Representing Organizational Competencies

Artur Caetano   João Pombinho   José Tribolet

Department of Information Systems and Computer Science,
IST, Technical University of Lisbon and
Organizational Engineering Center, INESC INOV.

INESC, Rua Alves Redol 9, 1000-029 Lisboa, Portugal.

+351 21 310 0258

{ artur.caetano | joao.pombinho | jose.tribolet }@inov.pt

ABSTRACT

Business activities within the organization are performed by a number of human or automated actors. For the organization to adapt to internal or external changes it must be able to understand how and why actors are related to and assigned to processes. This requires a consistent representation of the services required by the organization's processes and those provided by its actors. This paper focuses on defining the concepts that allow structurally aligning human actors and business processes through the description of the organizational competencies required to perform processes' activities. These structures can then be used dynamically within a marketplace-based model, supporting the management of actors and activities according to the supply and demand of competencies.

Categories and Subject Descriptors

H.1.0 [Models and Principles]: General
I.6.5. [Simulation and Modeling]: Model development

Keywords

Competency-based management, competency, skill, actor, scheduling, business process management.

1. INTRODUCTION

Competency-based management is becoming an object of growing interest as its importance is recognized from a strategic perspective. It aims leveraging the competitive advantage of knowledge-based organizations by observing individuals as strategic assets and explicitly integrating them into the organization's business strategy and processes. This approach has lead to the development of several methods and systems, often deriving from knowledge management related studies [4, 6, 7, 15].

Competency-based management provides individuals with a set of objectives to work towards and clear task definitions, linking organizational and personal objectives. It also facilitates process analysis and standardization inside and outside the organization's boundaries, making recruitment and compensation systems fairer and more open. However, competency related information is prone to become obsolete due to organizational changes, so it can

become expensive and time consuming to keep it updated [16]. Competency management support systems play a role in different organizational activities, such as expert finding, personnel recruitment and project management. These activities customarily involve human judgment to classify the skills held by workers, to evaluate the degree of competency and keeping up-to-date profiles. To facilitate the management of competencies and to minimize the subjectivity of human evaluation, several approaches put forward the use of information technologies to facilitate tasks such as contextualizing the communication between actors [5], managing skills and activities within teams [12] and using groupware to support information sharing [13].

Nevertheless, most of these solutions focus on supporting the operational phases of a business process and few assist the identification and selection of actors before the actual commitment to carry out its activities. This operational focus also means that skill management systems often do not relate the skill information pertaining to the actors to the activities' requirements as derived from the organization's process models. As a matter of fact, despite several approaches to business process modeling, such as BPMN [3], IDEF3 [14], Event Process Chains [1] and Role Activity Diagrams [10], focus on different modeling perspectives, such as activity coordination and resource modeling, they all lack the means to describe the competencies required to execute an activity and those provided by the organizational actors. Consequently, the information on the organization's competencies cannot be directly traced to its business processes. This hinders skill management from a process perspective and promotes mismatches within the description of the enterprise architecture.

This paper outlines a framework to overcome these issues by extending previous work on business process and actor modeling (v. [2, 9, 11]) with a set of concepts that allow representing and evaluating actors and competencies in the context of a business process. The remainder of this paper is structured as follows: section 3 describes our proposal on competency modeling, from a single competency structure to an integrated view of the modeling process, along with an example. Section 4 describes how to integrate this model into an organization using the concept of competency marketplace. Section 5 sets out the conclusions and describes current and future work.

2. RELATED WORK

Despite the research on competency-based management, the representation of competencies is not conveniently addressed in nearly all business process modeling standards. IDEFØ [21] is a method of modeling decisions, actions and activities of a given organization or system. It defines concepts such as function, input, output and control but it does not provide the means to repre-

sent how activities are performed by actors. BPMN [3] is a graphical notation that focuses on describing the activity flow within business processes. The elements of its diagrams are flow objects, swimlanes, artifacts and connecting objects. Participation in a business process is represented through swimlanes that relate activities to its performer, thus mapping responsibilities at a very high level of abstraction. There is also no means to specify the requirements an actor must fulfill to execute an activity. Other approaches, such as IDEF3 [14] and Role Activity Diagrams [10] focus on describing process flow but overlook the specification of human or system actor competencies as well as the requirements of an activity.

However, there is a large amount of research on competency modeling outside the scope of business process modeling. The "competency movement", credited to David McClelland [20], makes use of the concept of competency as the fundamental unit to classify human actors and to relate them with the ability of performing a task in a given environment. Spencer *et al.* [8] define competency as the "underlying characteristic of an individual that is causally related to criterion referenced effective and/or superior performance in a job or situation." *Underlying character-istic* means that the competency is part of a person and can be used to deduce behavior, being part of what Davenport classifies as tacit knowledge [22]. *Causally related* means that a compe-tency causes or can predict behavior and performance. *Criterion referenced* means that it is possible to quantify the level of per-formance of a competency in a task. This definition of compe-tency implies a causal relation between purpose and result. The performance criteria are critical since a characteristic or capacity cannot be considered a competency if it does not predict some-thing useful in the real world.

Defining the granularity of competency representation also plays an important role: on the one hand, a high-level, coarse-grained representation will not provide enough information. On the other, if it is too detailed, the entire representation process may become compromised, as it is effort and time-consuming [16]. Competen-cies are usually represented as hierarchical structures or compe-tency trees. An example of such structure is the "body of knowl-edge", in which competencies are classified according to specific areas of knowledge. A well-known example is IEEE's Software Engineering Body of Knowledge [23].

Pigneur *et al.* [18] also propose using a hierarchical competency structure to represent competencies since it simplifies processing and inference when compared to the mining of textual descrip-tions. Additionally, it is easier for people to identify their compe-tencies, the expectations of the organization as well as possible gaps. Pigneur proposes four competency categories as starting points for further classification: Enabling Technologies, Field Experience, Knowledge and Personal Traits. However, the seman-tics associated with these sub hierarchies is static by design, rely-ing on a functional domain specific classification to describe the competencies.

A more pragmatic approach can be found at CommOnCV [24]. This project aims at modeling *curricula vitae* during the recruit-ment process. Competencies are represented as sets of annotations derived from each *curriculum* and are formally represented using languages such as RDF [25] or DAML+OIL [26]. The annotations are based on a particular competency model, created within the scope of a specific knowledge area.

Although competency trees simplify competence representation they present some other problems related to flexibility and reus-ability. Since the trees are usually deep, with the functional cate-gorization specified in its upper levels, the description of a com-petency is not easily decomposable, leading to the dispersion of concepts between multiple nodes. As a result, the inference that can be done over those concepts is limited as they are not struc-tured and connected to each competency as individual compo-nents. Additionally, reusing competencies in different contexts is hindered as functional categorization is embedded at the tree's upper levels. Other issue is that relationships between different nodes, i.e. competencies, are not taken into account.

Next section proposes a competency representation framework that aims minimizing the issues here described.

3. COMPETENCY MODELING

Our goal is modeling the services provided by business actors and those required by the process activities. Such a framework makes possible to analyze the dynamics of human resource management and facilitates actor scheduling and team formation. It also en-ables assessing the alignment between actors and activities within the enterprise architecture.

The approach comprises three phases: definition, aggregation and analysis (v. Fig. 1). These phases aim at separating the definition of individual competencies from their hierarchy and aggregation rules as well as from the information that will instantiate the structures defined in the first two phases.
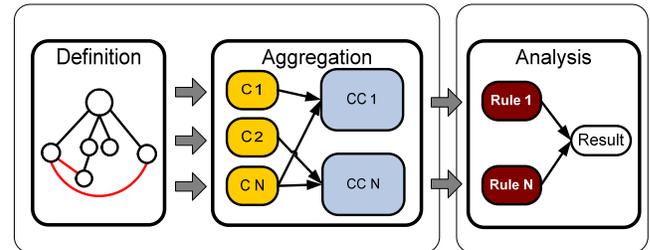

Fig. 1. Competency modeling phases.

The rationale behind separating the definition of a competency from its hierarchy is that it is not possible to anticipate the needs of an organization in terms of its competency categorization. If the definition and the corresponding hierarchy were entwined, then the competency would be defined in the scope of a single organizational context, making difficult specifying the same com-petency across multiple contexts. This is often the case where different organizational units use different classification schemes for the same competency. Thus, this separation aims at maximiz-ing the reusability and flexibility of the structure. The next three subsections detail each of the phases.

3.1 Phase 1 – Competency Definition

Defining a competency requires, first, specifying the concepts that are necessary for its description and, second, specifying the struc-ture that will hold these concepts and their relationships. Compe-tencies classify human actors according to their ability of per-forming tasks in a specific environment. They are the manifesta-tion of knowledge attained through the performance of an action. Creating a structure to relate competencies is not straightforward since it requires modeling a potential large set of disparate indi-vidual features that influence the performance of a given task,

such as knowledge, practical experience, psychological and social factors, context and motivation.

Before detailing how competencies are defined, it is important to disambiguate some fundamental concepts: skill, knowledge and competency. A *skill* is the ability to perform a particular activity, while *knowledge* is related to the information needed for enabling the skill. The improvement of skills and knowledge obtained through experience, training or education leads to *occupational competency* [19]. This means a single competency may encompass multiple skills. A *competency* relates to the behavior that individuals must perform as input into a situation while achieving a certain level of performance. It represents the association of knowledge and action. This means, on the one hand, that competencies are based on knowledge, and, on the other, that for a competency to be recognized as such, it is necessary to demonstrate the capability of giving use to that knowledge through an action that brings value to the related task.

A competency is an expression that is defined through the aggregation of elementary concepts, namely *nouns* and *verbs*. Examples of such expressions are "coding a breadth-first search algorithm in Lisp" and "coding a web service in C#". It is expected that the same elementary concept is often used on the definition of unrelated competencies. The concepts may also be related or functionally dependent (e.g. Lisp and C# are both programming languages although dealing with different programming paradigms). To deal with the representation of such features, we define a set of layered hierarchic structures that specify the competency's concepts regardless of their usage context. This type of structure allows concepts to be uniformly reused while defining different isolated competencies.

However, actors and activities relate to a number of competencies whose meaning depends on their usage context. Additionally, these individual competencies may be structurally related. To cope with this aspect, competencies are aggregated as coherent groups comprising individual competencies. Separating the aggregation of competencies from their definition enables their flexible rearrangement according to usage context without disrupting the corresponding definition. Groups of aggregated competencies can then be bound to organizational actors and activities, thus specifying the supply and demand of competencies within the organization.

3.1.1 Structuring Competencies

Competencies are often structured hierarchically to offset representation simplicity and expressiveness. However, most approaches rely on static classification schemes that prove difficult to adjust to the organization dynamics and its environment. Furthermore, hierarchical representations are defined in the scope of some functional context, so it is not straightforward to detach a structure from its context and keep its semantics. Therefore, as it is not reasonable to assume that functional contexts are shared even within the same organizational unit, let alone throughout the whole organization, the representation structure must be able to clearly separate competency definition from its categorization.

It is hard to tell apart the attributes that are intrinsic to actors and those that are context-dependent. Nevertheless, even if this separation is possible, a problem remains when identifying the granularity of the structure's elements. This problem surfaces because elementary competencies are not defined as unique throughout the

organization and are not structured and designed as atomic components, and, therefore, they are not easily composed nor reused.

To overcome these issues, we propose using a multi-dimensional tree structure, where each dimension relates the element with its usage context. This allows creating multiple views on the same object while keeping its uniqueness. Such structure can be perceived as a regular hierarchical tree where a number of contextual layers can be superimposed.

Nodes in the same dimension or layer are connected using intra-level links allowing relating elementary concepts. Inter-level links associate nodes from different dimensions allowing the specification of complex concepts. The semantics of this type of link is defined by ontological relations, such as "belong-to", "uses" or "enables."
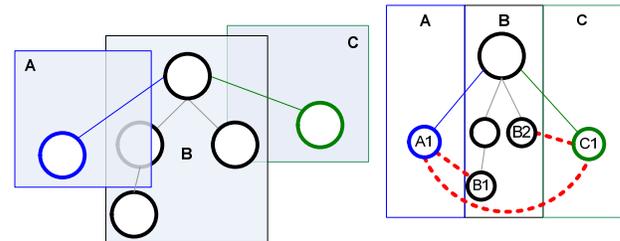
Fig. 2. Elementary concept structure (left). Competency definition (right).

The diagram in Fig. 2 depicts three layers (A, B, C), each defining elementary concepts in a given context. Using inter-level links (depicted as dotted arcs) to associate a number of elementary concept nodes in different contexts enables a competency defined as $\{\langle A1, B1\rangle, \langle A1, C1\rangle, \langle B2, C1\rangle\}$ to be specified.

To exemplify the proposed framework, we will model a number competencies related to software development. Fig. 3 shows a number of elementary concepts representing verbs and nouns.
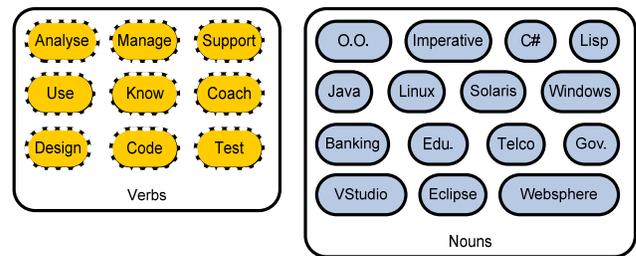
Fig. 3. Elementary concepts.

Competencies are defined by structurally relating elementary concepts using inter-level links (dotted arcs). Fig. 4 depicts the following set of competencies as a result of relating a subset of concepts:

§ Develop Software in the Banking Business Domain

§ Software Design using Imperative and Object Oriented Paradigms

§ Code Software in C# and Lisp Programming Languages

§ Know Java Programming Language

§ Administer Linux Operating System

§ Use any Operating System
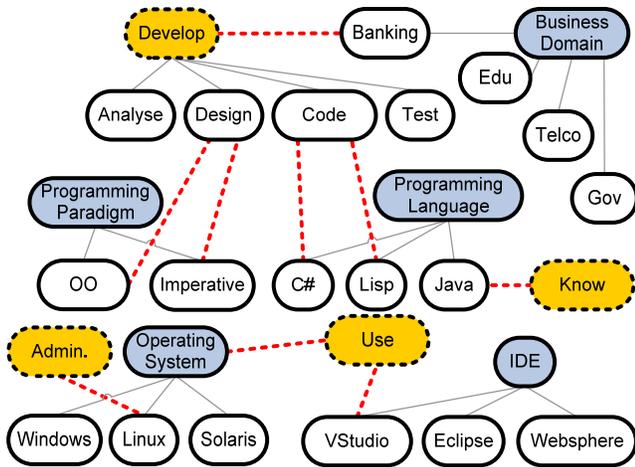
§ Use the Visual Studio IDE

Fig. 4. Competency definition.

3.2 Phase 2 – Competency Aggregation

Competency aggregation allows accommodating multiple contexts through the specification of different viewpoints. After the definition phase, each competency can be composed into a number of coherent sets so they can be handled and related as a unit. For instance, the meaning of the competency "programming" can change according to the perspective. One definition could be the knowledge of a given programming language. However, this definition could also be extended to take it as the aggregate knowledge in programming paradigms and programming languages. These two contexts can be modeled by specifying distinct viewpoints.

Fig.6 depicts an example of competency aggregation where four hierarchically aggregated competencies, C_1, C_2, C_3 and C_4 are described. Each of these individual competencies arise from the association of multiple concept nodes defined in the previous phase (e.g. C_1 is defined by the tuple $\langle A_{1,1}, B_{1,1} \rangle$).
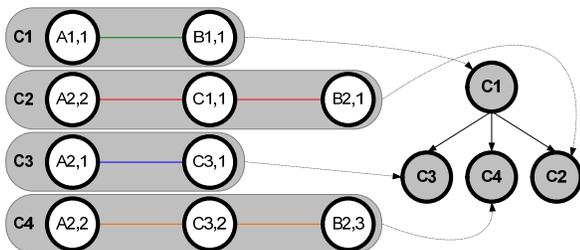

Fig. 5. Generic competency aggregation.

Different observations can be taken from the value of a competency element by changing the hierarchy while using the same set of inputs. It becomes necessary to ensure the independence between the viewpoint chosen to contemplate the individual's competencies in a given context and the competencies themselves. The latter cannot be dependent on the context and categorization chosen by the organizational unit: they are part of the individual and exist regardless of context. They can be adjusted and reorganized according to the needs of the situation, being invariant only if taken from a static perspective.

Therefore, we propose using two layers: one featuring objective information on the competencies and other with their classification and context information. This way, the hierarchy of upper

level nodes corresponds to categorization and views, allowing the creation of other structures that can reuse previously defined components. This approach is flexible since it enables the creation of a semantic level above elementary competencies by separating both and making individual's competencies context invariant.

Other aspect has to do with granularity: a monolithic structure is inadequate to representing competencies, so it is useful to create smaller trees, particularly in the case of straightforward relationships between nodes that arise from their own definition as concepts. As an example, take imperative, object-oriented and functional while programming paradigms. This classification can be considered context insensitive. These smaller trees are also known as "concept islands" [17], self-contained concept groups from sub-domains of a larger domain. Since these islands are interlaced, it makes sense to include them into multiple trees, which implies a complex network of node relationships. One approach is combining the concept islands into a single tree. This introduces complexity during ontology design. Worse, it limits the expressiveness and flexibility of the structure. Therefore, we propose representing the concept island as multiple trees and keeping them separately in a repository. This approach enables these structures to be used as modular building blocks when composing larger trees. It also supports the flexible rearrangement of blocks, without breaking existing connections, i.e., maintaining the connections traceable from runtime back to the atomic elements that form a competency.
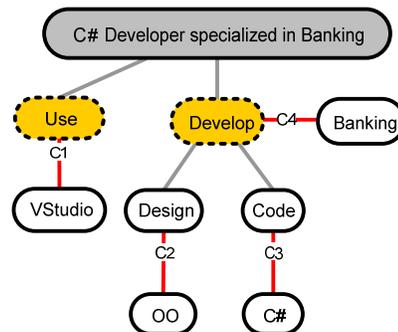

Fig. 6. Competency aggregation.

To illustrate the competency aggregation step, some of the competencies previously defined are now composed, thus specifying a new competency "C# Developer specialized in Banking", as shown in Fig. 6. Note that the original structure does not require revision. Aggregated competencies defined in this step may be used to map roles to competencies, for instance.

3.3 Phase 3 – Competency Analysis

The goal of the analysis phase is instantiating the competency structures according to the competencies actually provided by the organization's human actors. While the first two phases relate to the design of competencies, this phase relates to scheduling actors to activities actors, thus allowing the assessment of competency supply and demand.

Competency analysis requires a set of propagation rules to be defined between different hierarchy levels. Both the rules and the hierarchy itself are part of the classification scheme used by an organizational unit and should not be mistaken for the definition of the competencies. The latter should be independent of the

analysis phase as otherwise competencies will not be able to cross organizational boundaries.

Propagation rules describe how a node's information can be computed from its child nodes. They are defined bottom-up and are related to non-leaf nodes. Leaf nodes are instantiated with values that are propagated up to the root node using the rules on each intermediate node. A rule can include operations, such as logical expressions, weighted averages and threshold definitions. This enables real-time propagation of values, as opposed to static assignment where the connection to the original structure is lost from the moment that the first values are obtained.

When matching competency supply and demand there are important questions about the abstraction level or granularity used on both sides of the negotiation. If the supply is specified in more detail than the demand, simple bottom-up derivation is sufficient to match both sides. However, in the opposite case, it is necessary to request a further specification to the entity responsible for the offer. The updated specification will be included in the existing model as a new layer.
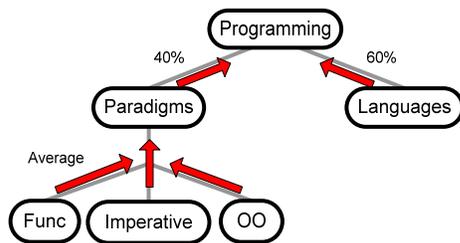

Fig. 7. Example of propagation rules.

While the most visible aspect of inference is performed at competency level, its application in that case is straightforward. The following example focuses on the application of rules to the relationships between elementary competency concepts. Fig. 7 depicts a scenario where a programming competency depends 60% on programming language knowledge and 40% on programming paradigms. As the programming paradigm is not a leaf node, another rule must specify how to compute its value. In this case, a non-weighted average of its descendents has been used.

A fundamental aspect of the proposed framework is the traceability across the three phases. It enables the identification of a specific atomic element that composes an aggregate competency instantiated at runtime. In turn, low-level syntactic translation mechanisms create the conditions to analyze competency elements that would otherwise be obscured by their macro description. The competencies are expandable down to leaf-level, therefore representing the most objective information possible about the organization's human actors. So, competencies can become as independent as desirable from their organizational classification and activity-specific arrangements. As a result, competency movement across different units implies ontological agreement only at the lowest hierarchical level, the most granular syntax of knowledge description.

4. THE COMPETENCY MARKETPLACE

The primary goal of representing competency supply and demand is allowing an organization to find, schedule and manage suitable actors or teams of actors to perform instances of activities. Our approach follows the paradigm of a marketplace transaction. This facilitates the management and tracking of competencies from an organizational perspective while promoting actors to develop their own competencies. It also facilitates correcting and analyzing the gaps between the actual requirements of an activity and those specified in the business process models through the usage of market rules.

The competency market comprises activities like scheduling, planning, recruitment, selection, training and compensation. These activities are usually supervised by the human resource management unit, which can be perceived as a mediator between the activity demand of competencies and actor supply. However, it is important that all parties in the marketplace share a common vocabulary and semantic model so that negotiation is possible. Such shared knowledge should be iteratively built and continuously updated instead of defined *a priori* in a static means. The market dynamics of competency management enable the evaluation of the negotiated work and done on a given activity, offering feedback for process improvement.

The implementation of a competency marketplace involves a specific process, which goes through the stages of Information, Negotiation and Settlement [18]. The Information stage consists in representing actual supply and demand. This is specified at during the Definition and Aggregation phases of the framework. The next step is finding the best candidate that fulfills the requirement of an activity. This is accomplished by searching the supply-side representations using inference and propagation mechanisms on the hierarchical competency structures, instantiated with the concrete proficiency values associated with each actor. The result is a set of actors and the corresponding quantitative evaluations computed from the similarity level between the required and provided competencies. After identifying a potential group of actors able to perform the task, the results must then be categorized according to the existing scheduling plans that track the assignment of actors to activities, thus limiting their availability. The results also need to consider valid business rules that may constrain the scheduling process (e.g. competency A must be 100% provided; actor B and C can not participate in the same activity instance) and the type of results (e.g. the result must be a team of two people). The information resulting from this step may be used to automatically propose a team or to assist the team formation process.

Team formation corresponds to the Negotiation and Settlement phase between the transaction participants. Here, each individual actor, the team as a whole and the activity owner set up the contract details, making explicit the assignment and scheduling terms and conditions. This phase ends when a contract is settled. The contract specifies the team and scheduling details and binds actors to activities through their competencies. Implementing the contract means performing the activity's tasks. Performance can be assessed using metrics whose goal is measuring how the contract is carried out by each individual actor and team.

This approach promotes an actor to be actively involved in the management of her own set of competencies. This influences the criteria of her evaluation and, as a result, how management observes her performance. It is in the best interest of the actor to be properly represented in the model since an accurate description will facilitate her scheduling to the tasks included on her competency pool. It will also enable the identification of gaps inhibiting good performance, calling for training on specific competencies.

5. CONCLUSIONS AND FUTURE WORK

The concept of competency is a fundamental to align actors with activities. It enables an organization to understand and evaluate what is actually required by its processes and what is provided by its actors. It also provides a means to make competencies visible to all process stakeholders, including its performers, designers and owners, allowing mismatches to be identified and continually adjusted. Provided the different stakeholders see the benefit of accurately representing the competencies they are responsible for, the marketplace will converge to a state that represents the actual supply and demand of competencies within the organization.

We have developed a web-based system that implements the concepts outlined in this paper, which is currently being evaluated in order to validate the approach. It allows employees to represent and advertise their competencies and process owners and designers to specify activity requirements and evaluate the performance of actors. The system also supports the dynamic aspects of actor scheduling, allowing for the search of competencies, assisting the process of team formation and evaluating the scheduling results.

Once matured, the competency marketplace will enable the extension of its processes beyond organizational boundaries. While one obvious usage is in the process of e-recruitment, it will also enable new resource management paradigms such as live-recruitment, involving real time resource allocation to serve the needs of ongoing processes. Outsourcing will be greatly improved by competency markets that will span organizational units and even different organizations.

The establishment of trust relationships is crucial and the integration into the Semantic Web will enable the creation of shared repositories of services, much as UDDI does today for Web Services. The unification between human competencies and technological services is subject of our future work. We are currently extending the framework to include the representation of the services provided by non-human actors, including information systems and other business process support systems. The goal is to unify the representation of the services supplied by human and those supplied by system actors. This will enable defining alignment metrics between the organization's service providers and its business processes.

6. REFERENCES

1. A.W. Scheer. *Business Process Modeling*. Springer Verlag, 1999.
2. A. Caetano, M. Zacarias, A. Rito Silva, J. Tribolet. A Role-Based Framework for Business Process Modeling. *Proc. HICSS 38*, IEEE, 2005.
3. Business Process Management Initiative. *Business Process Modeling Notation*. Version 1.0, May 3, 2004.
4. C. Woodruffe. *What is meant by a competency?* Leadership & Organization Development Journal, vol. 14, no. 1, 1993.
5. D. Yimam. *Expert finding systems for organizations: Domain analysis and the Demoir approach.* Beyond Knowledge Management: Sharing Expertise. MIT Press, 2000.
6. F. Stone. *New Definition of Corporate Competencies.* Management Review, vol. 84, no. 6, June 1995.
7. G. Hamel, C. Prahalad. *The core competency of the corporation.* Harvard Business Review, pages 79-91, 1990.
8. L. Spencer, S. Spencer, M. Signe. *Competency at work: models for superior performance.* John Wiley & Sons, Inc, New York, 1993.
9. J. Neves, A. Caetano, A. Vasconcelos, J. Tribolet. Integrating Knowledge Into Business Processes. *Proc. ICEIS 2001*, 2001.
10. M. Ould. Business *Processes, Modeling and Analysis for Reengineering and Improvement.* John Wiley & Sons, 1995.
11. M. Zacarias, A. Caetano, S. Pinto, J. Tribolet. Modeling Contexts for Business Process Oriented Knowledge Support. *Proc. Knowledge Management for Distributed Agile Processes, KMDAP'05*, Springer, 2005.
12. N. Gronau and M. Uslar. Integrating knowledge management and human resources via skill management. *Proc. IKNOW '04*, 2004.
13. R. Johansen, *et al.. Leading Business teams: How Teams Can Use Technology and Group Process Tools to Enhance Performance.* Addison-Wesley, 1991.
14. R. Mayer *et al.. IDEF3 Process Description Capture Method Report*. Knowledge Based Systems Inc., 1995.
15. S. Fletcher. *Competence-based assessment techniques.* Kogan Page, London, 1997.
16. Chartered Institute of Personnel and Development, UK. *Competency and competency frameworks*, 2005.
17. T. Lau and Y. Sure. Introducing ontology-based skills management at a large insurance company. *Proc. Modellierung 2002*, Tutzing, Germany, March 2002.
18. A. Lang and Y. Pigneur. Digital Trade of Human Competencies, *Proc. HICSS 32*, IEEE, 1999.
19. J. Jones and W. Bearley. *Developing 360º Leadership-Feedback Instruments*. Organizational Universe Systems, http://www.improve.org/ 360feedback.html
20. D. McClelland. *Testing for competence rather than for intelligence.* American Psychologist,. Jan;28(1):1-14, 1973.
21. National Institute of Standards and Technology. *Integration definition for function Modeling (IDEF0)*. Federal Information Processing Standards Publication 183. 1993.
22. T. Davenport and L. Prusak. *Working Knowledge: How Organizations Manage What They Know*, Harvard Business School Press: Boston, MA, 1997.
23. A. Abran and J. Moore. *Guide to the Software Engineering Body of Knowledge*. IEEE, 2004.
24. M. Harzallah and M. Lecrère, CommOnCV: Modelling the Competencies Underlying a *Curriculum Vitae. Proc. SEKE'02,* July 15-19, 2002, Ischia, Italy.
25. D. Brickley *et al. RDFS: Resource Description Framework Schema Specification* 1.0, 2000, http://www.w3.org/TR/2000/CR-rdf-schema-20000327/
26. J. Euzenat. *Research challenges and perspectives of the semantic web*. Technical report, EU-NSF Strategic Workshop, 2002, http://www.ercim.org/EU-NSF/semweb.html.