# Integration strategies in hospitals

Essay in "DIF8914 - Distributed information systems", December 2001

Gunnar Ellingsen
Gunnar.ellingsen@rito.no
Dept. of Computer and Information Science
Norwegian University of Science and Technology
N-7491 Trondheim

## Abstract

*This paper aims at responding to how two different strategies for integration (CORBA and Enterprise Application Integration) might influence comprehension of the quality of the information in the electronic patient record taken into account the problem dimensions on information systems integration (distribution, heterogeneity and autonomy. These issues are pursued through a study of electronic patient record systems (EPRs) in a large university hospital in Norway and the use of ICD-codes in different health related contexts. This essay also draws on the ongoing, large-scale Medakis project where the aim is to implement EPRs in the 5 largest hospitals in the country. Exactly as with Enterprise Resource Planning systems, EPRs in large hospitals are expected to integrate information and information based processes across departments, among different types of users and over time*

**Keywords:** integration, hospitals, CORBA, ERP-systems, Enterprise, Application Integration

### *1.1 Introduction*

Currently, hospitals typically have an abundance of poorly integrated information systems including patient administrative systems, laboratory systems, specialist systems of numerous kinds, a range of sensory/ graphical input devices such as X-ray, ultrasound, EKG and computer-tomographic images together with paper based records and indices (Grimson et al., 1998:124). Given this large number of partly overlapping, complementary and interdependent information systems, it is hardly surprising that considerable efforts have been poured into a tighter integration of these (Hartwood et al. 2001). Indeed, the integration of health information systems is currently something of a truism, a taken for granted ambition. As pointed out almost ten years ago:

> "The necessity for integration of systems and communication of information in [the health care] sector becomes evident when studying the variety of interested parties, the multitude of applications and their importance" (de Moor, 1993, p. 1).

Some scholars have ascribed the causes for this heterogeneity differences in database management, operating systems and design autonomy among component systems (Hasselbring 2000). A broader approach however suggests that incompatibility among information systems arises from differences in scope, use and responsibility across applications, sites and organizational units

Strategies and approaches to integration vary (Hasselbring, 2000) and includes technical solutions like federated database systems, World Wide Web and EDI (Grimson et al. 1998:124) as well as ERP-systems (Grimson, Grimson and Hasselbring, 2000). In addition, technologies such as CORBA and COM have emerged as a promising way to enable integration, as they are independent of programming language and operating systems (Emmerich, 2000:125). Also XML is considered for integration (Coulouris et al., 1994:23;

Bompani et al., 2000:305) in spite of insufficient reliability conformance (Emmerich, 2000:126). Defining technical integration strategies however does not solve the problem related to mutual *autonomy* between the components (Hasselbring, 2000; Sheth and Larson, 1990). In the following, this is a key aspect when I focus on the different integration approaches

This paper aims at responding to how two different strategies for integration (CORBA and Enterprise Application Integration) might influence comprehension of the quality of the information in the electronic patient record taken into account the problem dimensions on information systems integration (distribution, heterogeneity and autonomy. I pursue these issues through a particular instance of this problem, namely a study of electronic patient record systems (EPRs) and the use of ICD-codes in different contexts. Today, the ICD codes are stored in different hospital information systems as well as on a national level.

This essay draws on the ongoing, large-scale Medakis project supported by the Ministry of Social affairs in Norway establishing EPRs (dubbed DocuLiveEPR) at the 5 largest hospitals in the country. Exactly as with Enterprise Resource Planning systems, EPRs in large hospitals are expected to integrate information and information based processes across departments, among different types of users and over time (Hartwood et al. 2001; Grimson, Grimson and Hasselbring, 2000).

### *1.2 Medakis – an EPR project in the Norwegian university hospitals*

The Medakis project of the Norwegian health authorities has a long history (see Ellingsen and Monteiro (2001) for details), but has run since 1996 with Siemens as the privileged vendor. The DocuLiveEPR system, developed as part of the Medakis project is implemented to an interesting extent, especially at the regional hospital in Tromsø (RiTø) and the regional hospital in Trondheim (RiT). These two hospitals have used DocuLiveEPR for the shortest time (since February 1999), but nevertheless with the most widespread use among the five regional hospitals (Lærum, Ellingsen and Faxvaag, 2001)[1]. Both RiT and RiTø have completed their installations of DocuLiveEPR. Aligned with the development strategy, the users are currently primarily physicians and secretaries. Computer availability is reasonable for these two groups. The 5 regional hospitals together have about 6000 defined users of DocuLive EPR. On a normal working day, there are approximately 1400 concurrent users. A rough estimate suggests that there are about 2,5 million electronic documents, mostly physicians' notes but also sick notes and prescriptions. It is possible to import certain information from the patient administrative system (PAS) into DocuLiveEPR, but not the other way around[2]. With the upcoming (v5.0) version, the laboratory systems will be integrated and accessible through DocuLiveEPR. DocuLiveEPR also includes a workflow module that is strongly encouraged by management. A principal function here is the (digital) signing of notes by the physicians.

In spite of heavy deployment efforts the actual use of the EPR falls far below expectations. Physicians use the EPR for less than half of the tasks for which functionality is implemented

[1] The National hospital and Ullevål – one of the largest hospitals in northern Europe - are different with only 15% and 30%, respectively, installed.

[2] DocuLiveEPR imports demographic data like name, date of birth and addresses from PAS. In addition, the diagnostic and procedure codes may in principle be imported. This is seldom used, however, as these codes are normally recorded in the patient record before PAS. The integration mechanisms are not based on standards like CORBA or COM, but are rather hardware and language dependent solutions.

and a major part of the use is related to reading patient data from the EPR (Lærum, Ellingsen and Faxvaag, 2001).

### 1.3 Integration problem dimensions

Work on information system integration focuses to a great extent on three issues: autonomy, heterogeneity and distribution (Hasselbring, 2000; Sheth and Larson, 1990):

1. Distribution: hiding the geographical distribution of information systems.
2. Heterogeneity: hiding differences in platforms, programming languages and data models – as well as differences in perspective;
3. Autonomy: the extent to which the components are self-sufficient. The systems are autonomous in their design and meaning. Integration in under such circumstances normally require organisational change

Usually, SI aims at approaching the origin in this system of coordinates in Figure 1.

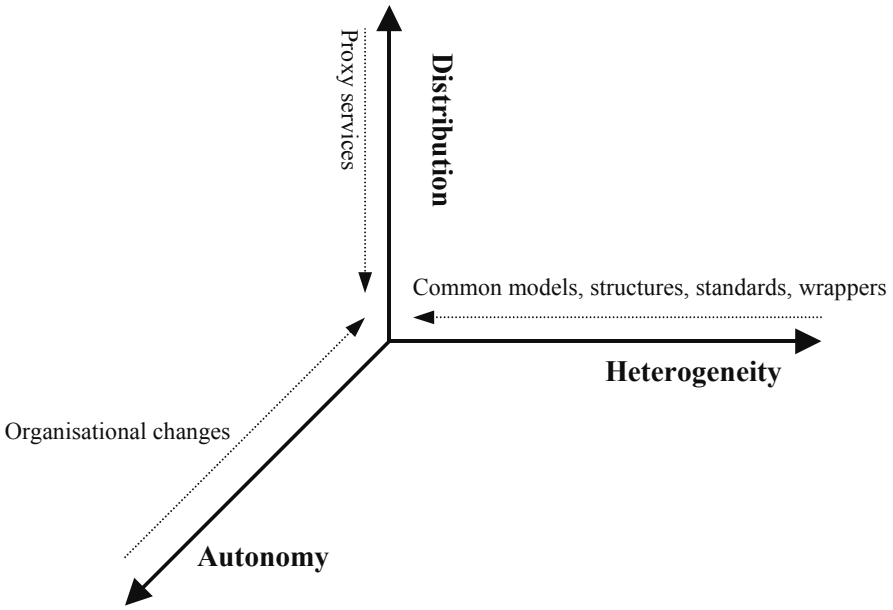


*Figure 1. Integration problem dimensions (Hasselbring, 2000)*

### 1.4 Achieving Integration through an ERP-system (EAI)

In this section I focus on Enterprise Resource Planning systems as the mechanism to achieve an integrated solution. In practice, this occurs by replacing existing systems. Enterprise Application Integration (EAI) in this context is in the following considered to be integration between modules that is encompassed by the ERP-application. In short, the idea is to replacing the existing IT-system with an ERP-system in which all the data is stored in one single database. This approach corresponds to the origin in the system of coordinates for the integration problem dimensions in Figure 1 (Hasselbring, 2000:54), which in practical terms means suppressing all of distribution, heterogeneity and autonomy.

An example of this is ASP solutions where ERP applications are central (Takahashi et al., 2000) and ERP-solutions promoted by SAP R/3, the dominant vendor of ERP systems and the world's fourth largest software company. An example of Enterprise Application Integration

(EAI) is the SAP R/3 module IS-H, which is deployed in a lot of German University hospitals. This module is specifically designed to support patient data management and some clinical functions and financing (Grimson et al., 2000:53). Integration with other R/3 modules can be achieved within the central database according to the general principle of SAP "everything in one database".

ERP systems appeal to health because they promise to provide and integrated application environment with seamless access to a single unified database (Grimson et al., 2000:53). In that sense, users relate to only one system and only one interface, which are assumed to enable easy learning and information availability as well as hampering fragmentation of information for the users and a more process-oriented organisation. Consider for instance the SAP website:

> "SAP R/3 overcomes the limitations of traditional hierarchical and function-oriented structures like no other software. [All the functions] are integrated into a workflow of business events and processes across departments and functional areas" (www.sap.com).

In the Medakis project the currently deployed solution in Norwegian university hospitals (DocuLive EPR) share several similarities with the EAI-approach. A key concern in the Medakis project has been the role of DocuLiveEPR in relation to the rich variety of other local, tailored and non-integrated information systems. In the requirements specification worked out in collaboration between the vendor and the hospitals, the presence of special purpose information systems is recognised:

> "There is a tendency that the specialist functions create their own information system to store and systematise data. These systems are only to a limited degree integrated or available in a uniform interface in Norwegian hospitals today" (Unified requirement specification, 1996)

This leads up to formulating a main goal of Medakis:

> "The EPR should replace many of the special purpose information systems that exist in the wards" (ibid.)

This is reinforced by the vendor's current strategy of subsuming specialist systems by replacing them by EPR controlled modules. As explained by a senior manager:

> "Instead of having many specialised systems you get a system that contains modules that can be added when needed. In that way you can let go of many specialised systems that have to be mutually integrated (...) [for instance] it is planned a PAS-module in the next version of DocuLiveEPR" (S1).

The visions about an all-embracing system for only one hospital is not the terminating point – the prospects of an Application Service Provider (ASP) solution is transparent in the currently concrete plans as is emphasised by a hospital manager:

> "Now, we challenge Siemens to elevate it one more level – above the hospital level to a concern level. We believe in installation of DocuLive for all our hospitals based on a common database. That implies that we wish to have the patient record for Odda, Stord or Voss in the same database as Haukeland (...) for us it is essential that the small hospitals shall have the same as the big – in that way we can over time include them in a complete concern" (M2)

### *1.5 Integration through CORBA*

CORBA is promoted by the Object Management Group (OMG) as a part of a strategy to develop an Object Management Architecture (OMA) for object-oriented computing (Blair and Stefani, 1998:34). OMG is a non-profit organisation where a major goal is to promote object technology for the development of distributed computing systems (Blair and Stefani, 1998:34-36). Currently, over 700 member companies support OMG (Szyperski 1998:178). Initially, OMG started out with a major goal in mind, "how distributed object-oriented systems implemented in different languages and running on different platforms interact?", and in that way it seems suited to respond to the call for more integrated IT-solutions in health care.

The purpose with the Common Object Request Broker Architecture (CORBA) is to enable interaction between objects running on different technologies and implemented in different languages. The core part is an Object Request Broker (ORB) that provides this communication. The ORB intercepts the call and is responsible for finding an object that can accept the request, pass it the parameters, invoke its method, and return the results (Hasselbring, 1997:194).

CORBA has essentially three parts (or components) (following Szyperski 1998:179):

- A set of *invocation interfaces*. They marshal an invocation's arguments such that the ORB core can locate the receiver object and the invoked method and transport the arguments
- The *object request broker* (ORB), which serves as a mediator between application objects across networks at runtime.
- A set of *object adapters* that unmarshals the arguments and invokes the requested method on the receiver object.

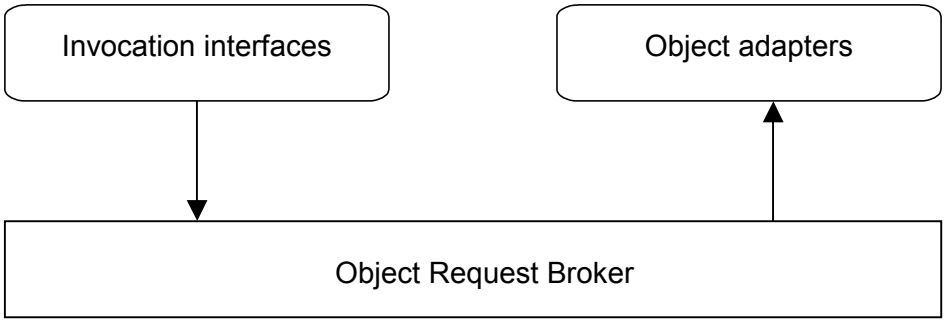


*Figure 2 The three parts of CORBA.*

The interface Definition Language (IDL) is used to describe the interfaces in a uniform manner, that is, language-independent. However the actual implementation of the object can typically be programmed in different languages like Smalltalk, Java and C++. Therefore a mapping (or binding) must be carried out between IDL and the language that is used to implement the object. IDL resembles C++, but with mechanisms to support distribution.

To enable efficient marshaling and unmarshaling of arguments, an ORB-specific OMG IDL compiler must be used to generate stubs and skeletons. A stub is responsible for forwarding all invocations through the ORB to the real target object. At the other side, a skeleton receives invocations, unmarshals arguments, and invokes the target method. This structure is presented in the figure below.

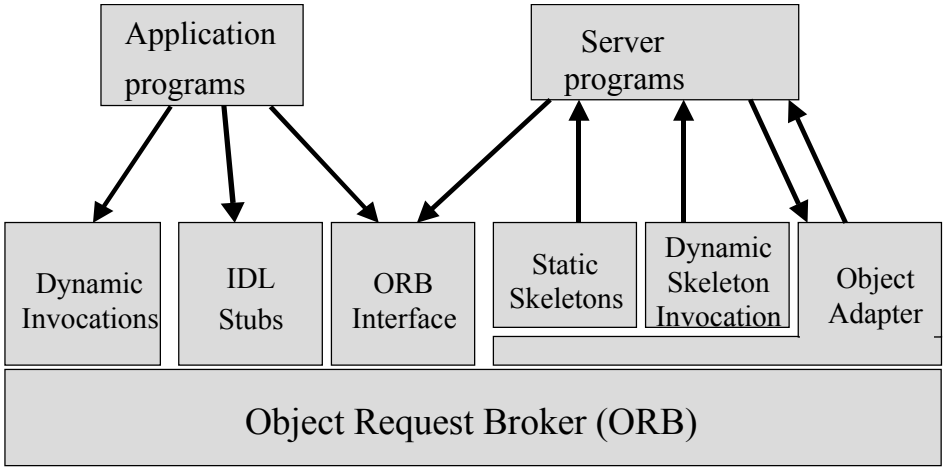


*Figure 3 The Object Request Broker and associated components*

### 1.5.1 OMA (Object Management Architecture)

Today OMG's major focus is OMA This represents a broader approach and includes the CORBA 2.0 specification and extended with several new areas of standardisation (Szyperski, 1998:182; Blair and Stefani 1998:39). According to the OMG website (2001) these are:

1. Object services
2. The Horizontal CORBAfacilities
3. Application objects

**The Object Services (CORBAservices)** are objects that offer basic services to the platform, for instance naming, concurrency control and transaction management as well as other domain-independent services. The Concurrency Control Service, for instance synchronize access between multiple clients towards shared objects, and the transaction service can ensure consistent updates of objects in accordance with a two-phase commit transaction protocol.

**The horizontal CORBAfacilities** offer higher-level services and sit between the CORBAservices and Application Objects. Horizontal facilities are still domain independent, but focus on specific application models. Unlike the Domain CORBAfacilities (describe next), these facilities are potentially useful across business domains. There are only four horizontal CORBAfacilities: The Printing Facility, the Secure Time Facility, the Internationalization Facility, and the Mobile Agent Facility (http://www.omg.org):

**The vertical domain CORBAfacilities** offer also higher-level services, but domain-specific or vertically oriented, like finance, health, e-commerce etc.

**Application objects**. The application objects are beyond the control of OMG as it is developed by outside vendors. The idea is that application objects represent domain-specific entities that can be plugged into component frameworks. An example of this is business objects, objects that directly represent abstractions used in specific businesses (også REFS Grimson). The figure below is from the OMB-website

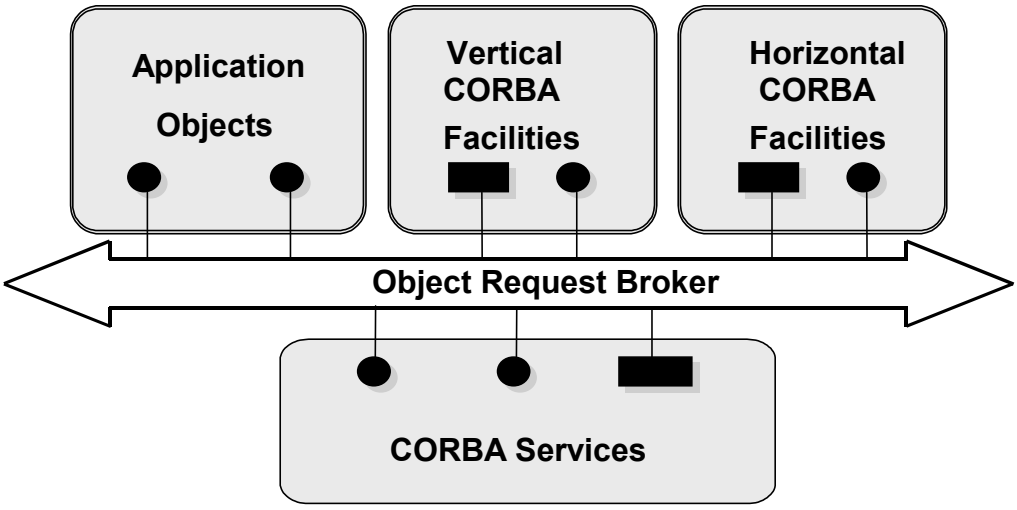


*Figure 4 The Object Management Architecture.*

OMG is also especially concerned about health by specifications through CORBAmed. CORBAmed's purpose is to improve the quality of care and reduce costs by use of CORBA technologies for interoperability throughout the global health care community. A set of services has been defined, like Clinical Observation Access Service, Clinical Image Access Service and Health care Resource Access Control. Nonetheless, progress is however slow – especially in Europe (ibid.). In that sense a record does not only consist of data, but also may include security, user validation services, patient identification etc, which constitutes added values for using CORBA (Grimson et al., 1998:137).

### *1.6 Integration in context*

The following discussion involve different topography of information systems, all supposed to include identical diagnostic and procedural codes based on the same patient from exactly on visit at the hospital. The involved IT systems is presented in the figure below (conform to a view that information technology enables decentralization of information (see for instance Blair and Stefani, 1998:15)):

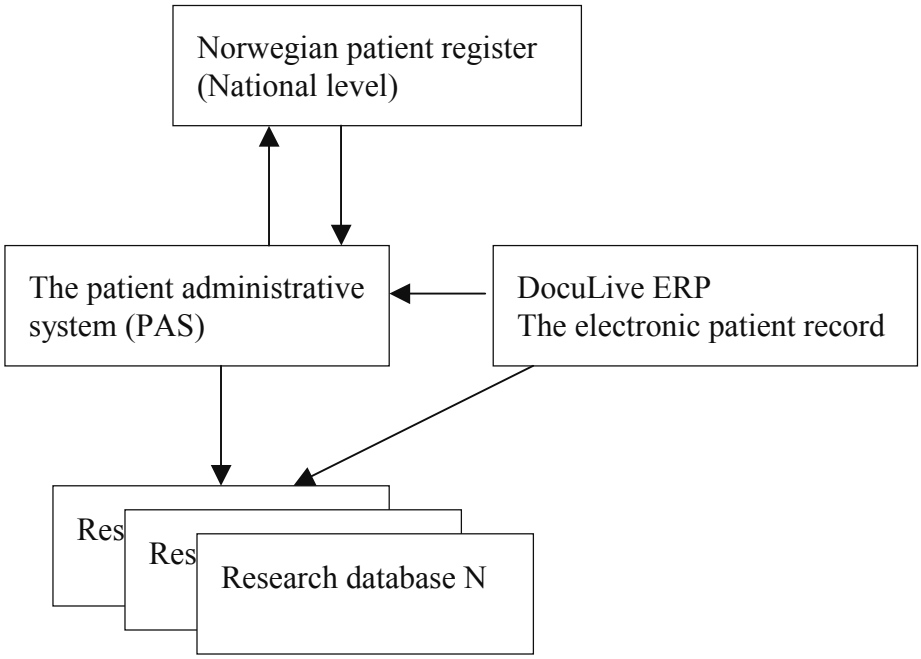


*Figure 5 The location of (ICD) diagnosis codes in Norwegian health care.*

As is indicated in the figure, the ICD[3] and NCSP[4] codes are stored in several systems among others in the patient record and in the PAS. But these codes take on quite distinct meanings in these two contexts of use. In the patient record, the ICD codes relate to the care and treatment of the patient and are the basis for subsequent discharge reports. In PAS, however, the ICD codes are used in a carefully designed way as the basis for governmental reimbursement according to the US adapted DRG[5] coding scheme. If the ICD coding fails to translate into the predefined DRG scheme, the hospital will not be reimbursed. The accounting department routinely checks the validity of the codes and reports back to the wards about errors or other reasons to make modifications. RiTø estimated a loss of 15MNOK in 1999 due to "incorrect" ICD/DRG coding.

### 1.6.1 Quality assurance through redundancy

Promises with ERP solutions (or centralised systems) fall into categories of streamlined business processes and improved data quality by entering and storing information only once. Taken the illustration in Figure 1 into account, such an argument is not surprising where the diagnosis codes are multiplied in several information sources. It is often argued that multiplied information in several information sources will result in a decline of the data quality because lack of consistency between the information sources (this is a challenge that is also recognised in the distributed systems literature, such as Blair and Stefani (1998:4-5); Allen et al. (1995:198).

---

[3] International Classification of Diseases as worked out by the World Health Organization (WHO).

[4] NCSP is an abbreviation for NOMESCO Classification of Surgical Procedures. The NOrdic MEdico-Statistical Committee was set up in 1966, following a recommendation by the Nordic Council. An aim of NOMESCO is to promote the coordination of health statistics in the Nordic countries.

[5] DRG is short for Diagnose Related Groups. The DRG system divides hospitalised patients into groups on the basis of diagnosis and treatment. Based on the hospital's operating costs, an expected price per patient discharged is estimated.

Blair and Stefani (1998:4-5) on the other hand argue that redundancy may increase availability and thus promote increased resource sharing. This are arguments commonly observed related to distributed systems environment. If one component fails, the others will continue to offer services.

Redundancy of components (or information sources) however may offer additional advantages directly influencing the comprehension of quality. For instance, a chief physician in the Dept of Medicine told that

> "The quality in the central IT-systems in keeping overview of performed Angiography and PTCA (heart examinations and out-blocking) is unsatisfactory. Therefore I maintain my own personal archive"

In that way he can compare the content of his personal archive containing information about heart examinations with the corresponding information in the PAS and EPR-systems and in such a way contribute to improve the quality in these systems that was more global in character.

Jet another example is from the dialysis section, dept. of Medicine:

> "We have dialysis-patients that come in regularly for treatment, and every time there are certain things that must be carried out (...) partly some extremely important computations which indicate if it is necessary to change modus and whether their medication is sufficient. (...) We paste the results into the discharge report."

In these cases, quality is related to a particular context and to the secretaries that do the actual writing of the discharge reports. This is the only place where these records are archived and only the secretaries in this section is trusted to write (based on a physician's dictating) these reports. Then quality in the EPR can be secured by support from other information sources.

### 1.6.2 How different integration strategies influence work

An EAI-approach implies uniquely defining ICD/DRG coding across the PAS and DocuLiveEPR. This will generate additional work for the physicians. If the ICD codes were identical in both systems, any changes of coding would have to be carried over to patient record as well. This includes, for instance, the discharge reports that are sent to the general physicians. This is by no means a straightforward task as nothing that is signed by a physician can be changed. The changed code would have to be written to DocuLive before sending new copies to the general physician with the updated codes. From a clinical point of view, this work is irrelevant as it only relates to economy.

The coding of ICDs is not necessarily "wrong", but has to be continuously updated according to new guidelines from the Ministry of Social affairs or misplaced main and secondary diagnoses. Often these aspects influence reimbursement. As the accountant explains:

> "We correct codes regularly (...) [and] the Ministry of Social affairs acknowledge that the hospitals can correct these codes going back the whole year. For instance, in September, new guidelines were issued that were supposed to be valid already from January 1. the same year".

An observed example from RiTø is the sequence of the two ICD codes C91.0 and Z51.1. The former expresses a type of cancer while the latter expresses appearance to chemotherapy related to the disease. Before the change, with C91.0 as the main diagnose and Z51.1 as the secondary, no reimbursement was generated. After correcting, i.e. swapping the sequence of the two, the reimbursement produced 11952 NOK.

### 1.6.3 Quality assurance through maintenance of different interpretations

Many have pointed out how medical work in general and patient record keeping in particular get caught up in an increasing number of roles (Berg and Bowker, 1997; Bowker, Timmermans and Star, 1995) (consider also the similarity with 'viewpoints' in the ISO RM-ODP object model (Blair and Stefani, 1998:27-28)). Beyond a resource for diagnostic purposes, it functions as a vehicle for coordinating work, as a source for cost- and income generation and become relevant for insurance companies**.** The use and meaning of diagnostic and procedural codes is an illustration of this.

Participants in clinical research often perceive that the quality of the diagnostic codes in PAS and even in the patient records are insufficient for their research projects and henceforth create their own clinical research databases where they maintain their own diagnose registers. These codes are also compared to the actual patient record text in order ensure the right quality. The former head of the clinical research department underscores this:

> "If you base clinical research solely on diagnose codes from the patient record [and PAS] then the validity will be challenged as it is well-known that the quality of these codes is poor (...) it was bad also before DRG was implemented[6], but has since deteriorated further (...) as the physicians are under pressure to code to maximise reimbursement"

This goes to the heart of their problems as it is essential to select patients belonging to certain diagnose categories in their research. This spawns local improvisation:

> "You may wonder, then, how it is possible to locate the patients when it is not possible to use the diagnose codes! We cope by using the laboratory results (…) [as index and subsequently] read through the patient record text in order to see what this is all about"

### 1.6.4 Legitimising heterogeneity through CORBA

Centralised approaches to integration of information systems ban systems adapted to more local use and wrongly considered a less complex solution compared to a distributed solution due to the erase of partly duplicated information sources. The reality however pulls in the opposite direction towards more decentralization of information:

> "Most modern-day organizations are very large and tend to consist of a number of separate institutions spread across a wide geographical area" (Blair and Stefani, 1998:15)

CORBA focuses relatively narrowly on technical heterogeneity (like heterogeneity in hardware, platforms, languages and management policies Blair and Stefani (1998:7-8). Generally this corresponds to the distribution aspect of Figure 1. Nonetheless, the focus on "content" (that is: no different perspectives of the content) however is missing which closely resembles the heterogeneity and autonomy aspect of Figure 1.

In spite of CORBA's relatively narrow focus on technical heterogeneity it carries a point of view that acknowledge that not only technical components are complex and heterogeneous, but also organisations, such as large hospitals as

> "modern organisations are highly distributed (...) [and therefore] a strong correspondence between distributed systems technologies and such organisational structurs" Blair and Stefani (1998:5).

---

[6] The DRG reimbursement schema was introduced in 1997.

Thus considering using CORBA as a mechanism for integration also indirectly recognise and acknowledge the existence of systems that play a key part in the respective organizations and that still have to do it in the continuation. For instance, that it is a similarity between distributed systems realised by CORBA and virtual organisation.

## *1.7 Conclusion*

For information systems, it is increasingly difficult to draw a line around an application system and say that you own and control it (Hasselbring, 2000). This is conform to the perspective on distributed systems that they shall be easy extendible.

It is argued that distributed systems (exemplified through CORBA) have a lot of benefits compared to centralised systems. In a medical context however a distributed system approach is especially attractive because the health care organisation can be considered a set of disparate users, performing diverse tasks on different technologies that have to base their daily work on information that to a high degree must be shared and integrated (Grimson et al., 2000:52). Therefore, these middleware solutions can prove to be the glue that connects the parts together.

## *1.8 References*


Allen, L.; Fernandez, G.; Kane, K.; Leblang, D.; Minard, D., and Posner, J. (1995): ClearCase MultiSite: Supporting Geographically-Distributed Software Development. In Jacky Estublier, edtor, *Software Configuration Management: selected papers/ICSE SCM-4 and SCM-5 workshops*, Springer Verlag LNCS 1005, pp. 194-214.

Berg, Marc and Bowker, Geoffrey (1997): The Multiple bodies of the medical record: Toward a Sociology of an Artifact. *The Sociological Quarterly*, vol. 38, no. 3, pp. 513-537.

Blair, G.S. and Stefani, J.B (1998): *Open distributed processing and multimedia.* Addison-Wesley.

Bompani, L.; Ciancarini, P. and Vitali, F. (2000): Software Engineering and the Internet: A Roadmap. In *Anthony Finkelstein (ED.): ICSE'2000 Roadmap papers*, pp. 305-315. IEEE CS Press.

Bowker, G., Timmermans, S. and Star, S. L. (1995): Infrastructure and organizational transformation: classifying nurses' work. In W. Orlikowski, G. Walsham, M.R. Jones and J.I. DeGross (eds.): *Information technology and changes in organizational work.* London: Chapman & Hall

Coulouris, G.H.; Dollimore, J. and Kindberg, T. (1994): *Distributed Systems: Concepts and Design*. Addison-Wesley, second editon.

De Moore, G. (1993). Standardisation in health care informatics and telematics in Europe: CEN 251 activities. In . De Moore, C. McDonald and J. Noothoven van Goor (eds.): Progress in standardization in health care informatics. IOS Press: Amsterdam.

Ellingsen, G. and Monteiro, M. (2001). Big is beautiful: electronic patient records in large hospitals in Norway 1980s – 2001. Submitted to *Methods of Information in Medicine*.

Emmerich, W. (2000): Software Engineering and Middleware: A Roadmap. In *Anthony Finkelstein (Ed.): ICSE' 2000 Roadmap papers,pages,*119-129, Limerick, Ireland, June 2000. IEEE CS Press.

Grimson, J.; Grimson, W.; Berry, D.; Stephen, G.; Felton, E.; Kalra, D.; Toussaint, P.; Weier, O.W. (1998): A CORBA-Based Integration of Distributed Electronic Healthcare Records Using the Synapses Approach. *IEEE Transactions on information technology in biomedicine*, vol. 2, no. 3, pp. 124-138.

Grimson, Jane; Grimson, William; Hasselbring, Wilhelm (2000): The SI challenge in health care. *Communications of the ACM*, vol. 43, pp. 48-55.

Hartswood, M.; Procter, R.; Rouncefield, M. and Sharpe, M. (2001): Making a Case in Medical Work: Implications for the Electronic Medical Record. To be published in *Journal of CSCW*.

Hasselbring, W (2000): Information system integration. *Communications of the ACM*, vol. 43, no. 6, pp. 32-38.

Hasselbring, W (1997): Federated integration of replicated information within hospitals. *International Journal on Digital Libraries*, vol. 1 6, pp. 192-208.

Lærum, H.; Ellingsen, G. and Faxvaag, A. (2001): A survey of physicians' use of electronic medical records systems in hospitals. *British Medical Journal* (http://bmj.com/cgi/content/full/323/7325/1344).

OMB website (accessed November 2001): http://www.omg.org/gettingstarted/specintro.htm

Sheth, A.P. and Larson, J.A. (1990): Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases. *ACM Computing Surveys*, vol. 22, no. 3, pp. 183-236

Szyperski, Clemens (1999): *Component Software: Beyond Object-Oriented Programming*. ACM Press New York, Addison-Wesley.

Takahashi, Kenji; Finkelstein, Anthony; Emmerich, Walfgan; Guerra, Sofia (2000): System Development Using Application Services over the Net. *Dif8914 Slides, cursory*

Unified requirement specification (1996): Unified requirement specification between SNI A/S and the University hospitals in Norway, date 09.05.1996.