A System for Pacemaker Treatment Advice[1]

Peter Lucas[2] and **Ruud Kuipers** and **Frederick Feith**[3]

Abstract. Previously it has been shown that the process of programming a cardiac pacemaker can be described in terms of the theory of diagnosis. A set-theoretical framework of diagnosis has been taken as the basis for the construction of a system for pacemaker programming that in its present form is capable of assisting cardiologists. The system has been made available commercially in 2003 by Vitatron for its C series pacemakers. In this paper, we discuss the practical requirement imposed by the clinical environment in which pacemaker programming takes place. The theory of diagnosis that has been used is briefly reviewed, after which we describe the capabilities and limitations of the implemented system. The paper is rounded off by some ideas for future development. As far as we know, this is the first system of its kind in the area of pacemakers commercially available.

1 INTRODUCTION

The rhythm and the coordination of the human heart are controlled by a complicated system consisting of excitatory and conductive tissue, which causes the heart to beat and to expel blood so that oxygen is being delivered to the cells and waste products are removed. Unfortunately, this system as well as the heart muscle can fail giving rise to various conditions collectively referred to as *arrhythmias* [2]. An example is abnormally low heart rate, called *bradyarrhythmia*, which occurs in many different forms. Signs and symptoms associated with arrhythmia vary from fatigue, dizziness, palpitations and fainting to death. There are a number of treatments available for arrhythmia including medical treatment by means of drugs, but these are mostly used to treat fast uncoordinated rhythms. Long-term treatment of bradyarrhythmia is best done by means of a cardiac pacemaker, in this case called a *bradypacemaker*. Bradypacemakers are produced worldwide by a number of companies, amongst others by Vitatron, a company that is part of Medtronic, a world leader in medical technology.

Modern pacemakers are sophisticated electronic devices, implanted in the patient's chest region, and capable of providing assistance on demand. Under normal conditions the patient is sent home after the implant of the pacemaker, with settings that are optimal at the time. Ingrowth of the pacemaker leads may give rise to changes, and after about 3 months the settings are again optimised. This is called the *pacemaker follow-up*, which is normally done every 6 months after the initial 3 months follow-up.

Pacemaker follow-up consists of the following steps:

1. patient 'welcoming' and connecting the ECG;
2. pacemaker interrogation, carrying out measurements, analysis of diagnostics, and reprogramming;
3. informing the patient, saying goodbye to the patient and updating the patient files.

The follow-up process takes approximately 30 minutes; Vitatron's aim is to reduce this to 5 minutes. Unfortunately, programming a pacemaker is not a straightforward task: it requires the availability of sufficient time and knowledge of pacemaker functionality and possible pacemaker therapy. It has been observed that due to a lack of one or both of these factors, in many patients a given pacemaker therapy is suboptimal: many implanted pacemakers are even kept in the settings with which they were originally shipped to the hospital. In addition, pacemaker technology is moving fast, yielding pacemaker devices that are almost annually enhanced in their capabilities. For example, the new C series of Vitatron pacemakers include a digital signal-processing unit that is able to automatically interpret signals from the heart. As a consequence, pacemaker equipment companies such as Vitatron are beginning to realise that some form of intelligent decision support is needed in order to let patients benefit from further advances in pacemaker technology.

The process of optimising programming of the pacemaker consists of observing signs and symptoms in the patient, collecting information of past and present electrical behaviour of the heart and pacemaker stored in the pacemaker device, taking into account stored pacemaker settings. Based on this information, advice can be given about desirable changes to pacemaker settings. The entire process has much in common with the process of diagnostic problem solving as was shown in a previous paper [5].

This paper reports on work that has been done subsequently to produce a practically useful advice system, called the *Therapy Advisor*, that is able to offer treatment advice about appropriate pacemaker treatment for patients. Several prototype models and systems have been developed during the past six years, with varying amounts of complexity. The advice system which is now being shipped with the C series pacemaker programmer is the result of this work. Although different representations and reasoning algorithms have been investigated in the course of time, at an abstract level all the approaches can be looked upon as special instances of the framework of diagnosis proposed in Ref. [4]. The design of the current system took into account a number of significant practical requirements, such as:

- the limited amount of time available to the clinician for usage of the advice system during a clinical consultation;
- routine use of the advice system should save time;
- the system should return with an advice within 2 seconds maximum;
- easy maintenance of the underlying model.

It is the first system of its kind that is being offered commercially.

[2] Institute for Computing and Information Sciences, Radboud University Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands, e-mail: peterl@cs.kun.nl
[3] Research and Development Department, Vitatron B.V., P.O. Box 5227, 6802 EE Arnhem, The Netherlands, e-mail: {Ruud.Kuipers, Frederick.Feith}@vitatron.com

Figure 1. The excitatory and conductive system of the heart.

Figure 2. C series pacemaker.

The structure of this paper is as follows. In the following section, we summarise some basic (patho)physiology of the heart, so far as needed for the understanding of this article. Next, in Section 3, the principles of the structure and function of pacemakers are briefly discussed. In Section 4, we present a theory of diagnosis used for pacemaker programming. Section 5 contains a description of the implementation of the current system. Finally, in Section 6 it is discussed what has been achieved by the current system and possible directions for future work are mentioned.

2 STRUCTURE AND FUNCTION OF THE HEART

The heart can be viewed as a pump, responsible for maintaining blood pressure and flow within the vascular system. Pressure and flow are the results of a rhythmic contraction of the cardiac muscle under control of specialised excitatory and conductive cardiac tissue. We shall briefly review the structure and function of this tissue, and mention some of the disorders associated with it.

2.1 The excitatory and conductive system of the heart

The excitatory and conductive system of the heart is shown schematically in Figure 1. The normal human heart beat is under control of the *sinoatrial node*, or *sinus node*, a small strip of specialised self-excitatory tissue, located in the wall of the right atrium. The sinoatrial node fires rhythmically at a rate of approximately 75 beats per minute. It generates an electrical impulse, called an *action potential*, that spreads through the atrial muscular wall, causing the atrial muscle to contract. Next, the impulse travels through a muscular fibre pathway to the right ventricle, where the impulse causes the *atrioventricular node* to fire. The generated impulse travels fast through the right and left bundle branches (indicated in the figure as two lightly shaded curves, originating in the atrioventricular node), which consist of specialised conductive tissue, to the muscular tissue of the right and left ventricles through so-called Purkinje fibres. The ventricular muscular tissue responds by a contraction.

2.2 Bradyarrhythmia

Bradyarrhythmia may be caused by a variety of disorders, such as *sinus node dysfunction*. The term *sick sinus syndrome* refers to a combination of symptoms, such as dizziness, fatigue, fainting and heart failure, due to sinus node dysfunction. Failure of the atrioventricular pathway to conduct electrical impulses from the atrium to the ventricle, called *atrioventricular block*, is a cause of bradyarrhythmia

that is rather common. In that case, only the frequency of ventricular contraction is decreased.

3 CARDIAC PACEMAKERS

A pacemaker is capable of taking over control of the rhythmic contraction of the cardiac muscle, thus replacing the function of the natural pacemaker, i.e. the sinoatrial node, or of parts of the conductive pathways.

3.1 Structure of a pacemaker

A pacemaker consists of a *can*, which contains a microprocessor, RAM, a battery and an impulse generator. Impulses are transmitted to the heart by means of a *lead*, which is attached to the can's connector. A lead is either unipolar or bipolar; a *unipolar* lead contains one insulated coil, whereas a *bipolar* lead contains two coils, separated by an inner insulation. The outer insulation shields a lead from the environment. The tip of a lead, which contains an electrode, is implanted into the inner surface of the heart; the actual location depends on the type of pacemaker. The pacemaker can is usually implanted in the chest region, with the lead running through the right subclavian vein to the internal surface of the heart. An example of a recent, advanced pacemaker is Vitatron's *C series* pacemaker (See Figure 2).

A pacemaker is programmed by means of a *programmer*, a computer with a special user interface for data entry and display, and with special software to control a magnetic programming head that communicates with the pacemaker. Figure 3 shows an example of a pacemaker programmer. The head is placed above the location of the pacemaker; information from the programmer to the pacemaker, and back, is transmitted by means of telemetry.

Figure 3. Pacemaker programmer.

3.2 Functions of a pacemaker

A modern pacemaker is not only capable of stimulating, or pacing, the heart, but also of sensing the intrinsic activity of the heart. Sensed activity is used as information for the pacemaker to adopt appropriate stimulating activity as therapy. Modern pacemakers are also capable of adapting their pacing rate, dependent on the patient's metabolic demands during exercise; this capability is called *rate responsiveness*.

A pacemaker, such as the C series pacemaker, collects and stores a lot of information, called *diagnostics*, that may be used to diagnose problems:

- *patient-specific information*, such as the patient's name, age, and date of implantation;
- *counters* collect information on the frequency of occurrence of certain events;
- *histograms* offer graphical information about the distribution of certain events;
- *holters* collect information about certain events over a particular period of time.

In addition to the diagnostics, the programmer also shows the programmed *pacemaker settings*, which determine the operation of the pacemaker.

3.3 Pacemaker problems

Cardiac signs and symptoms in a patient with an implanted pacemaker can be due to medical problems, inappropriate pacemaker settings, or pacemaker faults. In this section, we focus on two problems that occur due to inappropriate pacemaker settings in relation to atrial sensing and pacing. Synchrony between atria and ventricles is very important; the following two problems may affect this so-called *AV synchrony*:

- *atrial undersensing*: impulses generated by the sinoatrial node are not sensed by the pacemaker, e.g. because of an incorrect sensitivity setting. This may result in loss of synchrony and competitive atrial pacing;
- *atrial oversensing*: the pacemaker senses a signal, which, however, has not been generated by the sinoatrial node, but is, for example, the result of ventricular contraction (far field R-wave detection).

Both problems may give rise to what is called the *pacemaker syndrome*: the patient feels a beat in the neck, due to the regurgitation of ventricular blood through the atrioventricular valves back into the atria, caused by AV asynchrony.

4 PACEMAKER PROGRAMMING AND DIAGNOSIS

Pacemaker programming can be viewed as the process of finding appropriate values for pacemaker settings that avoid the occurrence of abnormal signs and symptoms in the patient. This process can be seen as a form of diagnostic problem solving: when there are particular signs or symptoms in the patient, indicating suboptimal pacemaker settings, pacemaker faults or a medical disorder, the possible causes should be determined and dealt with. The theory of diagnosis offers several ways in which such a diagnostic process can be described. We provide a brief overview of a set-theoretical framework of diagnosis (cf. [4]) that is used to describe the approach taken.
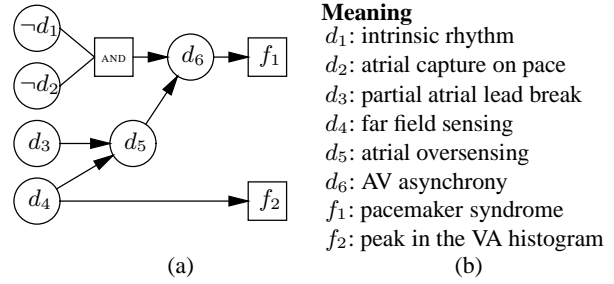

Meaning
d_1: intrinsic rhythm
d_2: atrial capture on pace
d_3: partial atrial lead break
d_4: far field sensing
d_5: atrial oversensing
d_6: AV asynchrony
f_1: pacemaker syndrome
f_2: peak in the VA histogram

(a)   (b)

Figure 4. Schematic representation of causal interactions.

4.1 The representation of interactions

Consider the following piece of pacemaker knowledge "absence of intrinsic heart rhythm combined with absence of atrial capture on pace leads to AV asynchrony, giving rise to the pacemaker syndrome in the patient; a partial atrial lead break or far field sensing may give rise to atrial oversensing, which in turn may cause AV asynchrony as well. Far field sensing can also be observed as a peak in the VA histogram." In Figure 4(a), a graph representing the causal knowledge embodied in this description is shown, where an arc denotes a cause-effect relationship. The meaning ascribed to the elements in the causal graph is indicated in Figure 4(b). Elements d_i represent defects, faults or disorders (they will be consistently called 'defects' in the following); elements f_i are observable findings. Note that, for example, the defects d_3 and d_5 are causally related to each other.

Interactions among defects, such as in the example above, can be captured more precisely by means of a mapping of sets of defects to sets of observable findings, yielding a diagnostic interpretation of this knowledge. Such a mapping will be called an evidence function. More formally, let $\Sigma = (\Delta, \Phi, e)$ be a *diagnostic specification*, where Δ denotes a set of defects, and Φ denotes a set of findings. Positive defects d (findings f) and negative defects $\neg d$ (findings $\neg f$) denote *present* defects (findings) and *absent* defects (findings), respectively. If a defect d or a finding f is not included in a set, it is assumed to be unknown. Let a set X_P denote a set of positive elements, and let X_N denote a set of negative elements, such that X_P and X_N are disjoint. It is assumed that $\Delta = \Delta_P \cup \Delta_N$ and $\Phi = \Phi_P \cup \Phi_N$. The power set of a set S is denoted by $\wp(S)$. Now, an *evidence function* e of a diagnostic specification Σ is a mapping

$$e : \wp(\Delta) \rightarrow \wp(\Phi) \cup \{\bot\}$$

such that: (1) for each $f \in \Phi$ there exists a set $D \subseteq \Delta$ with $f \in e(D)$ or $\neg f \in e(D)$ (and possibly both, which simply means that these findings may alternatively occur); (2) if $d, \neg d \in D$ then $e(D) = \bot$; (3) if $e(D) \neq \bot$ and $D' \subseteq D$ then $e(D') \neq \bot$. If $e(D) \neq \bot$, it is said that $e(D)$ is the set of *observable findings* for D (D is *consistent*); otherwise, it is said that D is *inconsistent*.

For the pacemaker knowledge depicted in Figure 4, it holds, among others, that:

$e(\{\neg d_1\}) = \varnothing$ $e(\{d_4\}) = \{f_1, f_2\}$
$e(\{\neg d_1, \neg d_2\}) = e(\{d_6\})$ $e(\{\neg d_2, d_3\}) = \{f_1\}$
$e(\{d_6\}) = e(\{d_5\})$ $e(\{\neg d_1, \neg d_2, d_3\}) = e(\{\neg d_1, \neg d_2\})$
$e(\{d_6\}) = e(\{d_3\}) = \{f_1\}$ $e(\{\neg d_1, \neg d_2, d_3, d_4\}) = e(\{d_4\})$

The property $e(D) \subseteq e(D')$ if $e(D), e(D') \neq \bot$ and $D \subseteq D'$, $\forall D, D' \subseteq \Delta$ expresses that the interaction between sets of defects
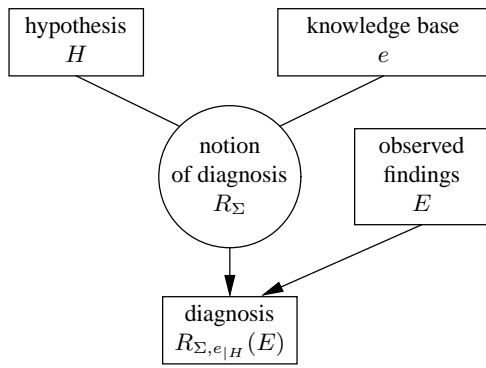
Figure 5. Schema of notion of diagnosis, diagnostic problem and solution.

is monotonic; the evidence function e is monotonically increasing. Nonmonotonic interactions may be due to masking of observable findings due to the presence of more than one defect in at the same time.

Local as well as global interactions between defects can be expressed readily in terms of evidence functions. A typical global property of evidence functions encountered in the literature is interaction freeness (cf. [6, 7]).

In the context of pacemaker programming, we know beforehand that particular sets of defects D cannot be present in the presence of other sets of defects D'. For example, 'retrograde conduction' (d) is only considered if 'atrial undersensing' (d') is not suspected. Assuming that $d \in D$, then we will not consider D' with $d' \in D'$ if it holds that $e(D \cup D') = \bot$.

4.2 Notions of diagnosis

An evidence function provides a semantic interpretation of a knowledge base in terms of expected evidence for the combined occurrence of (present or absent) defects; yet, it does not yield a diagnosis. To employ an evidence function for the purpose of diagnosis, it must be interpreted with respect to actually observed findings. Such interpretations will be called notions of diagnosis.

More formally, let $\mathcal{P} = (\Sigma, E)$ be a *diagnostic problem*, where $E \subseteq \Phi$ is a set of *observed findings*; it is assumed that if $f \in E$ then $\neg f \notin E$, i.e. contradictory observed findings are not allowed. Let R_Σ denote a *notion of diagnosis* R applied to Σ, then a mapping

$$R_{\Sigma, e_{|H}} : \wp(\Phi) \rightarrow \wp(\Delta) \cup \{u\}$$

will either provide a diagnostic solution for a diagnostic problem $\mathcal{P}$, or indicate that no solution exists, denoted by u (undefined). Here, H denotes a *hypothesis*, which is taken to be a set of defects ($H \subseteq \Delta$), and $e_{|H}$, called the *restricted evidence function* of e, is a restriction of e with respect to the power set $\wp(H)$:

$$e_{|H} : \wp(H) \rightarrow \wp(\Phi) \cup \{\bot\}$$

where for each $D \subseteq H$: $e_{|H}(D) = e(D)$. A restricted evidence function $e_{|H}$ can be thought of as the relevant part of a knowledge base with respect to a hypothesis H. An *R-diagnostic solution*, or *R-diagnosis* for short, with respect to a hypothesis $H \subseteq \Delta$, is now defined as the set $R_{\Sigma, e_{|H}}(E)$, where $R_{\Sigma, e_{|H}}(E) \subseteq H$ if a solution exists. In Figure 5, the idea underlying the definition of a notion of diagnosis R and diagnostic solution to a diagnostic problem is illustrated schematically.

A notion of diagnosis R provides the possibility to express interactions among defects and observed findings at the

level of diagnosis, which we call dependencies. We may also have that $R_{\Sigma, e_{|H \cup H'}}(E) = R_{\Sigma, e_{|H}}(E) \cup R_{\Sigma, e_{|H'}}(E)$, with $R_{\Sigma, e_{|H \cup H'}}(E) \neq u$, which means that the diagnostic solution with respect to the hypothesis $H \cup H'$ is obtained as the union of the solutions for the two separately examined hypotheses H and H'. This is called the *independence (or compositionality) assumption*. For many notions of diagnosis described in the literature, in particular for abductive diagnosis and consistency-based diagnosis, the independence assumption fails to hold.

To demonstrate how the definitions above can be employed, we consider a notion of diagnosis US (Unique Subset), such that $\text{US}_{\Sigma, e_{|H}}(E) = H'$ if it holds that H' is the only nonempty subset of H such that $e_{|H}(H') \subseteq E$; otherwise, $H' = u$. This notion of diagnosis expresses that a diagnosis consists of a set of defects which, on the one hand, can account for at least part of all observed findings, and, on the other hand, every finding associated with the set of defects that is taken as a diagnosis has been observed. Furthermore, there is only one such subset of the given hypothesis H. Some interesting diagnostic conclusions for the example in Figure 4 are: $\text{US}_{\Sigma, e_{|\{d_3, d_4\}}}(\{f_2\}) = \{d_4\}$, i.e. a peak in the VA histogram may be due to far field sensing, $\text{US}_{\Sigma, e_{|\{d_3, d_4\}}}(\{f_1, f_2\}) = u$, i.e. there does not exist a *unique* subset of H accounting for both a pacemaker syndrome and a peak in the VA histogram, and finally, $\text{US}_{\Sigma, e_{|\{d_4\}}}(\{f_2\}) = \{d_4\}$. In the first case, it is said that the hypotheses has been *adjusted*, in the second case, that the hypothesis is *rejected*, and in the last case, that the hypothesis has been *accepted*. This example demonstrates the flexibility of the approach.

4.3 Handling hypotheses by means of stored pacemaker data

The previous section defines a general framework that can be used to accommodate many different formal approaches to diagnosis. However, diagnostic problem solving normally also involves the selection of potential diagnoses using some extra constraints not covered above. As this feature is exploited in the context of pacemaker programming in a very specific fashion, we discuss this here as well.

The pacemaker stores lots of measurement data which appear to be useful for selecting diagnostic hypotheses based on their relevance with respect to what is known from such measurements. Let $M = M_P \cup M_N$ denote the set of all possible pacemaker measurement data. A *relevance function*

$$r : \wp(M) \rightarrow \wp(\Delta) \cup \{\epsilon\}$$

determines which of the hypotheses $H \subseteq \Delta$ are worth considering, i.e. only mappings $R_{\Sigma, e_{|H}}$ are considered for a given set of pacemaker measurements $S \subseteq M$ if $r(S') = H$, for $S' \subseteq S$ and $H \subseteq \Delta$, or none are considered if $r(S') = \epsilon$ for each $S' \subseteq S$.

5 RESULTING ADVICE SYSTEM

The Therapy Advisor is an integral part of the pacemaker-programmer (cf. Figure 3) software that is used at the pacemaker follow-up procedure. The first step of a follow-up that involves the programmer is the acquisition of pacemaker data. Secondly, the Therapy Advisor is activated after which the start-up screen, called the status screen, is presented to the user (cf. Figure 6). It displays the lead recordings of an ECG, the overall status of the pacemaker; at the lower part of the screen, messages generated by the Therapy Advisor are presented. The messages shown here are general in nature. A distinction is made between the following types of messages:
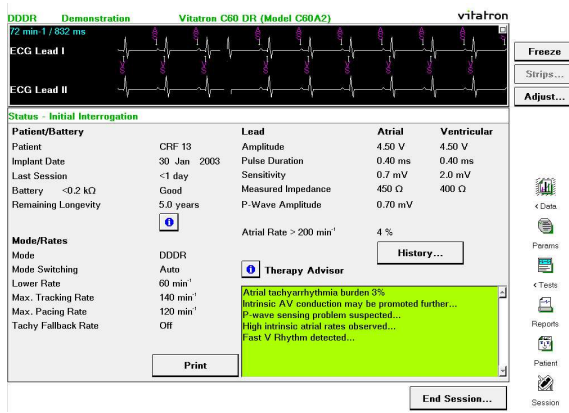
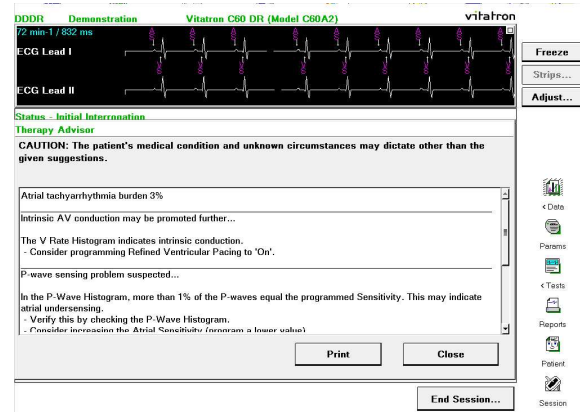**Figure 6.** Therapy Advisor start-up screen with status information.

Figure 7. Therapy Advisor screen with advice.

- pacemaker technical problems that result in the entire loss of diagnostics, e.g. pacemaker reset and battery depletion;
- pacemaker-patient interface problems. These problems arise when the technical pacemaker parameters, e.g. sensitivity or rate response, are not correctly aligned with the patient's needs and usually render parts of the diagnostics unreliable, e.g. far field R-waves distorted atrial sensing;
- clinical causes, possibly related to complaints or deterioration (e.g. atrial fibrillation and retrograde conduction);
- clinical effects, possibly related to complaints or deterioration (e.g. fast ventricular rhythm and low AV synchrony).

A separate screen is available to the user for more detail about the detected irregularities. This detailed information can consist of an explanation of the diagnostic observations, characterisation of the clinical cause (e.g. atrial fibrillation), therapy parameter advice, technical parameter advice and diagnostic parameter advice (cf. Figure 7).

The inference algorithm treats the model underlying the system as a forest of trees, where each tree can be interpreted as part of the definition of an evidence function e as discussed above. During problem solving, a search algorithm, implemented in C++, traverses the trees sequentially using the data from the pacemaker. As the main motivation for incorporating the Therapy Advisor into the pacemaker programmer software was to reduce time needed for follow-up, all data interpreted by the system come from the pacemaker. This also implies that information that is not available to the pacemaker, e.g. information obtained from the medical interview with the patient, is currently not taken into account by the Therapy Advisor in generating recommendations.

6 DISCUSSION

The development of the advice system described in this paper was preceded by the development of a number of prototype systems, which have been used to explore different research ideas. Even though the current system that is described here is now being shipped with the C series pacemaker programmer, it is likely that both model and underlying reasoning techniques will be enhanced in the future; at this stage we cannot speak of a 'definitive Therapy Advisor'. This is why it is important to base developments of such a system on a sound theory, as described in Section 4, as this allows making design decisions perspicuous.

Some of the earlier prototype systems were based on extensive causal models of abnormal behaviour, where in many cases the paths from initial effect to final cause were modelled in considerable detail (cf. [1]). Even though access to such causal knowledge may be important for explanation purposes, it is not really necessary for diagnostic problem solving (this is also clear from the theory summarised in Section 4). For example, special purpose algorithms were designed and implemented that were able to rewrite the causal network to a bipartite graph, which was then fed into another algorithm inspired by the set-covering algorithm by Peng and Reggia [6], that was modified to deal with multiple defects. However, strict upper limits on the acceptable response time, less than 2 seconds in the present version of the system, demanded disregarding multiple defects; in addition, the present model is relatively flat, and does not yet contain detailed knowledge of the causal mechanisms involved.

Quite a lot of effort has gone into the design of understandable messages, which in the model are linked to solutions produced by the system. The significance of this derives from the fact that the content of the messages determines whether the user will undertake the right actions.

It is currently still unclear whether the advice system will help improving the quality of the pacemaker therapy, whether in terms of percentage of patients with optimal therapy, time required for programming a pacemaker, or financially, e.g. because the patient's demands on clinical care are reduced. These are issues which need to be studied in the near future.

REFERENCES

[1] L. Console, D. Theseider Dupré and P. Torasso, A theory of diagnosis for incomplete causal models, in: *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, Los Angeles, USA (1989) 1311–1317.

[2] K.J. Isselbacher, et al., *Harrison's Principles of Internal Medicine*, 13th edition, McGraw-Hill, New-York, 1994.

[3] P.J.F. Lucas, Symbolic diagnosis and its formalisation, *The Knowledge Engineering Review* 12(2) (1997) 109–146.

[4] P.J.F. Lucas, Analysis of notions of diagnosis. *Artificial Intelligence* 105(1-2) (1998) 293–341.

[5] P.J.F. Lucas, A.M. Tholen and G. van Oort, An intelligent system for pacemaker reprogramming. *Artificial Intelligence in Medicine*, 17 (1999) 249–269.

[6] Y. Peng and J.A. Reggia, *Abductive inference models for diagnostic problem solving*. New York: Springer-Verlag, 1990.

[7] T.D. Wu, A problem decomposition method for efficient diagnosis and interpretation of multiple disorders. *Computer Methods and Programs in Biomedicine* 35 (1991) 239–250.