First experiences with a semantic modelling method to guide simulation-based PACS development

W.J.J. Stut jr.[1], M.R. van Steen[2], L.P.J. Groenewegen[3], A.R. Bakker[1]

[1] BAZIS, Central Development and Support Group Hospital Information System,
P.O. Box 901, 2300 AX Leiden, Netherlands

[2] TNO Institute of Applied Computer Science,
P.O. Box 214, 2600 AE Delft, Netherlands

[3] Department of Computer Science, Leiden University,
P.O. Box 9512, 2300 RA Leiden, Netherlands.

ABSTRACT

The design of PAC systems benefits from modelling and simulation. Within the scope of the Dutch PACS project several simulation studies have shown that simulation is a powerful tool to evaluate hardware architectures, image management strategies and clinical working methods before they are actually implemented. Besides using simulation techniques for obtaining insight into a PACS's performance and behaviour, they can also be applied to actively support the design of a PACS. In order to capture the full complexity of PAC systems in a simulation model, and to take full advantage of simulation as a design tool, we are currently developing new techniques for model construction. This paper summarizes the underlying principles of our modelling method, and describes our first experiences with it.

1. INTRODUCTION

BAZIS (the central development and support group hospital information system) has been actively involved in the PACS field for quite some time now. To support the development of hospital-wide integrated image information systems, in 1984 BAZIS initiated its IMAGIS (image information system) project. In close co-operation with physicians, industry and other PACS research centres, IMAGIS focusses on PACS research and software development. IMAGIS was embedded in the Dutch PACS project, a co-operation between BAZIS, the Utrecht University Hospital and Philips Medical Systems, and financially supported by the Dutch Ministry of Health (WVC).

World-wide, the realization of PAC systems is still in its initial stages. Mainly due to the vast quantity of data to be stored and transmitted, the construction of image information systems appears to be extremely difficult in practice. Therefore, BAZIS decided to support this design and construction process by computer modelling and simulation [1].
Simulation can be used to predict, analyse, and assist in solving performance problems which might arise from the introduction of PACS in real clinical environments. Furthermore, by projecting the practical experiences obtained at small-scale PACS experiments on a full scale PACS situation, simulation can be used to predict the performance of large-scale hospital-wide systems. Applied in this way, simulation reduces the possibility that we underestimate the complexity and quantity of data that are inherent to a full scale PACS realization.
Besides using simulation techniques for obtaining insight into a system's behaviour, they can also be applied to support the actual design of systems. After all, when modelling future PAC systems, design decisions have to be made. For example, when modelling an image prefetching algorithm, one has to specify the events by which it should be activated explicitly.

Although the results obtained with IMAGIS' current simulation software have shown to be very useful, some drawbacks were also revealed. One of these was that, due to the absence of an appropriate modelling method, it was difficult to maintain our simulation model. We found that, in order to capture the full complexity of PAC systems in a simulation model, and to take full advantage of simulation as a design tool, novel advanced techniques for model construction need to be developed. This paper presents ongoing research into the development of a new modelling method.

2. MODELLING METHOD

Our modelling method [2] is based on semantic data models [3] and parallel decision processes [4]. As a first step, a semantic data model is used to construct a high-level overview of the relevant objects and their interconnecting relationships. As a second step, dynamic aspects are to be included. For this purpose we use the PARADIGM modelling method, based on parallel decision processes. It is beyond the scope of this paper to discuss our modelling method in detail. Instead, both steps will be illustrated with simplified examples.

Our semantic data model has three main structural components: objects, functional relationships and ISA relationships (see Figure 1).

- Objects are used to model the components of a system. They are divided into passive and active objects, with the difference that the latter display autonomous behaviour. Objects are visualized by rectangles, at which active objects have one double-lined side. For example, examrequest is a passive object, technician is an active object.
- Functional relationships represent properties of an object. Two kinds of functional relationships are distinguished: passivity relationships referring to an object's static properties, and activity relationships conceptually describing activities that can be done by an object. Functional relationships are represented by diamond-shaped boxes, at which activity relationships have one double-lined side. For example, technician-has-worklist is a passivity relationship, technician-examines-patient is an activity relationship.
- ISA relationships provide the basis for the use of inheritance. When using inheritance in the definition of an object, a designer needs only specify what is new about the object in comparison with the objects from which it inherits properties. In the visual representation, the rectangle representing the object that inherits properties, is enclosed by the rectangle representing the object that provides properties. For example, an in_patient is a patient, and consequently inherits the patient's properties.

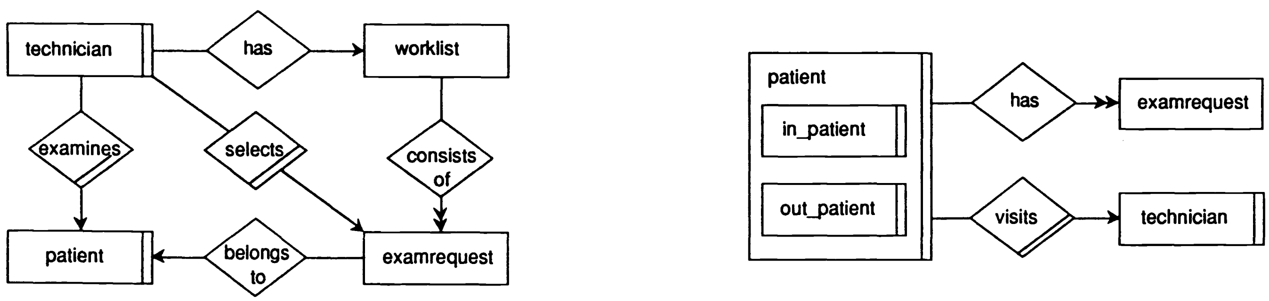

Figure 1. Two examples of semantic data models.

As semantic data models do not easily allow for the notion of time to be incorporated, special facilities have to be provided to model a system's dynamic behaviour. For that purpose we use the PARADIGM modelling formalism, based on decision processes, and providing the basic elements for modelling dynamic systems which exhibit parallel behaviour. Each decision process is used to model a system component which essentially exhibits sequential behaviour. Analogous to the underlying real-world system, decision processes may communicate.

Each decision process has a set of states, as well as state transitions (see Figure 2). At each moment the decision process is in exactly one state. The changes that occur in a system's component, are reflected by state transitions. For example, in state 2 the technician is calling the patient from the waitingroom. As long as the patient has not arrived, the technician stays in state 2. After arrival of the patient, the technician makes the transition from state 2 to state 3, and examines the patient. After the examination procedure, the quality of the images is checked. In the model the outcome of this quality check is simulated by stochastically selecting a transition which starts in the corresponding state 4. Once the quality is acceptable, the patient can be sent away.

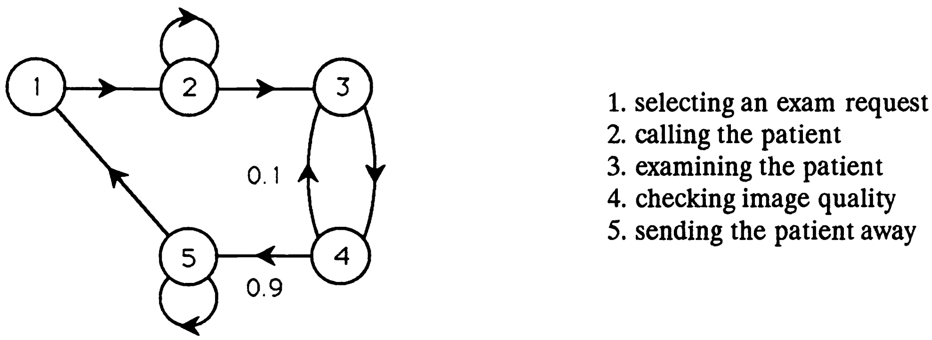

Figure 2. The decision process technician.

The starting-point in deriving a PARADIGM model is formed by active objects. Each active object is modelled as a decision process. The states and transitions are based on the object's activity relationships. By specifying the actual order in which states are to be visited, as well as the communication between decision processes, the dynamic aspects of a system (which are not addressed by the semantic data model) are incorporated.

After a PACS has been specified as a combination of a semantic data model and a PARADIGM model, modelling continues in a third phase mainly consisting of adding more details and by using a programming language (SIMULA) as the means for system specification. This eventually leads to a simulation program.

3. FIRST EXPERIENCES

In order to validate the new modelling method, we modelled a part of the working methods in the Utrecht University Hospital. During this validation the emphasis was on the modelling itself. Apart from some encountered difficulties (as discussed below), the validation showed that the method is very suitable to structure a real-world system, and that the obtained high-level structure can be directly translated into SIMULA code. As a result, the simulation program is a more or less direct representation of reality, and therefore easy to maintain. For example, the object patient as shown in Figure 1 can be directly translated to a SIMULA class, with a set of examrequests as attribute.

One of the problems encountered is the use of inheritance. Although the term inheritance is frequently used in the literature, precise analysis shows that there is no univocal interpretation of inheritance. Some object-oriented languages allow inheritance to be used in an ad hoc fashion, i.e. strictly as a code reusability mechanism. A special relationship (like ISA) between the giver and the receiver is not required as a basis: if one wants to build a new object that includes a capability already available in some other object, one can simply use inheritance to obtain the desired capability while ignoring those that are not desired. Our aim, however, is to use inheritance as a conceptual structuring mechanism, and to base the use of inheritance on ISA relationships (c.f. knowledge representation in artificial intelligence). For example, although a scheduled examination includes all properties of an examination request, in our model the former does not inherit those properties, since there is no ISA relationship with the latter.
A related problem is the use of multiple inheritance. Although our modelling method allows an object to inherit properties of several parent objects, SIMULA (as well as most other object-oriented languages) does not provide standard facilities to implement this construct.

The final semantic data model will generally contain many objects and relationships so that its visual representation will possibly be large and complex. In order to increase the understandability of a model, we decompose it into submodels by introducing for each object O a so-called fragment. The fragment of object O reflects its role in the complete semantic data model as viewed from O. To this aim, the fragment consists of those objects and relationships which are relevant to O. For example, the relationship technician-has-worklist is not included in the patient's fragment (see Figure 1), since it is not relevant to the patient.
The use of fragments promotes modularity. However, it is the modeller who determines which relationships are relevant to describe the role of object O, and it seems that there are no rigid guidelines for the construction of

fragments.

One of the current research topics is the transformation of a semantic data model to a decision process. Our aim is to provide the modeller with basic guidelines for obtaining decision processes with their respective states and transitions. Furthermore, if we are able to formalize this transformation, it can be partially automated. We emphasize that the dynamic aspects of a system (which are not addressed by the semantic data model) have to be explicitly specified by the modeller. For example, although a decision process' states and transitions are based on the object's activity relationships, the actual order in which states are to be visited still has to be determined.

4. CONCLUSIONS

In order to capture the full complexity of PAC systems in a simulation model, and to take full advantage of simulation as a design tool, we are currently developing a modelling method based on semantic data models and parallel decision processes. Although its theoretical basis has to be further extended, the method enables the construction of simulation models which are more or less direct reflections of the real system. As a result, these models are easy to understand and maintain, and actively contribute to the design of PACS. As the results of the first validation were promising, we are currently using the method to model the so-called ImageNet image network as developed by the Technical University of Aachen [5].

5. REFERENCES

1. W.J.J. Stut jr., J.P.J. de Valk, B.M. ter Haar Romeny, A.R. Bakker, "Modelling and simulation within the Dutch PACS project." In: Proceedings Medical Imaging III Conference, Schneider, R.H.; Dwyer III, S.J. (eds.), SPIE 1989, pp. 423-428.

2. W.J.J. Stut jr., A.R. Bakker, M.R. van Steen, L.P.J. Groenewegen, "A semantic modelling method to guide simulation-based system development." In: Proceedings of the 3rd European Simulation Congress, D. Murray-Smith et al (eds), September 1989, Edinburgh, pp. 101-106.

3. R. Hull, R. King, "Semantic Database Modeling: Survey, Applications, and Research Issues." ACM Computing Surveys, September 1987, Vol.19, No.3, pp. 201-260.

4. L.P.J. Groenewegen, G. Oosting, M.R. van Steen, "Modelling parallel phenomena." In Proceedings 2nd European Simulation Congress, Antwerp, Belgium, Vansteenkiste et al (eds), September 1986, pp. 45-51.

5. B. Fasel, F. Vossebürger, D. Meyer-Ebrecht, "High Speed Networks for PACS - Introduction and Comparison." In: Proceedings of the International Symposium CAR'89, H.U. Lemke et al (eds), June 1989, Berlin, pp. 577-580.