Adaptive Workflows for Healthcare Information Systems

Kees van Hee, Helen Schonenberg, Alexander Serebrenik, Natalia Sidorova, and Jan Martijn van der Werf

Department of Mathematics and Computer Science
Eindhoven University of Technology
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
{k.m.v.hee, m.h.schonenberg, a.serebrenik, n.sidorova, j.m.e.m.v.d.werf}@tue.nl

Abstract. Current challenges in Healthcare Information Systems (HIS) include supplying patients with personalized medical information, creating means for efficient information flow between different healthcare providers in order to lower risks of medical errors and increase the quality of care. To address these challenges, the information about patient-related processes, such as currently executed medical protocols, should be made available for medical staff and patients. Existing HIS are mostly data-centered, and therefore cannot provide an adequate solution. To give processes a prominent role in HIS, we apply the adaptive workflow nets framework. This framework allows both healthcare providers and patients to get an insight into the past and current processes, but also foresee possible future developments. It also ensures quality and timing of data communication essential for efficient information flow.

Keywords: adaptive workflows, EPR, medical protocols, Healthcare Information Systems.

1 Introduction

The recent study of the Netherlands Health Care Inspectorate [23] has established a number of serious shortcomings in the communication between healthcare providers that can cause risks for patient safety. The typical examples named in the study are a lack of communication between the anesthesiologist and the surgeon involved in the same surgery; repetitive overwritings of variable data, such as blood pressure, with no information when, in which circumstances and by whom the measurements were performed; the use of "no message — good message" principle in the communication between specialists; providing insufficient information for patients about their treatments and expected developments. Another tendency reported is the steady increase of the total amount of patient-related information available, which, on one hand, increases the survival rate, but on the other hand, complicates the overall picture due to the chaotic nature of the information. The chance that the surgeon is aware of all information relevant to the surgery being performed is estimated as very low. The problems mentioned indicate a great need in a new generation of Health Information Systems (HIS) that would satisfy new information and communication demands.

A. ter Hofstede, B. Benatallah, and H.-Y. Paik (Eds.): BPM 2007 Workshops, LNCS 4928, pp. 359–370, 2008.

At the same time, the patient should become a focal point of the new generation HIS. On April 30 of 2004, the European Commission adopted an action plan [7] aiming at making healthcare better for European citizens. As opposed to currently available provider-centered health systems, the action plan envisions citizen-centered health systems. A survey of Harris Interactive and ARiA Marketing conducted in 2000 [15] shows that more than 80% of the respondents are interested in obtaining on-line personalized medical information and electronic alerts specific to their medical histories, while 69% would like to have access to the charts that monitor the progress. Therefore, in this paper we aim at the creating a new concept of HIS that would increase the availability of personalized information in HIS both for care providers and for patients.

Currently employed HIS concentrate often on patient-related *data* rather than consider treatments as *processes* producing these data. By getting access to the process information, a patient can obtain a clear personalized picture of ongoing treatments, expectations and risks. Therefore, (s)he is more likely to become an informed decision-maker [8]—note that "participatory" decision-making model is recommended as the preferred model of treatment decision-making [4]. Moreover, shifting the focus of healthcare information systems from data to processes provides a practitioner with information on continuation of a treatment initiated by her and continued elsewhere. That is another reason why we advocate the move from data to processes.

To address the issues raised above, we propose a new approach to HIS. The key idea consists in associating a (number of) process(es) to each patient. By logging into the private web area, the patient gets access to the information related to processes associated to her. This information should include (a view on) data produced by previously performed treatments and a number of likely scenarios for the treatment continuation. Using the same information, a care provider can decide to alter an ongoing process, to abort such a process or to initiate a new one. New processes can either be suggested by a physician, or, more probably, they can be borrowed from a library of protocols, including e.g., hospital policies and medical guidelines.

Our approach relies on the framework of adaptive workflow nets [10,12]. Main advantages of the framework include adaptivity, adaptability and separation of concerns. By *adaptivity* we understand the ability of a process to modify itself as opposed to *adaptability*, which is the ability of a process to be modified by an external party. By *separation of concerns* we understand that every (sub)process has its owner, and the owner is the only actor responsible for performing the steps of the process. Still, the processes are interrelated and communicating.

In addition, we extend the currently used concept of the electronic patient record (EPR) with the notion of history. In other words, rather than overwriting the blood pressure in a database cell, we record the measurement every time it has been taken, by whom (the owner of the process), when and why (note that this information can be automatically generated since we know to which process this action belongs). This history is then used in the decision making process. It should be noted that the availability of history also enables the application of data and process mining techniques [2,17], which can lead to improving the quality of medical care.

2 Motivating Example

As a motivating example we consider the story of Saskia, a thirty-six years old Dutch preschool teacher, pregnant with her first child in the fifth week gestation. Saskia occasionally uses her home computer for chatting with her friends or looking up information on the Internet. During her pregnancy Saskia is assisted by a number of healthcare professionals: midwives, lab assistants, doctors, having (partial) access to Saskia's (electronic) patient record. Below we present the first steps of Saskia's pregnancy in the current situation and discuss possible improvements due to the implementation of our approach.

Currently. Saskia looks up the information on midwifes' offices in an on-line telephone directory. After consulting the web-sites of the offices, she selects one, and calls the office to arrange an appointment with a midwife. During the first visit, the midwife records Saskia's medical, social and gynaecological history and orders a number of standard lab tests: blood type, rhesus, iron deficiency, glucose level and urine. Moreover, since Saskia is thirty-six years old, the midwife briefly informs her on possible age-related risks for the baby and additional tests that can be performed. Upon receiving this information Saskia gets overanxious: she has no time to reflect on the matter and she is not able to decide whether she is willing to take these tests.

Getting home Saskia talks to her husband and looks for additional information available on-line. Based on this information she decides to undergo a nuchal translucency scan (NT screening). Saskia calls the midwifes' office to arrange another appointment at which she would order the test.

Blood tests show that Saskia is Rh(D) negative, while her husband is Rh(D) positive. Therefore, at twenty eight weeks gestation she gets an anti Rh(D) IgG immunoglobulin injection, which will be repeated after the delivery.

Desired. After selecting the office, Saskia fills an on-line form with her personal data. Being prompted what means of communication (e-mail, text messages, regular mail, phone call) does she prefer, Saskia chooses text messages. Indeed, Saskia is very excited about her pregnancy and wants to get all the information as soon as possible whether she is at home, at her office or out with friends. However, she is a working woman and not every moment might be appropriate for getting the information. Figure 1 presents the completed registration form.

Saskia receives a text message informing her which standard tests she should undergo. She also gets a link to a *personalized web-page* providing her with information on age-related risks and additional tests. Saskia takes her time to study the web-page together with her husband and decides to undergo an NT screening. She indicates this choice on her personalized web-page. A number of test times is proposed to her taking into account the scanner availability and that NT screening is performed between the eleventh and the fourteenth weeks gestation. The time selected and additional information are sent to the laboratory working with the midwifes' office.

Moreover, Saskia gets informed on possible future developments. For instance, she learns that if, according to the NT screening results, a chance of her baby being affected is high, she will be offered to undergo amniocentesis to obtain a conclusive evidence.

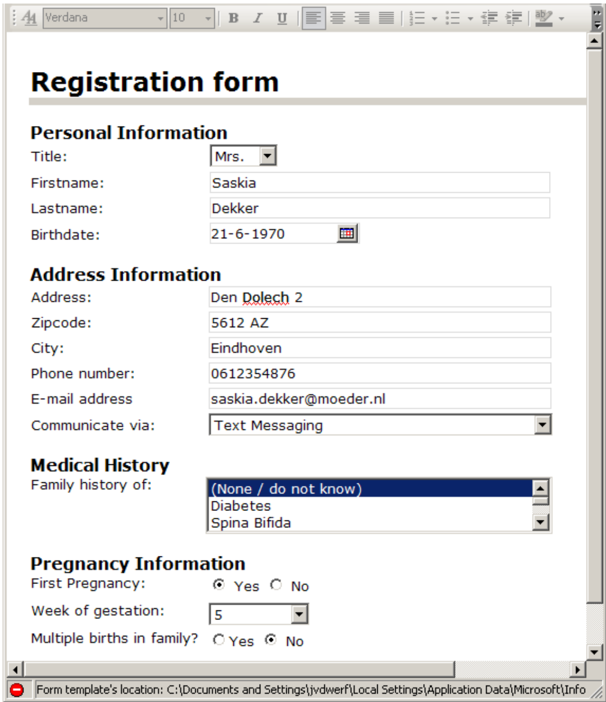

Fig. 1. Registration on-line form

The personalized web-page further includes a personalized pregnancy calender. All appointments arranged are automatically added and reminders are sent. Furthermore, depending on Saskia's personal data and outcomes of the preceding tests, new appointments to be scheduled are highlighted in the corresponding periods. For instance, anti Rh(D) IgG immunoglobulin injections will appear on the twenty eighth week and in her delivery procedure description. Note that healthcare providers have access to Saskia's pregnancy calender and get alerted, for example, if she fails to show up at her anti Rh(D) IgG immunoglobulin injection date.

Requirements imposed by the example. In order to support the desired scenario presented above a number of changes has to be introduced to current healthcare information systems. First of all, in order to provide Saskia with information on possible future developments, a healthcare information system should be made aware of the ongoing *processes* rather than only *data* involved. Recall that data relates to information till the current moment, while patients can be eager to know what can happen next. However, assuming one ongoing process is not realistic, as different healthcare providers have their own rules and processes. Sharing these processes might be superficial (a midwife should not necessarily be knowledgeable of the biochemistry of a blood test) or undesirable (lab assistant should not have access to Saskia's personal information). Therefore, rather than considering one process we envision *a series of interrelated but independent processes*.

Another important conclusion that we can draw from our example is that processes should be able to modify themselves on-the-fly. For instance, when Saskia's Rh(D) turns out to be negative, special treatment should be performed, i.e., a special process should be initiated in parallel with the ongoing one. We refer to the ability of a process to modify itself on-the-fly as *adaptivity*. Still, we do not intend to substitute medical stuff by an automatic decision making system. The choice of treatment protocols is certainly made by the care providers together with the patient. Therefore, processes should be not only adaptive but also *adaptable*.

3 Adaptive Petri Nets

As we could see in the Saskia example, HIS is a domain with a great need for adaptivity. In this section we present the theoretical foundations of our approach, so called *adaptive Petri nets*, and we illustrate the concepts introduced by means of a running example: a simplified version of the Dutch age-related prenatal diagnostics protocol, given in Figure 2. The owner of the considered process is the midwife.

A Petri net [18] is a bipartite graph whose nodes are called *places* and *transitions*. Transitions, graphically represented as rectangles, correspond to actions being taken. Places are represented as circles, and they are used to define the process flow. Given a transition we distinguish the *input places of the transition*, i.e., places with an arc going to the transition, and *output places of the transition*, i.e., places with an incoming arc coming from the transition. The Petri net presented in Figure 2 includes such transitions as "select age-related diagnostics, first trimester" and "communicate negative news and select further diagnostics". The name of a transition is written *under* the corresponding rectangle.

Workflow nets [1] are a special class of Petri nets, well-suited for modelling workflow processes. Workflow nets have exactly one place with no incoming arcs, called *the*

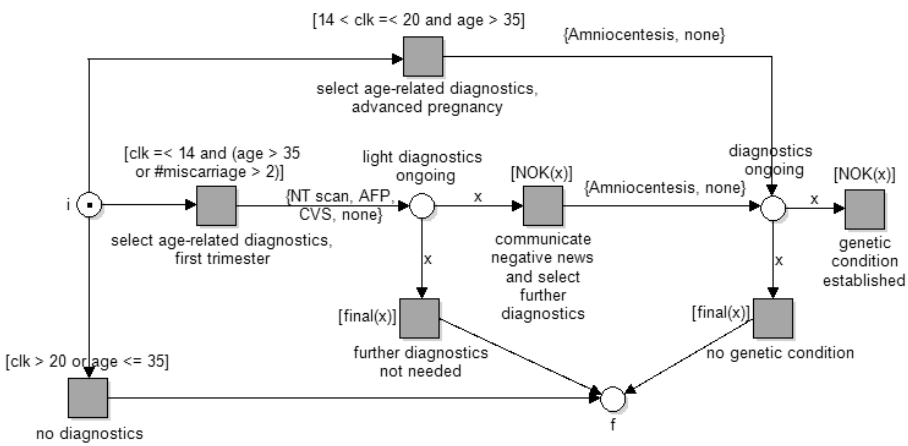

Fig. 2. Genetic condition related tests

initial place, and exactly one place with no outgoing arcs, called *the final place*. Moreover, every node in a workflow net is on a path from the initial place to the final place. *Extended workflow nets* [10] can be obtained from a workflow net by adding *exception transitions*, which are transitions with at least one input place and no output places. Exception transitions are used for modelling undesirable, abnormal or irregular events of such a nature that the process cannot decide itself how to continue and an assistance of a higher authority/layer is required. The Petri net in Figure 2 is an extended workflow net with the initial place i, the final place f and the exception transition "genetic condition established".

The state of a system is represented by means of *tokens*, drawn as black dots and residing in places. The initial state consists of a single token in the initial place i, while the final state consists of a single token in the final place f. The final state corresponds to the most likely, expected, termination of the process. Dynamics of a process is expressed by means of the *token game*: performing an action corresponds to a *transition firing* removing a token from every input place and adding a token to every output place of the transition. For instance, "select age-related diagnostics, first trimester" removes one token from the initial place i and produces one token in the output place *light diagnostics ongoing*. Figure 2 shows the state of the system before a firing of "select age-related diagnostics, first trimester". In the normal course of events, process terminates when we reach the state consisting of a single token on place f. Firings of exception transitions terminate the execution of the process independently of the process state, disregarding the fact whether there are still tokens left and transitions enabled.

Colored Petri nets [16] extend Petri nets by introducing data and time into the model, i.e., allowing to model a data flow in addition to a control flow. Due to historical reasons data types are commonly referred to as *colors*. Classical Petri nets are extended there by *guards* and *arc inscriptions*. The guard is a logical expression determining whether the transition may fire. Arc expressions at the outgoing arcs define data transformations. The firing of transitions become thus data dependent, and, moreover, transitions modify data.

Global history nets [12] further extend Petri nets by assuming the availability of a history record, registering all firings and the time of firing together with the information which process performed it. Transition guards can depend on the information contained in the history record. Guards are written between square brackets. For instance, "select age-related diagnostics, first trimester" has the following guard: "$\mathsf{clk} \leq 14$ *and* $(\mathsf{age} > 35$ *or* $\#\mathsf{miscarriage} > 2)$". This means that the corresponding EPR contains information that the gestation week is not later than 14 (clk denotes here the pregnancy clock), and either the maternal age exceeds 35 or the history record contains at least three previous miscarriages.

Adaptive nets[10] further extend the formalism by introducing a special color: nets. In other words, a token can be associated with another (adaptive) net, called a *token net*. Token nets are created by using suggested *library nets* or tailor-made nets. So, the firing of "select age-related diagnostics, first trimester" is in fact a procedure where the midwife together with the pregnant woman decides which prenatal test procedure will be carried out, if any at all. The suggested tests (NT scan, AFP, CVS and none) are indicated on the outgoing arc of the transition. By using "any" in the suggested set of token nets, we allow to take an arbitrary, possibly tailor-made, protocol. The considered example does not make use of these freedom of choice, since the set of tests is predefined. Note that the owner of the created token net is not necessarily the midwife. In this example, it will be a lab performing the selected test procedure.

Guards can be used to *synchronize* firings of the (upper layer) net with firings of a token net. The transition "communicate negative news and select further diagnostics" fires if the selected test procedure reports a high risk of genetic condition (NOK), for it is conditioned by the guard $\mathrm{NOK}(x)$, where x refers to the token consumed from the "light diagnostics ongoing" place. In other words, NOK in the selected test procedure and "communicate negative news and select further diagnostics" fire *synchronously*, which corresponds to the communication between the lab and the midwife. An additional form of synchronization is illustrated by "further diagnostics not needed". This transition can fire if the token net in the "light diagnostics ongoing" place has reached its final state, i.e., the only token present in the token net is residing in its final place.

4 Modelling Care Processes with Adaptive Nets

In this section we discuss how adaptive nets, discribed in Section 3 can be used to model care processes discussed in Section 2. First of all, we assume that (a view on) the electronic patient record (EPR) is available for all nets. EPR can be seen as a standard data record extended with the history log. Access to the EPR can and typically is further restricted depending on the task being performed: all tasks should be able to consult the pregnancy clock, but it is set during the registration and can be corrected by the midwife only.

Although the theoretical framework of adaptive nets does not restrict in the depth of nesting, we envision that the care applications will typically make use of three-layered processes. The *top-level* process represents the main process flow: registration, the choice of diagnostics and treatment protocols and closing the case. The main process flow can be associated with strategic goals. The second layer is normally a protocol layer. Processes of the second layer can can be carried out by the main care provider herself or delegated to other care providers. The operations can be seen as implementing the tactic goals of the process. The net at Figure 2 provides an example of a second level net. Finally, realization of medical protocols can require services of parties such as medical labs. These basic steps constitute the *third and the last level* of the process.

The deviations from the described process architecture scheme are of course possible: the midwife could for example demand Saskia to consult a cardiologist in case her blood pressure is repeatedly evaluated as too high. The cardiologist process, located at the second layer in this case, is however a strategic-level process, which is still a service for the midwife protocol. The nesting depth in this case could be greater than three.

The top-level process consists typically of the initialization (registration, analysis of the history available), a number of subprocesses taking care of adaptivity aspects, and a process termination (including a semi-automatic post-processing of the EPR to provide consistent data for future use). Figure 3 shows a pattern for a subprocess. The transition "initiate an additional protocol" is mostly often triggered by some event in the environment (e.g. an additional complaint of a patient). In response to this event, an additional

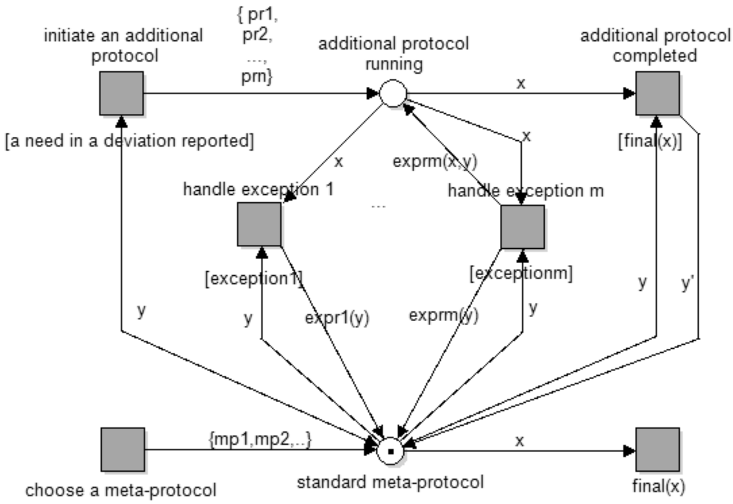

Fig. 3. Generic schema of a subprocess

protocol can be chosen and started (alternatively, or additionally, the token net on the place "standard meta-protocol" can be modified). Exceptions and the termination of this additional protocol are then processed by the corresponding transitions by initiating new processes of modifying the running protocol. Note that there are two kinds of exception-guards possible — the first one are guards demanding a synchronization with a corresponding exception transition in the token net, while the second one are guards specifying some external trigger.

The history record is composed from all actions taken in each protocol and it provides not only the information on the action taken but also specifies who (which process) took it. Also the initiations of additional protocols and modifications in the running protocols are logged. This allows to keep the (otherwise chaotic) history record well-structured, since we can always make a query to get all the information related to some treatment, abstract from unnecessary details (e.g. do not show actions of the third layer protocols), aggregate information related to a class of treatments (e.g. cardiological treatments), even if these treatments were initiated by different care providers (the midwife, the GP, the cardiologist).

Having the information on the processes running, we can construct a forecast view for both care providers and the patient by constructing and exploring a (partial) state space of the running protocols. Here different options are possible. The most simple one is constructing a forecast under the assumption that no exception will happen (e.g. the normal course of pregnancy, taking into account personal ERP information, like negative Rh(D)). More elaborated views allows to look into the future taking into account exceptions that are predefined in the running protocols. For instance, when the age-related diagnostics protocol in Fig. 2 is started for Saskia, it is possible to build a view informing that amniocentesis can be performed in case the NT-screening reports a high risk of a genetic condition.

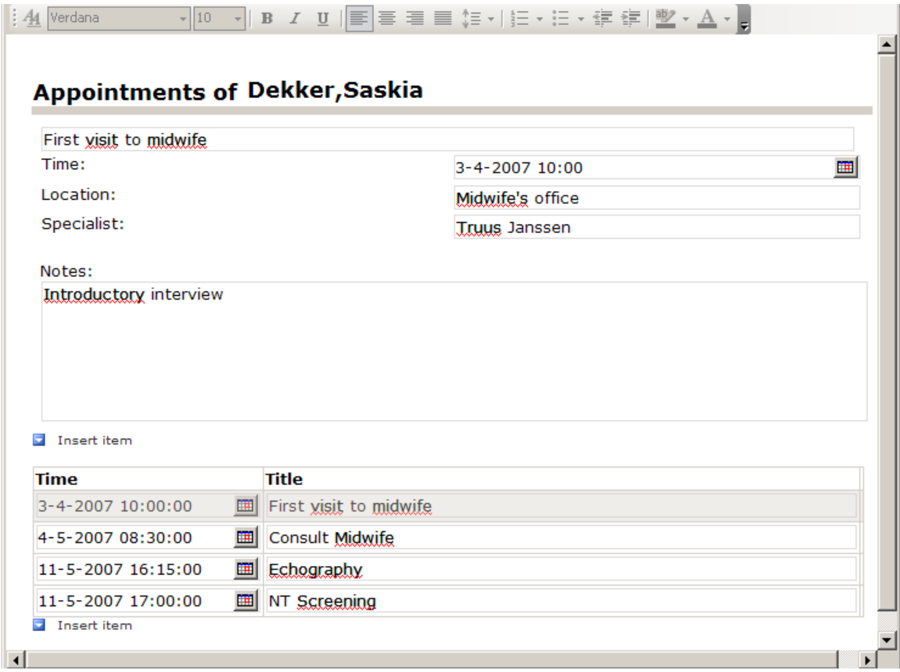

Fig. 4. Saskia's appointments calendar

The first type of a forecast view (the likely one) can also be used for such a pragmatic thing as planning appointments, tests and treatments. An important issue here is that the interprocess dependencies and restrictions can be handled, so that different protocols would not interfere or damage each other. Another trivial gain is e.g. providing a personalized pregnancy calendar for Saskia (see Figure 4) and providing her with reminders on appointments and tests supplied with additional information, like prerequisites of these tests.

5 Conclusion

The major contribution of our work consists in proposing a process-centered patient-centered framework for the new generation of HIS. While patient-centered systems attract more and more attention of the research community, the current proposals [5,21] concentrate mostly on implementation issues, such as data communication and heterogeneity of the application platforms, rather than on the conceptual ones, such as separation of concerns and adaptivity.

Our proposal is based on a solid theoretical foundation, namely, adaptive (nested) nets [11,10], a subclass of Petri nets. This allows us to perform a number of automatic correctness checks, like soundness and circumspectness. Soundness means that every process can terminate properly from any reachable state, while circumspectness means that every exception can be taken care of by the higher layer.

We do not expect the end user to know what Petri nets are and aim at providing a user-friendly web-interface instead. For this purpose, we developed the tool YasperWe [9] that is designed for prototyping IS. The tool integrates Yasper, which is a Petri net editor including a number of analysis options and compatible with some other analysis tools, with Microsoft Infopath.

Related work. Processes in healthcare are commonly documented in the form of *medical guidelines*. A guideline is not limited to doctors but also covers the workspace of nurses and paramedical personnel. There have been several attempts to formalize guidelines as *flowcharts* and *decision diagrams* and incorporate them into medical decision support systems.

Petri nets have been used for modeling of healthcare workflow, also known as careflow [13,19,20]. The guideline execution system GUIDE [20] translates formalized guidelines to a hierarchical timed colored Petri net. The resulting net can be run to simulate the implementation of the guideline in clinical setting. However, this formalism misses adaptivity and separation of concerns. The idea of adaptivity, i.e., controlled modification, in Petri nets has been considered in [6,14]. However, these approaches were able to model only processes involving two care providers, for instance, a midwife and a general practitioner, which is not sufficient for common healthcare processes.

Currently, there are several approaches that offer some degree of flexibility. Many share the idea of modeling with underspecification, i.e. a model where parts of the process are not given explicitly, but represented by a placeholder. At run-time the configuration is completed by binding the placeholder to a process from a repository, which is also known as late binding. Adaptive nets [10,11], worklets [3], pockets of flexibility [22] are modeling techniques that allow for late binding. Three main advantages of adaptive nets compared to the other approaches are the following:

– the ability to modify the structure of running processes in a controlled fashion;
– the possibility to define explicit synchronization between a token net and its owner;
– verification of a number of correctness properties for a subclass of adaptive nets.

Future work. For the future work, we are going to propose a number of process patterns to facilitate process modelling and a number of query patterns for the creation of different user views.

References

1. van der Aalst, W.M.P.: Verification of workflow nets. In: Azéma, P., Balbo, G. (eds.) ICATPN 1997. LNCS, vol. 1248, Springer, Heidelberg (1997)
2. van der Aalst, W.M.P., Weijters, T., Maruster, L.: Workflow mining: Discovering process models from event logs. IEEE Trans. Knowl. Data Eng. 16(9), 1128–1142 (2004)
3. Adams, M., et al.: Worklets: A service-oriented implementation of dynamic flexibility in workflows. In: Meersman, R., Tari, Z. (eds.) OTM 2006. LNCS, vol. 4275, pp. 291–308. Springer, Heidelberg (2006)
4. Charles, C., Gafni, A., Whelan, T.: Shared decision-making in the medical encounter: What does it mean (or it takes at least two to tango). Social Science and Medicine 44(5), 681–692 (1997)

5. Cheng, P., Yang, C., Chen, H., Chen, S., Lai, J.: Application of hl7 in a collaborative healthcare information system. In: 26th Annual International Conference of the Engineering in Medicine and Biology Society, 2004. EMBC 2004, Conference Proceedings, vol. 2, pp. 3354–3357 (September 2004)
6. Ehrig, H., Padberg, J.: Graph grammars and Petri net transformations. In: Desel, J., Reisig, W., Rozenberg, G. (eds.) ACPN 2003. LNCS, vol. 3098, pp. 496–536. Springer, Heidelberg (2004)
7. European Comission. Communication from the commission to the council, the european parliament, the european economic and social committee and the committee of the regions. e-health - making healthcare better for european citizens: An action plan for a european e-health area, SEC(2004)539 (April 2004)
8. Gerber, B.S., Eiser, A.R.: The patient-physician relationship in the internet age: Future prospects and the research agenda. Journal of Medical Internet Research 3(2), e15 (2001)
9. van Hee, K., Keiren, J., Post, R., Sidorova, N., van der Werf, J.M.: Designing case handling systems. In: Proc. of the International Workshop on Petri Nets and Software Engineering (2007)
10. van Hee, K., Lomazova, I., Oanea, O., Serebrenik, A., Sidorova, N., Voorhoeve, M.: Nested Nets for Adaptive Systems. In: Donatelli, S., Thiagarajan, P.S. (eds.) ICATPN 2006. LNCS, vol. 4024, pp. 241–260. Springer, Heidelberg (2006)
11. van Hee, K.M., Lomazova, I.A., Oanea, O., Serebrenik, A., Sidorova, N., Voorhoeve, M.: Checking properties of adaptive workflow nets. Fundamenta Informaticae 79(3–4), 347–362 (2007)
12. van Hee, K.M., Serebrenik, A., Sidorova, N., van der Aalst, W.: History-dependent Petri nets. In: Kleijn, J., Yakovlev, A. (eds.) ICATPN 2007. LNCS, vol. 4546, pp. 164–183. Springer, Heidelberg (2007)
13. Hoffman, K.: Run time modification of algebraic high level nets and algebraic higher order nets using folding and unfolding construction. In: Hommel, G. (ed.) Proceedings of the 3rd Internation Workshop Communication Based Systems, pp. 55–72. Kluwer Academic Publishers, Dordrecht (2000)
14. Hoffmann, K., Ehrig, H., Mossakowski, T.: High-level nets with nets and rules as tokens. In: Ciardo, G., Darondeau, P. (eds.) ICATPN 2005. LNCS, vol. 3536, pp. 268–288. Springer, Heidelberg (2005)
15. Q. Homan. Harris interactive and ARiA marketing healthcare satisfaction study (October 2000), Available at:
`http://www.harrisinteractive.com/news/downloads/HarrisAriaHCSatRpt.PDF`
16. Jensen, K.: Coloured Petri Nets. Basic Concepts, Analysis Methods and Practical Use. In: Monographs in Theoretical Computer Science, Springer, Heidelberg (1997)
17. Koh, H.C., Tan, G.: Data mining applications in healthcare. Journal of Healthcare Information Management 19(2), 64–72 (2005)
18. Petri, C.A.: Kommunikation mit Automaten. PhD thesis, University of Bonn, Bonn, West Germany (1962)
19. Quaglini, S., Panzarasa, S., Cavallini, A., Micieli, G., Pernice, C., Stefanelli, M.: Smooth integration of decision support into an existing electronic patient record. In: Miksch, S., Hunter, J., Keravnou, E.T. (eds.) AIME 2005. LNCS (LNAI), vol. 3581, pp. 89–93. Springer, Heidelberg (2005)
20. Quaglini, S., Stefanelli, M., Cavallini, A., Micieli, G., Fassino, C., Mossa, C.: Guideline-based careflow systems. Artificial Intelligence in Medicine 20(1), 5–22 (2000)
21. Ricci, F.L., Serbanati, L.D.: Mobidis: Toward a patient centric healthcare information system. Studies in Health Technology and Informatics 116, 557–562 (2005)

22. Sadiq, S.W., Sadiq, W., Orlowska, M.E.: Pockets of flexibility in workflow specification. In: Kunii, H.S., Jajodia, S., Sølvberg, A. (eds.) ER 2001. LNCS, vol. 2224, pp. 513–526. Springer, Heidelberg (2001)
23. van der Wal, G.: Rapport preoperatief traject ontbeert multidisciplinaire en gestandaardiseerde aanpak en teamvorming. Technical Report 2007-02, Staatstoezicht op de volksgezondheid. Inspectie voor de Gezondheidszorg (February 2007), http://www.igz.nl/15451/475693/2007-02_Rapport_Preoperatie1.pdf