alpha-Flow: A Document-based Approach to Inter-Institutional Process Support in Healthcare

Christoph P. Neumann and Richard Lenz

Friedrich-Alexander University,
Erlangen-Nuremberg, Germany,
{christoph.neumann|richard.lenz}@cs.fau.de,
http://www6.cs.fau.de/

Abstract. Inter-institutional collaboration requires clean task boundaries and the separation of responsibilities. In addition, healthcare processes are intrinsically fluid. Traditional activity-oriented workflow models or content-oriented workflow models do not provide adequate support for the paper-based working practice in healthcare. The *α-Flow* approach adopts electronic documents as the primary means of information exchange, fusing both paradigms into a combined workflow schema model, wherein workflow schemas are represented as documents which are shared coequally to content documents.

Topics: Process modeling in healthcare, workflow management in healthcare, context-aware healthcare processes, inter-institutional healthcare information systems, document-oriented integration

1 Motivation and Challenges

The patient treatment process increasingly changes from isolated treatment episodes towards a continuous process incorporating multiple organizationally independent institutions and different professions. Effective treatment of unclear symptoms or multimorbid patients increasingly leads to the need for establishing and managing dynamic teams of cooperating specialists. Independent electronic health records are discussed as a basis for inter-organizational cooperation, but despite of existing standards like openEHR, reality is still far away from this vision, and IT-support for inter-organizational patient treatment processes is an open issue. Today, IT support for healthcare processes is typically limited to intra-institutional approaches, and systems in different organizations are heterogeneous and rarely integrated.

Semantic scalability is an important requirement for a distributed IT application in healthcare. Therefore, we are looking for an evolutionary and decentralized approach to support inter-institutional processes in healthcare. The traditional approach to manage inter-institutional processes is based on documents with a dedicated semantics, such as a referral or a discharge letter. We pick

The original publication is available at www.springerlink.com (DOI: 10.1007/978-3-642-12186-9_55)

up this interaction paradigm and try to extend it to support more complex cooperation scenarios. The basic idea is to use electronic documents as self-contained units of information interchange which also carry process related information. As an illustrative application example one might consider disease management programs which are managed by paper-based documents that carry checklists.

2 Objectives

The goal of the *α-Flow* project is to develop a concept for an document-based workflow with loosely coupled heterogeneous systems at the participating sites. This particularly incorporates a meta-model for document-based process management, which provides the fundamental artifacts for process specification.

The *α-Flow* approach fuses the activity-oriented workflow paradigm with the content-oriented workflow paradigm into a combined model, wherein workflow schemas are represented as documents which are shared coequally to content documents. The intent is to allow access, viewing, and editing of the original content documents in standard ways like general editors without corrupting the workflow semantics. This paper will present the necessary artifacts and the resulting requirements to the infrastructure that is needed for an *α-Flow* implementation.

3 Background

The two basic aspects for collaborative activities are the support for content manipulation and the support for coordination. An information system traditionally focuses on the *manipulation of content*. Support is given for the gathering, the storage, access control, structuring, classification, and presentation of the information as well as the reaction to new information. Collaboration extends this focus with the concern of *coordination*. The information system must support the "articulation work" [1] as it is part of Computer Supported Cooperative Work (CSCW) [2]. Articulation support must enable cooperating actors to partition work into units, to divide it amongst themselves, and to schedule, mesh and interrelate their collective activities.

In *activity-oriented workflows* the central point is a task. Process definitions describe tasks with states and transitions (like Petri Nets) or with actors and activities (like BPMN[1]). At any given moment at run-time, the workflow is in a well-defined state and it moves to a different state when certain conditions are met. In Petri Nets the workflow engine enacts a set of actions during transitions, in BPMN the actions are internal part of an activity. Although each task is characterized by preconditions, postconditions, and possible exceptions, any required or generated artifacts, documents for example, are not necessarily considered by the workflow schema.

[1] Business Process Modeling Notation

Content-oriented workflow systems, in contrast, place a content artifact in the center of the workflow process, focusing on its creation and manipulation phases. Each workflow step alters the content object. At the end of each step, the state of the document reflects the step's result. The workflow definitions for content-oriented workflows have their origin in write-and-review processes in publishing companies.

Both approaches, the activity-oriented workflows and the content-oriented, are commonly based on predefined workflow schemas that can be instantiated by an enactment engine. For example, a content-oriented workflow schema will prospectively specify actors like "author", "reviewer", and "publisher" as well as any steps and content states like "private", "submitted", "reviewed", and "published". Ad-hoc workflows with an initially unknown set of actors and state/transitions are not considered, traditionally.

4 The *α-Flow* Approach

The *α-Flow* approach focuses on the relationship between content and coordination aspects of collaborative and large-scale environments. The collaboration is considered as a feature of the artifact and not of the application system. The *α-Flow* model adopts electronic documents, called *α-Docs*, as the primary means of information exchange and coordination.

4.1 Content vs. Coordination

In order to support heterogeneous systems, we need to decouple collaboration functionality from the application. Therefore we have to distinguish between content documents and coordination documents.

The *content documents* conduct medical information and are of arbitrary type, like Adobe PDF files, Microsoft Word documents, or HL7[2] CDA[3] documents. Content documents belong to the healthcare applications which generate them. The healthcare applications provide them to the distributed workflow engine by an export or by an universal resource identifier.

The *coordination documents* are independent of the application system and belong to the distributed workflow. They conduct information about actors, roles, and institutions, as well as system topology information, workflow reports or summaries, and control structures.

4.2 Artifacts: *α-Docs* and *α-Cards*

For applying a document-based workflow, the overall treatment process must be represented by a document. This document is a *treatment status artifact* and its state represents the overall state of the treatment process. Several documents,

[2] Health Level 7, `http://www.hl7.org`
[3] Clinical Document Architecture

both further coordination documents and arbitrary content documents, will accompany the treatment process. Later, we will elaborate on the relationship between the treatment status artifact, the supporting coordination artifacts, and the involved content artifacts. In this section, we will only focus on the structure of the α-*Doc* documents.

Motivating two granularity levels of document artifacts: α-*Docs* are the units of information interchange, but the unit of validation must be smaller. Parts of a singular α-*Doc* should be validated while others remain in a preliminary state and are filled as the workflow continues. For example, in a process to generate a report the report is represented by an α-*Doc*. The report consists of a form. The form will be filled by different organizations. It is necessary to structure the form schema and its later values into separate units such that it is possible to assign them to different organizations.

Inter-institutional healthcare example: The initial episode of breast cancer treatment is outlined in fig. 1. The goal of this treatment episode is to find out whether or not a knot in a breast is actually malignant cancer.

The treatment begins with a patient visiting her gynecologist, who writes the anamnesis documentation. After the anamnesis, the gynecologist conducts a sonography with an according report as its result. In fig. 1, the participant's superscript A stands for ambulant (office-based, primary care) in contrast to C for clinical (secondary care).

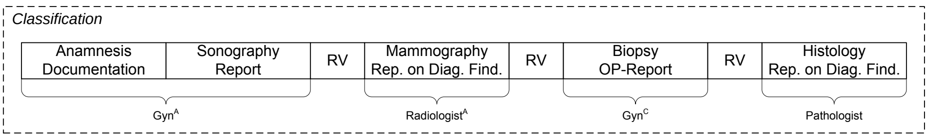

Fig. 1. The initial breast cancer treatment episode: classification

If the result is either malignant or dubious he/she will send the patient to a radiologist for mammography, using a referral voucher RV. After the radiologist's treatment,the mammography report on diagnostic findings is sent back to Gyn^A. The gynecologist evaluates the mammography report, primarily the medical indicator BI-RADS[4], and decides whether the patient has to be send to a hospital for a biopsy. The biopsy involves a clinic's Gyn^S, accordingly a referral voucher is created. The tissue is taken by Gyn^S and sent to a pathologist for histological diagnosis. The histology report is sent back from the pathologist to the clinical Gyn^S, who bundles the report with a short report about the biopsy operation and finally delivers the reports back to Gyn^A. The histology provides definite evidence, yet, the Gyn^A is the one who takes the histology result and is responsible for informing the patient. In the malignant case, another episode be-

[4] Breast Imaging – Reporting and Data System

gins now by sending the patient to a breast cancer treatment center for primary therapy.

For α-*Flow*, the traditional paper-based reports in fig. 1 can be considered as one report, that is successively filled by the participating institutions. Each such distributed treatment episode can be characterized by a common goal of the collaborating participants. Considering both the institutional and collaborative view on a treatment episode two kinds of granularity can be distinguished.

α-*Docs* versus α-*Cards*: The units of validation and organizational accountability are named α-*Cards*. One α-*Doc* consists of one or more α-*Cards*. On the one hand an α-*Card* is quite larger than single database fields. On the other hand the intended granularity of an α-*Card* is a more fine-grained one than the one that is experienced from paper-based working practice in healthcare. A single α-*Card* contains, for example, a diagnostic finding, clinical evidence, a diagnosis, a therapeutic measure, an order, or a prescription. This improves the structure of the patient files, and provides higher selectivity in retrieval and display. In fact, an α-*Card* is required to fulfill the fundamental feature of a document, to be self-contained on its own. To ensure this property of an α-*Card* rests with the institutional healthcare information systems.

Difference and coherence between α-*Docs* and α-*Cards*: In the document-based α-*Flow* approach for inter-institutional process support, each individual α-*Doc* is a collection of α-*Cards*. The α-*Docs* are the atomic units of information interchange. An α-*Card* is the atomic unit of validation, clearance, shared visibility, and cryptographic signatures.

α-*Cards* in the context of content vs. coordination: α-*Cards* are not only used for structuring content documents but also to consolidate coordination information. For example an α-*Card* might collect information about process participants, their institutional information, or treatment role models. Others provide information about system topologies or access control lists. The treatment status information is also consolidated in an α-*Card* conducting the overall workflow schema.

4.3 Fusing Activities into α-*Flow*

A distributed process as a structure of distributed activities is called an *episode*, or an α-*Episode* to accentuate the α-*Flow* context. One α-*Episode* is characterized by a particular goal and represented by one α-*Doc*. By using the term of an α-*Episode* in contrast to α-*Doc* we point out that the α-*Flow* approach is not blind to the necessary activities.

Yet, α-*Flow* tries to eliminate any modeling of activities in its coordination model. Activities are fused into the α-*Flow* approach by completely representing them by their results, the α-*Cards* as part of the α-*Doc*. This is necessary, because any decisions for process routing requires either a domain- and section-specific decision support system (e.g. based on rule-based artificial intelligence) or a human decision. No conditional model element like in activity-orientation is sufficient for most decision modeling and process routing in healthcare. Furthermore, most activities are human tasks or require a complex local health-care information system. Due to the heterogeneity and complexity of the existing systems, they are essentially factored out of the coordination layer but remain integrated by document-orientation.

The basic *α-Flow* assumption for inter-institutional workflows is that human or computer supported decisions can always be represented in a newly occurred demand for further information, e.g. patient-related information as it is well-known by the diagnostic-therapeutic cycle [3] in healthcare. An episode ends when no further information is required for the particular goal. Any decision that is made in the course of an *α-Episode* can be represented by the creation of a record keeping *α-Card*.

4.4 Workflow Progress by the Notion of Active Documents

An *active document* [4] is a document that allows a direct interaction with itself. Therefore, documents become active documents if they are assigned with *active properties*. In case of the *α-Flow* approach only the *α-Cards* are assigned with active properties. An *α-Doc* artifact becomes active if it contains at least one *α-Card* that has an active property. Not every *α-Card* necessarily has active properties but each can be assigned one.

The *"alpha"* in our artifact names relates to the active properties: By their active properties, *α-Cards* are triggers for workflow progress. Activities are not modeled explicitly but, instead, an newly demanded *α-Card* placeholder is created in the coordination list of the treatment status artifact. Workflow progress means successive fulfillment of requested content *α-Cards*. Workflow schema change means editing the list of required *α-Cards* and adopting the progress actions that occur at state change.

The active properties do not implement activities. They support the state change and exchange of the *α-Doc/α-Cards*. By propagating the *α-Doc* state change, the requests for *α-Card* fulfillment are delegated to the cooperating participants, indirectly triggering their institutional activities.

Active properties in regard to process, function, and data: In conclusion, the active properties' logic drive the progress of the *process*, but the *functions* that are equal to the notion of activities are subordinate to their results in form of *data*. The data is required in form of *documents* because *α-Flow* targets large-scale scenarios. Such require a document-oriented integration approach as we have detailed in [5]: Interface- or service-orientated integration approaches suffer from various shortcomings in order to provide a large-scale electronic health record infrastructure. In contrast, document-orientation is suitable to support the *deferred system design* principle [6] enabling evolutionary systems [7]. In *α-Flow*, a process definition basically consists of a set of *α-Card* documents and their control flow being expressed in active properties.

5 The *α-Flow* Meta-Model

The *α-Flow* approach is a dual workflow paradigm that aims at a unification of content-oriented workflows with activity-oriented workflows. In activity-orientation the activities' artifacts, either required or produced, are resistant to change and are inferior workflow elements. In content-orientation the information document is changed through collaboration with the focus on role models and notification mechanisms. The actual activities' tasks are commonly hard-coded, initially unknown sets of actors and state/transitions are not considered.

The *α-Flow* approach fuses both paradigms into a combined model. Workflow schemas are represented as *intelligent to-do lists*, in which each list item is an *α-Card* document that is either available or that is requested in order to progress the *α-Episode* process. The intelligent to-do list is represented by an *α-Card* generically named *treatment status artifact*. The intelligence is provided by its active properties. As outlined in sect. 4.2 the treatment status artifact is one of several coordination *α-Card* documents that are shared coequally to the involved content documents.

The state of an individual *α-Card* and an *α-Doc* has to be distinguished from the state of the treatment. The treatment process and its state will progress with the creation or the change of *α-Cards*, but each *α-Card* has its own properties independent of a treatment. Furthermore, it should be possible to use the original documents, which are basically the payload of an *α-Card*, in standard ways like general editors without corrupting the workflow semantics.

The state of an *α-Card* is based on what we call *adornment models*. Before we describe the adornment models in more detail, it is necessary to introduce what we call *collaboration resource models*.

5.1 The Collaboration Resource Models

Collaboration resources are illustrated as the who, where, with whom, and with what. The collaboration resource models contain the information about actors, roles, institutions, and systems. They are part of the coordination system and form a cross-cutting infrastructure being used by several of the *α-Card* adornment models.

5.2 The *α-Card* Adornment Models

The validity and visibility of *α-Cards* have to be considered separately. In traditional database-centric approaches, visibility is strictly coupled to validity. Information is only visible if it is committed, and the commit has to ensure the integrity constraints. In contrast, for document-centric approaches it is common to share documents preliminarily, by making them visible, although guarantees of validity are not provided just yet for the content.

The document-centric approach supports the separation between validity and visibility. The validity has to be distinguished into the *intent validity*, for expressing declarations of intent by humans being related to paper-based signatures, and *technical validity*, which is essentially defined by specifying how versions and variants are consolidated. Yet, providing electronic signatures for declarations of intent is subject to the local healthcare information systems. The α-*Flow* approach does not focus on PKI but is designed to fit into existing security frameworks like German eGK [8], PaDok [9] or IHE ATNA [10]. Contemporary workflow approaches, in regard to their artifact model, do not distinguish between the four aspects of intent validity versus visibility and versioning versus variants.

The introduction of α-*Cards* as an explicit unit of validation has been motivated by the need for flexibly dealing with intent validity and visibility. The *intent validity model* might simply consist of the classifiers "invalid" and "valid", whereas the *visibility model* might simply consist of the classifiers "private" and "public". Private α-*Cards* are for non-collaborative purposes or to prepare and configure collaborative purpose. An invalid public α-*Card* is interim information. A valid public α-*Card* is not allowed to change without versioning. Validity does not imply visibility: Valid private α-*Cards* are allowed, e.g. they are required if access control has yet to be configured for an α-*Card* before it is advertised by setting its visibility to public.

A *versioning model* is required both for content and coordination α-*Cards*. Versioning is mandatory for public and valid α-*Cards* because the individual systems require a global version for the tracking of changes. Any other α-*Cards* (in terms of visibility and validity) are equally allowed to use versioning as it seems appropriate. There is always exactly one current version of any α-*Card*.

A *variants model* is additionally required. In contrast to versions, there may exist several valid variants of an α-*Card* coequally at any given time. This concerns content α-*Cards*, especially but not solely public invalid α-*Cards*. It might as well be required for coordination α-*Cards*, for example if an ad-hoc medical consensus on further workflow activities is negotiated by different institutions by exchanging extended or modified variants of the treatment status artifact. Both the versioning model and the variants model contribute to the α-*Card* identity.

The *authentication model* and *authorization model* are augmenting the visibility with access control. The visibility is a general property of the α-*Card*. If access privileges have to be differentiated according to actors, roles, institutions, or arbitrary groups, a kind of access control mechanism is required. While the visibility state transition from private to public triggers a notification or a content advertising, the access control mechanism filters the audience of a notification or authorizes content access.

The *syntactic payload type model* describes the format of an α-*Card*. For example, PDF, Microsoft Word, or HL7 CDA for content artifacts. It would even be possible to exchange jPDL or BPEL files as coordination artifacts that are documenting intra-institutional processes. The MIME[5] types provide a common standard for the syntactic types.

With the *semantic payload type model* an α-*Card* is classified semantically. We distinguish the fundamental semantic type from the domain-specific semantic type, and eventually the user-specific semantic type. The *fundamental seman-*

[5] Multipurpose Internet Mail Extensions

tic type classifies artifacts into “content” vs. “coordination”. The *domain-specific semantic type*, for example, classifies content artifacts as “diagnostic finding”, “therapeutic measure”, or “prescription”. For coordination artifacts, there exist predefined semantic types like the “treatment status artifact”, the α-*Card* carrying the workflow schema, or the “treatment team artifact”, the α-*Card* carrying the information about participating collaboration resources.

The *subject model* describes the authors of an α-*Card*. For example a doctor for content α-*Cards* or a workflow-engineer for coordination α-*Cards*. In contrast, the *object model* describes the object of an α-*Card*. For example the patient for most content α-*Cards*, or a treatment process name for coordination α-*Cards*. Both the subject model and the object model contribute to the α-*Card* identity.

A *linkage model* is required to associate α-*Cards* arbitrarily among each other. The linkage model is also the basis for the navigation between the α-*Cards*. The linkage model could be based on XLink and XPointer technology, but is an adornment, allowing to associate even non-XML artifacts with other artifacts.

5.3 The α-*Doc* Adornment Models

An α-*Doc* is a named collection of α-*Cards*. Any progress of a treatment process will essentially change such a collection by the creation or change of α-*Cards* from distinct process participants and institutions. The *collection model* for α-*Docs* must provide an overview over all α-*Cards*. It enables the process participants to gain shared knowledge about each other’s activities.

A *transfer model* for the α-*Docs* is required. It references the transfer capabilities of particular institutions. Model elements are service endpoint declarations of the participating sites and applied communication protocols. The transfer model is based on the infrastructure that is provided by the collaboration resource models.

5.4 Active Property Model

The *active property model* has to provide mechanisms to assign active code to an α-*Card*. The active code of an active property is called a *progress action*. The active property model encompasses several sub-models: the ordering model, the activation condition model, and the evaluation phase model. Support for an *ordering model* of multiple active properties is required because there will exist several active properties for a single α-*Card*. The *activation condition model* must allow to describe conditions under which the active property is triggered, supporting both event-triggers and periodical triggers.

The *evaluation phase model* describes a three-phased evaluation cycle of active properties: The *verification phase* ensures the applicability according to any boundary conditions that are provided by the access control conditions of the adornment models or any conditions of the workflow model. The *operation phase* carries out the active code. The *finalization phase* carries out notifications and handles error or abort situations.

The modeling of a workflow schema can fulfill two very different intentions in loosely-coupled inter-institutional scenarios: a retrospective modeling for documentation and further delegation or a prospective modeling for enactment automation. For *retrospective modeling*, the focus lies on providing end-users with workflow schema editors to allow them to keep record of their latest process step and to allow them a process delegation to another institution. For *prospective modeling*, the focus lies on providing ad-hoc mechanisms for consensus finding if two participants have divergent notions of the treatment process articulated as variant treatment status artifacts. Both use-cases for workflow schema modeling have to be considered in the overall α-*Flow* approach.

6 Related Work

This section is separated into approaches related to active documents and approaches related to workflow models.

6.1 Active Document Technology

The X-Folders project [11] instruments WebDAV folders which can react to the insertion or modification of a document by starting a task. Because WebDAV folders can be distributed, multiple X-Folders can be combined to a site-spanning workflow. Yet, the guards and triggers are hard coded for each X-Folder, e.g. one folder to accept new forms, another folder with pending forms, and a last folder with accepted forms. The X-Folders project neither provides a formal representation of the workflow schema nor a dynamic adaptation; and it does not provide a distributed institution or role management.

The Placeless documents project [12] from Xerox PARC provides an infrastructure to implement active properties for arbitrary documents. It seemingly provides abstraction from existing document- and file-management interfaces.

With the DEUS mediated publish-subscribe infrastructure [5] we are implementing a distribution system for α-*Cards*. For implementing the α-*Flow* approach, we need a technical foundation that allows to assign active properties to an artifact like an α-*Card*. In the future, we will evaluate both the X-Folders and the Placeless documents system as partner to the DEUS platform to combine transfer features with active property features in α-*Flow*.

6.2 Related Workflow Approaches

There exist several related workflow approaches. The case-handling paradigm focuses on workflows like clinical pathways and requires semantic integration of medical data in form of data objects and forms. In [13], the authors acknowledge that the case handling creates an integration problem, because the state of a "case" is derived from "data objects" with a well-known schema which cannot be separated from the process. In addition, data objects are still product of a modeled activity, whereas *α-Flow* tries to separate an explicit model of intra-institutional activity from the inter-institutional model.

The object-aware workflow systems [14], focusing on write-and-review processes like job applications, the artifact-centric approach [15], and the data-driven process structures [16] all represent advanced solutions to the content-oriented approach. The *object*, respectively the *business artifact* or *data*, needs a structured and predefined content schema. All approaches allow to model life-cycle state-charts for the records. The coordination is provided by state-changes in the life-cycle model, as explained for content-oriented workflows which has been part of the inspiration for *α-Flow*. Yet, a fine-grained project-specific life-cycle and information model is required and cannot be changed ad-hoc. A comprehensive comparison of these types of workflow modeling is contained in [14].

From these approaches stem models to describe consistency between the process model and the object life-cycle. Both the artifact life-cycle language [17] and the work of Ryndina et al [18] seem promising and could eventually be adopted as formalism to *α-Flow*.

The primary boundary condition in the inter-institutional scenario is the integration of autonomous systems by loose coupling and respecting the manifold document standards which comprise arbitrary taxonomies and ontological standards for healthcare. The existing workflow approaches fail these conditions but provide sophisticated solutions for a homogeneous system environment with canonical information models. The core motivation of *α-Flow* is to support decentralized, large-scale scenarios in which semantically heterogeneous and even informal content types drive the distributed, collaborative workflow. Such requires the utter decoupling of content from coordination.

7 Conclusion

The *α-Flow* approach adopts electronic documents as the primary means of information exchange. The collaboration is considered as a feature of the artifact and not of the application system. Our artifacts themselves take on the active role in managing coordination. In order to support heterogeneous systems, we need to decouple collaboration functionality from the application.

This paper provides a systematic classification of the required elements for a document-based approach for inter-institutional process support in healthcare. Healthcare processes are intrinsically fluid and require clean task boundaries, separation of responsibilities, and multiple versions or variants of a document as well as initially unknown sets of actors and state/transitions. Neither activity-oriented workflow models nor traditional content-oriented workflow models provide adequate support for the paper-based working practice in healthcare. Only the fusion of both paradigms will enable a seamless enhancement of existing healthcare information systems with inter-institutional collaboration facilities.

References

1. K. Schmidt and L. Bannon. Taking CSCW seriously: Supporting Articulation Work. *Computer Supported Cooperative Work (CSCW)*, 1(1):7–40, 1992.
2. Irene Greif. *Computer-supported cooperative work: A book of readings*. Morgan Kaufmann, 1988.
3. J. H. van Bemmel, M. A. Musen, and J. C. Helder. *Handbook of medical informatics*. Bohn Stafleu Van Loghum, 1997.
4. 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 European Conference on Computer Supported Cooperative Work*, pages 1–20. Kluwer Academic Publishers Norwell, USA, 1999.
5. 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 Int'l Conf on Distributed Event-Based Systems*, July 2009.
6. N. V. Patel. *Adaptive Evolutionary Information Systems*. Idea Group Inc, 2002.
7. R. Lenz. Information Systems in Healthcare – State and Steps towards Sustainability. *IMIA Yearbook 2009 – Yearbook of Medical Informatics as a supplement of Methods of Information in Medicine*, pages 63–70, 2009.
8. Gematik. Einführung der Gesundheitskarte - Gesamtarchitektur - Version 1.5.0, Revision 2.3.4, September 2008.
9. Fraunhofer Institute for Biomedical Engineering (IBMT). PaDok – Patientenbegleitende Dokumentation. `http://www.ibmt.fraunhofer.de/fhg/Images/MT_padoknetzkonzept_de_tcm266-68980.pdf`, 2000.
10. ACC, HIMSS, and RSNA. IHE IT Infrastructure Technical Framework, vol. 1 (ITI TF-1): Integration Profiles, Rev 4.0, August 2007.
11. D. Rossi. Orchestrating document-based workflows with X-Folders. In *Proc of the ACM symposium on Applied computing*, pages 503–507. ACM, NY, USA, 2004.
12. P. Dourish, W. K. Edwards, J. Howell, A. LaMarca, J. Lamping, K. Petersen, M. Salisbury, D. Terry, and J. Thornton. A programming model for active documents. In *Proc of the 13th annual ACM symposium on User interface software and technology*, pages 41–50. ACM New York, NY, USA, 2000.
13. W. M. P. van der Aalst, M. Weske, and D. Grünbauer. Case handling: a new paradigm for business process support. *Data & Knowledge Engineering*, 53(2):129–162, 2005.
14. V. Künzle and M. Reichert. Towards Object-Aware Process Management Systems: Issues, Challenges, Benefits. *Lecture Notes in Business Information Processing*, 29.
15. David Cohn and Richard Hull. Business Artifacts: A Data-centric Approach to Modeling Business Operations and Processes. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*, September 2009.
16. D. Müller, M. Reichert, and J. Herbst. Data-driven modeling and coordination of large process structures. *Lecture Notes in Computer Science*, 4803:131, 2007.
17. C. E. Gerede and J. Su. Specification and verification of artifact behaviors in business process models. *Lecture Notes in Computer Science*, 4749:181, 2007.
18. K. Ryndina, J.M. Kuster, and H. Gall. Consistency of business process models and object life cycles. *Lecture Notes in Computer Science*, 4364:80, 2007.