

An expert system based sodium nitroprusside blood pressure controller

Citation for published version (APA):

Blom, J. A. (1990). An expert system based sodium nitroprusside blood pressure controller. In J. Cornelis, & S. Peeters (Eds.), *North Sea conference on biomedical engineering 1990, 2nd, Antwerp, 19-22 November 1990 : proceedings* Technologisch Instituut-K.VIV.

Document status and date:

Published: 01/01/1990

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

AN EXPERT SYSTEM BASED SODIUM NITROPRUSSIDE BLOOD PRESSURE CONTROLLER

J.A. Blom

Eindhoven University of Technology, Eindhoven, the Netherlands

ABSTRACT

When controlling a biomedical system, the control system needs to contain a significant amount of medical knowledge in order to be able to supervise and adjust the basic control mechanism, as well as to handle medically warranted exceptions to the normal control regime. Expert systems technology is a suitable foundation for such complex controllers.

Sodium nitroprusside (SNP) is the most frequently used drug when the blood pressure must be temporarily stabilized at a lower than normal level. Due to its characteristics, blood pressure control with SNP is not an easy task. Particularly the patient's sensitivity to SNP can vary so much that standard controllers cannot cope. On-line system identification is unsafe because the signal is often disturbed and because numerous other drugs and actions influence the blood pressure. A good approach proved to be to design a robust controller and to have a 'machine expert' monitor the controller's performance in order to readjust the control gain if it appears incorrect.

The expert system based SNP control system was tested during 30 cardiac surgery cases. It proved to be safe and could control all patients well.

Keywords: expert systems, adaptive control, arterial blood pressure, sodium nitroprusside

SODIUM NITROPRUSSIDE CONTROL SYSTEMS

The control of biomedical systems, including the control of blood pressure, is significantly more difficult to accomplish than many of the control engineering applications usually encountered in the regulation of man-made systems. In particular, hypotensive drugs act very rapidly and with high gain, the latter being patient dependent and showing great variance. They also have a very short half life (as they are being neutralized and excreted by the organism). The input-output delay time is relatively large compared to the half life. These half lives and delay times are often time varying, depending on both patient condition and the background level of other physiological variables and therapeutic agents.

Sodium nitroprusside (SNP), the most frequently used drug when the arterial blood pressure needs to be temporarily lowered, is an example. Due to its short half life, SNP is applied by infusion. In some patients manually controlled infusion is laborious because the flow rate has to be readjusted frequently; and since the response of the patient to the drug may change over time, the infusion's effect must be carefully monitored in *all* patients. Currently and in the past, many researchers [e.g. 1, 2, 3, 4, 5, 6, 7, 8, 9] have therefore designed automatic closed loop controllers.

As a control problem, the application of SNP is difficult. The problems are many: frequent artifacts in the blood pressure signal [10], a large delay time, a pronounced non-linearity of the response,

reflexes that fight the pressure decrease and that come into action after some time (the two latter factors result in a system which is not *identifiable* [11]), and in general a huge variability of the patient's characteristics, both between patients and for the same patient at different times [12, 13, 14]. Moreover, the exact pharmacology of SNP is unknown [15]. SNP is also toxic [16, 17]; overdosage has to be avoided at all costs.

Due to these problems, approaches based on classical and modern control theory are not good enough [11]. Although not always clearly stated, even modern adaptive controllers do not work correctly all the time [2, 4, 8, 18]. Clinical knowledge and experience seem to be necessary for a good control. Many physicians do as well as or better than automatic control systems, at least if they are constantly paying close attention. But that is exactly the problem: they do not; they have more and often more important things to do than adjusting infusion flow rates. Besides, they are prone to errors [19].

It is also unknown what constitutes 'optimal' control. There is no arbitrary level at which the blood pressure level is optimal. In the young and healthy, the usual level is 60-70 mmHg systolic, whereas in older individuals satisfactory conditions are attained at higher levels. Little is gained by depressing blood pressure more than necessary. The pressure is not necessarily kept at a fixed level but is allowed to fluctuate within reason, depending on surgical needs. The advantage of computer control may not necessarily be in improved control but rather the convenience and lack of susceptibility to distraction and fatigue.

Controllers must not only control induction of hypotension, they must also be able to handle pressure perturbations. Large 'noise' levels may occur in the blood pressure signal. In the ICU, this noise may be due to sudden changes in the patient's emotional state or level of activity. In the operating room, surgical stimulation during light anesthesia, bleeding, rapid fluid infusion, administration of other vasoactive drugs, and respiratory maneuvers can dramatically influence blood flow and arterial blood pressure. In all these cases the controller must take appropriate steps to regulate the blood pressure. However, routine procedures such as the taking of blood samples or the flushing of a catheter give readings which do not reflect the true blood pressure and must be disregarded.

So much medical knowledge is required in a reliable controller, that a purely algorithmic solution becomes difficult. The new blood pressure controller is therefore implemented as a SIMPLEXYS real time *expert system* [20], in which the basic controller is supervised by a 'machine expert' which also detects and handles medically warranted exceptions to the normal control regime. The main tasks of the machine expert are 1) to check when measurements are invalid, 2) to decide what to do when measurements are invalid, 3) to adjust the controller's characteristics so that control performance remains adequate and 4) to handle exceptions (unexpected transients in the signal).

KNOWLEDGE ACQUISITION

The more information is available about the object to be controlled, the better a controller can be designed. The first step in designing a good controller is thus to acquire as much information about the system as possible. This can be done either in a study especially devoted to this task, and/or on-line, while the controller is controlling.




Fig. 1. Typical dynamic (impulse response) characteristic of SNP in pigs. Both curves were obtained at the same mean arterial pressure, the upper curve about 90 minutes later than the lower one.




Fig. 2. Average normalized static (dose-response) characteristic of SNP in pigs. The lower curve was obtained at the beginning of the experiment, the upper one about six hours later.

Our experience shows that the blood pressure signal is contaminated with so much 'noise' due to spontaneous fluctuations and transient disturbances, that an on-line system identification or parameter estimation method will not be able to consistently produce reliable results. Moreover, the characteristics of the system (specifically its

time-varying non-linearity and the reflex mechanisms that elevate the 'zero flow' pressure) result in the fact that the system is not (or badly) observable [11], making accurate on-line estimation of the system's parameters an almost hopeless task.

To gain a better understanding of the variability of the response to SNP, experiments with Yorkshire pigs were done [14, 20]. Conclusions from this study were:

1. The shape of the dynamic (impulse or step) response to SNP was relatively constant for all subjects at all times; the response resembled that of a second order system with delay time, as shown in fig. 1.
2. The shape of the static (dose-response) curve was relatively constant for all subjects at all times; in most cases it showed a pronounced non-linearity (saturation at higher infusion flow rates resp. at lower pressure levels) which could be modelled well with an exponential curve, as shown in fig. 2.
3. The blood pressure that would result from stopping the infusion showed an upward shift at later times. This is evident in fig. 2.
4. The inter- and intra-patient variability of the sensitivity to SNP was the only parameter with a broad range (a factor of 80); fig. 1 shows an example of this variability. The sensitivity nor its variation could in any way be predicted from any of the patient's demographic or physiologic data.

SIGNAL VALIDATION

Many disturbances can cause the signal not to reflect the true blood pressure: blood clotting, air bubbles in the line, flushing the arterial line, sampling of blood, electrocautery etc. If these disturbances can cause an incorrect computation of the mean pressure, they, as well as transducer, catheter and amplifier failure, should be detected. A signal validation algorithm [20] computes, for each heart period, a number of features of the signal such as maximum, minimum and mean pressures, pulse pressure, period duration, systolic slope, etc. If all these features are within acceptability limits, the mean arterial blood pressure (MAP) value is considered valid and is passed to the control system together with a validity flag, which indicates that this value reflects the real MAP. If one or more of the signal's features are not within acceptability limits, the MAP value is considered invalid and cannot be used. Some of the acceptability limits are based on *a priori* knowledge of acceptable shapes of a period of the signal. Other acceptability limits are based on a signal model based on the average of some most recently validated periods; these limits are thus adaptive and track changes in the appearance of the signal.

During perfusion, the time during which the circulation of the blood is realized by a heart-lung machine, the pressure signal is almost flat. Since no heart rate period can be detected, an artificial period of one second is imposed. During perfusion, a different validation algorithm checks only the signal's mean value (the MAP), and the difference between the signal's minimum and maximum over a cycle.

The validation algorithm was shown to perform adequately: all significant artifacts were detected. A small percentage of acceptable periods was flagged as invalid, but not enough to significantly influence control actions. This is due to both the 5 second averaging algorithm which delivers a valid MAP as long as at least one period within the 5 second interval is valid, and to the fact that missing upto twelve successive 5 second average values has a negligible impact on the controller's performance.

A 'hold mode' solution, which simulations [21] showed to be safe, was chosen for the controller. Missing MAP values are supplied by the last valid MAP value and control continues

- during at most 30 seconds if the last valid MAP was more than 20 mmHg from the setpoint;
- during at most 60 seconds if the last valid MAP was less than 20 mmHg from the setpoint.

If the signal is lost for a longer time, the system gives an audible alarm and control returns to manual mode with a flow rate identical to the last valid flow rate given. To resume automatic control, the signal must be valid again; one key-press is then sufficient.

HANDLING OF TRANSIENTS

Whereas signal validation is a process to establish the physiological origin of the signal, transient detection is a process to establish whether a normal control regime can be maintained or not.

We define a transient to be a sudden large in- or decrease of the pressure which, even if not acted upon, disappears after a short time. When a large in- or decrease of the pressure suddenly appears, however, the control system has no way to 'look ahead' to establish whether it will be a transient or not. It is most prudent to temporarily assume a worst case situation.

Positive transients can be tolerated if they do not last too long. These positive transients are often caused by pain stimuli. Insufficient pain suppression may cause large increases in blood pressure that last as long as the pain stimulus occurs and subside when the pain stimulus is over or when pain suppression medication is increased. Such a blood pressure increase should not be suppressed by SNP, because SNP is not a pain killer. If it were suppressed by SNP, a large pressure undershoot might occur after the pain ends. Since a too low pressure is more dangerous than a too high pressure, a control strategy of, at least temporarily, not responding to a large sudden pressure increase seems best.

Negative transients are dangerous; the low pressure can indicate a state of shock. SNP infusion should stop immediately and should only be resumed when the blood pressure is approximately at the setpoint again. However, if the negative transient was due to an artifact of some sort, the after-effect of temporarily stopping the SNP infusion will often be a pressure overshoot.

The above mentioned *interpretations* of what causes transients, i.e. pain and shock, will often be incorrect, since transients may have other causes. The control strategy that is followed after such transients was chosen for the sake of patient safety and appears reasonable also if the cause should be different.

Transients need to be detected for more reasons. The controller simply cannot compensate for transients: they are too fast and too large. Moreover, an attempt to compensate for a transient would cause a large change of flow, but without much effect. And when the transient is over, the controller must recover and bring the flow back to the pre-transient level. It is therefore better to detect transients and control the flow in a different way until the transient is over. It may also be undesirable to compensate for transients; rapidly changing blood pressures are a diagnostic aid and should not be obscured.

PID CONTROLLERS

The traditional approach in biomedical applications has been to use a fixed-design PID controller to adjust the rate of drug infusion to the controlled system. These controllers are relatively robust, simple to implement, they do not require an elaborate model of the system, they produce zero steady-state error and when combined with logical

rules to avoid overdose, they can be successfully used for clinical therapy. But these controllers also possess several disadvantages. These include their tendency toward instability (for large integral-term gains), their excessive noise sensitivity (when large derivative gains are used in a noisy environment), and their inability to cope with patients whose sensitivity is either very small or very large.

Fixed PID controllers are too simple to adequately control such a complex variable as the human blood pressure. This is because no method exists to select the relative weights of the control parameters when the system description is unknown. Since physiological systems are often poorly characterized and may change over time, it is desirable to use controllers that automatically adapt their operation to changes in the system's characteristics. PID controllers can be *tuned* to the system they control if (some of) the system's characteristics are known [e.g. 22]. Tuning is impossible, however, if the characteristics of the system to be controlled change over time and if there is no reliable way to establish an estimate of those characteristics. A different, more robust approach is to keep track of the controller's performance in such a way that no parameter estimation needs to take place.

A PID-controller's performance can be established from observations in two ways [21]:

1. In a steady state condition (the pressure is or should be approximately constant), the controller's gain is too low if the offset (the difference between the actually measured pressure and the pressure setpoint) is, on average, significantly different from zero for too long a time; and the controller's gain is too high if the offset is, on average, zero, but shows oscillatory behavior.
2. Under dynamic conditions (the pressure changes or ought to change, e.g. due to a setpoint change), the controller's gain is too low if the pressure changes too slowly; and the controller's gain is too high if the pressure change is too fast.

Tests during 33 cases provided a verification of the animal data about the dynamic SNP response in patients; the dynamic SNP parameters are given in table 1, together with the controller's design values. The controller's nominal gain is - 0.2 mmHg/(mg/hr), equivalent to approximately - 0.04 mmHg/(μ g/kg/min) in an 80 kg patient, and total control gain is nominal gain times relative gain.

TABLE 1 Dynamic SNP parameters; extremes from 33 cases

	min found	max found	design range	
relative gain	.6	4.0	.1-9.0	
time delay	15	75	25-100	sec
time constant	30	100	30-120	sec

TABLE 2 Sensitivity classes and normalized control gain

sensitivity class	normalized gain
very insensitive	9
insensitive	3
normal	1
sensitive	1/3
very sensitive	1/9

Simulations were performed to establish the best (most robust) values for the control parameters; they showed that the derivative term should be set to zero, and that a constant integrative term could be chosen without a significant degradation of performance. The gain shows so much variation, however, that it needs to be adapted; no adequate control performance would otherwise be possible. The patient's sensitivity, as well as the corresponding normalized control gain, is classified into one of five classes, as shown in table 2.

The control formula which is used in the basic controller is written in such a way, that a flow *increment* is computed:

$$SNP_k = SNP_{k-1} + G * [I * (SET_k - MAP_k) + P * (MAP_k - MAP_{k-1})]$$

where

MAP	mean arterial pressure	[mmHg]
SET	setpoint pressure	[mmHg]
SNP	flow rate	[ml/hr] = [mg/hr]
G	adaptive control gain	(1/9, 1/3, 1, 3, 9)
I	I-parameter, .0960 to .0720	[(ml/hr)/mmHg]
P	P-parameter, .0056 to .0036	[(ml/hr)/mmHg]

The sample interval is 5 seconds. To improve safety and to ensure a quiet control signal, the flow increment is limited to 7% of the current flow or 1% of the maximum allowed flow, whichever is larger. Whereas simulations showed that, due to the expected noise, the D-parameter should be set to zero, the P- and I-parameters are allowed to vary because it became clear that in regulation different values were appropriate than in stabilization (a phase plane approach is described in [22] to realize a similar strategy). If the offset is 100 mmHg or more, the lowest values are used; at smaller offsets, the parameters proportionally grow to their highest values. This makes the controller non-linear, but since the offset will generally be less than 50 mmHg, the non-linearity is usually small.

For safety reasons, the system starts to control in the lowest gain class (1/9), as it initially assumes a very sensitive patient. In many cases, however, this control gain will be too low. Optimally, the controller's gain should be the inversely related to the patient's sensitivity.

In steady state conditions, the controller's gain is increased if the offset is, on average, significantly different from zero for too long a time; and the controller's gain is decreased if the offset, while on average zero, shows oscillatory behavior.

In dynamic conditions, after a setpoint change, the controller's gain is increased if the pressure changes too slowly; and the controller's gain is decreased if the pressure changes too rapidly.

Both offset level and pressure change can be monitored by constructing a number of *regions* around the setpoint and observing how long the actually measured pressure stays within such a region. The width of the different regions is chosen in such a way that if the control gain is correct, the MAP stays in each region for about 2 minutes. If the MAP stays in a region for longer than 4 minutes, the progress is considered too slow and the relative control gain is tripled, i.e. the patient's sensitivity is assumed to be one class lower. On the other hand, if the MAP moves through a region within a period of one minute, the gain is considered too high and is divided by three. Progress is thus determined by the time the MAP stays in a region, or rather by a *slope* (MAP change per time period). Direct computation of a slope in a noisy signal is not reliable; to

reduce noise influences, the slope should be computed either over a certain time or over a certain MAP change. The gain adaptation mechanism uses both: if the MAP changes slowly, the gain will be increased when after a certain time the MAP is still found to be in the same region; if the MAP changes quickly, the gain will be decreased as soon as a region border is crossed.

It was decided to limit the infusion flow rate to a 'safe' 2 µg/kg/min, a factor 2.5 to 4 lower than the maximum flow rate mentioned by others [1, 23], but in agreement with what we observed to be standard practice. This implies that very insensitive patients cannot be controlled. In such cases standard practice was to accept a higher setpoint; due to the non-linearity of the response (see fig. 2), such a new setpoint usually needed to be only slightly higher.

THE CONTROL SYSTEM

The controller was designed as a SIMPLEXYS real time expert system [20]. It runs on an IBM-PC-AT compatible computer and has a cycle time of 5 seconds. Its input, the arterial pressure, is obtained from the monitoring equipment, sampled by a LabMaster AD-converter, and validated. The beat-to-beat output of the validation algorithm is averaged over a 5 second interval; invalid values are, of course, not included in the average. If in the 5 second interval no valid beats were found, a 'pressure not valid' flag is set by the algorithm. Thus the input to the expert system controller consists of, at each 5 second interval, a mean blood pressure value and a validity flag. The output, a new SNP flow rate, is communicated every 5 seconds to an IMED 929 infusion pump. Extra provisions are displays of MAP, setpoint and SNP flow rate, both numerically and as a trend graph over the last 30 minutes.

A diagram of the blood pressure control system is shown in fig. 3. The system communicates with the outside world in three ways: it acquires the analog arterial blood pressure signal from the AD-converter, it communicates with the infusion pump through an RS-232 interface, and it communicates with the user through the computer's keyboard and video display.

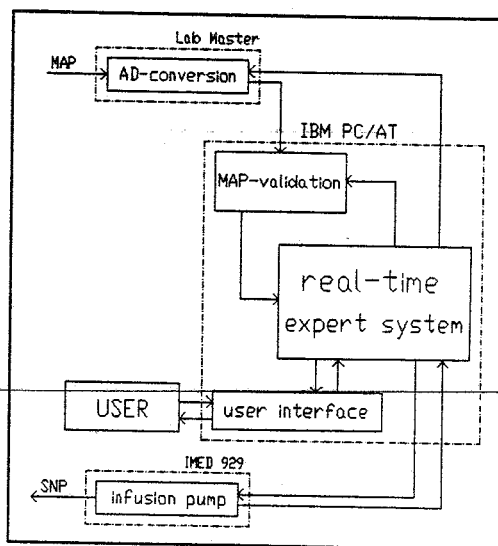


Fig. 3. The arterial blood pressure control system

The controller has two control modes: manual and automatic. The controller starts up in manual mode with a zero flow rate. Through the keyboard's function keys, the setpoint can be in- or decreased in steps of 1 mmHg. In manual mode, the infusion flow rate can be in- or decreased in steps of 0.1 ml/hr (approximately 0.02 µg/kg/min for an 80 kg patient). A safety key allows stopping the infusion immediately.

THE CONTROLLER'S PERFORMANCE

The control system was tested during cardiac surgery procedures, mainly bypass cases. Clinical tests proceeded in three phases. During the first phase all signal processing and display functions were tested, including the data validation [24]. During the second phase (33 cases) the system was tested under 'open loop' conditions; although control was still manual, a flow rate was computed by the expert system. The first and second phases also provided extra information on the ways in which clinicians controlled the SNP flow rate. The second phase was mainly used to debug the system, to correct the knowledge base where necessary, to gain sufficient confidence in the system's performance, and to familiarize the clinicians with the system. The third phase (30 cases), closed loop control, was started when we thought the system could be trusted under all conditions; during this phase the knowledge base remained unchanged.

Bypass surgery has three stages. During the first stage access to the heart is made, and blood vessels are resected from a leg. During the second stage, these vessels are used to replace defective coronary arteries. The heart is inoperative and blood circulation and oxygenation is provided by a heart-lung machine (perfusion); the resulting blood pressure signal is almost flat. The MAP is mainly controlled by the heart-lung machine and is at an exceptionally low level (around 40 mmHg). During the third stage the heart is reactivated, the chest is closed and the MAP is allowed to rise again to its pre-perfusion level.

The controller was generally switched to automatic soon after the arterial pressure measurement became available, and remained in control before, during and after perfusion. Occasionally the clinician considered the controller's response to be too slow; in such cases manual control was entered, the correct flow rate was set, and automatic mode was returned to.

A major aspect of the controller's performance can be assessed by measuring the offset, the difference between MAP and setpoint, over time. Manual control during a total of 33 cases (85889 5 second MAP averages) was compared with automatic control during 30 cases (60970 5 second MAP averages; manual control episodes are not counted) [24]. In the interpretation of the results, several factors are important to consider. In the first place is it, both for the clinician and for the controller, impossible to control the pressure if the setpoint is above the MAP (no negative flows can be given), or when the imposed maximum value limits the flow rate. In the second place were clinicians during manual control sometimes too preoccupied with other matters to promptly mention the new setpoint that they considered appropriate in a new situation. Third, the numbers of cases compared are too small to allow definitive pronouncements.

Averages and standard deviations of the offset distributions are presented in table 3 for both manual and automatic control. On average, both clinicians and controller kept the MAP close to the setpoint. The controller's standard deviation is smaller, however, indicating that in automatic control the MAP is less often far from the setpoint.

TABLE 3 Manual and automatic control regime statistics

	manual	auto	
offset	- 0.8	- 1.1	mmHg
S.D.	18.1	12.6	mmHg

Another way to present the results is by noting how often the offset stayed within a zero-centered band (table 4). We notice that automatic control is consistently better; 3 to 8% of the offset values *more* fall in bands closer to zero.

TABLE 4 Percentages of the offset in zero-centered bands

offset band	manual	auto
± 5 mmHg	30%	33%
± 10 mmHg	55%	61%
± 15 mmHg	70%	78%
± 20 mmHg	81%	88%

In other respects, too, the controller proved to be adequate and safe. It could handle all cases well. All invalid measurements were rejected. Some valid measurements were rejected as well, mainly following sudden changes of the mean pressure level, but this had no significant impact on the controller's performance. Transients, both positive and negative, were recognized and processed correctly. Gain adaptation worked correctly as well but was overly cautious in some cases: when the flow rate had been (almost) zero for some time, the system decreased the gain, assuming that an increase of the patient's sensitivity could have occurred. This assumption often proved to be incorrect. This type of gain decrease action is probably unnecessary in most, maybe in all, cases; its removal would make the controller's response somewhat faster.

The overall conclusion is that automatic control is slightly better than manual control.

A KNOWLEDGE ENGINEERING APPROACH TO CONTROL

The SIMPLEXYS real time expert systems toolbox [20] offers a knowledge engineering programming language that is especially well suited for the implementation of systems like the one described. As an example, the SIMPLEXYS language offers the *protocol* construct to denote and link chunks of time-sequencing knowledge. The knowledge about the patient's sensitivity resp. the control gain is maintained in a sensitivity protocol, part of which is shown in fig. 4. The boxes represent a *state* the patient or the system can be in, and the directed links between the boxes represent *conditions* (triggers) that can make another state active. Connected to a state are one or more *goals*, evaluations to be performed and/or tasks to be executed. In the example, a change of state corresponds with a decrease or increase of the patient's estimated sensitivity.




Fig. 4. Part of the sensitivity protocol

The SIMPLEXYS language encourages modular and top-down programming. The control system as a whole is constructed in a number of ever more 'symbolic' layers:

- interrupt routines for the AD conversion, keyboard handling and communication with the infusion pump;
- the validation routine that delivers the MAP;
- simple MAP-based calculations such as MAP filtering, MAP slope determination and calculation of the SNP flow rate by the basic control algorithm;
- basic rules that test data, parameters, flags and key presses;
- evaluation rules that combine basic rules and/or other evaluation rules, including rules that form a 'safety net' around the basic controller;
- the protocol.

Higher layers always refer to lower ones, except the highest two, which refer to each other. States and their goals and triggers represent the system at its highest level of abstraction. Design can start at this level, and subsequent refinements successively make the system more and more functional. A major advantage of the expert systems approach to control systems is thus the design methodology that is offered.

CONCLUSIONS

SIMPLEXYS proved to offer a satisfactory set of tools for the design of this control system. It offered the possibility for a top-down design that allowed testing in all design stages. The expert system based controller runs on an IBM-PC-AT computer system and yet is fast enough for this fairly complex real time task.

The controller proved to be safe and could handle all cases well. Invalid measurements which could cause incorrect behavior of the controller were rejected. Transients, both positive and negative, were recognized and processed correctly. The gain adaptation worked correctly as well, although it was overly cautious in some cases.

REFERENCES

1. Slate JB. Model based design of a controller for infusing sodium nitroprusside during postsurgical hypertension. PhD Thesis, Univ of Wisconsin in Madison. 1980.
2. He WG, Kaufman H, Roy R. Multiple model adaptive control procedure for blood pressure control. *IEEE Trans Biomed Engng* 33:1, 10-19. 1986.
3. McInnis BC, Deng LZ. Automatic control of blood pressures with multiple drug inputs. *Annals of Biomed Engng* 13, 217-225. 1985.
4. Millard RK, Hutton P, Pereira E, Prys-Roberts C. On using a self-tuning controller for blood pressure regulation during surgery in man. *Computers in biology and medicine* 17:1, 1-18. 1987.
5. Reid JA, Kenny GNC. Evaluation of closed loop control of arterial pressure after cardiopulmonary bypass. *British Journal of anaesthesia* 59:2, 247-255. 1987.
6. Rosenfeldt FL, Chang V, Grigg H, Parker S, Learns R, Rabinov M, Xu WG. A closed loop microprocessor controller for treatment of hypertension after cardiac surgery. *Anesthesia and Intensive Care* 14:2, 158-162. 1986.
7. Stern KS, Chizeck HJ, Walker BK, Krishnaprasad PS, Katona PG. The self tuning controller, comparison with human performance in the control of arterial pressure. *Annals of Biomed Engng* 13, 341-347. 1985.
8. Meline LJ, Westenskow DR, Somerville A, Wernick RT, Pace NL. Evaluation of two adaptive SNP control algorithms. *J Clin Monit* 2:2, 79-86. 1986.
9. Westenskow DR, Meline L, Pace NL. Controlled hypotension with sodium nitroprusside: anesthesiologist versus computer. *J Clin Monit* 3:2, 80-86. 1987.
10. Rampil IJ. Intelligent detection of artifact. In: Gravenstein JS, Newbower RS, Ream AK, Smith NT (eds). *The automated anesthesia record and alarm systems*. Butterworths, Boston. 1987.
11. Genderingen HR van. Some aspects of system identification and control theory in the design of an adaptive blood pressure controller (in Dutch). Master's Thesis, Dept of Electr Engng, Eindhoven Univ of Technology. 1984.
12. Wood M, Hyman S, Wood AJJ. A clinical study of sensitivity to sodium nitroprusside during controlled hypotensive anesthesia in young and elderly patients. *Anesthesia and Analgesia* 66:2, 132-136. 1987.
13. Blom JA, Bruijn NP de. Peroperative estimation of sodium nitroprusside sensitivity. *Proc IEEE Southeastcon*, Sandestin, 564-566. 1982.
14. Hutchinson WF, Hollway TE. An interesting effect of sodium nitroprusside [letter]. *Anesthesia* 40:11, 1128. 1985.
15. Patel CB, Laboy V, Venus B, Mathru M, Wier D. Use of sodium nitroprusside in post-coronary bypass surgery; a plea for conservatism. *Chest* 89:5, 663-667. 1986.
16. Butler AR. Further investigations regarding the toxicity of sodium nitroprusside. *Clinical Chemistry* 33:4, 490-492. 1987.
17. Chizeck HJ. Adaptive control theory and applications to drug delivery. 1986 American Control Conf, Vol 2, 871-873. 1986.
18. Hopking BDA. Hazards and errors in anesthesia. Springer. 1980.
19. Blom JA. The SIMPLEXYS Experiment: real time expert systems in patient monitoring. PhD Thesis, Eindhoven Univ of Technology. 1990.
20. Lammers JO. Knowledge based adaptive blood pressure control: a SIMPLEXYS expert system application. AIO Thesis, Dept of Electr Engng, Eindhoven Univ of Technology. EUT report 90-E-236. 1990.
21. Krijgsman AJ, Verbruggen HB, Bruijn PM. Knowledge based real time control. In: Rodd MG, Suski GJ. *Artificial Intelligence in real-time control*. Pergamon. 1989.
22. Packer JS, Mason DG, Cade JF, McKinley SM. An adaptive controller for closed-loop management of blood pressure in seriously ill patients. *IEEE Trans Biomed Engng* 34:8, 612-616. 1987.
23. Zwart R. Implementation and evaluation of a robust adaptive blood pressure controller (in Dutch). MSc Thesis, Dept of Electr Engng, Eindhoven Univ of Technology. 1990.