A Light-Weight System Extension Supporting Document-based Processes in Healthcare

Christoph P. Neumann and Richard Lenz

Friedrich-Alexander University,
Erlangen-Nuremberg, Germany,
christoph.neumann@cs.fau.de,
http://www6.cs.fau.de/

Abstract. Inadequate availability of patient information is a major cause for medical errors and affects costs in healthcare. Traditional information integration in healthcare does not solve the problem. Applying the classic diagnostic-therapeuthic cycle as the model for a document-oriented information exchange protocol allows to foster inter-institutional information exchange in healthcare. The goal of the proposed architecture is to provide information exchange between strict autonomous healthcare institutions, bridging the gap between primary and secondary care, following traditional paper-based working practice. The combination of a RESTful architecture with a distributed light-weight workflow model provides minimized requirement for participating systems.

Topics: Process oriented system architectures in healthcare, process interoperability & standards in healthcare, context-aware healthcare processes, inter-institutional healthcare information systems, document-oriented integration

1 Introduction

In a systems analysis of adverse drug events, 18% of the medical errors were associated with inadequate availability of patient information [1]. The problem of inadequate availability of patient information as a major cause for medical errors is aggravated by the rise of healthcare networks and the increasing number of healthcare parties that are involved in a treatment: The aging of western society affects the public health sector, chronic diseases and multimorbidity become the focus of interest, and the cost pressure rises. For example, cancer, diabetes, asthma, or cardiac insufficiency require more healthcare parties than common diseases. Coevally, the rapid advance in medicine leads to an advancing specialization of physicians that is an additional cause for the increasing number of involved parties regarding a single patient's treatment. For improving the treatment quality and in order to avoid unnecessary costs, an effective information and communication technology is vital for the support of inter-institutional patient treatment.

In order to foster the continuity of care, the inter-institutional cooperation needs to bridge the current gap between institutions of the primary and sec-

The original publication is available at www.springerlink.com (DOI: 10.1007/978-3-642-12186-9_54)

ondary care. Such effort must not instrument regional standards, as it is done in *regional healthcare information networks* (RHIN) [2], but *transregional standards*. Accomplishing information exchange in distributed healthcare scenarios requires the integration of heterogeneous and *strict autonomous IT systems*. To allow for *inter-institutional cooperation* the support of distributed and seamless flow of information is required, thus changing paradigms from closed and hegemonic to open and distributed architectures. The proposed architecture adheres to these boundary conditions.

2 Idea and Objectives

This paper outlines the goal to focus on a document-oriented paradigm [3] for healthcare system integration following the paper-based clinical work practice as reference model. There are two prime objectives of the proposed solution. The first is the abdication of any central server, like joint databases, transaction monitors, and central context managers, as adherence to the strict autonomy of the institutions. The second objective is the application of document-oriented integration with lightweight interfaces instead of service-oriented integration with semantically rich interfaces. Document-orientation favors local autonomy by adhering to the design goal of loose coupling.

A subsequent design objective is to aim for minimal standards in order to yield minimal requirements to the participating systems. Favoring local autonomy over central hegemony requires, for example, that distribution of information will not be enforced, but is voluntary and process participation can be supplemented on demand. Platform independence and the avoidance of vendor lock-ins require that the basic architecture is decoupled from any specifically instrumented middleware and components off-the-shelf. Loose coupling, as desirable property of the proposed system extensions [4], particularly means that it should be possible to add and remove participants without any modification of other participants. Thus, without previously interconnecting two participants, it should be possible to interchange information[1].

A risk in instrumenting a central content storage, like German D2D or Google Health[2], is an information leak that potentially involves all patients. This is not comparable to any possible abuse scenario in today's paper-based infrastructure: No current healthcare institution hoards information about so many patients as will do any centralized solution for inter-institutional scenarios. The distributed approach mirrors the current state in paper-based working practice and provides *information locality*: The patient information is available only to the directly involved healthcare systems. As a result, the consequences of a security breach are limited to a fraction.

[1] Excluding considerations for a federated, large-scale security infrastructure which might still impose coupling on certain levels.

[2] `http://www.google.com/health`

3 Methods

This section will motivate the document-oriented integration approach with its capacity to support loose coupling and deferred system design. For implementation, a REST[3] architecture will be instrumented which provides document-orientation naturally.

3.1 Foundations of Semantic Compatibility

Exchange of patient information among institutions requires data compatibility. *Data integration* achieves data compatibility either by common standards or by data transformation. Data integration for medical processes requires standards for medical terminology that have to deal with volatile medical concepts [5]. Over the intervening years numerous standards for medical ontologies have been created on *type level* for system implementers at design-time and on *instance level* primarily for end-users as a semantic reference at run-time.

At instance level, standards like ICD , SNOMED , and LOINC exist which unremittingly evolve over time. The HL7[4] v2 is a well-established standard for clinical message specification [6]. It is a standard on type level, and incorporates coding schemes and terminologies on instance level. The HL7 v2 standard allows for the specification of self-defined messages, which has lead to incompatibilities. The relatively new HL7 v3 standard is based on the HL7 v3 *reference information model* (RIM) and is radically different from the v2 standard: It allows for new types to be derived from a limited number of core classes, enabling RIM-based systems to handle even unknown message-types in a generic way.

Electronic medical record (EMR) and *electronic health records* (EHR), e.g. [7], are popular approaches to share patient related information among institutions. They typically contain data that can be extracted on demand. Yet, it is unclear how these systems scale and how direct communication between institutions can be effectively supported in large-scale scenarios. A conceptual change from messages and records to documents is provided by the HL7 v3 *clinical document architecture* (CDA). CDA provides a framework for XML-structured medical documents. EMRs and EHRs fit in the notion of our approach by specifications like the *continuity of care document* (CCD), a U.S.-specific standard, which is a constraint on HL7 and focuses on document-oriented medical content types. In Germany, based on HL7 v3 CDA, the SCIPHOX[5] working group develops specific document content types for German healthcare; for example, referral vouchers and discharge letters[6]. The CCD and SCIPHOX standards do not consider process history or coordination information. Any new standards should respect the ones

[3] REpresentational State Transfer

[4] Health Level 7, `http://www.hl7.org`

[5] Standardized Communication of Information Systems in Physician Offices and Hospitals using XML

[6] Particularly "eArztbrief SCIPHOX CDA R1" and its advancement "eArztbrief VHitG CDA R2"

already in practice for backwards-compatibility and to achieve and maximize acceptance. Therefore, the exchanged documents of the proposed system extension are motivated by the increasing importance of HL7 v3 CDA. Yet, the proposed infrastructure does not depend on CDA.

3.2 Interface- vs. Document-Orientation

System integration in healthcare is traditionally based on interface-orientation. Three-tier network-based architectures with remote procedure calls are yet the dominant architectural style for information systems. The most common technological occurrence of remote invocations is based on SOAP with WSDL. The *interface-oriented integration* focuses on available functionality, and the integration method affects semantically rich interfaces. An invocation uses parameters to detail its synchronous service request to a target system. In interface-oriented integration the information being passed is not necessarily viable on its own but often in the context of the service request only.

Even if a service is triggered event-oriented using asynchronous messaging, like it is done in HL7 v2-based systems, such parameters or messages essentially represent transient fine-grained information that is assimilated by the targeted system. The three main problems in information integration projects, including healthcare systems, are insufficient synchronization of redundant data, problems with data consistency, and functional overlapping [8]. Therefore, interface-oriented and message-oriented integration between distinct institutions is complex and custom-designed.

In contrast, documents are coarse-grained, self-contained, and viable. A document can exist independently from the system it stems from. Changes are not propagated by update information, but by creating an updated document that replaces its predecessor. The *document-oriented integration* focuses on available information, and the integration method affects the semantic scalability of document models, using standardized and minimal interfaces for hand-over. Redundancy in data distribution is not critical with documents because, due to the self-containedness, a synchronization in the classical sense is not required. Instead, document versioning and variant management solutions are effective. Likewise are traditional data consistency checks confined to the scope of the document, inconsistencies between documents represent logical errors or divergence in opinion on such semantically high level that a conflict can only be detected or solved by specialized decision support systems or humans.

3.3 Loose Coupling and Deferred System Design

The *deferred system design* principle of evolutionary systems [9] requires semantic decisions not to be frozen in an interface schema because they are hard to revise. Applying a document-oriented approach improves the adaptability of the information systems by deferring schema decisions from design-time to deploy- or run-time.

HL7 v3 CDA provides semantic scalability for healthcare documents, both because this has been an inherent feature of the underlying RIM and because CDA is particularly structured in three levels of semantic abstraction: CDA level 1 is the unconstrained CDA specification. CDA level 2 applies section-level templates. CDA level 3 applies entry-level templates. For example, CDA level 1 simply ensures the ability to display a document like a PDF file. Any CDA document can be accepted without immediate support for processing. Advanced semantic processing support of CDA level 2 or 3 can be added to the system, seamlessly enhancing the information value of already stored CDA documents. HL7 v3 CDA supports deferred system design by its semantic scalability. The proposed solution applies CDA as primary document type, extending the CDA header with information exchange protocol and process participant information.

As an architectural style which implies minimal requirements to be supported by participating systems, the REST paradigm is applied in DMPS. The REST architectural style is the generalization of the architecture of the web, proposed by Fielding [10], the co-author of the original HTTP[7] with Berners-Lee. REST provides a paradigm for decentralizing applications in which applications are decomposed into resources with various representations and links between them. The RESTful approach does not require an additional marshaling layer as do interface-oriented remote invocation approaches. The focus of REST lies in the explicit modeling of the representation; in the interface-oriented approach the representation is often generated implicitly by vendor-specific development environment tools. The benefits of a REST architecture are its minimal requirements and its coarse-grained resource/representation approach which compels loose-coupling and allows for a document-oriented architecture.

4 Proposed Solution

The proposed exchange protocol and system extension is named *distributed medical process support* (DMPS). DMPS targets a document-oriented process support between strictly autonomous institutions. It follows the paper-based work practice as reference model, focusing on referral vouchers and result reporting by letters of referral. The DMPS exchange protocol is deduced from the traditional diagnostic-therapeutic cycle [11] in fig. 1.

The basic technological DMPS adoption of the diagnostic-therapeutic cycle will provide remote information exchange for the edges of the cycle. It is based on the REST architectural style, requiring only HTTP and support of HL7 v3 CDA documents as well as the implementation of the procotol statechart as it will be described in sect. 4.2.

In contrast to clinical environments, in which the focus of process support often relates to decision support (which is e.g. based on rule-based artificial intelligence) and process control (to guarantee process quality), the inter-institutional

[7] HyperText Transfer Protocol

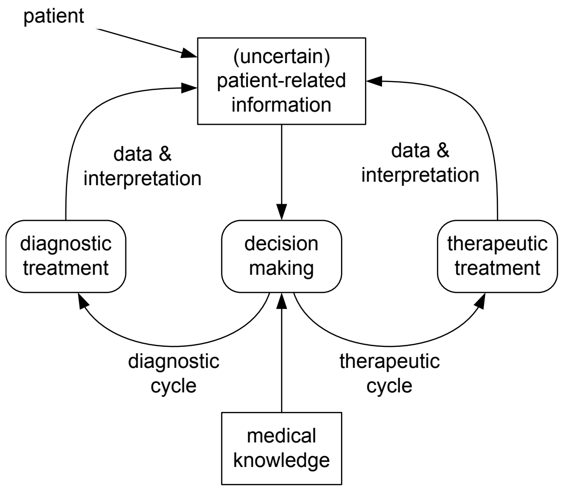

Fig. 1. The traditional diagnostic-therapeutic cycle

process support only supplements coordination information and enables coordination guidance instead of control. At the moment our approach focuses on sharing a distributed process identifier and managing a merged process history of participating institutions. To provide context awareness in form of a process history is one of several most important challenges for inter-institutional information system integration in healthcare [12].

4.1 Scenario

The basic scenario of the intended DMPS system extension is outlined in fig. 2. Document transportation can be done *online* by REST access or *offline* by transporting documents with the aid of an external medium like an eGK smart card or a flash drive. Another dimension of transport classification is the *counterpart accessibility*: the delivery can be done by *direct* communication or by a *mediated* approach. The direct approach delivers documents directly from the source DMPS node to the sink DMPS node using online or offline mechanisms, whereas the mediated approach uses a third DMPS node as intermediator, for example a DMPS node that hosts DMPS accounts for patients. The third classification dimension, which is not visualized in fig. 2, is the *counterpart identity availability*: the counterpart can be known at the time of document shipping (*addressed* communication) or the counterpart can yet be unknown (*unaddressed* communication).

The DMPS scenario considers three communication variations: *direct/online*, *direct/offline*, and *mediated*. The mediated approach is a composite of directed communications, and arbitrarily uses online or offline transport for each of its atomic edges. The distinction due to the counterpart identity availability is motivated by most basic examples: Letters of referral are addressed communication scenarios, but referral vouchers are unaddressed ones because the voucher

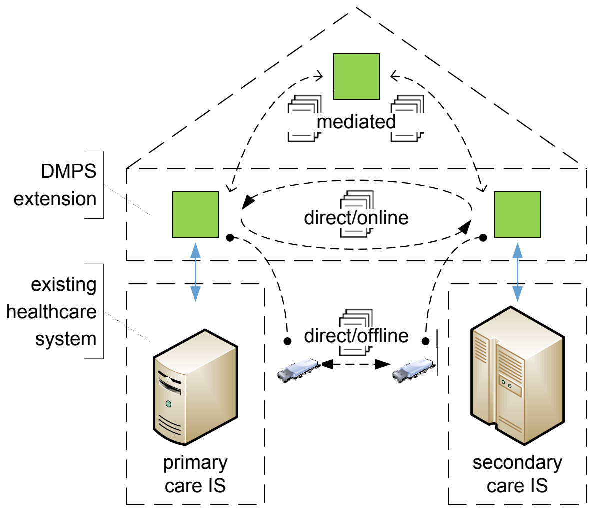

Fig. 2. The DMPS scenario

lists only the medical specialty, while the patient can arbitrarily choose the actual institution and medical specialist. A referral voucher or a prescription can, for example, be delivered in a mediated/online fashion using a patient-centered third-party DMPS node which hosts patient DMPS accounts, enabling the patient to collect documents inside his or her account history and to delegate initially unaddressed vouchers or prescriptions to a healthcare institution of his or her choice. If the patient does not want to use an own DMPS account, e.g. because the necessary IT education is not present, the referral voucher or prescription can still remain unaddressed using a direct/offline fashion, e.g. instrumenting an eGK smart card. A mutual exclusion exists only between direct/online and unaddressed because every direct/online communication requires a known delivery endpoint.

Introductorily, the direct/online approach is outlined: Each existing *healthcare information system* (HCIS)[8] that participates in the DMPS information exchange posts documents that are to be delegated to another institution to its local DMPS extension using a REST/HTTP endpoint of the DMPS. The DMPS extension manages the process instance and delivery protocol by a statechart implementation and forwards the documents to the DMPS extension of the receiving HCIS (called "downstream" as it is outlined in fig. 3). The delivery instruments an intra-DMPS REST endpoint. An independent local process instance is created and managed by the downstream DMPS extension and the documents are delivered to the local HCIS by a third REST endpoint, which has to be provided by the local HCIS. Once the diagnostic or therapeutic treatment has been accomplished, the

[8] The information systems in primary care are abbreviated either HIS (hospital) or CIS (clinical). To avoid any clash, healthcare information systems in general (of the primary and the secondary care) have been abbreviated as HCIS.

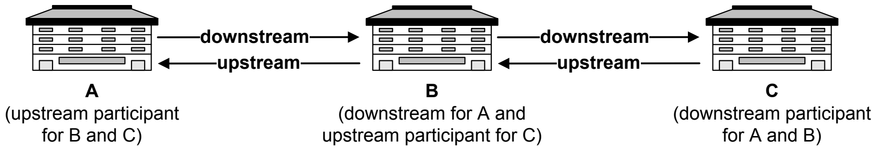

Fig. 3. Downstream and upstream relationships

downstream HCIS reports its result documents to its DMPS extension which will return them to its upstream DMPS correspondent. The upstream DMPS extension delivers the result documents, e.g. CDA-based letters of referral, pictures, or PDF documents, to its local HCIS.

In the whole process, each HCIS is free to delegate diagnostic-therapeutic treatments to one or many downstream institutions. The DMPS subsystem creates a document that contains the global process information, or propagates this process document to the downstream DMPS correspondents. The process status information consists of a shared process ID and the process history with any involved institutions. The process history provides information about the pre-treatment or mutual treatment providers.

The process status information that is sent downstream or upstream by the DMPS and can be configured to be filtered by each institution individually: It is neither necessary to inform downstream institutions about preceding institutions nor to inform the upstream institution about self-conducted delegations. The local DMPS extension manages all available historic information that has been provided by the upstream and downstream institutions together with its own actions, but can reduce or eliminate this information for each own delegation. This is necessary because DMPS also targets generalized scenarios where in complex hospital environments there exists a DMPS endpoint for information exchange with external institutions, but in addition the internal delegation between hospital subsections is just as well supported by intra-institutional DMPS extensions. Therefore an arbitrary control of a HCIS over the globally observable process history is required.

4.2 Leight-Weight Protocol for Inter-Institutional Exchange

The statechart that is deduced form the traditional and essential diagnostic-therapeutic cycle (fig. 1) to facilitate inter-institutional treatment delegation is shown in fig. 4 and named *pandiagnostic-pantherapeutic protocol*[9].

There are two entry points, the first is the leadoff patient entry: the patient visits the first healthcare professional which internally cycles diagnostic-therapeutic phases until a delegation to an external institution is decided. The second entry point into the statechart is an incoming treatment request from an upstream institution that is accepted. The documents are subsequently delegated to the local HCIS while the process waits for the decision either to reply a

[9] Greek “pan-” as a prefix: “of everything”, “involving all members” of a group

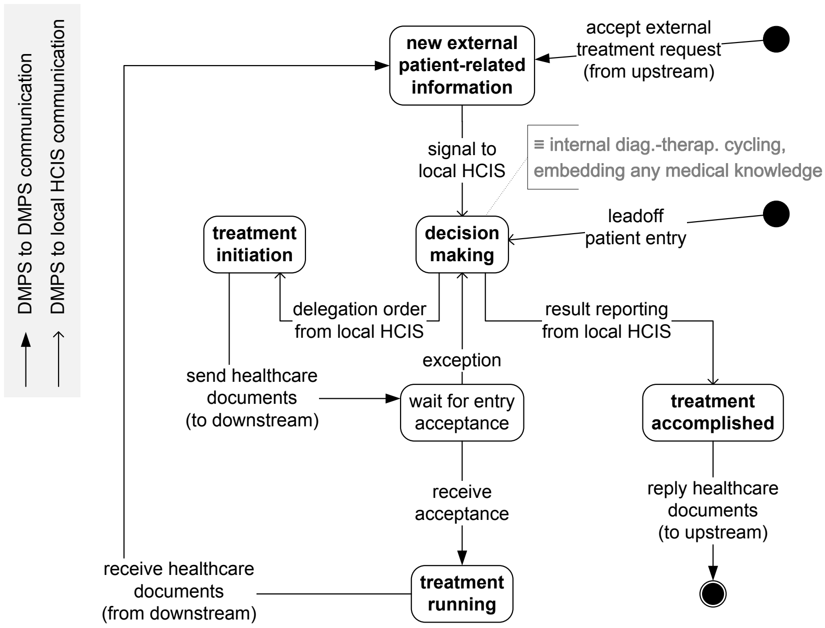

Fig. 4. The pandiagnostic-pantherapeutic protocol

result report from its local HCIS to the upstream DMPS or to initiate a successive delegation to another downstream institution. The statechart comprises the four combinations of from/to and upstream/downstream.

Encapsulating the exchange and process support mechanism into a distinct component allows to extend DMPS by a public-key infrastructure (PKI). In terms of security, cryptographic measures have to be distinguished into the ones for expressing declarations of intent (equal to paper-based signatures) and technical ones for documenting information origin (securing exchange independently of human signatures). Signatures for declarations of intent are object to the HCIS, documents that are delegated to the DMPS extension for delivery are either signed or not. Integrating PKI mechanisms for technically securing the information origin could be plugged into the REST/HTTP endpoints. The DMPS project does not focus on PKI but is designed to fit into existing security frameworks like German eGK [13], PaDok [14] or IHE ATNA [15]. Yet, in order to build authentication and authorization during the establishment of DMPS-to-DMPS relationships a generalized PKI component integration is required, being applicable independently of national PKI specifications.

In DMPS a broad spectrum of document exchange scenarios is supported. First and foremost DMPS is compliant to traditional paper-based working practice, provides process history, and enables process support. It additionally intends to allow for patient-centered document management, fostering cross-sectional life-long patient-centered healthcare documentation.

5 Related Work

Existing protocol standards for information exchange in healthcare focus on hospitals of the secondary care, commonly instrumenting centralized system functionality for tailor-made integration purposes: For example, the *cross-enterprise document sharing* (XDS) [15] standard from IHE allows for distributed document repositories and access delegation. In order to find documents in such a repository a single central document registry is specified, reusing ebXML registry methodology to provide a centralized method of indexing documents. The central registry is a global system node that allows queries and that delegates the access to referenced documents to the original document repositories. Such architecture targets complex hospitals with associated ancillary systems and is even applicable to regional integration efforts, but fails for nationwide application due to its centralized approach. Most of the non-standardized tailor-made regional integration efforts are based on a central database system with distributed transaction systems and communication middleware. Even wide-area RHIN architectures like HYGEIAnet [16] on the island Kreta are tightly-coupled by their complex middleware.

Standards for electronic information exchange between the practice management systems of the primary care and the hospitals and institutions of the secondary care are rare. No universal exchange protocol and format exists for interchange of referral vouchers and discharge letters. In Germany, the governmental project "Elektronische Gesundheitskarte" has not provided a solution for the issue since the project's outset in 2002. Effective platforms like D2D[10] require a central server for document handover.

6 Future Work

In DMPS, the shared information is the global process ID, being created during leadoff delegation and propagated among the involved DMPS nodes. The patient identity has to reside in the transported documents which are not evaluated by DMPS. Providing distributed process support of multiple process instances for an individual patient in the context of a patient identity requires a distributed master patient index. Traditional federated master patient index systems like IHE PIX[11] or OMG[12] PIDS[13] instrument hierarchical federation with central system nodes and are not applicable in distributed environments like the DMPS scenario. Therefore, a loosely-coupled distributed patient identification service for inter-institutional purpose is required.

Whereas DMPS supports unidirectional transport for traditional healthcare supply chains, closely cooperating dynamic teams require mechanisms for team

[10] Doctor to Doctor, `www.d2d.de`, based on PaDok cryptographic infrastructure [14]
[11] Patient Identifier Cross-Referencing
[12] Object Management Group
[13] Patient IDentification Service

publication. With the DEUS mediated publish-subscribe infrastructure [17] we are implementing a distribution system for document-oriented integration purposes. At the moment, neither the process identifier nor the merging of process history is integrated into DEUS but both efforts will converge into a unified platform.

The document-oriented information exchange is a foundation for inter-institutional process support. At the moment, the DMPS extends the transported documents with process history information within its distributed light-weight workflow layer, and provides an exchange protocol that adheres to the diagnostic-therapeutic cycle as coarse-grained intuitional reference from working practice. In DMPS, a workflow status model for cooperative but distributed medical treatment processes is not yet available. Such workflow model must further formalize the activities that take place inside each institution, being represented by the central state "decision making" in the pandiagnostic-pantherapeutic protocol. Yet, extending documents with process information can be considered as a preparation to achieve such workflow support in form of *active documents* [18].

7 Conclusion

For inter-institutional process support, there exists a semantic gap that is not covered by standards, concerning the functional integration of autonomous healthcare information systems. The initial goal of the proposed DMPS architecture is to foster the availability of patient information in order to bridge the gap between primary and secondary care. The prime objectives in design are the document-oriented integration approach and the abdication of any central servers. The essential argument for document-oriented integration over interface-oriented integration lies in its capacity to support deferred system design. Deferred system design supports demand-driven system evolution which is needed for healthcare information systems.

The DMPS architecture achieves a document-oriented process support between strict autonomous institutions following the paper-based work practice as reference model. The document exchange includes propagation of the process history which provides information about the pre-treatment or mutual treatment providers. It is oriented at traditional healthcare communication directly between healthcare institutions without patient involvement. Additionally, it intends to allow for patient-centered document management by optionally mediating document transport through a patient DMPS account. The combination of a REST architecture with a distributed light-weight workflow model provides a minimal set of requirements to be supported by participating systems.

References

1. L. L. Leape, D. W. Bates, D. J. Cullen, J. Cooper, H. J. Demonaco, T. Gallivan, R. Hallisey, J. Ives, N. Laird, and G. Laffel. Systems analysis of adverse drug events. ADE Prevention Study Group. *JAMA*, 274(1):35–43, 1995.

2. M. H. Williams, G. Venters, and D. Marwick. Developing a regional healthcare information network. *Information Technology in Biomedicine, IEEE Transactions on*, 5(2):177–180, 2001.
3. Mario Beyer, Klaus Kuhn, and Richard Lenz. Potential der CDA in verteilten Gesundheitsinformationssystemen. In Richard Lenz, Ulrich Hasenkamp, Wilhelm Hasselbring, and Reichert Manfred, editors, *EAI*, volume 141 of *CEUR Workshop Proceedings*, Marburg, Germany, July 2005. CEUR-WS.org.
4. R. Lenz. Information Systems in Healthcare – state and steps towards sustainability. In Antoine Geissbuhler and Casimir Kulikowski, editors, *IMIA Yearbook of Medical Informatics*, Stuttgart, 2009. Schattauer. Accepted for publication.
5. C. J. McDonald, J. Marc Overhage, P. Dexter, B. Takesue, and J. G. Suico. What is done, what is needed and what is realistic to expect from medical informatics standards. *International Journal of Medical Informatics*, 48(1-3):5–12, 1998.
6. HL7v2. ANSI/HL7 V2.6-2007. `http://www.hl7.org/Library/standards_non1.htm`.
7. J. Powell and I. Buchan. Electronic Health Records Should Support Clinical Research. *Journal of Medical Internet Research*, 7(1), 2005.
8. R. Lenz and K. A. Kuhn. A strategic approach for business-IT alignment in health information systems. *Lecture notes in computer science*, pages 178–195, 2003.
9. N. V. Patel. *Adaptive Evolutionary Information Systems*. Idea Group Inc, 2002.
10. Roy T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD thesis, University of California, 2000.
11. Richard Lenz and Manfred Reichert. IT Support for Healthcare Processes. In Wil M. P. van der Aalst, Boualem Benatallah, Fabio Casati, and Francisco Curbera, editors, *Business Process Management*, volume 3649, pages 354–363, 2005.
12. R. Lenz, M. Beyer, C. Meiler, S. Jablonski, and K. A. Kuhn. Informationsintegration in Gesundheitsversorgungsnetzen. *Informatik-Spektrum*, 28(2):105–119, 2005.
13. Gematik. Einführung der Gesundheitskarte - Gesamtarchitektur - Version 1.5.0, Revision 2.3.4, September 2008.
14. Fraunhofer Institute for Biomedical Engineering (IBMT). PaDok – Patientenbegleitende Dokumentation. `http://www.ibmt.fraunhofer.de/fhg/Images/MT_padoknetzkonzept_de_tcm266-68980.pdf`, 2000.
15. ACC, HIMSS, and RSNA. IHE IT Infrastructure Technical Framework, vol. 1 (ITI TF-1): Integration Profiles, Rev 4.0, August 2007.
16. M. Tsiknakis, D. G. Katehakis, and S. C. Orphanoudakis. An open, component-based information infrastructure for integrated health information networks. *International Journal of Medical Informatics*, 68(1-3):3–26, 2002.
17. C. P. Neumann, F. Rampp, M. Daum, and R. Lenz. A Mediated Publish-Subscribe System for Inter-Institutional Process Support in Healthcare. In *Proc of the 3rd ACM International Conference on Distributed Event-Based Systems*, July 2009.
18. A. LaMarca, W. K. Edwards, P. Dourish, J. Lamping, I. Smith, and J. Thornton. Taking the work out of workflow: mechanisms for document-centered collaboration. In *Proc of the 6th conference on European Conference on Computer Supported Cooperative Work*, pages 1–20. Kluwer Academic Publishers Norwell, USA, 1999.