QUALITATIVE SPATIAL AND TEMPORAL REASONING IN CARDIAC ELECTROPHYSIOLOGY.

Jim Hunter

University of Aberdeen †

Nick Gotts and Ian Hamlet

University of Sussex ‡

Introduction

We believe that there are persuasive arguments for investigating the application of qualitative modelling techniques within medicine[1]. Although the work reported here is concerned with the specific domain of cardiac electrophysiology, we expect that some of the techniques will be more widely applicable both within medicine and elsewhere.

Basic Cardiac Electrophysiology

We have space here to give only the briefest outline of the processes we are trying to model; the interested reader should consult one of the standard texts[2].

The heart consists of four chambers: two atria and two ventricles. Electrically, however, the two atria can be considered as one, and we shall refer only to *the atrium*; likewise for the two ventricles.

In essence the function of the cardiac electrical system is to generate impulses at the correct rate, and to deliver them to other parts of the heart so that they arrive at the correct times. Impulse conduction can be likened to a row of falling dominoes. Cells that are polarised have a negative potential across their walls (the upright dominoes). If they remain in this state for long enough, they will spontaneously depolarise with the wall potential becoming positive (the domino will fall over). However cells in different parts of the heart have different spontaneous depolarisation rates. Under normal circumstances the cells in the Sino-Atrial (S/A) node have the shortest spontaneous depolarisation time; this region therefore serves as the primary generator of impulses. Any polarised cell next to a cell which depolarises will itself depolarise (one domino knocks over the adjacent one). In this way, the impulse is propagated throughout the conducting system and hence to the heart muscle. After some time in the depolarised state, a cell will repolarise (the domino picks itself up) ready to repeat the cycle.

As we have said, in the normal heart impulses are generated the S/A node, which is located in the upper atrium. They are then conducted through the atrium, causing the atrial muscle to contract and giving rise to the *P-wave* observed on the surface electro-cardiogram (ECG). The impulse is then conducted from the atrium to the the ventricle through a single channel called the Atrio-Ventricular (A/V) node located between the two. The main functions of the A/V node are to delay the propagation of the impulse (to allow the atrium to contract fully) and to control the rate of impulse transmission to the

Extended abstract submitted to the European Group on Qualitative Physics for the Second Workshop on Qualitative Physics, Paris, July 26th-28th 1988.

† Department of Computing Science, University of Aberdeen, Kings College, OLD ABERDEEN AB9 2UB, UK; email: jhunter@uk.ac.abdn.cs

‡ Artificial Intelligence in Medicine Group, School of Engineering, University of Sussex, BRIGHTON BN1 9QT, UK.

ventricles should the atrium pathologically present it with too rapid a sequence of impulses. When the impulse reaches the ventricle and depolarises the ventricular muscle, we observe the so-called *QRS-wave* on the surface ECG.

The conducting system is subject to all manner of malfuctions which give rise to abnormal behaviours (arrythmias); for example:

- impulses may be generated in the wrong part of the heart;
- impulse pathways may be blocked;
- parts of the system may take longer to repolarise than they should;
- there may be additional pathological pathways.

Many of these defects may co-exist, making the number of potential behaviours very large indeed.

Objective

Our objective in this work is to model the operation of the conducting system at the above level of physiological detail, to the point where a surface ECG can be given an interpretation in terms of a behaviour of the model. The model is basically a qualitative one, although we do make use of approximate numerical information when we believe such information to be used by clinicians; however we are a long way from *full* numerical simulation[3]. We are not concerned with detailed signal processing of the surface ECG but rely on human intervention to segment the input signal and to attach (albeit tentative) labels (e.g. P-wave) to the segments.

We were inspired by the work of Bratko et al. on the Kardio system[4]; however it should be noted that our aim is to work at a much more detailed level. Kardio dealt with descriptions of the ECG which were averaged over time; we aim to explain each feature of the signal. Furthermore, Kardio's model did not operate at the level of physiological mechanism (polarisation, depolarisation, repolarisation, etc.) which we represent.

This work brings together two existing strands of work in the Artificial Intelligence in Medicine Group: that on qualitative spatial modelling (Gotts) and that on temporal reasoning (Hamlet).

State Representation and Transitions

We describe more completely elsewhere a formalism for describing the possible electrical states of a *region* of the conducting system[5]. A region can be as large or as small a portion of the system as is desired depending on the degree of resolution required. At present we are using a three region model (atrium, the A/V node, and ventricle). Regions are divided into *lobes* where a lobe is a continuous portion of tissue whose cells are all in the same electrical state - all polarised (P) or depolarised (D). Thus the state of a region is described in terms of its lobes and their relationship to one another - a region with two lobes of depolarised tissue and one polarised being represented by the graph:

D-P-D

Again one can consider arbitrarily complex descriptions; we have limited ourselves at present to regions with one or two lobes, having four possible state descriptions:

P D-P D P-D

These correspond to the situations where the region is totally polarised, is depolarising, is totally depolarised, and is repolarising. Our formalism allows us to derive a set of transition rules for the states of the region. In this example they would be:

P -> D-P -> D -> P-D -> P -> D-P ...

We can express the connectivity of the regions in the forms of rules. For example:

> **If** the A/V node is in the state P,
> **and** the atrium makes a D-P -> D transition,
> **then** the A/V node will make a P -> D-P transition.

This expresses the fact that if the A/V node is polarised, the arrival of a wave of depolarisation from the atrium will trigger a wave of depolarisation in the A/V node.

We now have a way of describing the states of different components (regions) of the system, and have rules which define the allowable transitions between these states. We can aggregate these individual region-state descriptions together into a global state vector for the entire system, for example:

{D-P P P }

i.e {atrium A/V node ventricle}

Given this mechanism we can now explore the possible behaviours of the system from a given starting state; for example:

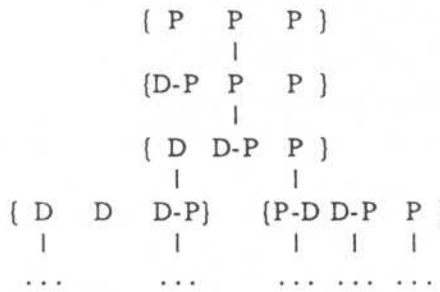

This is similar to the way in which Kuipers' QSIM grows a tree of possible behaviours[6], and would suffer from the same fate as QSIM, namely a branching behaviour which is very difficult to control, were it not for the fact that there is a further set of constraints that we can bring to bear, namely constraints on the duration of particular region-states. For example, we know approximately the time that the ventricle will stay polarised before spontaneous depolarisation takes place. Likewise we can put upper and lower limits on the repolarisation time of the A/V node. Any behaviour which violates these constraints can be pruned.

Temporal reasoning

Temporal constraint management is handled by a general purpose temporal reasoner which is described in greater detail elsewhere[7]; it is similar to the temporal database described by Dean and McDermott[8]. We maintain a database of temporal intervals in which different types of information can be stored:

- Allen's 13 primitive temporal relationships (after, before, during, meets, etc) between intervals[9];
- start times of intervals in absolute time units; imprecision is handled by expressing these in [min,max] pairs; likewise for end times and durations of intervals;
- times between the start and/or end points of intervals (again using [min,max] pairs).

The temporal reasoner maintains consistency between all of this information, and informs the higher level problem solver if any inconsistencies are detected. Searching is made more efficient by the incorporation of an ATMS[10], which allows different branches of the search space of behaviours to be explored easily by changing context.

The temporal reasoner is used in qualitative simulation by mapping individual region-states into temporal intervals. We can then express the state transition rules using an appropriate sub-set of Allen's relationships. Consider, for example the D-P state of the atrium. We know that this state is *met by* the previous P state of the atrium. It also *meets* the following D state of the atrium and if the A/V node is

in a P state at that time, then the A/V node will enter its D-P state (which is *met by* the atrial D-P state):

```
<---- P ---|----- D-P ----|---- D ---->     atrium
                <---- P -----|--- D-P ----->  A/V node
```

In addition, as we have said, we can impose constraints on the durations of individual states. For example, in a normal heart, the A/V node takes between 120 and 200 msec to depolarise (D-P state). Thus as we proceed along a particular branch of the tree of behaviours, we may come to a point where we cannot satisfy all of the temporal constraints. That branch is then pruned.

So far we have been exploring *all* of the realisable behaviours from a particular model. However, if we have a segmented ECG signal, we can go further and attempt to interpret those observations in terms of a particular behaviour (or behaviours) of the model. We know that a given surface feature on the ECG corresponds to a particular internal state; thus the P-wave is the external manifestation of the depolarisation of the atrium (D-P state). We can therefore assert that normally an interval corresponding to an atrial D-P state must match exactly (Allen's *equals*) a P-wave interval.

Modelling Pathological Conditions

Our model is specified in terms of:

- the level of aggregation we chose for each region
- the degree of complexity of lobes we allow for each region;
- the temporal constraints we place on the durations of states;
- the rules which govern the ways in which a state transition in one region affect state transitions in others.

It is possible to alter this model in a number of ways in order to represent different pathologies; for example:

- we can introduce new regions to model additional abnormal conducting pathways;
- we can alter state durations to model different electrical properties of a region.

Use of the Model

At present we generate a hypothesised (possibly faulted) model externally, and then run it to see if it is capable of explaining the observed ECG. We would hope to move to a situation where a short list of hypotheses might be derived from a fairly simple set of rules applied to the description of the ECG - a surface, associational rule-based system generating hypotheses to be tested by a deeper, model-based system. As a further refinement, the generate, test, and debug paradigm of Simmons and Davis[11] is attractive.

It is not our intention to offer such an approach as the basis for a system to interpret raw ECG signals; such systems are available commercially and achieve high rates of accuracy. However, as far as we are aware, they rely purely on signal processing and pattern matching, and have no underlying model. When we observe our collaborating cardiologist we find that he uses such models extensively in teaching, both to paramedical staff (nurses and ambulance drivers) and to junior doctors. It would be our ultimate goal to integrate our techniques into a tutorial/critiquing system.

Conclusion

We believe that the technique of applying temporal constraints to control excessive branching in qualitative simulation has a wider application than the specific domain of cardiac electrophysiology. It should in principle be applicable to any simulation methodology in which behaviours are described inherently in terms of temporal intervals, and we are currently investigating its use in conjunction with the QSIM algorithm.

Acknowledgements

We are grateful to all of the members of the Artificial Intelligence in Medicine group for providing a stimualting environment in which to carry out this work, and in particular to Richard Vincent, our clinical collaborator. Nick Gotts has been supported by SERC/Alvey grant IKBS/044; Ian Hamlet has been supported on an SERC studentship.

References

1. J.R.W. Hunter, "Qualitative Models in Medicine: An Artificial Intelligence Perspective," in *Measurement in Clinical Medicine*, pp. 47-52, Institute of Measurement and Control, London, 1986.
2. L. Schamroth, *An Introduction to Electrocardiography (Sixth Edition)*, Blackwell Scientific Publications, 1982.
3. M. Malik, T. Cochrane, and A.J. Camm, "Computer Simulation of the Cardiac Conduction System," *Computers and Biomedical Research*, vol. 16, no. 5, pp. 454-468, 1983.
4. I. Bratko, I. Mozetic, and N. Lavrac, "Automatic Synthesis and Compression of Cardiological Knowledge," *Machine Intelligence*, vol. 11, 1986.
5. N.M. Gotts, "A Qualitative Spatial Representation for Cardiac Electrophysiology," in *AIME-87: Proceedings of the European Conference on Artificial Intelligence in Medicine*, ed. J. Fox, M. Fieschi and R. Engelbrecht, pp. 88-95, Springer-Verlag, 1987.
6. B.J. Kuipers, "Qualitative Simulation," *Artificial Intelligence*, vol. 29, pp. 289-338, 1986.
7. I. Hamlet and J.R.W. Hunter, "A Representation of Time for Medical Expert Systems," in *AIME-87: Proceedings of the European Conference on Artificial Intelligence in Medicine*, ed. J. Fox, M. Fieschi and R. Engelbrecht, pp. 112-119, Springer-Verlag, 1987.
8. T.L. Dean and D.V. McDermott, "Temporal Data Base Management," *Artificial Intelligence*, vol. 32, no. 1, pp. 1-55, 1987.
9. J.F. Allen, "Maintaining Knowledge about Temporal Intervals," *Communications of the ACM*, vol. 26, no. 11, 1983.
10. J. De Kleer, "An Assumption-Based TMS," *Artificial Intelligence*, vol. 28, no. 2, pp. 127-162, 1986.
11. R. Simmons and R. Davis, "Generate, Test and Debug: Combining Associational Rules and Causal Models," *IJCAI-87*, pp. 1071-1078, 1987.