Further Progress in Qualitative Modelling of Cardiac Electrophysiology

Ian Kirby and Jim Hunter

Department of Computing Science, University of Aberdeen, King's College, OLD ABERDEEN, AB9 2UB, Scotland, UK. †

Abstract: A qualitative model of the electrical conducting system of the heart is developed which attempts to represent basic physiology and pathophysiology. State descriptions and state transition rules are used to generate a tree of possible behaviours from a given start state. We demonstrate that without any further knowledge, explosive combinatorial branching occurs, but that qualitative information on state durations can be employed to reduce very significantly the number of behaviours. An initial indication is given as to how the approach could be used to explain observed electrocardiograms (ECGs) in terms of the underlying physiological processes, and the advantages of generating envisionments rather than simulations are presented.

1. INTRODUCTION

This paper explains and discusses an approach to the qualitative modelling of the electrophysiological changes which take place in the human heart as it beats either normally or abnormally; it builds on work reported in [1] and [2].

Our ultimate objective is to interpret electrocardiograms (ECGs) in terms of the underlying cardiac electrophysiological processes that give rise to them. We will not cover all aspects of ECG interpretation - only those that relate to the diagnosis of *arrhythmias* (disturbances of rhythm). This excludes, for example, the interpretation of features of the ECG that enable cardiac enlargement and myocardial infarction (heart attack) to be diagnosed.

Qualitative modelling [3] describes the state of the process or entity being modelled in non-numerical terms. A common approach, which we take here, is to define a formalism to represent the state of the system at a moment of time or over a interval, and to develop a set of constraints (or rules) which specify what transitions between states are possible.

In subsequent sections we will review some necessary background knowledge in cardiac electrophysiology; we will then develop our formalism for describing electrical state, and derive sets of rules which govern transitions between these states. The use of these rules on their own to develop a tree of behaviours from a given starting state leads to an explosive growth in the number of states explored. We show how knowledge about state duration, expressed qualitatively, can provide very significant constraints on this growth. We then dicuss how such models can be used for ECG interpretation. Simulation proves to be inappropriate in a domain where we are trying to represent continuous cyclic or quasi-cyclic activity, and we demonstrate the advantages of producing envisionments. We conclude by comparing the somewhat better known approach of KARDIO [4] with ours.

This is an updated version of the paper "Using Quantitative and Qualitative Constraints in Models of Cardiac Electrophysiology" which appeared in *Artificial Intelligence in Medicine*, 3 (1991) 41-61.

† Phone: (+44 224) 272287; FAX: (+44 224 487048); email: ikk@cs.abdn.ac.uk, jhunter@cs.abdn.ac.uk

2. BASIC CARDIAC ELECTROPHYSIOLOGY

The physiological function of the cardiac conduction system is to coordinate the contraction and relaxation of the heart, so that it pumps blood into the arteries in an efficient fashion, and responds appropriately to changing levels of activity by speeding or slowing the heart rate. As used here, the term 'cardiac conduction system' covers both the cells of the heart muscle (myocardium), and those of the *specialized* conduction system running through it. It is the cells of the myocardium that are directly responsible for cardiac contraction; the reversal of electrical polarisation across the surface of myocardial cells (depolarisation) causes contraction of the muscle fibres to which they belong, thus forcing the blood in the required direction. When the cells repolarise, the fibres relax again.

The cells of the specialized conducting system do not belong to fibres that are able to contract; their role is to ensure that the myocardial cells depolarise, and hence contract, in the correct sequence. Depolarisation is a 'domino' process; depolarised cells influence their still-polarised neighbours to depolarise in turn. A depolarisation wave begins when one or more cells depolarise spontaneously. The wave then spreads through the conductive tissue of the heart, as long as it continues to encounter cells ready to depolarise; normally its progress will be continuous. Repolarisation, on the other hand, is not a domino process; the time at which a cell repolarises is determined not by the repolarisation of its neighbours, but by the time it depolarised, its intrinsic properties, and the general physiological conditions in the local extracellular environment. Nevertheless, there generally *appears* to be a coordinated repolarisation 'front', with nearby cells repolarising close together in time. This is a result of the smooth distribution of the intrinsic rates of recovery of cells in the normal heart.

2.1. Anatomy and Function of the Cardiac Conducting System

Anatomically, the heart consists of four muscular 'chambers': the left and right atria, and the corresponding ventricles. The two sides of the heart are separate pumps working in parallel, the left side taking oxygenated blood from the lungs and pumping it out to the tissues, the right collecting deoxygenated blood from the tissues and returning it to the lungs. When considering the electrical processes of the heart, however, the left/right division is not very important; the myocardium can mostly be considered as consisting of two parts, the atrial myocardium and ventricular myocardium. The atrial and ventricular parts of the myocardium are electrically separated from each other by a ring of insulating tissue, which in the normal heart is pierced only by a narrow 'atrioventricular bundle' or AV node. This is one part of the specialized electrical conducting system, depicted in Figure 1.

The other important parts are the sinus node, which is a small structure in the wall of the right atrium where the wave of depolarisation begins in the course of a normal heartbeat, and the ventricular conducting system, which carries the depolarisation wave from the AV node to the entire ventricular myocardium.

The ventricular conducting system can be subdivided into the bundle of His (not shown) and the left and right bundle branches which carry the depolarisation impulse to their respective ventricles.

In some abnormal hearts there are extra electrical connections between the atrial and ventricular myocardia in addition to the AV node. The extra electrical connections which can be provided by these additional structures allow rhythms that cannot arise in an anatomically normal heart.

2.2. Events In The Cardiac Conducting System During A Normal Heartbeat

Since the cardiac conducting system (the myocardium plus the specialized conducting pathways) is composed of cells, it is useful to look at the events within individual cells over the course of a heartbeat - although it turns out that for clinical purposes the division into cells can largely be ignored and the conducting tissue regarded as a continuous, solid conducting medium.

Each cell in the cardiac conducting system goes through repeated cycles of depolarisation and repolarisation. As already noted, an electrically-conducting cardiac cell can be caused to depolarise by the depolarisation of its neighbours; once a cell has depolarised, repolarisation follows automatically after a short period.

We can distinguish two basic types of electrically-conducting cell in the cardiac conducting system: pacemaker cells and non-pacemaker cells. A pacemaker cell will depolarise spontaneously a certain

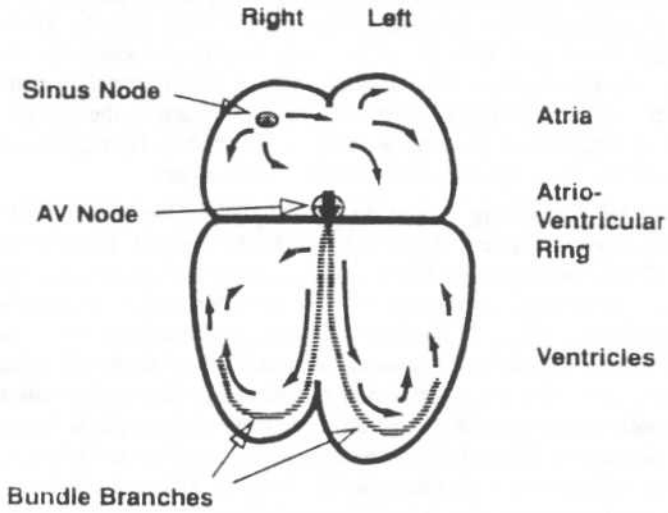

Figure 1: Anatomy and Function of the Cardiac Conducting System

time after its previous depolarisation, if it has not already been depolarised by its neighbours. Pacemaker cells are concentrated in the sinus node, and in the lower part of the AV node. Normally the sinus node cells initiate the heartbeat, the AV nodal pacemakers forming a backup if the sinus impulse fails. Non-pacemaker cells normally have a passive role, passing on depolarisation impulses that they receive from neighbours. However, non-pacemaker cells may develop pacemaker properties under certain pathological circumstances, when their spontaneous depolarisation time decreases (increased *automaticity*).

The physiological properties of cells vary according to the anatomical region and may be altered by disease processes. These properties influence the speed of conduction, the automaticity, and the length of time necessary before a region can transmit another impulse (*the refractory period*).

In a normal heartbeat, as already noted, depolarisation begins in the sinus node, which contains cells that are particularly quick to depolarise spontaneously. Depolarisation then spreads through the atrial myocardium to the AV node. The depolarisation wave then travels through the AV node, where conduction is slower, particularly in the upper part of the node. The delay in the AV node ensures (when the heart is working correctly) that the contraction of the atria 'primes' the main, ventricular pump; that is, it ensures that the ventricles are full of blood when they begin to contract. Either too short or too long a delay in the AV node can reduce the efficiency of the heart.

Reaching the bundle of His, the depolarisation wave speeds up again, then splits first into two as the left bundle branch separates itself off, then into many parts as it travels out into the ventricular myocardium.

2.3. The ECG - What Can Be Seen Of Cardiac Electrical Activity

A standard 'surface' ECG is made by attaching one or more electrodes to the skin and translating the electrical impulses that they detect into deviations from straight, horizontal lines drawn by a pen on squared paper.

On a standard ECG, only three physiological events normally show up as deviations from the baseline: the depolarisation of the atrial myocardium (which produces a deflection called the 'P-wave'), the

depolarisation of the ventricular myocardium (producing the 'QRS-complex'), and the repolarisation of the ventricular myocardium (producing the 'T-wave'). These types of event show up on surface ECGs because they involve a sufficient mass of tissue undergoing electrical change in a coordinated fashion; the repolarisation of the atria, and events in the specialized conducting system, fail to show up. In an ECG of a normal heartbeat, the types of deflection corresponding to the three types of event appear repeatedly in the order: P, QRS, T - each deflection normally being separated from the previous one by a flat part of the trace. The flat part of the trace between each P-wave and the subsequent QRS-complex corresponds approximately to the passage of a depolarisation wave through the AV node. Figure 2 shows a normal ECG, and Figure 3 shows a schematic version of the ECG produced by a single typical heartbeat.

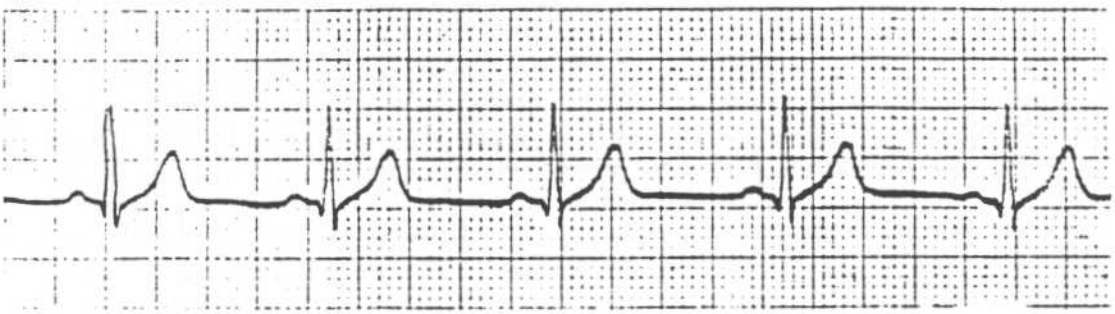

Figure 2: ECG Trace of Normal Sinus Rhythm

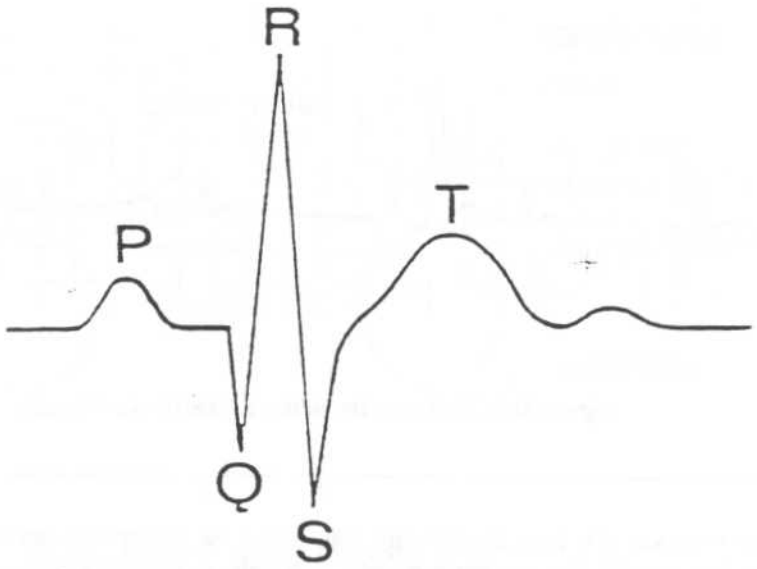

Figure 3: Schematic Diagram of ECG Trace of Normal Sinus Rhythm

2.4. Arrhythmias

Arrhythmias are caused by defects either of impulse conduction or generation and manifest themselves as ECGs which do not correspond to the pattern of Figure 2. Defects can occur in combination, but that it is not possible to derive the ECG for the combined arrhythmia from the individual ECGs of the single arrhythmias. Currently, our model can handle around 15 different arrhythmias; here we describe two which will be used as examples in subsequent sections.

2:1 Block: There are a number of conduction defects which can arise in the AV node; in 'second degree block', the AV node does conduct, but takes longer than normal to recover. This has the effect that the next wave of depolarisation arriving from the atria finds the AV node still in a refractory state, and that atrial beat fails to be conducted to the ventricles. By the time the third beat arrives, the AV node has recovered and the beat is conducted normally. One variant is **2:1 block**, which gives rise to a characteristic pattern in which every second atrial beat is blocked - as shown in Figure 4.

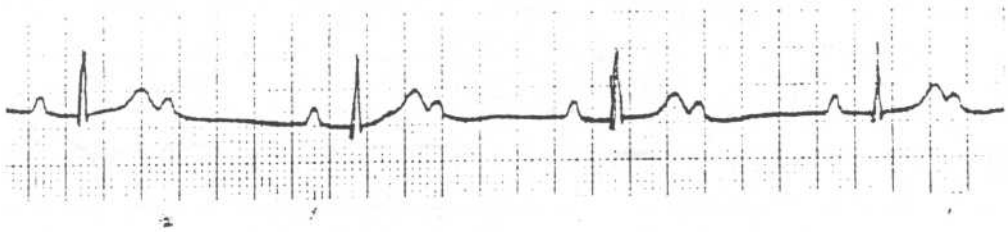

Figure 4: ECG Trace of 2:1 Block

Figure 5: ECG Trace Including a Ventricular Ectopic

Ventricular ectopics: We have already said that, under the appropriate (possibly pathological) circumstances, it is possible for any group of cells within the conducting system to act as a pacemaker. This has its positive side, since if the sinus node fails, another part of the system can take over and keep the heart going. However, it can happen that other pacemakers can become active even when the sinus node is functioning normally. Either the new pacemaker fires regularly at a higher rate than the sinus node, in which case it takes over its function completely, or else it does so occasionally, giving rise to occasional beats - so-called 'ectopics'. If the sporadic pacemaker is located in the ventricles, such a beat is called a **ventricular ectopic**; an example is given in Figure 5.

3. STATE REPRESENTATION AND STATE TRANSITIONS

Our current work, as discussed in this paper, is concerned with modelling the operation of the conducting system at the level of physiological detail described in the previous section. We start by considering a normal beat; only then will we consider the introduction of abnormalities. In a normal beat, depolarisation is initiated in the sinus node. The wave of depolarisation passes through the atrial myocardium, and when it reaches the boundary with the AV node initiates depolarisation in the latter in a similar fashion. The wave eventually reaches the boundary with the ventricles, and initiates depolarisation there. Once depolarised, the different parts will then repolarise in their normal way.

We have developed a *class* of models, the **Ticker** models, which represent the space in which the processes of interest take place as a set of one or more anatomical *regions*; regions are disjoint, and together cover the whole of the space concerned. They are connected to one or more neighbouring regions across *region boundaries*. We say *models* since we obtain a specific member of the class when we define the set of regions, and their connectivity; these form the fixed aspect of the representation of space in a region model. In our models, neither the regions nor the region boundaries are assigned any spatial properties - other than connectivity.

We have worked with a number of models. In this paper we will refer to:

- a four region model, in which the regions are: sinus node, atria, AV node, ventricles (Figure 6A);
- a three region model, in which the sinus node and the atrial myocardium have been collapsed into one region which we will call the atria (Figure 6B); whenever we refer to the 'atria' in connection with this model, this will be taken to include the sinus node;
- a seven region model in which the ventricles are expanded into: the His bundle, the left and right bundle branches, and the ventricular myocardium (Figure 6C).

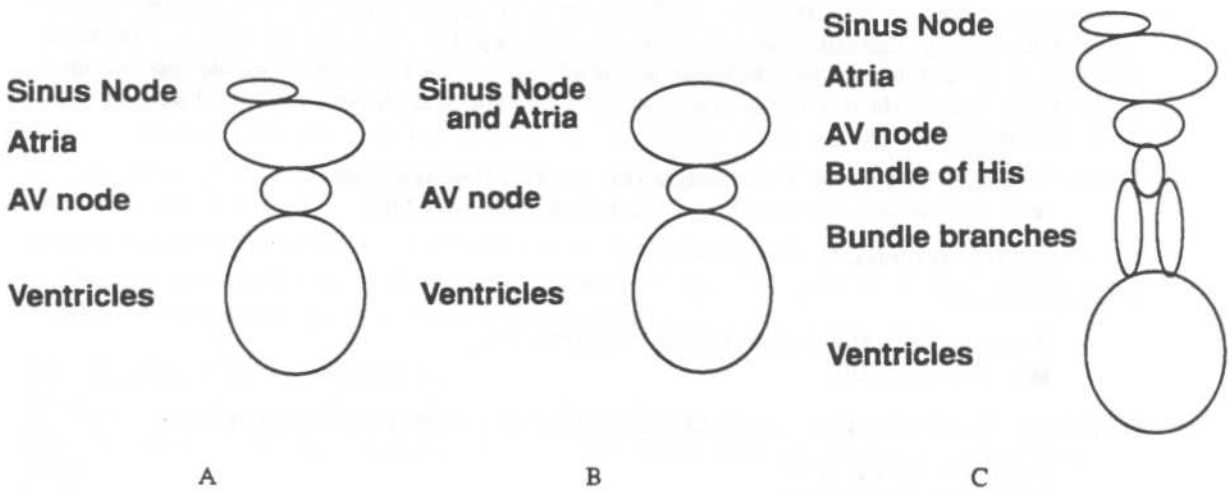

Figure 6: Ticker Anatomical Models

The state of a region at a particular time is refered to as a *region state*. Our implementation for the normal heart is based on a simple set of descriptors for those region states:

P the region is completely polarised,

DP the region is depolarising - partly polarised and partly depolarised,

D the region is completely depolarised, *or* the region is repolarising (i.e. partly depolarised and partly polarised).

A specific region state will be indicated by putting the name of the descriptor, in parentheses, after the name of the region: thus atria(DP) means that the atria are depolarising.

Which region states can be succeeded by which others (what *region state transitions* are permitted) can be represented as a set of rules, each rule being of the form:

if <region> is in <state1>
then <region> can make a <state1> → <state2> transition

or in a more compressed form:

if <region> (<state1>) **then** <region>(<state2>)

We can therefore express the natural cycle of depolarisation and repolarisation within a region by the rules:

if <region>(P) **then** <region>(DP)
if <region>(DP) **then** <region>(D)
if <region>(D) **then** <region>(P)

These rules, however, do not properly model the normal function of the heart in that they imply that all regions can make the P → DP transition spontaneously. This is only true for pacemaking regions and we must therefore remove the P → DP rules for all others.

The P → DP transition can be triggered by the arrival of a wave of depolarisation at the boundary with a connecting region; such an arrival is represented by a DP → D transition in the connecting region. However the wave of depolarisation only crosses the boundary if the receiving region is in a receptive state, which, in our simplified representation, means a P state; if the region receiving the depolarisation is not in that state, conduction will not take place and the impulse is blocked. To model this, we need rules of the following type:

if <region-2> is in state P **and** <region-1> makes a DP → D transition
then <region-2> will make a P → DP transition at the same time

where <region-1> is connected to <region-2>

Or, in our notation:

if <region-2>(P) **and** <region-1>(DP) → <region-1>(D)
then <region-2>(DP)

For example, for our three region model shown in Figure 6B, we have the following rules:

if atria(P) **then** atria(DP)
if atria(DP) **then** atria(D)
if atria(D) **then** atria(P)

if AV_node(DP) **then** AV_node(D)
if AV_node(D) **then** AV_node(P)

if ventricles(DP) **then** ventricles(D)
if ventricles(D) **then** ventricles(P)

if AV_node(P) **and** atria(DP) → atria(D) **then** AV_node(DP)
if ventricles(P) **and** AV_node(DP) → AV_node(D) **then** ventricles(DP)

The first seven rules express the autonomous state transitions in a region. The last two are consulted to see if a change in one region will affect another i.e. conduct an impulse.

A *global state* for this three region model is built up of the three individual region states and we can construct a global state vector as follows:

[state_of_atria state_of_AV_node state_of_ventricles]

For example, a state in which the atria are depolarised (D), the AV node is depolarising (DP), and the ventricles are still polarised (P) would be written:

[D DP P]

For simplicity we assume that two independent region state transitions cannot occur simultaneously.

Given a particular starting state - say [P P P], we can simulate all possible sequences of global states (behaviours). Because there are usually several possible global state transitions available at any point, these behaviours form a tree in which any particular behaviour forms a branch. We will often refer to global states as *nodes* in the context of a tree of behaviours.

This generation of a tree of behaviours is similar to Kuipers' QSIM [5] and suffers the same fate, that of an explosive branching which is very difficult to control. The unconstrained tree grown to depth 13 (corresponding to about 2 heartbeats) will have 13,780 nodes. As described in the following sections, considerable pruning can be achieved by applying knowledge of temporal constraints.

4. TEMPORAL CONSTRAINTS

We consider region states as corresponding to temporal intervals (c.f. Allen [6]). Figure 7 shows possible relationships between intervals for a normal heartbeat in a three region model. Our use of temporal constraints is based on the representation of the durations of these intervals.

Approximate numerical data on the durations of region states can be extracted from textbooks, clinicians, etc. and initially we attempted to represent them as [min max] pairs; for example, in a normal heart, the AV node takes between 80 and 120 msec to depolarise (DP). Although this approach achieves substantial pruning of the tree of behaviours, we felt that there were a number of weaknesses:

- clinicians seem to reason more qualitatively than this; either they don't know these figures, or if they do, they don't use them;
- minor changes to the numbers chosen may produce significant variations in results;
- we cannot tell whether two similar parts of the tree are significantly different; we would like to 'collapse' equivalent parts giving us a graph - possibly cyclic.

4.1. Qualitative Temporal Constraints

We have chosen to represent the qualitative knowledge that doctors have about durations by generating a partial ordering of the region states according to their relative durations. Thus for the normal heart we have:

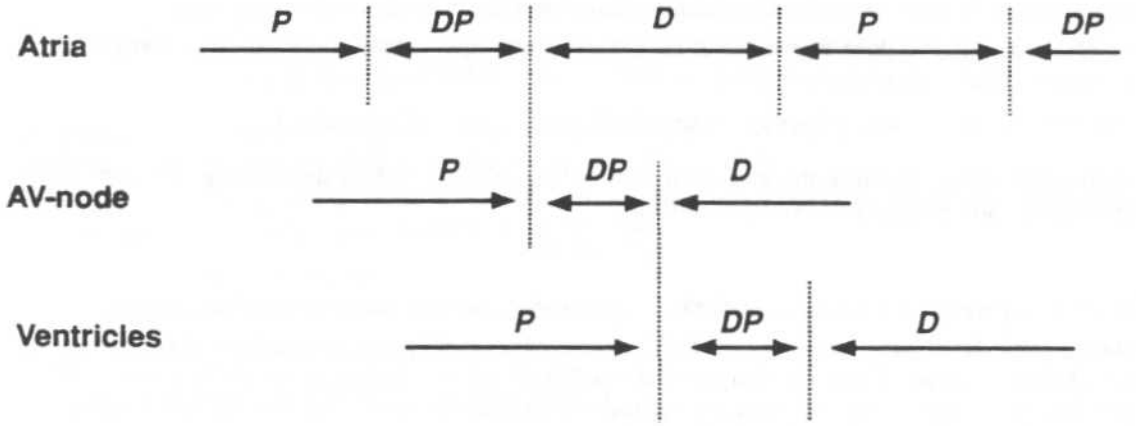

Figure 7: Relationships between Intervals for a Normal Heartbeat

ventricles(P)	↑	*longest*
atria(P) AV_node(P)		
AV_node(D)		increasing duration
ventricles(D)		
atria(D)		
atria(DP) AV_node(DP) ventricles(DP)		*shortest*

Using these qualitative relative durations we can again constrain the generation of our tree of states. We can assert that one region (in a particular state, at a particular time) must end (make a transition) before, after or at an unknown time in relation to another. These *end point relations* can be written for regions A and B:

A < B - A must make a transition before B
A > B - A must make a transition after B
A u B - either may make a transition first

For completeness we should include the transition:

A = B - A and B make a transition at the same time

We do not do so because for a physiological system we have never found the need to specify that two durations are exactly equal; this is also consistent with our previous assumption that two independent region state transitions can not occur simultaneously.

These constraints can be asserted in three cases:

1) *To relate two regions making a transition at the same time*

If two regions were to enter new states at the same time we could read from the above table which of them must make a transition first. For example, if atria(D) and AV_node(DP) are entered at the same time, this partial ordering adds the constraint that the AV node must transition before the atria, because AV_node(DP) is shorter than atria(D) - as illustrated in Figure 7.

transition	end point relation
[DP P *any*] → [D DP *any*]	atria(D) > AV_node(DP)

2) *To relate one region making a transition to another which is not*

If the atria are already in the D state, and the ventricles make a transition to P, we can say that the atria must make a transition before the ventricles. However, if the ventricles are already in P and the atria enter D the relation is unknown, since we cannot say how long ventricles(P) has existed and therefore cannot say how long before we might expect a transition.

transition	end point relation
[D *any* D] → [D *any* P]	atria(D) < ventricles(P)
[DP *any* P] → [D *any* P]	atria(D) u ventricles(P)

3) *To relate two regions neither of which is making a transition (but another did)*

When one region makes a transition, the relation of the two other regions can be carried forward. If the AV node makes a transition, the relation between the atria and the ventricles will be the same.

transition	end point relation
[D *any* P]	atria(D) < ventricles(P)
↓	
[D *any* P]	atria(D) < ventricles(P)

Consider Figure 8 which shows a tree of behaviours starting with the myocardium relaxed and all regions in the P state. The relationships for each of the nodes is given in Table 1.

The most general qualitative constraints we can apply are that all relations are 'u' so that for node 1, all relationships are undefined.

Applying the 'normal heart' transition rules, the only rule which succeeds is atria(P) → atria(DP) to give node 2. Although atria(DP) is shorter than AV_node(P) and ventricles(P), we cannot say anything about the relations between the atria and the AV node, and the atria and the ventricles because AV_node(P) and ventricles(P) may have started a long time ago and may now be due to make a transition. All relations are still therefore undefined for node 2.

However between nodes 2 and 3 we have simultaneous transitions for the atria and the AV node which allows us to establish that atria(D) > AV_node(DP). Thus node 4, which has the atria making a transition before the AV node is pruned. Other nodes which are pruned in a similar manner are indicated by the faint arrows.

Because we have a finite number of descriptors, it is possible to find identical nodes and to 'collapse' the tree to form a graph so that only unique nodes are expanded. In the worked example we have been discussing, it can be seen that two of the nodes (11 and 13), on the last level shown, are identical and the branches can be collapsed giving an acyclic graph. In fact, collapsing can take us even further since if the worked example had been continued, we would have seen [P P P] configurations appearing and shortly afterwards (at a depth 4 greater) all of the branches of the tree would have generated nodes equivalent to nodes at the first few levels of the tree. Thus we produce a cyclic graph representing the cycle of events in the normal heartbeat.

The operation of our temporal constraint mechanism can be seen as performing one of the functions of Williams' temporal constraint propagator [7], where the information about temporal relations consists of the partial ordering of durations. Weld also uses temporal constraints to constrain the proliferation of

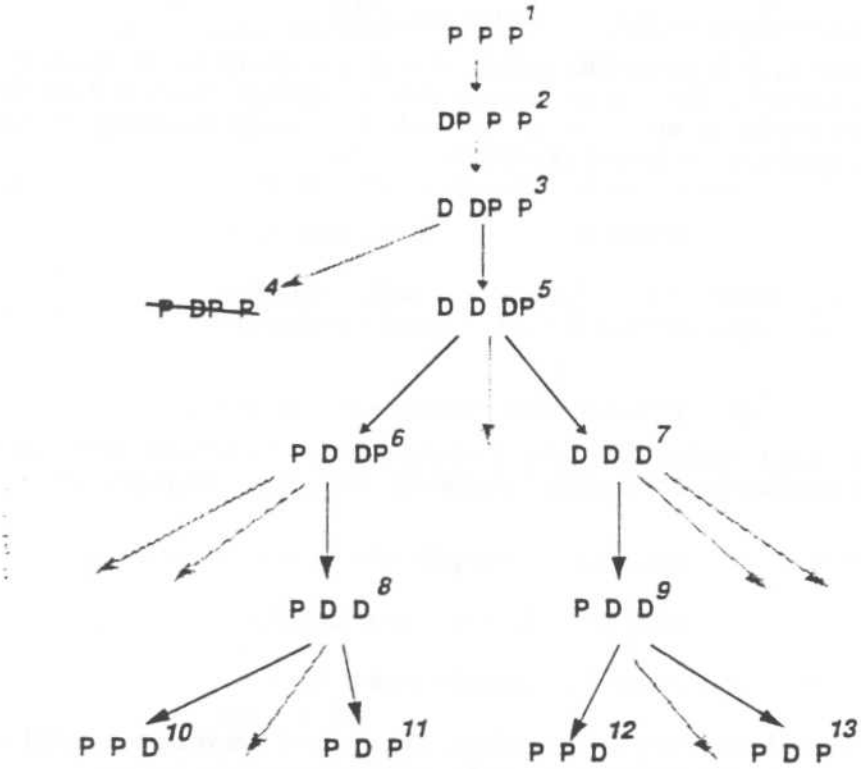

Figure 8: Application of Qualitative Constraints

states in HR-QSIM [8]; the main differences between his work and ours are: (i) his durations are derived from the model - in our case they represent external inputs (derived from our expert cardiologist); (ii) he is concerned with the relationships between infinitessimal, finite and infinite values - all of our values are finite.

Our group has also considered the application of externally obtained temporal constraints to QSIM [9].

4.2. Modelling Pathological Conditions

Our model is specified in terms of:

- the number of regions represented,
- the level of detail used to represent the electrical state of a 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 affects state transitions in others.

Node	atria - AV node	atria - ventricles	AV node - ventricles
1	u	u	u
2	u	u	u
3	>	u	u
4	>	u	u
5	<	u	>
6	>	>	>
7	<	<	u
8	>	u	u
9	>	>	u
10	u	u	>
11	>	<	<
12	u	>	>
13	>	<	<

Table 1: Application of Qualitative Constraints

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

- we can introduce new regions to model additional abnormal conducting pathways;
- we can introduce new descriptors of electrical state (fibrillation, flutter, and abnormal depolarisation);
- we can alter the temporal constraints to model different electrical properties of a region;
- we can specify a region to be blocked (incapable of transmitting impulses);
- we can specify a region to have become a possible pacemaker.

One example is the 2:1 block we described in section 2.4; recall that this is a conduction defect of the AV node in which it takes a longer time than normal to repolarise. As a result the next atrial depolarisation occurs *before* the AV node has fully recovered; the AV node therefore fails to conduct the impulse. For the third atrial discharge, the AV node is ready and conducts normally. We model this by moving the duration of AV_node(D) up the partial ordering:

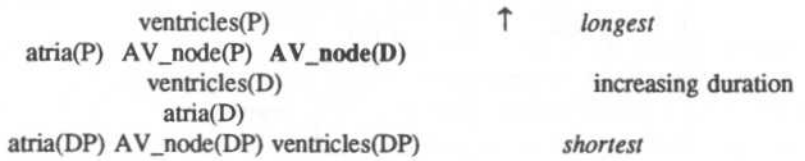

Another example discussed in section 2.4 is that of ventricular ectopics, where the automaticity of the ventricle is increased, allowing it to act occasionally as the pacemaker. We implement this in the model by allowing the ventricles to be a pacemaker, enabling spontaneous ventricles(P) → ventricles(DP) transitions.

5. ECG INTERPRETATION

So far we have been describing models of different levels of sophistication and discussing how to explore all of the realisable behaviours from a particular model. A closely related problem is that of *measurement interpretation*. In this context, 'measurement' means the features of the observable ECG. By 'interpretation' we mean the identification of a model(s), probably pathological, which is capable of

generating behaviours which explain the observed ECG. There exist commercial pattern-matching systems for ECG interpretation but these do not represent or explain what is going on at the physiological level; they use mainly statistical methods to identify individual deviations from the baseline and then to produce a diagnostic hypothesis; there is no attempt to model the events responsible for the ECG.

We know that a given ECG feature corresponds to a particular internal state; for example, the P-wave is the external manifestation of the depolarisation of the atria - atria(DP). We show in Table 2 how observable ECG features match a region state, and thus define a vocabulary for the qualitative descriptions of ECGs. The state of the AV node is not directly observable on the surface of the body and activity in the ventricles masks any other activity.

atria	AV node	ventricles	Notation	Meaning
DP	*any*	P	-P-	P-wave
any	*any*	DP	-QRS-	QRS-complex
P *or* D	*any*	D	-T-	T-wave
P *or* D	*any*	P	--	Nothing observable
DP	*any*	D	-P+T-	P wave falling on T

Table 2: Relating Internal States to the ECG

Any behaviour can thus be translated to provide a description of the corresponding ECG sequence. For example, consider the ECG sequence shown in the first column of Table 3 which could arise from 2:1 block. One possible 'explanation' for this sequence is provided by the behaviour in the second column which is derived from the 2:1 block model described in section 4.3. Note the AV node has not yet recovered, and therefore 'blocks' in the transition: [DP D P] → [D D P]

ECG description	**Ticker-R3 description**
--	[P P P]
-P-	[DP P P]
--	[D DP P]
-QRS-	[D D DP] [P D DP]
-T-	[P D D]
--	[P D P]
-P-	[DP D P]
--	[D D P] [D P P] [P P P]
-P-	[DP P P]
etc.	*etc.*

Table 3: Example of ECG Matching for 2:1 Block

6. ENVISIONMENT

One of the problems in working with simulation is the need to identify a particular state from which to start. In a continuous cycling system such as the heart, it is meaningless to try to find such a state. One of our objectives is to characterise the classes of behaviour which a given model can give rise to, and it is easier to see how this might be done if all the behaviours were available at the same time, rather than being generated incrementally in simulation.

For these reasons we have been recasting our model to generate envisionments; according to Forbus, an envisionment "represents all possible qualitative states a system may take on, and all legal transitions between them" [10]. One of the problems with generating such envisionments is the potentially large number of distinct states (nodes). The maximum number of distinct states global states for Ticker models is $d^r o^c$

where: d: is the number of state descriptors
r: is the number of regions
o: is the number of relation operators (<, >, or u)
c: is the number of pairs of regions, i.e. $r(r-1)/2$

For our three region model with three state descriptors, the maximum number is $3^3 \times 3^3 = 729$. For a model with four regions and six state descriptors, this rises to $6^4 \times 3^6 = 944{,}783$ nodes.

However, some of these states are self-contradictory, as for example, the following set of end point relations:

$$A > B \quad B > C \quad C > A$$

Others are implicitly self-contradictory, e.g.

$$A > B \quad B > C \quad C \; u \; A$$

From the first two statements it is obvious that A > C. Of the 27 combinations of end point relations in the three region model, 8 are illegal in this way.

The number of nodes can be further reduced if we assume that all behaviours must be cyclic; nodes with no predecessors and nodes with no successors can be pruned.

Even with such reductions, the potential for unacceptable demands on computational resources exists, unless appropriate steps are taken. We will not describe these steps here (they involve the use of compressed representations, saving of intermediate computations, and careful consideration of the order in which nodes are generated), but we have been successful to the point where most four region envisionments can be generated in a few minutes on a SUN 4/60 with 16 Mbytes of physical memory. Our implementation is in POP-11, an AI language which combines the list-processing and pattern-matching facilities of LISP with a procedural syntax and semantics.

Table 4 shows how many nodes there are in the envisionments of various arrhythmias; the meanings of the other columns will be explained later.

We have discussed earlier the need to compare the predictions of our model (as expressed by the envisionments) with the surface ECGs. One important step towards this is to abstract the envisionment to the level of the ECG. The resulting graph will represent a class of ECGs - those ECGs which are consistent with the model which was used to generate the envisionment. Any specific ECG from that class can be generated as a path (probably involving cycles) through the graph. The technique is simple, and involves aggregating together any two nodes which are neighbours, and which represent the same ECG feature. Figure 9 presents an artificial example. The process of abstraction brings a further reduction in the size of the graph as shown in Table 4.

Given that we are interested in cyclic behaviour, it is of interest to identify the cycles in these abstracted graphs. We have used an algorithm to identify unique elementary cycles: *elementary* in that no cycle contains the same node twice (in which case it could be split into two simpler cycles) and *unique* in that cycles passing through the same nodes in the same order are recognised as being the same. Table 4 shows the number of such cycles for the different models.

Arrhythmia	Maximum size of envisionment	Actual size of envisionment	Actual size of abstraction	Number of unique cycles in abstraction
Sinus rhythm	59,049	161	30	25
Sinus bradycardia with junctional rhythm	186,624	244	56	301
Sinus rhythm with 2:1 AV nodal block and and ventricular ectopics	186,624	301	77	660
Sinus node disorders, atrial ectopics, and ventricular fibrillation	455,625	106	42	9

Table 4: Results for Four Region Envisionments

Using such abstracted graphs to decide whether the originating model could give rise to a specific ECG involves trying to fit the ECG to the graph in a fairly obvious way. Such a technique is similar to the method of interpreting observations of physical systems as described by Forbus [11].

With such a complex set of models, it is clearly necessary to establish a mechanism for validating them. The most obvious approach is, for each arrhythmia, to get a clinician to describe the class of ECGs to which it should give rise. We would then set up a Ticker model corresponding to this arrhythmia, create an envisionment, abstract it, and then compare the abstraction with the ECG class description; see Figure 10. In order to achieve this, we are currently working on the following problems:

- how to define a formal language for describing classes of ECGs;
- how to design an algorithm for making the comparison.

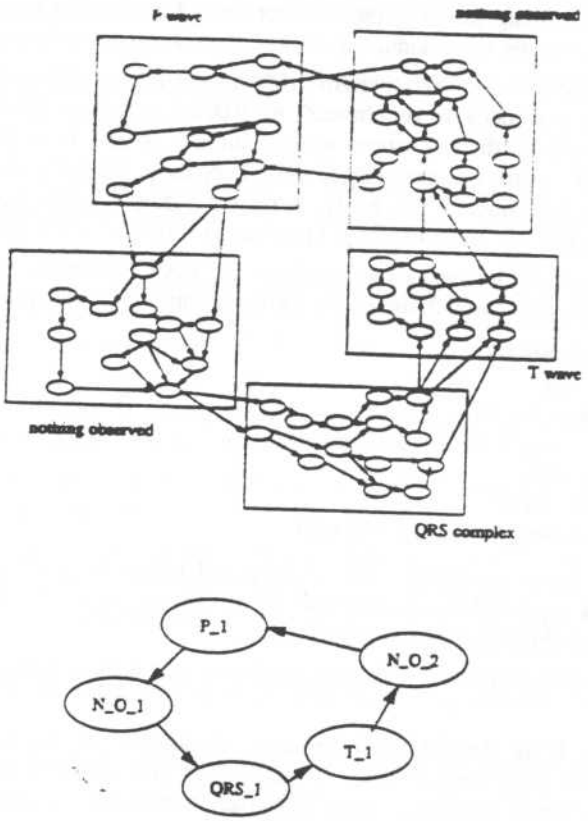

Figure 9: Artificial Example of Abstraction

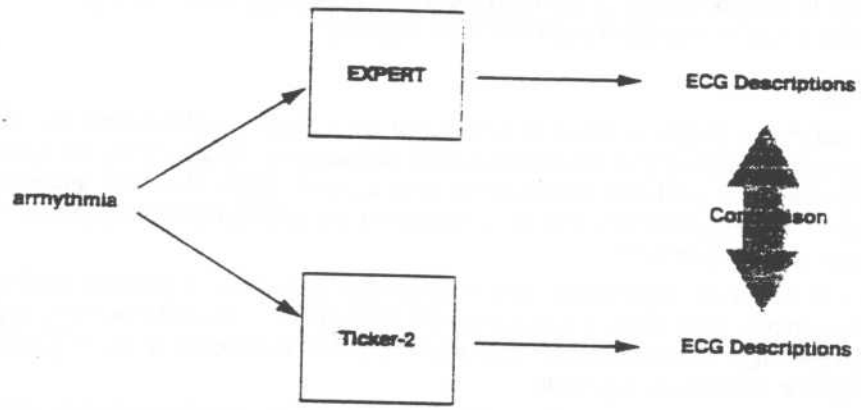

Figure 10: Validating the Ticker Models

7. KARDIO

We have been considerably influenced in this work by the qualitative model of the heart developed as part of the KARDIO project [4]. This is used to generate 'shallow' knowledge in the form of a knowledge-base of rules describing the ECG traces which should correspond to particular arrhythmias. Although KARDIO operates in the same domain as Ticker, there are three main differences.

Firstly, the approaches are at different levels. KARDIO does not follow anatomical structure and physiological function as closely as does Ticker. Instead, KARDIO has a more functional orientation: the anatomy is secondary. Sometimes these functions would clinically be said to be performed by the same anatomical component of the heart. For example, KARDIO separates the AV node into a generator and a conductor, although it is anatomically one entity. A second example is that KARDIO has structures where two impulse descriptions are combined (e.g. in the atria) which have no physical analogue.

Secondly, whereas Ticker models moment by moment events, KARDIO reasons about *averaged* or typical behaviour. For example, in KARDIO, the ECGs corresponding to 2:1 block are described thus:

rhythm_QRS = regular
dominant_P = normal
rate_of_P = between_60_100
relation_P_QRS = after_P_some _QRS_missing
dominant_PR = normal
dominant_QRS = normal
rate_of_QRS = between_60_100 or under_60

Unlike Ticker, KARDIO is not concerned with the individual events of depolarisation and repolarisation. As a result it can not explain, for example, how a particular atrial ectopic will reset the timing pattern of subsequent P-waves generated from the sinus node.

Finally, KARDIO does not represent basic physiological concepts such as *refractory period*, *automaticity*, etc.

Experiments have indicated that the KARDIO programs are accurate in diagnosing arrhythmias from ECGs and predicting ECGs given arrhythmias. KARDIO must be considered as one of the most successful and complete qualitative models to date. However, it seems to us that KARDIO and Ticker must not be seen as being in competition, but rather as different complementary levels in the spectrum of possible representations.

In connection with our work on validating the Ticker models (see previous section) we have been considering using KARDIO as the 'expert' in Figure 10. However we are not completely satisfied with KARDIO's language to describe classes of ECGs. There are certain ambiguities in its semantics, and it does not handle a few specific cases that we consider to be important.

8. CONCLUSIONS

Cardiac electrical activity is complex, and it is certain that the concepts explored here will not be sufficient to represent this complexity in its entirety. Indeed, although we are attempting to represent a physical system about which a good deal is understood, there are many areas where this understanding fades into speculation. If representing well understood systems is not simple, representing poorly understood systems is correspondingly harder!

Nevertheless, we feel that some progress has been made in clarifying some of concepts involved. In particular, when we compare what we have undertaken with KARDIO, it is clear that we have explored another 'level' of knowledge. Representing the same physical system at different levels of detail is an activity that will become increasingly important.

The technique of applying temporal constraints to control excessive branching in qualitative simulation has been explored by a number of workers; we hope that this practical demonstration will be of interest outside our specific domain.

Likewise we feel that the use of envisionments and their abstractions to characterise classes of behaviours (particularly those involving cycles) has more general applicability; we intend to develop this aspect further in the future.

Acknowledgements

Some of the research reported here was funded by the Alvey/SERC grant: GR/D 18271 (IKBS/099) - *An Intelligent Model Based System for Diagnosis in Cardiology.* Ian Kirby has been supported by an SERC studentship. We are grateful to Nick Gotts, Ian Hamlet and Richard Vincent for significant contributions.

We would like to thank our two anonymous referees for their helpful suggestions.

References

1. N.M. Gotts, J.R.W. Hunter, I. Hamlet, and R. Vincent, "Qualitative Spatio-Temporal Models of Cardiac Electrophysiology," AUCS/TR8903, Department of Computing Science, University of Aberdeen, 1989.
2. J.R.W. Hunter, N. Gotts, I. Hamlet, and I. Kirby, "Qualitative Spatial and Temporal Reasoning in Cardiac Electrophysiology," in *AIME-89: Proceedings of the Second European Conference on Artificial Intelligence in Medicine*, ed. J.R.W. Hunter, J. Cookson and J. Wyatt, Springer-Verlag, 1989.
3. D. S. Weld and J. de Kleer, (Editors), *Readings in Qualitative Reasoning about Physical Systems*, Morgan Kaufmann, San Mateo, CA, 1990.
4. I. Bratko, I. Mozetič, and N. Lavrač, *KARDIO: A Study in Deep and Qualitative Knowledge for Expert Systems*, MIT Press, 1989.
5. B.J. Kuipers, "Qualitative Simulation," *Artificial Intelligence*, vol. 29, pp. 289-338, 1986.
6. J.F. Allen and H.A. Kautz, "A Model of Naive Temporal Reasoning," in *Formal Theories of the Commonsense World*, ed. J.R. Hobbs and P.C. Moore, pp. 251-268, Ablex, Norwood, New Jersey, 1985.
7. B.C. Williams, "Doing Time: Putting Qualitative Reasoning on Firmer Ground," *AAAI-86*, vol. 1, pp. 105-112, 1986.
8. D.S. Weld, "Exaggeration," *Artificial Intelligence*, vol. 43, pp. 311-368, 1990.
9. I. Arana and J. Hunter, "The Use of Temporal Knowledge in QSIM," Department of Computing Science Report, University of Aberdeen, 1990.
10. K.D. Forbus, "The Logic of Occurrence," University of Illinois Report No. UIUCDCS-R-86-1300, 1986.
11. K.D. Forbus, "Interpreting Observations of Physical Systems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 13, pp. 350-359, 1987.