

Implantable Medical Devices as Agents and Part of Multiagent Systems

Touby Drew
Medtronic Neurological, Inc.
800 53rd Avenue NE; MS N240
Columbia Heights, MN, 55421-1200
touby.drew@medtronic.com

Maria Gini
University of Minnesota
4-192 EE/CSci, 200 Union St SE
Minneapolis, MN 55455
gini@cs.umn.edu

ABSTRACT

The consideration of medical implants as an increasingly important population of isolated agents is a valuable perspective that should not be ignored by the agent community. Implanted Medical Device (IMD) applications are complex, naturally distributed, and could benefit from such attention. This paper explores implantable medical devices and their attributes in an agent context and terminology. It submits that an increasing body of IMDs should be considered agents and that there are opportunities for incorporating these implantable agents into multiagent systems (MAS). This will include: (i) Discussion of several IMDs in traditional agent terms. (ii) Discussion of trends and issues in IMDs related to their potential role in MAS. (iii) Experimental exploration of some potential MAS applications in the problem of medical monitoring. (iv) Broader discussion of the value of framing IMDs and applications involving them in the agent paradigm.

Categories and Subject Descriptors

I.2 [Computing Methodologies]: Artificial Intelligence

General Terms

Medical Monitoring, Implant, Medical Device, Agent, Multiagent System, Implantable Agent

1. INTRODUCTION

Currently there is a wide range of different implantable medical devices (IMDs). Although the importance of stents, bio-implants, valves, joint replacements and other “passive” IMDs is clear, this paper focuses largely on those which contain more sophisticated processing capability (sometimes called ‘active’ IMDs in industry). Still, in this category there are quite a number of IMDs made by a range of companies. A variety of pacemakers, Implantable Cardioverter

Defibrillators (ICDs), neurostimulators, programmable drug pumps, and other devices are currently being used by patients around the world. These devices are created by companies including Medtronic¹, Guidant, St. Jude Medical, Boston Scientific, NeuroPace, and many others.

1.1 Implants and the Agent Literature

There is already discussion in the literature of both medical monitoring and systems interacting with implantable devices [2, 6, 9, 10, 15, 19, 24]. As part of this there has even been the limited suggestion that IMDs responsible for monitoring patients “might also benefit from natural language and artificial agent technologies” [6]. To the knowledge of the authors, however, no one has discussed implants as agents or explored them as participants in agent-based systems. The broader discussion of what is and what is not an *agent* is already present in the literature and provides an important context in which we consider this material in Section 2.

1.2 Brief History of Pacemakers

About a decade after the first external pacemakers were designed and built, Medtronic began to produce early implantable pacemakers² and continues to manufacture their complicated descendants today. The earliest pacemakers were fairly simplistic as they supported no real sensing, processing, or communication capabilities. After a few years pacemakers could be recharged and switched to different rates of pacing. Development of wireless techniques for communication emerged. Pacemakers were eventually made to be “responsive” to their environment in a closed loop fashion so that they could, for example, sense the level of activity of their host and change the pacing rate. Many modern devices involve significant processing to facilitate sensing, monitoring, detecting, responding, communicating, and other activities. Several of these support so called “downloadable firmware” in which at least some software code can be sent to the device for it to run in addition to or as a replacement

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS’06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

¹This company employs the primary author of this paper, is the largest IMD maker with a broad range of cardiovascular, neurological, and other devices.

²The earliest work was done by a Swedish team on a patient who ultimately went through 22 pacemakers and 5 lead systems over about 43 years until he died in 2001 from a problem not related to his pacemaker systems [11].

for existing code.^{3 4}

2. IMPLANTS ARE AGENTS?

We submit that some IMDs might be considered to meet many of the criteria and completely satisfy some of the notions set for agency in AI tradition and literature. By reviewing specific IMDs and agent concepts we hope to explore and help elucidate the place of implants in the spectrum and taxonomy of agents. In a longer version of this paper [4] we address these questions in greater detail, including agent properties such as persistence, action-perception relationship, autonomy, social ability, mobility, benevolence and veracity, etc. Here we will discuss environments, sensors, and effectors as key aspects of agents which are “anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors.” [23]

2.1 Implants’ Environments

IMDs transition through a few environments (e.g. in the sequence Manufacturing, Implant(ed), and finally Explant(ed) with the majority of device life in an implanted environment). Consider the SynchroMed IITM (S2) implantable drug pump as an example. It may be useful to think of an implant as a very high reliability robot, but one which is never turned off and is almost always capable of wireless communication. Generally, within a minute of power a S2 has completed some initialization, diagnostics and logging; enabled wireless communications; sounded an audible alarm, and depending on its configuration (or lack thereof) begun to pump drug and sense its capability to do so. After it is completely assembled, a S2 is communicated with wirelessly to ask it to configure itself and perform diagnostics. Eventually, a clean-room worker would put the implant into a sterile package where it would remain until taken out in the operating room. Eventually the device would be put into its implant(ed) environment. At this time, a doctor uses the physician’s programmer to inform the implant of this change. The device will do some processing and make some changes to its behavior. As time passes and different events occur the device eventually will indicate that the physician should consider replacing it (Elective Replacement Indicator or ERI) and may change its behavior (especially if allowed to reach End Of Service or EOS). At some point, the device will be removed from the patient (explanted) and discarded or returned to Medtronic for analysis.

2.2 How Implants Sense Their Environment

There is a great variety of sensors employed in IMDs to allow them to better and more appropriately sense their environment. The S2 device, for example, senses its electromagnetic environment for the presence of wireless telemetry communications. The S2 also monitors itself, checking for stall, measuring its battery voltage, tracking the passage of time, self-diagnosing numerous other issues that might arise and responding appropriately based on its understanding of its environment, capabilities, status, history, and configuration. During testing each S2 pump expels fluid while moni-

³This is presumably valuable for confronting issues related to software/firmware which occur in the field but more fundamentally illustrates the flexibility of the devices.

⁴The capability of IMDs varies; typical clock speeds 32kHz to 3MHz, memory (RAM,ROM,EEPROM,Flash,...) 8kB to 2MB, and miscellaneous custom circuitry.

toring itself. If the pump were to stall (e.g. in the presence of a strong magnet or damaged part) the S2 would detect this and indicate it by audible alarm and indication in a following telemetry session. For a more subtle example consider that the S2 tracks how much volume it has pumped and if it believes it is implanted when this exceeds some threshold it will trigger an audible alarm and consider itself in an ERI or EOS state. Furthermore, if at that time a physician asks the S2 to return to its “implant behavior” it will not (even where the condition causing it to leave its implant behavior was transient such as recognition of a temporary dip in battery voltage). Even so, relative to many other IMDs the S2 is relatively sensor poor.

In the broader pool of IMDs more sensors have been proposed and put to use. We mentioned earlier that cardiac pacemakers make use of sensors to determine patient activity. The determination of activity, posture, etc is accomplished for example using 1, 2, or 3 axis accelerometers in a range of devices like the Medtronic Kappa 900TM. For the most part, cardiac devices also monitor the electrical activity of the heart (ECG) using leads to propagate the signals. Recently some devices including Medtronic’s RevealTM Insertable Loop Recorder and Reveal Plus have enabled “leadless” or subcutaneous sensing of the ECG. Yet other cardiac devices are sensing oxygen, pressure, or other⁵ levels [20].

Furthermore, though traditionally limited in their sensing capabilities, neurostimulators and other IMDs are bringing traditional and novel sensing modalities to new areas of the body; for example the NeuroPace Responsive Neuro-StimulatorTM (RNS) which measures the electrical activity of the brain (EEG) to detect epileptic seizures [18]. Temperature sensors have also been used (e.g. to address tissue heating with recharge). Some devices utilize sensors to simplify patient communication (e.g. using accelerometers to detect a patient tapping on the implant or respond to the presence of a simple magnet). In the future, research may lead to direct “Brain Computer Interfaces” [13]. For shorter life IMDs, a number of other sensors, including notably pH sensors, like those in Medtronic’s BravoTM, can be used. In the near future, IMD sensors may sense a range of new things including chemicals like “hormone levels, anesthetic agents and other factors within the human body such as ... drugs” [16].

2.3 How Implants Affect Their Environment

The S2 devices, used as an example previously, affect their environment by delivering drug to it. With the ability to kill or incapacitate a patient by delivering drug too quickly (overdose; e.g. with morphine in pain therapy) or too slowly (under-dose; e.g. with Baclofen in spasticity therapy), the significance of this effector can hardly be understated. Although more rarely utilized, another effector is the S2’s built in audible alarm which can generate a range of tones and patterns to alert the patient, caregivers, or other user of different conditions. Also of great importance, but more subtle is the device’s ability to generate wireless fields (for the purpose of communicating).

While also able to communicate wirelessly, cardiac devices have slightly different effectors that often include the ability to electrically stimulate the heart. In the future

⁵For example respiration or lung fluid levels via impedance measures.

cardiac and other devices may include both electricity and drug delivering effectors [8]. Furthermore, stimulation is not targeted only to the heart, for example, neurostimulators like Medtronic's ActiviaTM device for Parkinson's therapy and NeuroPace's RNS for epilepsy both use stimulation targeted to the brain, and those like the Cyberonics device deliver stimulation through cuff electrodes to the vagal nerve. Cochlear implants, various sight related implants (including direct interface to the visual cortex, subretinal implant, etc.) [26], diaphragm stimulation [3], and a number of other devices and therapies reflect the diversity of stimulation targets. Even beyond electrical stimulation, drug delivery, wireless communication and audible alarm there are some "other" and emerging technologies for implantable effectors. Certainly an exhaustive list of sensors or effectors is beyond the scope of this paper, but consider for example potential implantable thermoelectric devices, to be used for example in cooling the brains of epilepsy patients [12].

3. IMPLANT TRENDS AND POTENTIAL ROLES IN MULTIAGENT SYSTEMS

The IMD industry is large, competitive and growing quickly as it reflects the developing market for the devices it produces. IMDs themselves are becoming smaller, longer lasting, more broadly applicable and better reimbursed both in response to and as a fuel for their growing industry. As medical devices and other technologies keep patients alive longer the aging population further supports the industry with replacement and combination therapy needs⁶. Important results of these positive feedback systems are that IMD systems are becoming more prevalent, flexible, complex, capable and accessible to potential customers and users. As this continues, we submit that the potential role and importance of IMDs in multiagent systems and as implantable agents themselves is also growing.

While there is strong pressure to add intelligence, capability and flexibility to implants, these things often bring complexity. Even more so than in other systems, there is also pressure to keep implantable components as simple as possible so as to increase their longevity⁷, simplify testing, reduce the likelihood of error, reduce device cost⁸, and speed time to market. With these forces there is pressure to move some system capability outside of the implantable device and into instruments which communicate with it from outside of the body. In practice there has been an increase in communication capabilities, range and complexity in instruments like the patient and physician programmer devices that are designed to communicate with implantable devices.

This growing role of instruments can be seen not only with traditional, life-sustaining uses of implants (e.g. pacemakers), but also as IMDs are increasingly used for non-

⁶It has been well documented in agent literature that the aging population presents a significant target for agent applications [21].

⁷Implants often do as little processing as feasible to reduce the power consumed by the microprocessor and other circuitry. By reducing power, the device can last much longer with the same size battery.

⁸As with any highly regulated, mass produced, embedded system where space and power is at a premium, the cost of memory and other miniature electronic components, even read only memory, can be high and have an effect on production cost.

life-sustaining therapy, combined therapy and monitoring, exclusive-monitoring and other applications. As this trend continues, the volume and complexity of the data available from these devices also grows, presenting the primary problems of medical monitoring with implantable devices. This results in further pressure for increasing memory, automation, decentralization, and prominence in the role of external instruments and in many cases has affected a shift to a more patient centric system involving patient programmers, telemedicine, and transferable electronic medical records. Despite the complexity more patients, physicians and caregivers are using instruments which interact with medical devices and the data they gather and as this continues the role of human computer interaction and "human factors" becomes increasingly important [14].

3.1 Implantable Agent Interaction

Traditionally IMDs have been fairly solitary agents, for the most part interacting only with their environment and the humans that communicate with them almost directly (through simple programmers and the applications that run on them). Consider that in the future this could change, and implants may begin to interact with other agents to a greater extent. This could take the form of implant to implant communications, implant to wearable agent communications, or implants communicating with agents that are transiently (or ubiquitously) available in the environment.

3.2 Implanted Networks

As some of the trends like increasing device prevalence continue, the likelihood of finding multiple implantable devices within communication range increases. This enables implant to implant communication in the field. Although the idea of the IMDs in two acquaintances or even strangers utilizing their transient link is fascinating, the simplest and most tangible cases to consider might be between implants concurrently implanted in the same person. There is some precedent for this, concern, and even design relating to this in existing IMDs. Despite this concern and ongoing work, there is significant potential for and interest in implant-to-implant communication within a patient's body [1, 7]. Some of the motivations for this might include surgical implantation issues, device coordination issues, robustness⁹, as well as resource and capability sharing¹⁰.

3.3 Implantable-Wearable Systems

One of the simplest and most obvious uses of implantable-wearable systems is one in which a single wearable component regularly interacts with a single implantable compo-

⁹In the simplest case, the implants could regularly check on each other, and if for some reason one consistently does not respond appropriately, the other could respond (e.g. alert the patient, log troubleshooting information, etc.). More practically implants could propagate messages to each other increasing both the reliability and size of their telemetry cloud.

¹⁰For example, only one device need have an audible, vibratory or other means of alerting the patient and other devices could use it. In another embodiment, a primary closed loop therapy (e.g. neurostimulation for epilepsy) may recognize that it is being unsuccessful in suppressing a life threatening event and call in the services of a secondary device (e.g. implantable drug pump delivering a potent anti-epileptic drug).

ment. The most fundamental uses of such systems are valuable for supplementing basic implantable capabilities (like processing, sensing, and memory). Consider for example that modern cochlear implants generally include an external device which senses and processes audible inputs and an implanted device which stimulates the auditory nerve. Conversely, as an example, implantable memory is not sufficient to store the large amounts of data modern cardiac (and other) devices monitor and so these devices can be used with an external holter monitor like device to record much greater amounts of data (which may be particularly useful in tuning therapy). For the Medtronic RestoreTM Neurostimulator system, there is currently a wearable system for recharging the implantable. Over a period of time the Restore recharger provides a strong field in short bursts enabling inductive recharge. In between bursts, the recharger measures and models the temperature and then decides if and how long to wait before continuing the recharge (to avoid implant heating issues). Another potential embodiment for such a system, and one perhaps more in line with the tradition of telemedicine, might be one where wearable devices are not the primary link between implantable agents and the rest of the system. Instead, implantable agents may interact directly with agents which are only transiently available (e.g. not worn or implanted in the same body-area network).

3.4 Implantable Communication with Transiently Accessible Devices

Telemedicine systems which occasionally interact with implantable devices are already achievable and in place. Consider for example that a relatively small, privately held company NeuroPace has developed and tested the NeuroPace Data Transmitter (DTR) system. A recent abstract on this system for “on demand monitoring in the Outpatient Environment” explains: “On an intermittent basis, the DTR is connected to a phone line and data is transmitted to the secure server where it can be viewed by the treating physician over any Internet connected computer.” [22] To the knowledge of the authors, it is not yet discussed in the literature whether findings in more traditional cardiac centered telemedicine systems can be generalized to systems and therapies like the NeuroPace DTR. With this caveat in mind, there is evidence that such systems are practical to implement and provide clear potential for cost savings [5]. These systems are not perfect with respect to usability, patient compliance, data loss and other concerns, but Medtronic’s CareLinkTM system, with over 50,000 patients, is an example of their viability at a fairly large scale. These traditional telemedicine systems involving implants require the patient to consciously interact with the system on a regular basis. We submit that the interaction between a transiently available telemedicine system and implants could be further automated and its users better served using a multiagent approach as we will discuss in Section 5.

Beyond traditional telemedicine motivations, there are other applications and embodiments of implantable communication with transiently accessible devices some of which might benefit from agent technologies. We leave these for future work, but will suggest that these may involve distributed clinical research, anonymization, authentication and identification, patient based electronic medical records, device update, maintenance and support, data integration and summarization, emergency response and localization, trauma

and triage, automatic fleet management, improved software platforms for interacting with medical devices, auditing, reporting, novel physiologic computer interfaces, etc.

4. IMPEDIMENTS TO IMPLANTABLE AGENT APPLICATIONS

4.1 Business Case

Although it is fairly clear that IMDs could technically be used as part of a multiagent system, the business need and benefits of such a system are, perhaps, less obvious. As discussed elsewhere in this paper, the use of a MAS approach could improve data management, instrument usability and possibly other aspects of IMD systems, but the cost of their implementation, validation, deployment, is a significant business hurdle. Even a relatively large telemedicine system like CareLinkTM would have a hard time overcoming these costs in the sale of its services. Consider, however, that this is not the business model that really matters. Revenue from telemedicine systems like this is very small in comparison to that from device sales. When viewed as a feature that could differentiate a product or company and potentially grow market share and increase implant sales the business case and potential is much more apparent. Potential future uses in distributed clinical and device research as well as new device applications (e.g. requiring collaboration or improved adaptability) may also become business drivers.

4.2 Technical Stability and Complexity

It is certainly true that some implants are already complex, real-time systems, difficult to upgrade or fix without replacement, performing critical functions, and need to function reliably for many years. For these and other reasons, keeping the implants simple is important because in many ways, simpler is safer. On the other hand, as an underlying premise of this paper we submit that some active implants are already capable of acting as agents there would be little or no added complexity *in the implant*. Interestingly and although not widely discussed in agent terms, this existing design is in large part due to the desire to provide flexible but robust and autonomous systems and is part of what makes it transparent to the implant whether it is communicating with another agent or simply a program.

As a broader MAS including non-implantable agents, there would also be added complexity and thus cost, perhaps especially in testing. The related issues would almost certainly mean slow adoption and limits to functionality initially. It is almost certain in the near future that the connection of IMDs with all other agents and resources of a MAS will be limited. Throughout this paper it is assumed that implants would play a limited, intermittent role in MAS’s but could still have significant potential benefits and in some part it is this very requirement for some independence and automation that makes IMDs and the components that interact with them natural candidates for AS and, we submit, MAS applications.

One interesting subtlety worth noting is that instrument software which communicates with implants currently has an extremely long life cycle as it is expected to be used and maintained or replaced from launch date until the last date at which any device of the supported type is still implanted (which could be more than a decade after its implant) or

on the order of 15 years. Over this time tens of millions of dollars are spent on migrating applications to new platforms as operating systems and other once plentiful and well supported things become obsolete. For this reason portability, extensibility, testability and longevity are especially important. To address some of these and other concerns, some have proposed that applications interact with a single set of autonomous, persistent agents (e.g. for telemetry or data management) which are discovered and interacted with at run-time using appropriately flexible languages.

4.3 Legal, Regulatory, and Security

Though the specifics of these and other important concerns are well beyond the expertise of the authors, it is worth mentioning explicitly their existence and some initial opinions. First of all whether certain IMDs are viewed through the agent lens or discussed with agent terminology or even designed with certain agent properties in mind is not of direct importance to regulatory agencies like the FDA. Instead, the clear and potential risks, mitigations and benefits of such designs are of primary importance. Furthermore, we submit that certain products already approved by such agencies are already reasonably considered agents and in agent terminology by ourselves and some of our peers. On the other hand, this clearly does not change the fact that solid verification, validation, and consideration of risks and alternatives needs to happen for any medical product, especially IMDs. Privacy (e.g. in HIPAA compliance), robustness in the face of misuse, intellectual property clearance, and potential liability issues are at least in large part the responsibility of the manufacturer and essential items to address in potential products (AS, MAS, or not).

5. MAS IN MEDICAL MONITORING WITH IMPLANTS

There are significant motivations for medical monitoring. The most basic approaches to medical monitoring with implantable monitors center on having the patient periodically come in to a clinic where a caregiver manually uses an external tool (e.g. programmer) to off-load data from the IMDs. Telemedicine systems are becoming another important approach, but one which requires patients to manually use a device in their home to send data to physicians electronically (e.g. indirectly through the CareLinkTM system or on electronic medium recorded by a patient programmer). Both of these approaches have positive and negative aspects. For example, a traditional clinical visit may be more easily reimbursed, does not require a patient-device interaction (e.g. CareLinkTM base-station upload session or patient programmer downloading session) and might better facilitate a verbal dialog or other complex, high-fidelity interaction (e.g. visual inspection) which may have positive effects on patient perception as well as outcome. On the other hand, telemedicine approaches may reduce travel costs, increase the speed and ease with which data can be accessed and reviewed, automate electronic documentation, and allow both the patient and physician to participate asynchronously and outside of the clinic. These are generalizations, but we believe that there may be advantages to a medical monitoring system based on a multiagent framework in which an implantable agent interacts with at least one external agent. Some of these benefits could be realized by a MAS as sim-

ple as a basic telemedicine system in which the traditional base-station software has been changed to accommodate an agent which autonomously propagates information from an implantable agent to the physician (or other part of the system).

It is difficult to assess a MAS approach to medical monitoring because such a system did not exist at the start of this project and to test it appropriately would require a clinical study. Furthermore, to generalize the results of such a study the patient population would have to be sufficiently sizable and reflective of the many different device technologies, patient types, and other factors over a sufficient period of time. However, as a small first step on the theoretical, confidence-building journey to this end we provide the results of some simple simulations to explore some of the issues discussed.

5.1 Methods

Using the Swarm agent based simulation framework simulations were created to explore the medical monitoring domain and especially the data loss problem. The simulation was parameterized to allow easy configuration of settings related to simulation type (traditional, telemedicine, or MAS), range over which the rate of data accumulation is set for the population, range over which time elapses between office visits, frequency with which data-offloading is attempted (by the patient in the telemedicine style simulation or by the external agent in the MAS style simulation), etc. Experiments were defined by giving values to these parameters. The amount of data lost and amount of data off-loaded was calculated for each of the experiments described later to achieve the average of 10,000 agent-runs with each agent being simulated for 500 days.

5.2 Experiments

The most basic experiment explored simulation of the traditional approach to medical monitoring with IMDs in which data is accumulated (at a rate randomly distributed between 0 and 30% of capacity each day) until the implant's memory is full and from that point continually lost until the next office visit (set 180 days apart). To explore simulation of the modern telemedicine approach a second simulation alters this such that users attempt (with between 80 and 100% success) to off-load their IMD's data regularly (every 14 days) at home as well as during office visits. We also explored potential implementations of multiagent approaches to medical monitoring including one aimed exclusively at improving the success rate of home off-loading sessions (resulting in between 90 and 100% success).

Finally we explored a more realistic, though at the time not yet realized, multiagent approach to medical monitoring with IMDs that is aimed at approaching continuous monitoring. That is, it simulates agents which not only off-load their data during office visits but also interact with external agents to off-load their data at home. The key difference in this approach is that because the external agents and IMD agents autonomously off-load the data, they can attempt to do this at least daily and keep trying until they are successful (perhaps more than a day later) or determine no data need be off-loaded¹¹.

¹¹In some respects, the interactive communications to determine the need to download data or to avoid downloading could be the most interesting. Telemetry with implantable devices affects their longevity through power drain to carry

Parameter Selection

In real clinical settings, the frequency of office visits, success rate of data-offloading attempts and particularly range of data accumulation rates vary significantly. The choice of 180 days between follow-ups reflects the conservative end of common practice with pacemakers (about every 6 - 12 months) and many other devices. The success rate of telemedicine offloading is also at the conservative end of that reported for CareLink [24]. Though many IMDs sense and process continuous real-time physiological signals and this data would be valuable for various retrospective uses, to our knowledge, there is currently no IMD which is capable of storing all of this for more than a short period. Instead IMDs tend to use tricks including lossy compression, saving only short data snippets surrounding detected events, and using priority schemes so the lowest priority data is overwritten first. The application of these approaches and interpretation of how much of the data is actually valuable varies widely, but we submit that our choice is conservative and would be so conservative for an external device recording the same signals that it would be laughable. It is also worth noting that in our experience the simulation results were not highly sensitive to relatively small changes in these choices.

LATITUDE and CONEXUS

The practical use and implementation of such a MAS may be hard to recognize, but could improve patient care in many ways. The industry is just now taking the first steps in this direction; consider that since the start of this work, Guidant has released the LATITUDETM Communicator system. This allows automated daily monitoring via a bedside device that operates in a manner largely transparent to the patient and is part of what Biotronic, Inc. spokesmen are saying is “definitely the wave of the future.” [25] This system also includes an agent scale component which weighs the patient and gives them spoken feedback and adaptive instructions. Further, since the first submission of this paper Medtronic issued a press release for the first clinical trial of its new line of IMDs with 2 to 5 meter range “Conexus Wireless Telemetry” which “will enhance efficiencies at device implant and during in-office follow-up visits, as well as enable *automatic*, wireless data transmission from the patient’s device to a home monitor.” [17]

5.3 Results

Indicator Functions and Metrics

In the experimental scenarios described previously in this paper, only data which has been off-loaded is retained to be accessed outside of the implant. With this in mind, the amount of data that has been off-loaded is a valuable metric in determining success. Furthermore, the loss of data in medical monitoring is generally undesirable. The value of absolute data retention, or minimal to moderate levels of data loss depends on the intended applications, but may well not be linear. For example, some cases, like the retrospective analysis of continual real-time algorithms on the

out a full off-load. Interactions between IMD and external device could limit or terminate an off-load if the IMDs power reserves are low or if the data buffer is not yet sufficiently full to warrant data collection. In principle, this type of MAS-like interaction could optimize IMD performance across several parameters.




Figure 1: Discrete Loss: Experiments 1 - 4.

gathered data, may require continuous or near continuous data for very long periods to accurately reflect performance. In the experimental scenarios described in this paper, any data which is gathered by an implant full of not-yet-off-loaded data is lost. In volume this is equivalent to the loss of the oldest data not yet off-loaded in a full implant and can alternately be thought of this way if desirable.

The mathematical relationship between the data lost by an agent pair (IMD and external home monitoring agent) is equal to the difference between the data on the device and the sum of the data off-loaded and the total data acquired by the IMD (see equation 1).

$$\begin{aligned} \text{dataLostFrom}(\text{agent}) &= \text{totalDataAcquiredFrom}(\text{agent}) - \\ &\quad [\text{dataOffloadedFrom}(\text{agent}) \\ &\quad + \text{dataOnDeviceOf}(\text{agent})] \end{aligned} \quad (1)$$

We created a metric allowing better understanding of the data in graphical analysis called *discrete loss* or *dloss*. In this context, *dloss* is defined for each agent as 1 if the agent has not lost data since the last off-loading and 0 if it has (see equation 3). The discrete loss of the agent population is thus the average across agents (see equation 2) and reflects the portion of the population that has not lost data since their last off-loading. Thus for n agents data loss is defined as:

$$\text{dloss} = \frac{\sum_{\text{agent}=1}^n \text{dloss}(\text{agent})}{n} \quad (2)$$

where for each agent

$$\text{dloss}(\text{agent}) = \begin{cases} 1 & \text{for : dataLostSinceOffload} = 0 \\ 0 & \text{for : dataLostSinceOffload} > 0 \end{cases} \quad (3)$$

This indicator is useful in recognizing and understanding cases in which only a fraction of the agents are responsible for all of the data loss.

5.4 Discussion

Comparing the results of the simulations, it is clear that, given our assumptions, there are some potential advantages to be had with the application of multiagent techniques over traditional and telemedicine approaches in terms of data loss and data off-loading.

The simulations clearly reflect the data loss relationship (equation 1) and the parameters chosen relating to device memory capacity, data acquisition rate and device off-loading rate. As discussed, the parameters chosen reflect our subjective experience with medical monitoring implants but under key assumptions which might be further explored in clinical studies and other future work. In speaking with medical and other experts who work with data and electronics, the diminutive amount of memory available in long-term implantable devices is often surprising.¹² In contrast consider that a single channel of EEG or ECG signal sampled at 256Hz with 2-byte precision might generate about 30kBytes per minute, filling all of memory for even extreme IMDs in only 1 or 2 hours.

The intuitive benefits of telemedicine over traditional approaches was apparent in our results with more than an order of magnitude improvement over the simulated traditional approach (Experiment 1) in terms of data off-loaded (and a factor of about 2.3 reduction in data lost). This improvement is due to the additional data-offloading opportunities afforded by the supplementary telemedicine approach.

Also as one might expect, the multiagent system (MAS) simulation aimed only at improving the off-loading success rate showed a small improvement over the simulated telemedicine approach as a result of the increase in off-loading success rate (gains in data off-loading and data loss rates of about 7 and 10% respectively). Considering only these differences, this rudimentary approach provides only incremental improvement that would likely not justify the cost of designing, implementing and validating it. On the other hand, it may be worth noting that this approach allows the patient to not worry about the process or their implant and perhaps even forget that it exists. This may be highly beneficial for patient usability and compliance (especially hard with elderly, cognitively impaired, or non-tech savvy), and perhaps in other ways as well. For example, it has been noted that patients who feel that their role is to recharge or otherwise service an implantable device may do so too frequently. This can negatively affect the device (e.g. depleting the battery) and reflect lost time and increased anxiety. Similarly, patient confusion and misuse of telemedicine equipment can result in costs to their healthcare and service providers by burdening call centers and technical service providers.

The final experiment, aimed at exploring a more realistic MAS approach demonstrated an improvement of almost 60% in data off-loaded and about 600% reduction in data loss (over the rudimentary MAS approach aimed only and increasing off-loading success rate). At the same time benefits in usability and transparency of the rudimentary MAS approach are still present. This simulation reflects a significant potential advantage of thinking of the medical monitoring system as a MAS. Though the findings may be intuitive, this is key evidence that suggests a MAS approach is appropriate because it hinges on autonomy and avoids patient compliance as a major issue. In other words, as a more complete paradigm shift it approaches a level of autonomous communication and cooperation between implant and telemedicine agent that is not practical in traditional

approaches where asking the patient to constantly measure the contents of their implant, make off-loading decisions and initiate appropriate off-loads would be unreasonably burdensome and far from transparent.

Implantable Agents in Multiagent Systems

Traditionally the implantable medical device market has been focused very heavily on physical implantable technologies and products (as these drive revenue for companies like Medtronic). Telemedicine, however, underscores the increasingly important role of instruments and software. Furthermore, as IMD-based systems and software continue to develop, the need for them to be distributed as highly robust, portable and increasingly automated systems is clear. We submit that agent based approaches can be applied to address this need. The first step may be thinking of implants as agents. IMD-wearable, implanted network, and transient IMD interaction with agents in the environment are emerging interaction patterns that can be used as lenses through which opportunities for IMD systems to benefit from the agent paradigm can identified and developed.

The experimental work in this paper is meant only as an example, not as a conclusive prescription for implementation. Taking the paradigm shift to agency and thinking of implants as potential parts of MAS allows for other, and perhaps better, solutions to this and many other problems. For example, we recognize that the implantable-wearable solution to medical monitoring discussed previously has not been considered adequately in our simulations.

One potential use of implantable-wearable systems is in providing access between IMDs and other external entities. Currently the external components of such systems are primarily simple tools to enable patient interaction with implantables (e.g. a patient programmer device which can be worn on the belt). In the future, however, small wearable telemetry booster devices could act to connect the short-range, proprietary wireless systems in modern implants to remote systems via longer distance wireless schemes. Were such a system to come into place it could create potential for a very large system in which implantable agents, hospital systems, emergency systems, company systems, and perhaps even patients' personal systems are intertwined. The role of agents in this system is not certain, but their advantages in automation, etc. could well shape this potential into what may become a huge and heterogeneous multiagent system with an important fraction of its constituents taking on an implantable form. This type of approach could open up new opportunities for a richer, more reliable medical monitoring systems addressing not only continuity of data issues, but also reporting, data priority, caching, analysis automation, data integration, distributed anonymous clinical research, and other important areas. Although some of the applications of these systems could well address these various goals and concerns, we leave this for future work but examine some examples in the longer version of this paper [4].

6. CONCLUSIONS AND FUTURE WORK

The use of IMDs is a large, complex, and growing domain which naturally lends itself to applications of distributed artificial intelligence. This field is particularly appropriate for autonomous agent and multiagent system consideration as a result of its vast base of existing autonomous, distributed, largely solitary agents and emerging needs and

¹²Implantable memory takes up valuable space and current in an IMD. In contrast to commercial memory, implantable memory is specially designed and validated to be highly reliable and draw minuscule amounts of power.

capabilities for these agents to work as part of multiagent systems in medical monitoring and therapy. Though work in this nascent area of overlap between AAMAS and IMDs has already begun¹³ it is not thought of in these terms and could benefit significantly from more explicit consideration.

7. REFERENCES

- [1] M. A. Arent and I. A. Sun Microsystems. EP1022035A1: Communication network and devices to be implanted within a subject.
- [2] T. Bachmor, J. Schochlin, and A. Bolz. Transmitting patient and device data via gsm - central management for decentral mobile medical devices. *Biomedizinische Technik. Biomedical Engineering*, 47(1):346–349, 2002.
- [3] A. F. DiMarco. Restoration of respiratory muscle function following spinal cord injury. Review of electrical and magnetic stimulation techniques. *Respir Physiol Neurobiol*, 147(2-3):273–287, Jul 2005.
- [4] T. Drew and M. Gini. Implantable Medical Devices as Agents and Part of MAS. November 2005.
- [5] L. Fauchier, N. Sadoul, C. Kouakam, F. Briand, M. Chauvin, D. Babuty, and J. Clementy. Potential cost savings by telemedicine-assisted long-term care of implantable cardioverter defibrillator recipients. *PACE*, 28(2):S255–S259, January 2005.
- [6] G. Flammia. The Web: A communication medium for healthcare. *IEEE Intelligent Systems*, 17(2):88–89, 2002.
- [7] H. D. Funke and I. A. Medtronic. Patent: US4987897: Body bus medical device communication system.
- [8] S. D. Harpstead, L. M. Otten, T. R. Prentice, and I. A. Medtronic. Patent: US5697951: Implantable stimulation and drug infusion techniques.
- [9] J. Huang, N. R. Jennings, and J. Fox. An agent architecture for distributed medical care. In *ECAI-94: Proceedings of the Workshop on Agent Theories, Architectures, and Languages on Intelligent Agents*, pages 219–232, New York, NY, USA, 1995. Springer-Verlag New York, Inc.
- [10] V. Koutkias, I. Chouvarda, and N. Maglaveras. Multi-agent system architecture for heart failure management in a home care environment. *IEEE Computers in Cardiology*, 30(3):383–386, 2003.
- [11] B. Larsson, H. Elmqvist, L. Ryden, and H. Schuller. Lessons from the first patient with an implanted pacemaker: 1958–2001. *Pacing Clin Electrophysiol*, 26(1 Pt 1):114–124, Jan 2003.
- [12] R. P. Lesser, W. R. S. Webber, Y. Motamedi, Gholam K. and Mizuno-Matsumoto, and T. J. H. U. (Assignee). Patent: US20050149123A1: Techniques using heat flow management, stimulation, and signal analysis to treat medical disorders.
- [13] S. P. Levine, J. E. Huggins, S. L. BeMent, R. K. Kushwaha, L. A. Schuh, E. A. Passaro, M. M. Rohde, and D. A. Ross. Identification of electrocorticogram patterns as the basis for a direct brain interface. *J Clin Neurophysiol*, 16(5):439–447, Sep 1999.
- [14] C. Lewis. Emerging Trends in Medical Device Technology: Home Is Where the Heart Monitor Is. *FDA Consumer*, May 2001.
- [15] N. Maglaveras, G. Gogou, I. Chouvarda, V. Koutakias, I. Lekka, D. Adamidas, C. Karvounis, G. Louridas, and E. Balas. Using contact centers in tele-management and home care of congestive heart failure patients : The CHS experience. *IEEE Computers in Cardiology*, 29(02):281–284, 2002.
- [16] I. Medtronic. Carmel Biosensors Collaborates With Medtronic In Research And Development Of Innovative Biosensors. [Online; accessed 07-Nov-2005].
- [17] I. Medtronic. Medtronic Initiates “Concerto AT” Study In United States. [Online; accessed 01-February-2006].
- [18] I. NeuroPace. Implantation and testing of responsive neurostimulator (RNS) system for epilepsy. *Annual Meeting of the American Society for Stereotactic and Functional Neurosurgery*, May 2003. [Online; accessed 07-November-2005].
- [19] Z. Obrenovic, D. Starcevic, E. Jovanov, and V. Radivojevic. An agent based framework for virtual medical devices. In *AAMAS ’02: Proceedings of the 1st International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 659–660, New York, NY, USA, 2002. ACM Press.
- [20] A. Ohlsson, S. H. Kubo, D. Steinhaus, D. T. Connelly, S. Adler, C. Bitkover, R. Nordlander, L. Ryden, and T. Bennett. Continuous ambulatory monitoring of absolute right ventricular pressure and mixed venous oxygen saturation in patients with heart failure using an implantable haemodynamic monitor: results of a 1 year multicentre feasibility study. *European Heart Journal*, 22(11):942–954, June 2001. Clinical Trial.
- [21] M. E. Pollack. Intelligent Technology for an Aging Population: The Use of AI to Assist Elders with Cognitive Impairment. *AI Magazine*, 2005.
- [22] E. K. Ritzl, E. H. Kossoff, G. K. Bergey, P. S. Coe, and D. S. Gupta. Complementing the responsive neurostimulator system with a patient operated data transmitter on demand monitoring in the outpatient environment, 2005. [Abstract; accessed 11-Nov-2005].
- [23] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2003.
- [24] M. H. Schoenfeld, S. J. Compton, R. H. Mead, D. N. Weiss, L. Sherfese, J. Englund, and L. R. Mongeon. Remote monitoring of implantable cardioverter defibrillators: A prospective analysis. *PACE*, 27(1):757–763, June 2002.
- [25] J. Van. Remotes keep tab on heart devices: Wireless tracking offers piece of mind. *Chicago Tribune*, November 2005.
- [26] P. Walter. [Electronic visual prostheses]. *Klin Monatsbl Augenheilkd*, 222(6):471–479, Jun 2005.

¹³Consider previous examples like the RestoreTMIMD that works with a wearable recharger under heating constraints.