

Adaptive Version Clocks and the OffSync Protocol

Christoph P. Neumann, Andreas M. Wahl, Richard Lenz
Computer Science 6 (Data Management)
Friedrich-Alexander University, Erlangen, Germany
christoph.neumann@cs.fau.de

Abstract—The α -OffSync project offers a synchronization concept for α -Flow, an electronic process support in heterogeneous inter-institutional scenarios in healthcare. A distributed case file is provided by α -Flow to represent workflow schemas as documents which are shared coequally to content documents. α -OffSync allows the detection and resolution of global concurrent modification conflicts on process artifacts based on locally available information by using logical timestamps. Further mechanisms are included to ease the management and dynamic extension of the group of participating actors.

Keywords—Logic Clocks, Synchronization, Optimistic Concurrency Control

I. INTRODUCTION & BACKGROUND

The research project α -Flow, as described in [1], [2], [3], provides an approach to electronic process support in healthcare following the case handling paradigm [4] while focusing on distributed environments and document-oriented systems integration. The traditional paper-based interaction paradigm in medicine is picked up and further extended. Both the often initially unknown amount of participating actors and the dynamic order of treatment steps in healthcare processes are considered to offer benefit to the cooperating physicians.

Relevant medical documents are gathered in a distributed electronic case file called α -Doc. Each α -Doc represents one α -Episode in which the actors collaborate in the treatment of a specific patient. An α -Doc consists of multiple α -Cards which are the units of validation, organizational accountability and subject to synchronization. Changes in the treatment process are reflected by the creation or modification of α -Cards. Each α -Card is decomposed into a descriptor with process-relevant status attributes and a payload carrying information about medical results.

Unlike passive paper-based case files, an α -Doc offers a variety of additional active properties [5]. The α -Doc, as an active document [6], embeds all facilities for direct interaction with itself and the embedded pieces of information. It does not depend on another application, but is a light-weight application itself. Among others, these active properties include automatic synchronization with the remaining actors and administration of the

group of collaborating participants. Actors only need a file copy of the α -Doc in order to participate. The α -Doc replicas represent nodes in an overlay network for synchronization.

α -Docs are files on the participant's desktop and active only when opened. In terms of networking, all nodes have an offline characteristic and usually no two peers are online at the same time. The objective is to provide generic concepts for synchronization that is offline-capable such that locally conducted synchronization operations provide global consistency across all physically distributed but logically centralized replicas. Technical integration of these concepts into the existing system architecture of α -Flow is also covered.

II. METHODS & OBJECTIVES

The communication channel is supposed to use store-and-forward to deliver messages once recipients are reachable. Non-FIFO behavior of the channel is tolerated as messages may be delayed in transit or arrive out-of-order.

Due to these preconditions, common synchronization mechanisms, based on mutual exclusion or on other techniques that require reaching a global consensus between participants, are not sufficient for the α -Flow system. A suitable protocol must detect global conflicts, but for reconciliation a local decision must suffice: Further communication is not possible because all other nodes are offline. A decision about a valid version must be derived instantly otherwise the local human actor is blocked. Local reconciliation must ensure global consistency among all nodes: it provides a globally unified partial order of the versions.

The synchronization adopts lists of logical timestamps, inspired by vector clocks [7] and version vectors [8]. Further synchronization approaches such as Independent Updates or Timestamped Anti-Entropy had been evaluated; a survey is best gained by [9].

III. RESULT

An *adaptive version clock* (AVC) has been implemented as a new data structure for dynamic lists of logic clocks and an offline-capable *OffSync protocol* for synchronization of dynamic nodes in a totally partitioned

network (cf. [10]). Each AVC holds timestamps in an associative array, which contains one key for each locally known node in the overlay network. Associated counters indicate how often which actors have modified an α -Card from which network nodes. Each modification increments the counter of the change originator. Applying logic clocks is necessary to establish an order on change events of an artifact. This order is based on element-wise comparing the AVCs received from other nodes with the locally known AVCs.

The OffSync protocol requires a version control system (VCS) that supports artifact-specific logical timelines together with history manipulation capabilities. If the comparison indicates that an incoming version properly succeeds the latest local one, it is classified as a valid version. Gaps in an artifact timeline can be computed from AVCs and indicate the existence of further versions being delayed in transit. In both cases the incoming version is persisted at the appropriate position in the VCS history.

Concurrent modifications are global conflicts. Thus, the OffSync protocol queries the local version repository and computes the last globally non-conflicting version (LGNCV). All versions between the LGNCV and the conflict-causing one are invalidated; this means they are moved to a conflict branch within the VCS for data provenance purposes. The LGNCV is selected as the reconciled version to be reinserted into the timeline; its AVC is then set to the element-wise maximum of the locally known preceding timestamps. Global consistency is achieved by discrete local reconciliation without further communication.

To enable dynamic change of actors, nodes and AVCs (now in terms of cardinality), a join protocol forms a second part of the OffSync protocol. New copies of an α -Doc introduce themselves to their peers and automatically synchronize by mutually exchanging information about locally known α -Cards based on their AVCs. New nodes are incrementally updated to the latest state. Each participant gains knowledge about newly joined actors.

The outlined concepts have been integrated into the existing prototypical system architecture of the α -Flow project. SMTP and IMAP are preliminarily used for store-and-forward communication. In the future, they are to be substituted by other protocols being more reliable in terms of guaranteed delivery, such as XMPP. Hydra [11] is used as VCS implementation.

IV. DISCUSSION

Unique to the introduced synchronization approach is the ability to establish a shared view on the process state among all participating actors in totally partitioned networks where no guaranteed assumptions about the reachability of any network nodes can be made.

Global conflicts can be detected and reconciled without additional online message exchange for determining a globally valid version. α -Flow is enabled to facilitate the management of dynamic groups of participants by minimizing the effort for inviting new actors and joining an ongoing treatment episode to provide the necessary flexibility for heterogeneous processes in healthcare.

REFERENCES

- [1] C. P. Neumann and R. Lenz, "The alpha-Flow Use-Case of Breast Cancer Treatment – Modeling Inter-Institutional Healthcare Workflows by Active Documents," in *Proc of the 8th Int'l Workshop on Agent-based Computing for Enterprise Collaboration (ACEC)*, Jun. 2010.
- [2] C. P. Neumann, P. K. Schwab, A. M. Wahl, and R. Lenz, "alpha-Adaptive: Evolutionary Workflow Metadata in Distributed Document-Oriented Process Management," in *Proc of the 4th Int'l Workshop on Process-oriented Information Systems in Healthcare*, Aug. 2011.
- [3] C. P. Neumann and R. Lenz, "The alpha-Flow Approach to Inter-Institutional Process Support in Healthcare," *International Journal of Knowledge-Based Organizations (IJKBO)*, vol. 2, no. 3, 2012, accepted for publication.
- [4] W. M. P. van der Aalst, M. Weske, and D. Grünbauer, "Case handling: a new paradigm for business process support," *Data & Knowledge Engineering*, vol. 53, no. 2, pp. 129–162, 2005.
- [5] A. Todorova and C. P. Neumann, "alpha-Props: A Rule-Based Approach to 'Active Properties' for Document-Oriented Process Support in Inter-Institutional Environments," in *Lecture Notes in Informatics (LNI) Seminars 10 / Informatiktage 2011*, L. Porada, Ed. Gesellschaft für Informatik e.V. (GI), Mar. 2011.
- [6] P. Dourish, W. Edwards, A. LaMarca, J. Lamping, K. Petersen, M. Salisbury, D. Terry, and J. Thornton, "Extending document management systems with user-specific active properties," *ACM Transactions on Information Systems (TOIS)*, vol. 18, no. 2, pp. 140–170, 2000.
- [7] C. J. Fidge, "Timestamps in message-passing systems that preserve the partial ordering," *Australian Computer Science Communications*, vol. 10, no. 1, pp. 56–66, Feb. 1988.
- [8] D. S. Parker Jr, G. J. Popek, G. Rudisin, A. Stoughton, B. J. Walker, E. Walton, J. M. Chow, D. Edwards, S. Kiser, and C. Kline, "Detection of mutual inconsistency in distributed systems," *IEEE TSE*, vol. 9, no. 3, pp. 240–247, 1983.
- [9] R. Lenz, *Adaptive Datenreplikation in verteilten Systemen*. Stuttgart: Teubner, 1997, PhD dissertation.
- [10] A. M. Wahl, "alpha-OffSync: Verteilte Datensynchronisation in Form von IMAP-basiertem Mailtransfer," Bachelor Thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, 2011.
- [11] S. A. Hady, "alpha-VVS: An integrated Version Control System as a Component of Process Support based on Active Documents," Master's thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, 2011.