HematoWork: A Knowledge-based Workflow System for Cancer Therapy

R. Müller[1], B. Heller[2], M. Löffler[2], A. Winter[2], R. Speer[2], H. Berger[2], M. Klöss[2], L. Mantovani[3], F. Brümmer[1], E. Jödecke[1], U. Neubert[1], F. Fiebig[1], E. Rahm[1]

1 Objectives

The domain of hemato-oncology is characterized by a complex patient management and the involvement of several institutions (e.g. oncological ward, private practitioner, external review panels) in the context of protocol-directed trials. Current research efforts in this domain (e.g. [1-3]) focus on specialized subtasks such as chemotherapy dose calculation and toxicity monitoring, but fail to support inter-application data flow and coordination aspects which have been identified as essential for integration in heterogeneous and distributed clinical environments (e.g. [4,5]). At the University of Leipzig, we are presently developing the distributed workflow system HEMATOWORK, which has explicit knowledge about the intended oncological treatment strategy and the associated communication paths between the involved institutions. In particular, HEMATOWORK intends to support the following basic tasks:

- ***Treatment Plan Functionality:*** This core functionality of HEMATOWORK covers therapy management and diagnostic monitoring, and is achieved through specialized applications (e.g. for calculating chemotherapy dosages) and databases coordinated by HEMATOWORK.
- ***Intra-hospital Communication Functionality:*** As every specialized medical workflow system inherently requires services of other local sections and departments, HEMATOWORK supports processes for communication and material transfer between the oncological site and other local departments via the Hospital Information System (e.g. appointment management with the radiological section, finding requests etc.).
- ***Inter-hospital Communication Functionality:*** This covers the communication and material transfer between the treating hospital and external expert panels and central commissions. It includes sending medical reports and test material from the hospital to external specialists and commissions.
- ***Tracking Functionality:*** This part of HEMATOWORK is located at the central commissions, and manages the tracking of the patients participating within the clinical trials. This involves periodical checks whether all reports have been received for a particular patient, and includes generating reminders and admonitions if reports are missing. Furthermore, incoming data are checked with respect to medical plausibility and correctness.

2 Methods

For workflow modeling, an extended Petri Net (PN) formalism has been chosen [6,7,8]. Besides the classical PN constructs such as transitions, places, firing conditions, and tokens, the following extensions are used:

[1] Institut für Informatik, Universität Leipzig
[2] Institut für Medizinische Informatik, Statistik und Epidemiologie, Universitätsklinikum Leipzig
[3] Medizinische Klinik und Poliklinik II, Abteilung für Hämatologie/Onkologie, Universitätsklinikum Leipzig

- To allow flexible data modeling and to facilitate an implementation of the PN workflow model within a heterogeneous and distributed environment, tokens are *objects* conform to the CORBA-model [9], being organized in class hierarchies and providing IDL[4]-interfaces.
- Places and transitions can be subclassed; instances of subclasses usually are constrained concerning their token input/output and induce a specific behavior of the net interpreter. An important transition subclass of the oncological model is, for example, *DetermineAnd-ExecuteMedicalPlan*. Instances of this class are placeholders for diagnostic and therapeutic plans, and induce a request to a knowledge base to determine an appropriate medical plan, depending on the specific cancer type and the results of previous therapy procedures applied to the patient. Similar to tokens, places and transitions have an IDL-interface.

3 A Petri-Net Workflow Model for Distributed Cancer Therapy

3.1 Organizational Processes

First of all, HEMATOWORK provides a set of generic workflow nets expressing the organizational structure of distributed hemato-oncology and abstracting from particular oncological diseases and therapeutic procedures. Processes of this organizational class mainly cover disease-independent aspects such as patient admission and discharge, or document management and material transfer between the involved departments. Fig. 1 shows the global communication paths and abstract high-level transitions of distributed hemato-oncology.

As it is expected that additional hospitals and specialist practices will participate in oncological trials in the future, this organizational model does not make any assumptions about the number or geographic location of the participating organizations (*organizational scalability*). For example, the diagnostic panel and the central commission for a particular hematological cancer type may be located at the same site, or in different towns. To facilitate a site-specific implementation, the model specifies all data requirements in CORBA/IDL.

3.2 Disease-Specific Process Refinement

Disease-specific aspects are linked to an organizational net by cancer-type specific token classes (such as leukemia specific *Report* token subclasses derived from a generic *Report* class), or by refining the net during run-time through disease-specific nets. To separate organizational from disease-specific

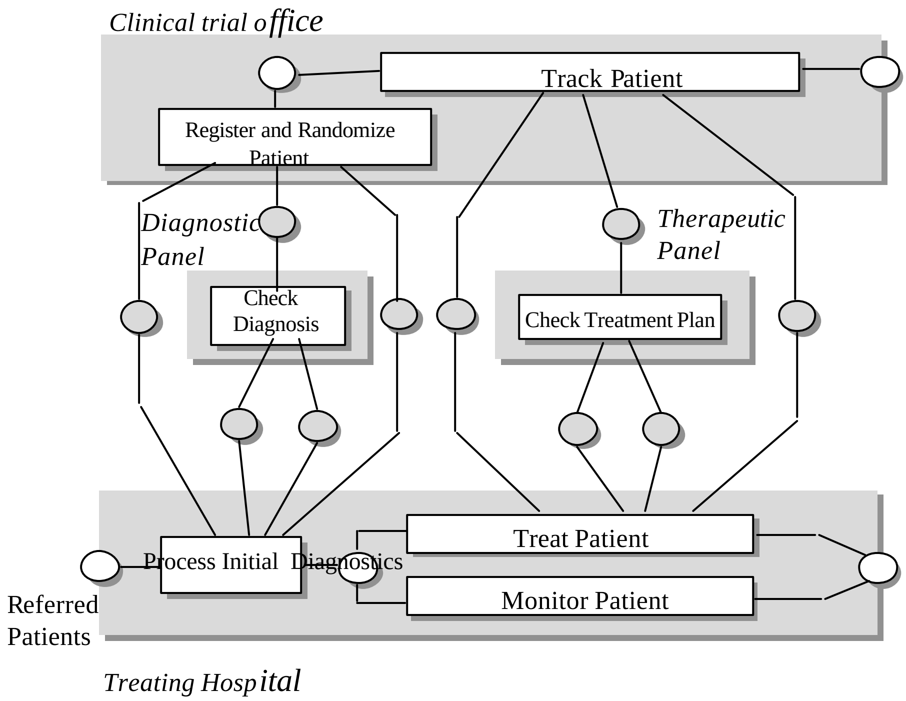

Fig. 1: High-level PN describing the data flow between the departments.

[4] Interface Definition Language

processes and to support multiple usage of medical knowledge (as knowledge is used not only by the workflow system, but also by other applications), this disease-dependent knowledge is encapsulated in a knowledge base. Beside terminological information and knowledge about the structure of the basic diagnostic and therapeutic procedures, this knowledge base contains "procedural" knowledge about the disease-specific conditional sequences of diagnostic and therapeutic steps. The knowledge of this procedural layer is then used to refine a generic workflow at run-time in the following way (also see Fig. 2):

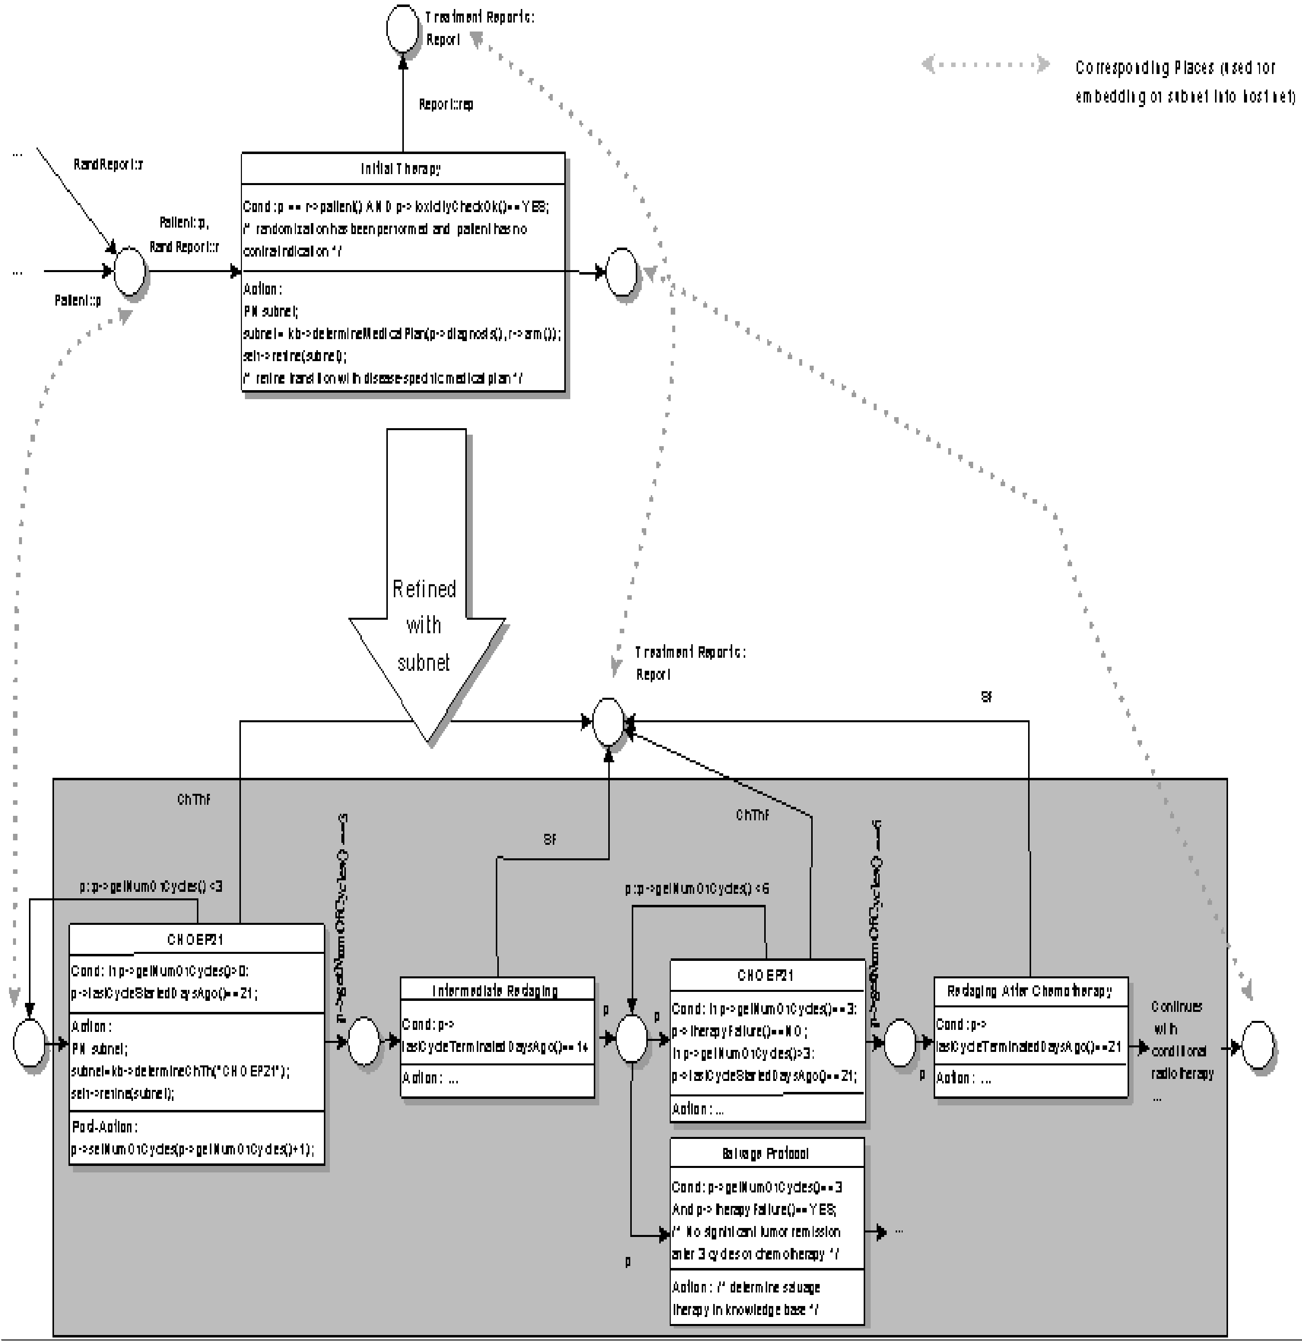

Fig. 2: Refinement of the abstract transition named *Initial Therapy*, which is an instance of the transition class *DetermineAndExecuteMedicalPlan*. The transitions named "CHOEP21" (a chemotherapy of the Non-Hodgkin lymphoma treatment) are instances of the class *DetermineAndExecuteChemotherapy* (which is a subclass of *DetermineAndExecuteMedicalPlan*); "kb" is a global CORBA reference to the knowledge base.

If the workflow engine processes a transition of the type *DetermineAndExecuteMedicalPlan*, this triggers the activation of a so-called *Plan Refinement Manager*, which identifies the patients specific situation (via inspecting the input *Patient* token), and then communicates with the knowledge base to determine the specific medical (sub-)plan representing the appropriate medical procedures to apply next to the patient. The identified plan is then transformed to a PN-notation and used to refine the *DetermineAndExecuteMedicalPlan* transition in a patient- and disease-adapted manner. The refined PN is then further processed by the workflow engine (which may detect further transitions of class *DetermineAndExecuteMedicalPlan*, inducing the same refinement process on a level of finer knowledge granularity). A more detailed description of the knowledge refinement process and the underlying knowledge representation approach can be found in [10].

4 Architectural and Implementation Issues

Currently, a prototypical implementation of the workflow system HEMATOWORK is realized at Leipzig University, based on IONA ORBIX and COSA Workflow, a PN-based workflow management system. To achieve independence and autonomy of the involved departments, the net in Fig. 1 has been split up into 4 parts (indicated by the gray rectangles), each part having its own workflow server and communicating with the other workflow systems. Communication between geographically distributed departments is realized with IONA OrbixWeb. A chemotherapy calculator, a patient database and a report database have been implemented with ORACLE/DELPHI, the oncological knowledge base is realized with O_2.

References

(1) Musen, M.A., Tu, S.W., Das, A.K et al.: EON: A Component-Based Architecture for Automation of Protocol-Directed Therapy. Journal of the American Medical Informatics Association vol. ... (3), 1996: 367-388.

(2) Séroussi, B., Morice, V., Dreyfus, F. et al.: Control theory as a conceptual framework for intensive care monitoring. Artificial Intelligence in Medicine, 7, 1995: 155-177.

(3) Müller, R., Sergl, M., Nauerth, U. et al.: TheMPO: A Knowledge-Based System for Therapy Planning in Pediatric Oncology. Computers in Biology and Medicine vol. 27(3), 1997:177-200.

(4) Brigl, B., Ringleb. P, Steiner, T. et al.: An Integrated Approach for a knowledge-based Clinical Workstation: Architecture and Experience. Methods Inf. Med. 37, 1998: 16-25.

(5) Taboada, M., Marín, R., Mira. J. et al.: Integrating Medical Expert Systems, Patient Data-Bases and User Interfaces. Journal of Intelligent Information Systems 7 (3), 1996: 261-285.

(6) Aalst, W.M.P. van der: Three Good reasons for Using a Petri-net-based Workflow Management System. In: Navathe, S. and Wakayama, T. (eds.): Information and Process Integration in Enterprises, *Kluwer Academic Publishers*, Norwell, 1997.

(7) Lakos, C.A.: From Coloured Petri Nets to Object Petri Nets. In: Proceeding of the 16th International Conference on Application and Theory of Petri Nets, Turin, June 1995 (LNCS 935): 278-297.

(8) Schöf, S., Sonnenschein, M., Wieting, R.: Efficient Simulation of THOR Nets. In: Proceeding of the 16th International Conference on Application and Theory of Petri Nets, Turin, June 1995 (LNCS 935): 412-431.

(9) Ben-Natan, Ron, B.-N.: CORBA : A Guide to the Common Object Request Broker Architecture. New York: McGraw-Hill, 1995.

(10) Müller, R., Heller, B.: A Petri Net-based Model for Knowledge-based Workflows in Distributed Cancer Therapy. To appear in: Proceedings of the Workshop on Workflow Management Systems of the Conference for Extended Database Technology (EDBT98), Valencia, Spain, March 1998.