VETO!
A Peer-to-Peer Paradigm Based Medical Knowledge Management System

Claus Eikemeier

University of Bremen, Faculty of Mathematics and Informatics,
Bibliotheksstr. 1, 28359 Bremen, Germany
`cei@informatik.uni-bremen.de`
`http://www.eikemeier.ch.vu`

Abstract. Insufficient knowledge of the patient history leads to severe medical, ethical and economical problems [5]. Adverse drug reactions (ADR) or repeated examinations are well-known results. The paper presents a blueprint of applying the *VETO!* principle to medical knowledge management. This helps to reduce such errors.
VETO! is a system that exploits distributed knowledge of the patient record while keeping high medical security standards. The data is uncovered in a Clearinghouse where possible medical risks are identified. If problems are anticipated, a "Veto!"-message is released to the originating entity. In severe cases additional information can be provided. Unlike other systems, *VETO!* is an evolutionary P2P system: the resulting quality increases with the number of contributors. It exploits several beneficial characteristics of Peer-to-Peer systems.

1 Introduction

Knowledge plays an important role in medicine. Decision support systems (DSS) can generate case and domain specific advice [8, chap. 16] from the provided knowledge. DSS are part of artificial settings or can help human agents. They are categorized into autonomous systems on one hand and solicited or unsolicited advice systems on the other. In settings where human users are involved, each system must handle inconsistencies, data gaps, errors and other problems that occur in real life gracefully. The system design and its implementation has to reflect these requirements to provide a convincing performance. Peer-to-Peer (P2P) systems provide the required characteristics: superior performance of the service based on massive parallel contributions of the peers, on the similarity of the nodes and on non-determined dynamics of parts of the individual peers [1]. Beneficial effects increase if more peers participate and contribute (network effects).

The field of application puts special restrictions on the implementation of the DSS. Medical applications, for example, have to deal with restrictions like law (e.g. obligation to store medical data) or data security. On the other hand
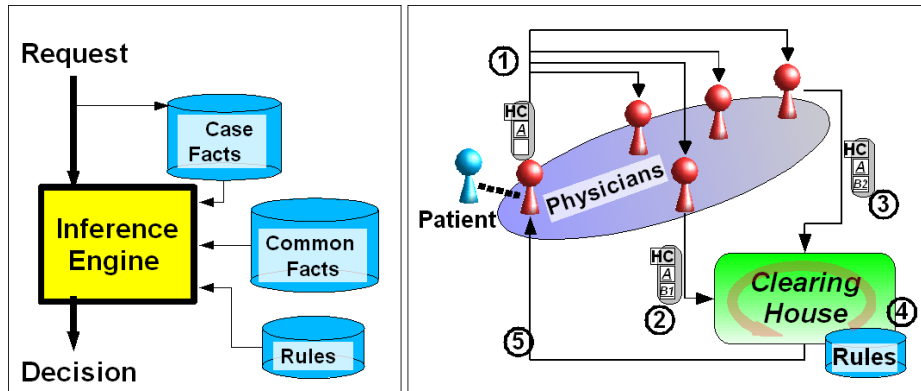
Fig. 1. (left) Building blocks of an expert system

Fig. 2. (right) Overview of VETO!-System

such environment provides beneficial aspects like a controlled vocabularies or common application and communication interfaces (e.g. HL/7 [2]).

There still remain problems because of the traditional, unconnected "data islands" in healthcare. The missing integration of information in the healthcare domain is reason for the afore mentioned problems [3, 7]. A surveillance service that considers the medication and treatment status of the patient reduces these problems. In real world settings, the distributedly stored patient record prevents from accessing all data when it is needed.

This paper explains the ideas of a P2P based decision support system: the *VETO!* system gathers data from the peer community (here: healthcare providers) and checks it against domain specific knowledge of experts. *VETO!* works unsolicited, automatic and in near real-time. Because of the possible incompleteness of the knowledge base, the system is not able to prove a certain decision, but it is able to return a negative judgement (a 'veto') if the decision is alarming or even harmful.

2 Description of the system

The main idea of *VETO!* is similar to falsification of a hypothesis: if there is a contradiction, the hypothesis must be discarded, if not, we do not know. If a decision is based on distributed data, all available information should be taken into account as each additional piece of data enhances the decision. Besides available case facts (e.g. start of treatment), intelligent systems also use an inference engine, rules (expert knowledge) and common pharmaceutical facts (cf. Fig. 1).

The case related part of the knowledge base changes over time as more and more information arrives from different data providers. The amount of knowledge increases as each examination or illness contributes to the underlying patient record. The patient record changes over time. One can briefly think of an

erroneous and incomplete facts base: there is a need for gathering of the actual patient data in real-time and for dealing with incomplete and inconsistent data.

The reasoning rules represent the expert knowledge from previous experience with drugs or the treatment or guidelines (medical, economical or other). As example, an ADR rule can be written as:

$$factA \land factB \to ADR$$

with $factX = takes\_in\_drug\_X$ and $ADR = degree\_of\_possible\_negative\_effect$ Besides simple logics, temporal logics and sometimes even spatial logics is needed to reason over the provided knowledge. E.g. drug interactions commonly only occur during some time after intake of another drug:

$$ADR \land (timeB - timeA < criticalTime\_A\_B) \to warning$$

If a rule is fired, a possible negative effect can occur. This information is delivered to the physician who initiated the process.

The treating physician (lef in Fig. 2) triggers the surveillance process by issuing a prescription to the patient. Therefore some information is spread to his colleagues (1). The message itself (A) is encrypted (PKI with key of Clearinghouse), only a hashcode that identifies a group (!) of patients, is visible. The choice of the hashcode is founded by demographics of the population: the group size per hashcode should be big enough to prevent from identifying a single patient, but needs to be small enough to not produce too much unnecessary work for the Clearinghouse (CH). For example, the hashcode consists of one or two letters of the name and the month of birth. If patients are assigned sequential numbers with a patient card ID, a modulo operation (e.g. mod 100) can generate the group membership. So neither the message (A encrypted) nor the hash code (it identifies a group) unveils the identity of the patient. Each healthcare provider who has information that matches the hashcode adds its encrypted information (2,3) and sends this information to the CH. The CH decrypts the different message parts A,B (4). Iff the patient identity in both parts of the message match and the rules predict a harmful outcome, a notification (warning) is issued to the physician (5). He checks if the additional information is of relevance for his former decision.

The A and B part of the message can only be decrypted by the CH. The CH is the only part that needs to be operated by an official authority to ensure data safety. In general, the CH performs the same tasks as a standard inference engine in other expert systems. Therefore existing treatment surveillance systems and tools from the AI domain match the technical needs of the *VETO!* system.

3 Related work

Two competing approaches are based on different premises. The card approach stores the data 'near' the person, e.g. gathered on the patient card. Card solutions demand for a similar control logic because of the increased amount of data to be

considered. Another solution includes a central "global" repository to contain all patient data. The NHS system is of the second type [4]. Both are revolutionary "big" solutions that change the involved healthcare system significantly.

4 Summary and Discussion

Ethical and economical reasons demand for an enhanced medical quality. *VETO!* is an automatic, reactive DSS that issues warnings related to medical treatment quality. The distributed patient record remains unchanged, information is gathered in realtime. This ressembles to P2P systems.

Events like a drug prescription or starting a new treatment trigger the automatic surveillance process. The initiating peer queries the medical community for information on the patient. A Clearinghouse processes the collected data. Identity information is unveiled only there. On principle only negative effects of the intended treatment can be detected. In the case of indicated danger, further information is reveiled to the physician, who can react accordingly. This helps the physician to avoid errorneous decisions and thus increases the quality of the treatment.

The *VETO!* system is not yet implemented, currently it is seen as a model how medical services can be implemented when new technological possibilities are available. As healthcare telematic platforms are currently under development, the idea of *VETO!* can initiate seminal discussions of how to softly implement the needed services. This is in contrast to "revolutionary" approaches that demand for significant changes of the infrastructure.

References

1. Fattah, H.M.: P2P: How Peer-to-Peer Technology is Revolutionizing the Way We Do Business. Dearborn Trade Publishing, Chicago, Il., 2002.
2. Health Level Seven Consortium *HL/7 homepage* online at URL: http://www.hl7.org (access on 2004-02-17)
3. Lazarou, J., Pomeranz, BH., Corey, Pn. *Incidence of adverse drug reactions in hospitalized patients: a meta-analysis of prospective studies* JAMA, 1998-04-15, 279(15), pp.1200-5
4. National Health System news *Care records contract* online at URL: http://www.nhs.uk/nhsupdate/news.asp?newsid=772 (access on 2004-03-29)
5. Shiff. G. et al.: Linking Laboratory and Pharmacy. Arch Intern Med, Vol 163, pp. 893-900, 2003
6. Schollmeier, R.: A Definition of Peer-to-Peer Networking for the Classification of Peer-to-Peer Architectures and Applications. online at personal homepage (www.lkn.e-technik/tu-meunchen.de) (access on 2002-07-04).
7. Stieler, W.: *Elektronische Gesundheitskarte soll Milliarden einsparen* (Electronic health card shall save billions of Euros), Heise News, online at URL: http://www.heise.de(access on 2004-03-27)
8. Van Bemmel, J., Musen, M.A.: *Handbook of Medical Informatics* Springer, Heidelberg 1999