Combining Rate-Adaptive Cardiac Pacing Algorithms Via Multiagent Negotiation

Francesco Amigoni, *Member, IEEE*, Alessandro Beda, *Student Member, IEEE*, and Nicola Gatti, *Student Member, IEEE*

***Abstract*—Simulating and controlling physiological phenomena are notoriously complex tasks to tackle and require accurate models of the phenomena of interest. Currently, most physiological processes are described by a set of partial models capturing specific aspects of the phenomena, and usually their composition does not produce effective comprehensive models. A current open issue is thus the development of techniques able to effectively describe a phenomenon starting from partial models. This is particularly relevant for heart rate regulation modeling where a large number of heterogeneous partial models exists. In this paper we make the original proposal of adopting a multiagent paradigm, called anthropic agency, to provide a powerful and flexible tool for combining partial models of heart rate regulation for adaptive cardiac pacing applications. The partial models are embedded in autonomous computational entities, called agents, that cooperatively negotiate in order to smooth their conflicts on the values of the variables forming the global model the multiagent system provides. We experimentally evaluate our approach and we analyze its properties.**

***Index Terms*—Cooperative negotiation, multiagent systems, multisensor pacemakers, rate-adaptive cardiac pacing.**

I. Introduction

PHYSIOLOGICAL processes are a class of extremely complex systems to study and to model [1], [2]. A physiological process usually emerges from the interaction of several elements belonging to an intricate network of relationships, where each element is involved in more processes [3]. Currently, a large number of physiological phenomena are described by a set of related partial models that individually describe the single elements. The relationship among these partial models has not yet been fully explored. The lack of comprehensive models of physiological processes negatively affects the design of simulation and control systems for these processes [2]. For instance, in the adaptive cardiac pacing field, several partial models have been developed in the attempt to mimic the natural regulatory system of heart rate, but their reliability and accuracy are so poor that their use in commercial pacemakers is still limited [4]–[6]. Thus, the study of how to combine partial models represents a fundamental issue to improve the accuracy of modeling physiological processes within biomedical engineering. What is sought is a paradigm able to combine several partial models to effectively represent the observed phenomenon. Such a paradigm should provide a flexible tool to support different combination techniques for comparatively evaluating their effectiveness. In the literature on rate-adaptive pacing, several techniques have been adopted to combine partial models, including weighted average, cross-checking, and overdrive [5]. All these techniques exhibit some drawbacks: they are rather *ad hoc*, not easily extensible, and their overall performance is modest.

We already developed a paradigm, based on multiagent systems and called *anthropic agency* [7], able to integrate a number of partial models of a physiological phenomenon in order to globally produce a comprehensive model of the phenomenon. We applied the anthropic agency paradigm to the modeling of the glucose-insulin metabolism, a field in which a large number of models have been proposed. In our approach, partial models, describing the insulin response to food ingestion and to physical activity, are embedded in *agents* that are autonomous computational entities [8] behaving as decision makers. The global model emerges from the interaction of these agents in a distributed decision process. More specifically, the agents perform a *cooperative negotiation* in order to find an agreement on the values of the variables that are "shared" among the models they embed. The anthropic agency approach provides a flexible infrastructure to analyze different model combinations, allowing the dynamic insertion and removal of agents (and of associated models).

In this paper, we further build on the anthropic agency paradigm and we present an improved cooperative negotiation protocol—the core of partial model combination—and the application of the paradigm to adaptive cardiac pacing. The innovations of the negotiation protocol with respect to that in [7] are a formal definition that allows us to analytically prove its stability and the introduction of a number of parameters to tailor the protocol to the characteristics of a specific person. The application of the anthropic agency paradigm to adaptive cardiac pacing is important because of the partiality of the available models of heart rate regulation. A desirable solution would be to combine the partial models in a way that each one prevails just when it provides a good emulation of the normal sinus activity. The cooperative negotiation mechanism presented in this paper is an effective and flexible approximation of such solution.

The original contributions of this paper are 1) a cooperative negotiation protocol that integrates partial models of physiological processes, 2) the application of this protocol to heart rate regulation modeling, and 3) the experimental comparison of a global model of heart rate regulation obtained by our cooperative negotiation technique with those obtained by traditional

Manuscript received August 2, 2004; revised December 4, 2004 and March 2, 2005.

F. Amigoni and N. Gatti are with the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy (e-mail: amigoni@elet.polimi.it; ngatti@elet.polimi.it).

A. Beda is with the Institute of Sound and Vibration Research, Southampton University, Highfield, Southampton SO17 3RZ, U.K. (e-mail: ab1@isvr.soton.ac.uk).

Digital Object Identifier 10.1109/TITB.2005.855564

techniques presented in the literature. The immediate target of our system is the offline analysis of different combinations of partial models of heart rate regulation, while the long-term target is its possible use in implantable pacemakers. This paper does not aim to find the best set of partial models to be combined to solve the problem of cardiac pacing.

This paper is structured as follows. Section II describes the state of the art in the fields of modeling via multiagent systems and of heart rate regulation systems for cardiac pacemakers. Section III introduces the proposed negotiation protocol for combining partial models of heart rate regulation. Section IV presents the experimental evaluation of our approach. Section V discusses the results obtained. Section VI concludes the paper.

II. Related Works

A. Model Combination via Cooperative Negotiation

Two main aspects are discussed in this section: the motivations behind the choice of cooperative negotiation to address partial model combination, and the adoption of the multiagent paradigm to support cooperative negotiation.

Physiological processes are complex systems [9] which are traditionally represented by using different models at different scales or operating contexts [10]. The decentralized optimization approach [11] has recently emerged as the most effective technique to combine different models. In this case, modeling can be formulated as a decentralized multiobjective problem where n independent decision makers, each one embedding a local model, cooperate to achieve a common goal. However, the solution of this problem is often computationally expensive. Approximate techniques that allow a considerable reduction of the computational complexity are based on economics [12]. Each decision maker i operates a local optimization of its utility function $\mathcal{U}_i$ (its local model), and a global optimum is achieved via a cooperative negotiation. The optimization activity is thus performed at two levels: at the individual decision maker level and at the overall system level. At this latter level, cooperative negotiation [13] tries to solve the conflicts, derived from the different goals the decision makers try to achieve, in order to optimize the social utility $\mathcal{U}$ of the system. Since we apply the optimization process both to the individual models and to their combination, our approach can be considered as an extension of that described in [2].

The determination of the optimal global states that should be achieved and that maximize $\mathcal{U}$ is not trivial in the case of decision makers embedding partial models of a physiological process. Such knowledge can be derived by monitoring the process. For example, in [7], where we addressed the problem of modeling (and regulating) the glucose-insulin metabolism, we adopted a social utility function $\mathcal{U}$ defined as the absolute value of the difference between the insulin values estimated by the cooperative negotiation and the real ones. In that case, $\mathcal{U}$ has to be minimized. In this paper, we define a similar social utility function $\mathcal{U}$ relative to the heart rate regulation problem.

The use of the multiagent paradigm to implement a model combination via cooperative negotiation is somehow a "natural" choice [8]. By using such paradigm, each partial model is embedded in a separate entity (i.e., an agent), sensibly reducing the problems related to the heterogeneity of the partial models to combine (e.g., based on differential equations or on logic, continuous or discrete). The cooperative negotiation is then implemented as a negotiation between the agents that embed the partial models. Furthermore, a multiagent system can be dynamically reconfigured and structured by inserting and removing agents. Such property implies that we can use the same platform to evaluate and compare virtually every possible combination of partial models.

B. Heart Rate Regulation Systems for Cardiac Pacemakers

Since the 1960s, permanent cardiac pacing has been used as an effective and reliable therapy for specific pathological alterations of the cardiac rhythm (arrhythmia) [4]. Rate-adaptive pacing, in which pacing frequency is determined as a function of some physiological parameters measured by sensors, aims at developing a multisensor heart rate regulatory system that effectively emulates the natural regulation of heart rate in healthy subjects. Different combinations of physiological parameters, sensors, and control algorithms have been proposed over the years [4]–[6]. Some studies [4] have shown that single-sensor pacemakers can emulate effectively the normal sinus activity only under specific conditions that depend on the sensor and the algorithm used. For example, most of the commercial pacemakers set the pacing frequency according to body movements detected by an accelerometer. In many cases, activities with completely different metabolic requests correspond to similar accelerometer readings (e.g., walking up or down a slope) [4]. Thus, it seems rather unlikely that a single-sensor pacemaker will ever be able to regulate heart rate effectively under every possible physiological condition and requirement imposed by activities of daily life. In order to overcome these limitations, different strategies to combine control algorithms have been pursued, including overdrive (the pacing frequency is set by the model that proposes the highest value for heart rate), cross-checking (one model is used as a reference to avoid abrupt outcomes of a "controlled" model), and weighted average (the pacing frequency is a weighted average of the values proposed by the models) [5]. To the best of our knowledge there are no established comparative results on the effectiveness of these combination techniques.

To summarize, in the heart rate pacing community, there is an evident need for an instrument that allows comparision of different combinations of heart rate regulation algorithms. This instrument should be functional, versatile, and effective, and should use a general technique to combine control algorithms, including as special cases overdrive, cross-checking, and weighted average. The anthropic agency approach presented in this paper is an answer to this need.

III. A Cooperative Negotiation Protocol for the Combination of Heart Rate Regulation Models

This section describes the innovative approach proposed in this paper for combining heart rate regulation models (Fig. 1). The partial models of heart rate regulation are embedded in agents. They receive at time 0 the values of some parameters

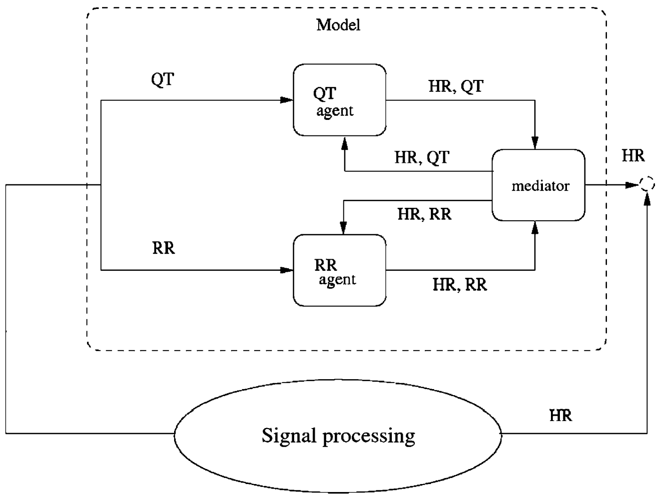

Fig. 1. Overview of our modeling system.

measured by sensors and independently operate an individual optimization (*agent optimization*). Then, each agent proposes its optimal state to a *mediator* that calculates the agreement, that is an optimal state, for the whole system, or agency. A cooperative negotiation, namely an interleaved succession of proposals and counter-proposals, goes on until the proposals of the agents converge to a stable agreement or a timeout expires (*agency optimization*). The process restarts when the agents receive new values of parameters at time T_s, where T_s is the sampling period of the system (2s in our implementation).

Considering the issue of how to generate the inputs for our system, given the early stage of the project and the limited knowledge we have on the behavior of our system, we rejected the option of animal and human testing as premature. Moreover, we are not aware of any closed loop simulator for cardiac pacing applications (a simulator able to mimic the reaction of the human body to imposed heart rates during different types and levels of physical activity). Hence, recalling that our aim is to preliminary validate our approach, we used an open loop simulator (described in Section IV). Such a decision has the strong limitation of assuming that the heart rate does not affect the behavior of the cardiovascular system (and, more generally, of the organism), but nevertheless allowed us to carry out initial tests on our system, and to illustrate its promising potential and properties.

A. Agent Optimization

In our experimental activities, we have considered two partial models relating the heart rate HR (measured in beats per minute, bpm) to the length of QT interval (expressed in ms) and to the respiration rate RR (expressed in cycles per minute), respectively. The models we adopted are feed-forward; the use of more complex models (e.g., feed-back regulated [14]) appeared unjustified given both the evident limitations of the open loop approach we used for the simulator and our aim to illustrate the potential of our approach rather than finding the best set of partial models for heart rate regulation. The two partial models we considered significantly illustrate the potential of our approach since they share the same output space (HR). Although the possibilities of the proposed approach could be even more evidenced by combining partial models that also share (part of) the same input space, the simple partial models we have chosen have input values (QT and RR) that can be easily extracted from the electrocardiogram (ECG) signal by simulated sensors [15]. For experimental activities, this means that we can use long ECGs of healthy subjects, acquired using holter devices, which are likely to cover the full range of activities of daily life. We do not need the availability of complex correlated data (such as a synchronous acquisition of ECG, accelerometer, and thoracic impedance data).

More specifically, the first partial model relates HR to the length of QT interval according to [16]:

$$\mathrm{HR} = -q_1 / \log\left(\frac{q_2 - \mathrm{QT}}{q_3}\right). \tag{1}$$

The second partial control model relates HR to RR, according to a slight refinement of the model [17]:

$$\mathrm{HR} = r_1 \cdot \arctan\,(r_2 \cdot \mathrm{RR} - r_3) + r_4 \tag{2}$$

where q_i and r_i are parameters that depend on the specific patient. Following our anthropic agency approach, we embedded in a first agent, called *QT agent*, the model (1), and in a second agent, called *RR agent*, the model (2).

Since we want the agents to initially perform a local optimization, utility functions $\mathcal{U}_{\mathrm{QT}}$ and $\mathcal{U}_{\mathrm{RR}}$ must be defined. In particular, $\mathcal{U}_{\mathrm{QT}} : A_{\mathrm{QT}} \rightarrow \Re$ (where $A_{\mathrm{QT}} \equiv \mathrm{HR} \times \mathrm{QT}$) and $\mathcal{U}_{\mathrm{RR}} : A_{\mathrm{RR}} \rightarrow \Re$ (where $A_{\mathrm{RR}} \equiv \mathrm{HR} \times \mathrm{RR}$). Hence, the state of the agent QT can be described by a pair of values $\langle hr, qt \rangle$ and the state of the agent RR by a pair $\langle hr, rr \rangle$. We call $A \equiv A_{\mathrm{QT}} \bigcup A_{\mathrm{RR}}$; thus, a vector $\boldsymbol{p} = \langle hr, qt, rr \rangle \in A$ describes the state of the entire system. In our applicative scenario, the values qt and rr are extracted from (simulated) sensors and hr is estimated by negotiation (Fig. 1). Given a point $\boldsymbol{q} \in A_{\mathrm{QT}}$ and calling $d(\boldsymbol{q})$ the distance of $\boldsymbol{q}$ from the curve (1), $\mathcal{U}_{\mathrm{QT}}(\boldsymbol{q})$ is calculated as:

$$\mathcal{U}_{\mathrm{QT}}(\boldsymbol{q}) = \begin{cases} \frac{d(\boldsymbol{q})^2}{10} & d(\boldsymbol{q})^2 < 20 \\ 20 + 10 \cdot \log \frac{d(\boldsymbol{q})^2}{10} - 20 + 1 & d(\boldsymbol{q})^2 > 20 \end{cases}. \tag{3}$$

In a similar way, $\mathcal{U}_{\mathrm{RR}}(\boldsymbol{r})$ (with $\boldsymbol{r} \in A_{\mathrm{RR}}$) is calculated. Fig. 2 shows $\mathcal{U}_{\mathrm{QT}}$ and $\mathcal{U}_{\mathrm{RR}}$. The formulas (3) have been experimentally determined in order to assign high values to the points that are far from the optimal curves (1) and (2).

The utilities $\mathcal{U}_i$ are defined as potential functions where the points with minimum potential (equal to zero) satisfy (1) and (2); namely these points represent the combination of values for which the heart has an optimal functioning (at least according to the partial models we use). Thus, in our approach, the most desirable situation corresponds to the minimum values of $\mathcal{U}_i$. Indeed, from the medical literature, the optimal curves (such as (1) and (2)) are usually the only information available on a physiological phenomenon. We can reasonably assume that, when a state is far from an optimal curve, the performance of the

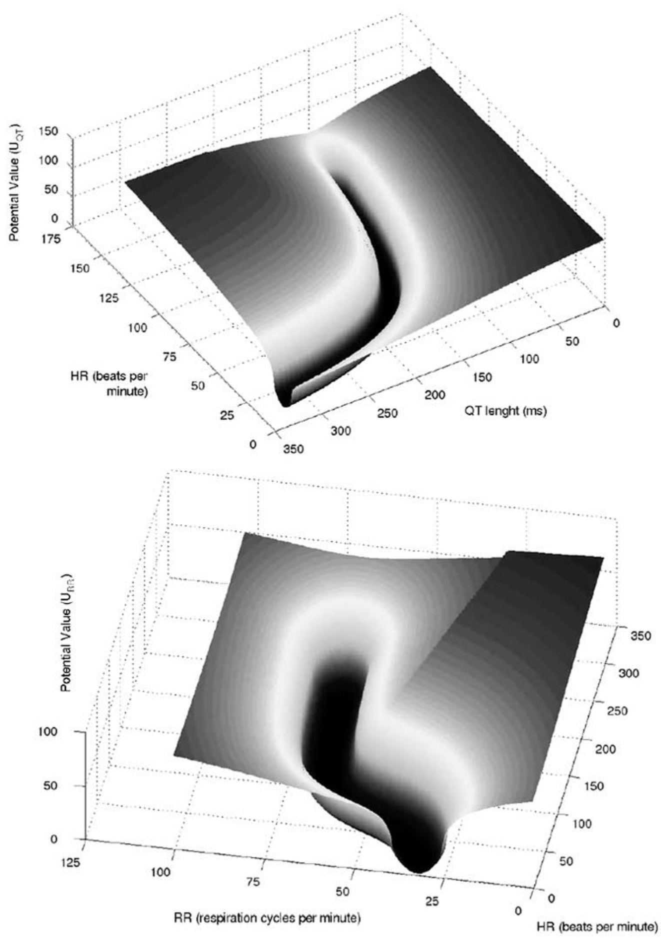

Fig. 2. Utility (potential) functions for QT [16] (top) and RR [17] (bottom) models.

heart in that state is far from optimal; this aspect is captured in the definition of the previously described steep potential functions. The goal of an agent is to change the heart rate in order to reach a minimum potential state near to the current state, according to its model. Note that, starting from optimal curves describing a physiological process, it is usually possible to build utility functions that do not have local minima.

We remark that, although in our implementation we used simple feed-forward models, our approach and, in particular, the process of defining utility functions, can be extended to a broader class of partial models. For example, for feed-back models represented as $\mathrm{HR} = \mathcal{F}(\boldsymbol{x})$, where the components of the vector $\boldsymbol{x}$ can include past values of HR, the definition of an utility function using formulas similar to (3) is straightforward. The possibility of using our approach both with open and closed loop control models is linked to the fact that cooperative negotiation is not a control algorithm, but a sort of control *meta*algorithm that harmonizes a number of control algorithms.

Given a vector $\boldsymbol{p} = \langle hr, qt, rr \rangle$ representing the current state of the global phenomenon at time 0 (for example, as read by sensors) and its projections $\boldsymbol{p}_i$ on the spaces of the agents, the agents independently perform a local optimal search with objective function $\mathcal{U}_i(\,\cdot\,)$ (for instance they could use a hill-climbing technique) under specific constraints f_i:

$$\boldsymbol{p}_i^0 = \mathrm{argmin}_{\boldsymbol{p}_i} \mathcal{U}_i(\boldsymbol{p}_i),$$
$$\text{subject to } f_i(\boldsymbol{p}_i^0, \boldsymbol{p}) = 0 \quad \forall i \in \{\mathrm{QT}, \mathrm{RR}\}$$

$\boldsymbol{p}_i^0 \in A_i$ is thus the "desired" state of agent i. In our application, we assume the following constraints: $f_{\mathrm{QT}}(\boldsymbol{p}_{\mathrm{QT}}^0, \boldsymbol{p}) = 0$ iff $\Delta\mathrm{QT} = 0$ and $f_{\mathrm{RR}}(\boldsymbol{p}_{\mathrm{RR}}^0, \boldsymbol{p}) = 0$ iff $\Delta RR = 0$; in other words, the values of QT and RR cannot change during the local search. This is because the QT and RR values are given and the agents have to determine the HR values. We can safely assume that hill-climbing techniques efficiently find the global minimum of $\mathcal{U}_i(\,\cdot\,)$ since local minima do not exist by construction. In conclusion, from the agent optimization algorithms, two vectors, $\boldsymbol{p}_{\mathrm{QT}}^0 = \langle hr_1, qt \rangle$ and $\boldsymbol{p}_{\mathrm{RR}}^0 = \langle hr_2, rr \rangle$, are produced.

B. Agency Optimization: The Negotiation Protocol

Once the agents have performed individual local optimizations, their partial views are harmonized via cooperative negotiation. (The following description is partially taken from [18].)

A *negotiation session* is a sequence of interleaved proposals of the agents to the mediator $(\boldsymbol{p}_{i\to e}^t)$ and counter-proposals of the mediator to the agents $(\boldsymbol{p}_{e\to i}^t)$, starting at time 0 and possibly ending at time $\tau \leq T_s$. For example, the portion of a negotiation session involving agent i can be represented as: $\boldsymbol{p}_{i\to e}^0 \succ \boldsymbol{p}_{e\to i}^0 \succ \boldsymbol{p}_{i\to e}^1 \succ \cdots \succ \boldsymbol{p}_{e\to i}^\tau$.

Each agent proposes its offer by sending the following pair to the mediator: $\langle \boldsymbol{p}_{i\to e}^t, \mathcal{U}_i(\boldsymbol{p}_{i\to e}^t) \rangle$. The initial proposal of an agent i is $\langle \boldsymbol{p}_i^0, \mathcal{U}_i(\boldsymbol{p}_i^0) \rangle$, with $\boldsymbol{p}_i^0$ calculated as discussed in the previous subsection. The mediator receives the offers of all the agents at time t and calculates their weighted average:

$$\boldsymbol{m}^t = \frac{\sum_{i=1}^n \boldsymbol{p}_{\boldsymbol{i}\to e}^{\boldsymbol{t}} \cdot (1 + \mathcal{U}_i(\boldsymbol{p}_{i\to e}^t))}{\sum_{i=1}^n (1 + \mathcal{U}_i(\boldsymbol{p}_{\boldsymbol{i}\to e}^{\boldsymbol{t}}))}$$

$\boldsymbol{m}^t$ is the *agreement* reached in the negotiation session at time t. Note that $\boldsymbol{m}^t$ is a point in A; in the above formula, the sum at the numerator is intended to sum only the corresponding elements of the vectors (for example, $\langle hr_1, qt \rangle + \langle hr_2, rr \rangle = \langle hr_1 + hr_2, qt, rr \rangle$). The mediator communicates to the agents its counter-proposals $\boldsymbol{p}_{e\to i}^t$, that are the projections of $\boldsymbol{m}^t$ on the A_i, by sending to each one of them the value $\langle \boldsymbol{p}_{e\to i}^t \rangle$.

After receiving the counter-proposal, every agent i calculates its new proposal to the mediator as

$$\boldsymbol{p}_{i\to e}^{t+1} = \boldsymbol{p}_{i\to e}^t + \left\| \boldsymbol{p}_{e\to i}^t - \boldsymbol{p}_{i\to e}^t \right\| \left(\alpha_i(\boldsymbol{p}_i) \boldsymbol{u}_i^{t+1} + \beta_i(\boldsymbol{p}_i) \boldsymbol{w}_i^{t+1} \right) \tag{4}$$

where $\alpha_i(\,\cdot\,)$ and $\beta_i(\,\cdot\,)$, called *negotiation parameters*, are two functions on A_i to $\Re$, $\|\cdots\|$ is the vector norm, and $\boldsymbol{u}_i^{t+1}$ and $\boldsymbol{w}_i^{t+1}$ are two versors defined below. (In our implementation of the negotiation protocol, the proposal $\boldsymbol{p}_{i\to e}^{t+1}$ is forced equal to $\boldsymbol{p}_{i\to e}^t$ when their difference is below a given threshold δ_i.)

Considering the vector connecting $\boldsymbol{p}_{e\to i}^t$ to $\boldsymbol{p}_{i\to e}^t$ and normalizing it results in

$$\boldsymbol{u}_i^{t+1} = \frac{\boldsymbol{p}_{e\to i}^t - \boldsymbol{p}_{i\to e}^t}{\|\boldsymbol{p}_{e\to i}^t - \boldsymbol{p}_{i\to e}^t\|}$$

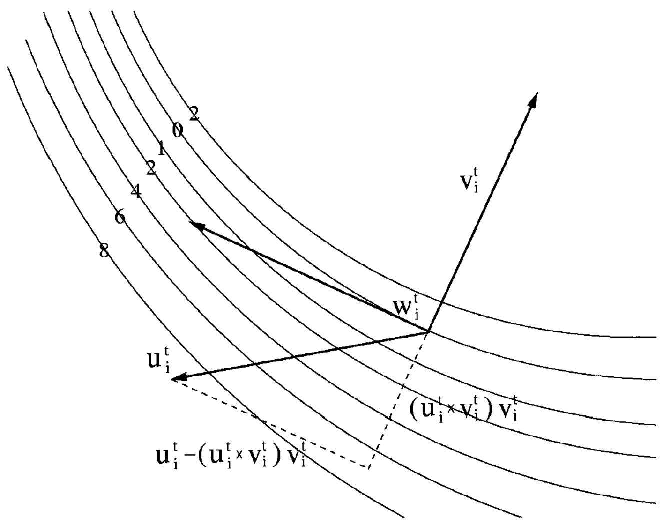

Fig. 3. $\boldsymbol{u}_i^t$, $\boldsymbol{v}_i^t$, $\boldsymbol{w}_i^t$ in a 2-D space.

This versor is headed towards the agreement (in A_i) at time t (Fig. 3). From (4), it can be seen that the next proposal of the agent i gets closer to the last agreement of the negotiation by a quantity $\alpha_i(\boldsymbol{p}_i)\|\boldsymbol{p}^t_{e\to i} - \boldsymbol{p}^t_{i\to e}\|$, proportional to the distance between the last proposal of agent i and the last agreement.

Considering the vector along the direction of the gradient of $\mathcal{U}_i(\,\cdot\,)$ in $\boldsymbol{p}^t_{i\to e}$ and normalizing it, we obtain the following versor:

$$\boldsymbol{v}_i^{t+1} = \frac{\nabla\mathcal{U}_i\left(\boldsymbol{p}^t_{i\to e}\right)}{\left\|\nabla\mathcal{U}_i\left(\boldsymbol{p}^t_{i\to e}\right)\right\|}.$$

It points toward the direction of maximum increasing of $\mathcal{U}_i(\,\cdot\,)$ and every vector orthogonal to $\boldsymbol{v}_i^{t+1}$ is tangent to an equipotential curve (gray lines in Fig. 3). We define the following versor that is orthogonal to the gradient direction, and thus tangent to the level curves of the potential space (Fig. 3):

$$\boldsymbol{w}_i^{t+1} = \frac{\boldsymbol{u}_i^{t+1} - \left(\boldsymbol{v}_i^{t+1}\cdot\boldsymbol{u}_i^{t+1}\right)\boldsymbol{v}_i^{t+1}}{\left\|\boldsymbol{u}_i^{t+1} - \left(\boldsymbol{v}_i^{t+1}\cdot\boldsymbol{u}_i^{t+1}\right)\boldsymbol{v}_i^{t+1}\right\|}.$$

From (4), it can be seen that the next proposal of an agent tries to keep constant the potential value in the direction of the last agreement, by moving in the direction of $\boldsymbol{w}_i^{t+1}$ by a quantity $\beta_i(\boldsymbol{p}_i)\|\boldsymbol{p}^t_{e\to i} - \boldsymbol{p}^t_{i\to e}\|$.

To summarize, the next proposal of an agent takes into account two components: the first one contributes to accomodating the tendency of the society, while the second one contributes to keeping the individual agent close to its optimal curve. The values $\alpha_i(\boldsymbol{p}_i)$ and $\beta_i(\boldsymbol{p}_i)$ determine the relative weights of these two components. For example, if $\alpha_i(\,\cdot\,)$ is small (namely, the weight of the component that pushes towards the negotiation agreement is small), the model embedded in agent i has high confidence and its proposal tends to move only slightly from its optimal curve. Since in the experiments presented below the values of $\alpha_i(\,\cdot\,)$ and $\beta_i(\,\cdot\,)$ have been kept constant, we call them α_i and β_i. Finally, note that the parameters α_i and β_i can be tailored to the specific patient whose heart rate is being estimated; this possibility is exploited in Section IV. (Implementation details of the proposed cooperative negotiation protocol are discussed in [19].)

C. *Some Properties of the Proposed Negotiation Protocol*

The cooperative negotiation protocol previously presented generalizes, with an appropriate choice of the negotiation parameters α_i, β_i, and of the utility functions $\mathcal{U}_i$, the techniques commonly adopted to combine partial models (and related control algorithms) of heart rate regulation. For example, when $\alpha_i = 1$, $\beta_i = 0$, and $\mathcal{U}_i = \omega_i$ (where ω_i is the weight of the model embedded by agent i), a weighted average combination is applied. As another example, when $\beta_i = 0$ and either (*a*) $\alpha_i = 1$ and $\mathcal{U}_i = 0$ if the HR value proposed by agent i is less than that proposed by other agent j or (*b*) $\mathcal{U}_j = 1$ otherwise, an overdrive mechanism is present. In a similar way, it could be easy to show that our protocol generalizes cross-checking: the weight of the "controlled" model is 0 when its output differs more than a threshold from the output of the "controlling" model. This makes our anthropic agency approach suitable for testing different model combination techniques.

The proposed negotiation paradigm has been proved to converge to an agreement [18], namely to reach a "stable" value of HR for each negotiation session, if the negotiation parameters satisfy

$$\begin{aligned}\forall i\colon\ & (0 < \beta_i < 1 \wedge \beta_i < \alpha_i < 2 - \beta_i)\\ & \vee (-1 < \beta_i < 0 \wedge -\beta_i < \alpha_i < 2 + \beta_i)\,. \qquad (5)\end{aligned}$$

IV. Results

A. *Experimental Setting*

We used five 24-h-long ECGs taken from [20] and relative to patients 16265, 16273, 16420, 17052, 17453. These patients have been randomly selected out of those in the database who do not present persistent artifacts on the complete 24-h ECG signals or excessively large signal-to-noise ratios. From each ECG signal we extracted the temporal sequences of the beat-to-beat HR values and the QT interval length values using the algorithms in [21], and the temporal sequence of the ECG-derived respiratory signal (EDR) values using the algorithms in [15]. The RR values have been calculated by using a zero-crossing algorithm with a 20-s window on the EDR sequence. The QT values have been calculated by averaging the values of the QT interval length in a 10-s window. In summary, for each patient, we have three temporal sequences of values of HR, QT, and RR.

In order to best tune the parameters of the models (1) and (2) on the specific patients, we extracted samples (uniformly distributed and covering the 10% of the length of the ECG track) of the three temporal sequences relative to HR, QT, and RR. We sampled over the full sequences because no information is available about the activities performed by the patients during ECG recording. We applied a nonlinear least-squares technique [22] to determine the values of parameters q_i and r_i that minimize

TABLE I
PARAMETERS OF THE MODELS

Patient	QT model			RR model			
	q_1	q_2	q_3	r_1	r_2	r_3	r_4
16265	8.20	1930.4	1734.7	26.98	0.0995	2.5947	95.02
16273	9.00	1930.7	1745.0	26.80	0.0970	2.5945	95.12
16420	9.50	1930.5	1744.1	27.30	0.0990	2.5950	94.56
17052	9.20	1930.3	1732.0	30.40	0.0846	2.6250	92.75
17453	10.00	1930.7	1736.4	28.40	0.0960	2.5950	90.32

TABLE II
WEIGHTS AND NEGOTIATION PARAMETERS THAT MINIMIZE $r_{\rm tr}$

Patient	$\bar{\omega}_{\rm QT}$	$\bar{\omega}_{\rm RR}$	$\bar{\alpha}_{\rm QT}$	$\bar{\alpha}_{\rm RR}$	$\bar{\beta}_{\rm QT}$	$\bar{\beta}_{\rm RR}$
16265	5.5	1	0.020	0.350	0.020	0.030
16273	3.2	1	0.050	0.070	0.010	0.055
16420	5.2	1	0.020	0.170	0.000	0.020
17052	2.4	1	0.050	0.080	0.030	0.065
17453	7.8	1	0.020	0.200	0.005	0.100

TABLE III
EXPERIMENTAL RESULTS EXPRESSED AS $r_{\rm tr}$ AND $r_{\rm te}$ (BOTH IN BPM)

Patient	QT model	RR model	Weighted Average		Cooperative Negotiation	
	r_{te}	r_{te}	r_{tr}	r_{te}	r_{tr}	r_{te}
16265	7.1	11.2	4.2	6.7	3.6	4.1
16273	5.2	7.1	3.0	4.5	2.9	3.2
16420	11.2	14.1	8.1	10.2	5.9	6.3
17052	11.9	17.8	9.6	10.1	6.8	7.1
17453	11.8	16.4	7.0	11.3	6.2	6.5
Averages	9.4	13.3	6.4	8.6	5.1	5.4

TABLE IV
DIFFERENT DIVISIONS FOR THE TRAINING AND TESTING SETS
($r_{\rm tr}$ AND $r_{\rm te}$ ARE IN BPM)

Training set	Testing set	Weighted Average		Cooperative Negotiation	
		r_{tr}	r_{te}	r_{tr}	r_{te}
first 40 minutes	last 20 minutes	4.2	6.7	3.6	4.1
first and last 20 minutes	middle 20 minutes	4.5	6.6	3.8	4.0
last 40 minutes	first 20 minutes	4.4	6.6	3.6	4.1
first 30 minutes	last 30 minutes	4.0	7.7	3.2	4.8
last 30 minutes	first 30 minutes	3.8	7.5	3.0	4.6
first half of the track	second half of the track	3.1	9.7	2.8	5.1
second half of the track	first half of the track	4.3	9.8	3.4	5.2

the error between the HR values produced by the model (1) and (2), respectively, and the real values read from the ECG. Table I shows the values obtained from this calibration process performed on each patient.

We embedded the QT and RR models relative to a patient in two agents, as explained in the previous section, obtaining five instances of the anthropic agency, each one representing the heart rate regulation models of a patient. The input of the system is sampled at a fixed frequency (every 2 s) from the sequences QT and RR.

B. Evaluation of the Negotiation Protocol

We compared our negotiation-based approach with other combination techniques that are currently adopted in rate-adaptive pacing applications. We considered the five patients in the following four scenarios: the QT and RR partial models individually taken (as examples of single-sensor pacing systems), their weighted average, and their combination by our negotiation protocol. Each scenario corresponds to a different global model for HR. We evaluated the effectiveness of a modeling technique by the root mean squared error $r = \sqrt{E[e^2]}$ (where $E[\,\cdot\,]$ is the average over time). The error e is defined as the difference, at a given time, between the real value of HR (as read from the ECG) and the value of HR estimated by the model (Fig. 1).

For each patient, we used a portion of the sequences relative to the patient as *training set* to tune weights (in weighted average) and negotiation parameters (in cooperative negotiation), and the remaining portion as *testing set*. The training set is composed of the first $2/3$ of the data in every 60-m window, while the remaining $1/3$ of the windows compose the testing set. This method of selecting the training and testing sets is justified by the fact that we do not know how the statistical properties of the signal change during the 24-hours. Hence, it is reasonable to select samples homogeneously distributed on the track. We calculated r both for the training set ($r_{\rm tr}$) and for the testing set ($r_{\rm te}$). The weights of the partial models $\omega_{\rm QT}$ and $\omega_{\rm RR}$ have been selected in order to minimize $r_{\rm tr}$ (namely, in order to produce the *best* weighted average on the training data). To determine these best weights ($\bar{\omega}_{\rm QT}$ and $\bar{\omega}_{\rm RR}$) we set $\bar{\omega}_{\rm RR} = 1$ and we made an exhaustive search with $\omega_{\rm QT}$ changing from 1 to 10 with step 0.1. Results are reported in Table II. The negotiation parameters α_i and β_i have been chosen in order to minimize $r_{\rm tr}$ (namely, in order to produce the *best* negotiation-based model combination on the training data). The negotiation parameters $\bar{\alpha}_i$ and $\bar{\beta}_i$ that minimize $r_{\rm tr}$ for the five patients are reported in Table II and have been determined by using an exhaustive search within the space bounded by constraints (5) and variable quantization intervals between 10^{-1} and 10^{-3} (heuristically reducing the quantization interval around the values that provide the best performance). The performances of the different combination techniques are shown in Table III.

In order to investigate some interesting issues that emerged from the extended experimental activity, in the following we report results about patient 16 265; the considerations about other patients are similar. We evaluated the effects of adopting different divisions for the training and testing sets. The results (for patient 16 265) are showed in Table IV (the first five rows are relative to 60-m windows). The negotiation performs better than the weighted average for every choice of training and testing sets we tried.

To analyze the role of β_i in our cooperative negotiation protocol we compared the performance of a negotiation that uses the $\boldsymbol{w}_i^t$ component with respect to a negotiation that does not use it (i.e., with $\beta_i = 0$). For patient 16 265, we determined, keeping $\beta_{\rm QT} = \beta_{\rm RR} = 0$, the values of $\alpha_{\rm QT}$ and $\alpha_{\rm RR}$ that minimize $r_{\rm tr}$. Then, we used these values (still keeping $\beta_{\rm QT} = \beta_{\rm RR} = 0$) to calculate $r_{\rm te}$. We obtained $r_{\rm tr} = 3.7$ bpm and $r_{\rm te} = 5.6$ bpm. Comparing these results with Table III we note that the presence of the $\boldsymbol{w}_i^t$ component reduces $r_{\rm te}$ of 1.5 bpm and $r_{\rm tr}$ of a negligible amount.

Although general transformation formulas that relate negotiation parameters to weights are not trivial to find, informally we note that models with ω_i are embedded in agents with small α_i (i.e., with high confidence). Moreover, we found that $r_{\rm tr}$ and $r_{\rm te}$ are more sensitive to the variations of the weights than of the negotiation parameters. We compared the two approaches (for patient 16 265) varying the parameters of the QT model (i.e., $\alpha_{\rm QT}$ and $\omega_{\rm QT}$) around their optimal values $\bar{\alpha}_{\rm QT}$ and $\bar{\omega}_{\rm QT}$

reported in Table II and keeping constant the (optimal) values of $\bar{\alpha}_{\mathrm{RR}}$, $\bar{\omega}_{\mathrm{RR}}$, and $\bar{\beta}_i$. More precisely, given the optimal ratios $\bar{R}_\alpha = \bar{\alpha}_{\mathrm{QT}}/\bar{\alpha}_{\mathrm{RR}}$ and $\bar{R}_\omega = \bar{\omega}_{\mathrm{QT}}/\bar{\omega}_{\mathrm{RR}}$, we varied α_{QT} and ω_{QT} in order to vary the ratios around $\bar{R}_\alpha$ and $\bar{R}_\omega$ in the ranges $[0, 2\bar{R}_\alpha]$ and $[0, 2\bar{R}_\omega]$ in 100 uniform steps, respectively. We obtained that r_{te} increases quasilinearly when we get far from the optimal ratios. The value of r_{te} increases up to 14.6% (up to 4.7 bpm) and 4.8% on average in the negotiation case, and up to 67.0% (up to 11.2 bpm) and 18.7% on average in the weighted average case.

With a similar approach, we evaluated the sensitivity of the negotiation to small variations of β_i. For patient 16 265, given the optimal ratio $\bar{R}_\beta = \bar{\beta}_{\mathrm{QT}}/\bar{\beta}_{\mathrm{RR}}$, we varied β_{QT} (keeping constant to their optimal values all the other negotiation parameters) in order to vary the ratio around $\bar{R}_\beta$ in the range $[0, 2\bar{R}_\beta]$ in 100 uniform steps. We obtained that r_{te} increases quasilinearly, getting far from the optimal ratio, up to 11% and 4.5% on average.

Finally, we performed a preliminary test to examine the possibility of implementing our negotiation approach in an embedded device for application in real pacemakers. Specifically, we used SIMIT-ARM [23] to simulate the QT and RR agents and the mediator on an ARM720T chip (32 bit RISC processor with 8 KBytes cache) with 0.09 μm technology. Although cooperative negotiation is by far much more computationally complex than weighted average, we obtained an (over)estimated power consumption of about 50 μW for the former and of about 10 μW for the latter. The value is low enough to consider the possibility in the future of using our approach in implantable pacemakers. However, these issues require substantial further investigations.

V. DISCUSSION

The experimental results (Table III) indicate that negotiation improves single models and weighted average approaches. However, these results are not always significant because the partial models we adopted are very simple. Moreover, the present feasibility study is based on a limited number of patients; a more rigorous investigation on larger data sets is planned. Even if our practical implementation of a heart rate control system is not really significant from a clinical point of view, our aim was to demonstrate that our approach can effectively constitute a common platform to test and compare different combinations of control models. It is at this research level that our contribution is appreciable, since we propose an effective and innovative tool to develop clinically-significant control systems.

It is worth noting that although we kept α_i and β_i constant during all the negotiation sessions in our experiments, their values could change according to the confidence level of the partial models they refer to. Ideally, this should be done at runtime since physiological conditions are usually unknown *a priori*. However, we have not yet explored this possibility. Along this direction, the development of a learning module to automatically calibrate the negotiation parameters on a given patient could drastically reduce the time required for this activity.

It will be also interesting to consider other combinations of partial models for estimating HR; in particular the combination of a model based on physical activity (measured by an accelerometer) and of a feedback model based on respiratory or cardiac indexes (e.g., ventilation, respiratory frequency, stroke volume, and preejection period). In this way, as suggested in [14], the accelerometer open loop control is integrated as an additional and quick forward actuator in a closed loop system developed using the second model. We also plan to refine our approach by changing the properties of the modeling system when the patient is sleeping, and by estimating the optimal delay between atrial and ventricular stimulation (for pacemakers with double pacing mode).

VI. CONCLUSION

In this paper, we have described how a multiagent cooperative negotiation paradigm can be used to combine partial models of heart rate regulation for cardiac pacing applications. The proposed approach involves individual agent (model) optimization and agency optimization, performed according to our negotiation protocol, by which the agents achieve a global agreement. The negotiation protocol presented here generalizes and, as experimental results suggest, improves traditional techniques for combining models in cardiac pacing applications. The adoption of a multiagent paradigm enables us to use the same platform to effectively evaluate and compare virtually every possible combination of partial models of heart rate regulation.

ACKNOWLEDGMENT

The authors would like to thank the two reviewers for their insightful comments.

REFERENCES

[1] A. Goldberger and B. West, "Chaos in physiology: Health or disease," in *Chaos in Biological Systems*, A. Holteon and L. F. Olsen, Eds. New York: Plenum, 1987, pp. 1–5.
[2] R. Rao, B. Aufderheide, and B. Bequette, "Experimental studies on multiple-model predictive control for automated regulation of hemodynamic variables," *IEEE Trans. Biomed. Eng.*, vol. 50, no. 3, pp. 277–288, 2003.
[3] F. Amigoni and N. Gatti, "On the simulation of multiagent-based regulators for physiological processes," in *Multi-Agent Based Simulation II*, J. S. Sichman, F. Bousquet, and P. Davidsson, Eds. New York: Springer-Verlag, 2002, pp. 142–154.
[4] C.-P. Lau, "The range of sensors and algorithms used in rate adaptive cardiac pacing," *Pacing and Clinical Electrophysiology*, vol. 15, pp. 1177–1211, 1992.
[5] M. Glikson and D. Hayes, "Cardiac pacing. A review," *Med. Clin. N. A.*, vol. 85, no. 2, pp. 369–421, 2001.
[6] T. Harvey, *Design of Cardiac Pacemakers*. Piscataway: IEEE Press, 1995, ch. 13, pp. 290–304. rate-adaptive pacemakers.
[7] F. Amigoni, M. Dini, N. Gatti, and M. Somalvico, "Anthropic agency: A multiagent system for physiological processes," *Artif. Intell. Med.*, vol. 27, no. 3, pp. 305–334, 2003.
[8] G. Weiss, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. Cambridge, MA: MIT Press, 1999.
[9] P. Antsaklis, "Defining intelligent control. report of the IEEE control systems society task force on intelligent control," *IEEE Control Syst. Mag.*, vol. 14, no. 3, pp. 58–66, 1999.
[10] D. Meyer, "Towards the global: Complexity, topology, and chaos in modeling, simulation, and computation," in *Proc. Conf. Complex Systems*, Nashua, USA, 1997, pp. 397–409.
[11] K. Astrom, P. Albertos, M. Blanke, A. Isidori, W. Schaufelberger, and R. Sanz, *Control of Complex Systems*. New York: Springer-Verlag, 2001.

[12] P. Heiskanen, "Decentralized method for computing pareto solutions in multiparty negotiations," *Eur J. Oper. Res.*, vol. 117, no. 3, pp. 578–590, 1999.
[13] A. Roth, *Game-Theoretic Models of Bargaining.* Cambridge, U.K. Cambridge Univ. Press, 1985.
[14] J. Werner, M. Hexamer, M. Meine, and B. Lemke, "Restoration of cardio-circulatory regulation by rate-adaptive pacemaker systems: The bioengineering view of a clinical problem," *IEEE Trans. Biomed. Eng.*, vol. 46, no. 9, pp. 1057–1064, 1999.
[15] ECG-Derived Respiration (2004). [Online]. Available: http://www.physionet.org/physiotools/edr/
[16] J. Sarma, R. Sarma, M. Bilitch, D. Katz, and S. Song, "An exponential formula for heart rate dependence of qt interval during exercise and cardiac pacing in humans: Reevaluation of bazett's formula," *Amer. J. Cardiology*, vol. 54, no. 1, pp. 103–108, 1984.
[17] P. Voukydis and J. Krasner, "A physiologically regulated cardiac pacemaker," *Brit. J. Exp. Pathol.*, vol. 48, no. 1, pp. 118–123, 1967.
[18] N. Gatti and F. Amigoni, "A cooperative negotiation protocol for physiological model combination," in *Proc. Int. Conf. Autonomous Agents and Multiagent Systems*, New York, 2004, pp. 656–663.
[19] A. Beda, N. Gatti, and F. Amigoni, "Heart-rate pacing simulation and control via multiagent systems," *Proc. Work. Agents Applied in Health Care*, Valencia, Spain, 2004.
[20] The Research Resource for Complex Physiologic Signals (2004). [Online]. Available: http://www.physionet.org.
[21] QRS Detection and Waveform Boundary Recognition Using ECG-Puwave (2004). [Online]. Available: http://www.physionet.org/physiotools/ecgpuwave/
[22] J. Morè, *Numerical Analysis*, ser. Lecture Notes in Mathematics. New York: Springer-Verlag, 1978, ch. The Levenberg-marquardt Algorithm: Implementation and theory, vol. 630, pp. 105–116.
[23] W. Qin and S. Malik, "Flexible and formal modeling of microprocessors with application to retargetable simulation," *Proc. Design, Automation and Test in Europe Conf.*, pp. 556–561, Munich, Germany, 2003.

Francesco Amigoni (M'98) received the Laurea degree in computer engineering from Politecnico di Milano, Italy in 1996, and the Ph.D. degree in computer engineering and automatica from Politecnico di Milano in 2000.

From December 1999 to September 2000, he has been visiting scholar at the Computer Science Department of Stanford University, Stanford, CA. Since February 2002, he has been an Assistant Professor at the Dipartimento di Elettronica e Informazione, Politecnico di Milano. His main research interests include agents and multiagent systems, mobile robotics, and the philosophical aspects of artificial intelligence. He is the author of about 40 papers published in international journals, books, and conference proceedings.

Alessandro Beda (S'04) received the Laurea degree in biomedical engineering from the Politecnico di Milano, Italy, in April 2003. Since November 2003 he has been a Ph.D. student at the Institue of Sound and Vibration Research of the University of Southampton, U.K. His research is focused on modeling the autonomic regulation of the cardiovascular system.

Nicola Gatti (S'03) received the Laurea degree in biomedical engineering in 2001 from Politecnico di Milano, Italy, where he is currently pursuing the Ph.D. degree in information engineering.

He is the author of about 15 papers published in international journals, books, and conference proceedings. His main research interests include agents and multiagent systems, cooperative negotiation, and robotics for disabled people.