acquired pneumonia rate is calculated at the end of the month and provides an intermediate measure that needs to trend in the right direction for the process to have face validity. However, the definition of ventilator-acquired pneumonia varies between sites because of inclusion and exclusion criteria. Although it has continued to step down if you do formal statistics to compare sequential 6-month blocks, we have not found a meaningful benchmark. What we feel are more important are the summative measures, which took another year to be able to see. These are the mortality rate, length of stay, and cost for patients who are on the ventilator for longer than 96 hours, enough time for ventilator management to really make a difference. Those summative measures have continued to come down almost in a straight line and can be benchmarked. So I think that this idea is to combine immediate process measures, whatever clinical measures we can get intermediately, and longer-term measures that really focus on survival, cost and morbidity to evaluate closed-loop control. That seems to work.

DuBose, Winston Salem: Thank you, Dr. Stead, for this very interesting presentation. I wanted to mention the issue of chronic disease management within the context of the patient-centered medical home to which you alluded, because I think it's quite clear that one of the things we do very poorly in the health-care system in this country is coordinate care. It seems to me that your model and the application hold promise for what has been extended beyond the patient-centered medical home, and that's the concept of the patient-centered medical home neighbor, where there is coordination of care between a primary-care physician and some specialist. I wonder if one of the things that's occurring right now within subspecialty organizations is modeling of their favorite chronic disease, chronic kidney disease for nephrology, and heart failure, for cardiology, and some of the conditions you mentioned for some of the other subspecialties. But as we go forward, there is clearly a lack of the workforce needed in primary care to address all

of these problems. There must be better coordination, and it seems to me that a system like this holds promise for better control of the health-care delivery system and for allowing both subspecialists and primary-care physicians to work as a team to manage these problems in a more systematic matter. And I just wondered what your comment might be with that application.

Stead, Nashville: I think your observation is correct. I don't think that we will succeed in taking care of the patient until we have a holistic plan for the patient that crosses all of the care specialties the patient needs, where each of the specialists actually interacts with that overarching plan and we work out what we want to do so that we are, in fact, prioritizing the plan. Otherwise, the things any one of us could want to do, while they might make sense within our piece of the problem, will not make sense or be doable for the patient as a whole. And so I think your point is very valid.

Shannon, Philadelphia: Thank you for that very, very interesting presentation. I wonder whether you've looked at either patient or provider satisfaction in this system-driven model, and if you can tell us a little bit about what the docs think.

Stead, Nashville: Everybody who is participating in the proof-of-concept is ecstatic. That is the only way I can describe it. The patients are getting up in church and giving testimonials. Admittedly we have held the physicians harmless of impact on compensation and revenue, but they've got the support to focus their time and what is now being called practicing at the top of their license. They are not having to do other things. They can depend on each one of them doing their part because they have a common access to common information, and they are working against a common plan. They all like it. What we elected to do as a medical center was to get a proof-of-concept and get the data that really could prove that we could get to 80% physiologic control. We set that as our bar. Anything short of that and, we will stop and iterate the model until we get there. If in fact we can do that, and if we get the workflow and utilization data together, then we are actually going to have to sit down and calculate how to reimburse people to make this work, but we took that out of the critical path of what will be proof of concept.

Cohen, Washington, D.C.: I want to congratulate you on a very interesting approach. There is a lot of talk, as you well know, about accountable-care organizations being one of the pathways to improvement in the implementation of the accountable care act. In fact, there is going to be a lot of money available for modeling accountable-care organizations. What is yet to be done is to figure out what it means to be accountable. A lot of work needs to be done on that score, and it strikes me that the approach you are taking could provide a crucial template or model for accountability for these organizations. If they implement these kinds of rational control systems in managing the chronic diseases they treat, it seems to me that that could be a very potent measure of accountability for these organizations. Would you agree or not?

Stead, Nashville: Yes I would agree. Accountable-care organizations are largely being thought about from the legacy of managed care and other forms of risk taking, and I think we've got a lot of evidence that the health-care industry doesn't really know how to do that. What we are hoping we can do is to actually get the care right and take out all unwanted variability. What the control system does, this visual system that I described to support the care team, is to leave the team in charge. They can override it every time they want to, but all the variation we have is intentional, so it takes the noise out and it's easy then to see the effects on outcomes. If we could get this kind of control system to get the care right, get the noise out, then I think it would be easy to accept risk for the remaining variability, but I think we have to flip it around and first get the care right.