

Landscape Maps for Enterprise Architectures

L. van der Torre¹, M.M. Lankhorst², H. ter Doest², J.T.P. Campschroer³, F. Arbab⁴

¹ University of Luxembourg, Luxembourg

² Telematica Instituut, Enschede, the Netherlands

³ Ordina, the Netherlands

⁴ CWI, Amsterdam, the Netherlands

Abstract. Landscape maps are a technique for visualizing enterprise architectures. They present architectural elements in the form of an easy to understand 2D 'map'. A landscape map view on architectures provides non-technical stakeholders, such as managers, with a high-level overview, without burdening them with technicalities of architectural drawings. In this paper we discuss the use of and techniques for landscape maps. A formal model for landscape maps is introduced as the basis of visualization and interaction techniques. Moreover, we show how a landscape map can be generated from its underlying model. Finally we show several interaction techniques, for example to build a landscape map from scratch, independently of an underlying model, or to change a landscape map together with its underlying model.

1 Introduction to landscape maps

The IEEE 1471-2000 standard [11] promotes the use of viewpoints for architectural description, and it presents as examples the structural, behavioural, physical connect, and the link bit error rate viewpoint. Moreover, to relate to other standards, it includes discussions on the decomposition and allocation, enterprise, information, computational, engineering, and technology viewpoint. Many other viewpoints have been proposed. Also others, such as Finkelstein *et al.* [5], and Lassing *et al.* [9], and Nuseibeh *et al.* [10], have advocated the use of viewpoints for describing architectures.

In this paper we discuss so-called landscape map viewpoints used in decision support of, e.g., information planning. Decision support viewpoints help managers in decision making by offering insight into cross-domain architectural relations. Typically, this is accomplished through projections and intersections of underlying models, but analytical techniques also play a role in construction of landscape maps. Such manipulations of architectural models typically result in lists, tables, matrices and reports. As such, decision support viewpoints create high-level, coherent overviews of enterprise architectures, providing the 'big picture' required by decision makers.

Landscape map viewpoints are used for example to publish an overview for managers and process or system owners, or they are employed by architects as a convenient tool for the analysis of changes or to find patterns in the allocation of resources. A landscape map, as defined by Van der Sanden and Sturm [14], is a matrix that depicts a three-dimensional coordinate system representing architectural relations. Figure 1 is an example of a landscape map that shows which information systems support the operations of an insurance company. The vertical axis represents the company's business

functions; the horizontal axis shows its insurance products. An application rectangle covering one or more cells means that this particular function/product pair is supported by the application, e.g., contracting of a legal aid insurance is supported by the legal aid backoffice system.

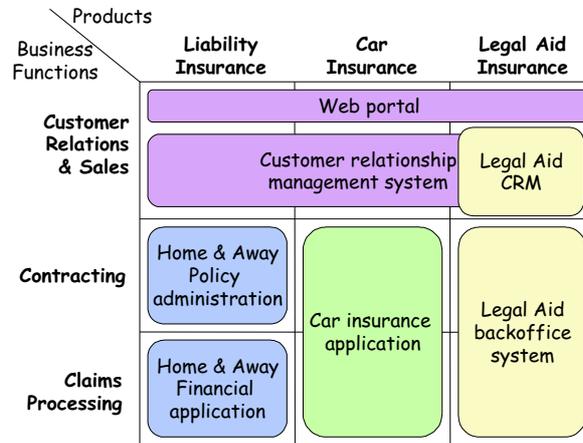


Fig. 1. Example of a landscape map

The dimensions of the landscape maps can be freely chosen from the architecture that is being modeled. In practice, dimensions are often chosen from different architectural domains, for instance business functions, products and applications, etc. In most cases, the vertical axis represents behavior like business processes or functions; the horizontal axis represents ‘cases’ for which those functions or processes must be executed. These ‘cases’ can be different products, services, market segments, or scenarios. The third dimension represented by the cells of the matrix is used for assigning resources like information systems, infrastructure, or human resources. The value of cells can be visualized by means of colored rectangles with text labels.

We propose to use landscape maps as a presentation format (modality) of enterprise architecture models expressed in the ArchiMate language [6, 8]. However, our approach is not restricted to this particular modeling language. In this paper we illustrate how ArchiMate models can be mapped to landscape maps, and how landscape maps can be used as an interactive medium for architecture design. For instance, the landscape map in Figure 1 relates business functions (customer relation and sales, etc.) and products (liability insurance, etc.) to systems (web portal, etc.). The relation between business functions and products is directly supported by the assignment relation. The relation between products and systems is indirectly supported: products are assigned to processes (or functions), which in turn use systems.

The layout of this paper is as follows. In Section 2 we introduce our running example, in Section 3 we give our formal model of landscape maps. In Section 4 we discuss the visual aspects of landscape maps, and we show how a landscape map can be constructed from an underlying model. In Section 5 we discuss interaction with landscape maps, where we distinguish between editing a landscape map from scratch, without reference to an underlying model, and changing the landscape map while simultaneously changing its underlying model as well.

2 Running example

To illustrate the concept of landscape maps, we introduce an example to be used in the remainder of this paper. Our example involves ArchiSurance, an imaginary though reasonably realistic insurance company. ArchiSurance, originally in the business of home and travel insurance, has merged with PRO-FIT (car insurance) and LegallyYours (legal aid). As a result of this merger, the company's main products are now in home, travel, car, liability, and legal aid insurance. To create high-level insight in ArchiSurance's primary operations, the company is described in terms of its main business functions: Customer Relations & Sales, Contracting, Premium Collection, Claims Processing, and Document Processing. Post-merger integration is in full swing. The first step in the integration process has been the creation of a single department for Customer Relations and Sales. However, behind this front office are still three separate back offices:

- Home & Away: this department was the original pre-merger ArchiSurance, responsible for home and travel insurance.
- Car: this department is the core of the old PRO-FIT and handles car insurance, including some legal aid.
- Legal Aid: this is the old LegallyYours, except for the part that has now moved to the Customer Relations & Sales department.

As in many recently merged companies, IT integration is a problem. ArchiSurance wants to move to a single CRM system, separate back-office systems for policy administration and finance, and a single document management system. However, Home & Away still has separate systems for claims handling, premium collection, and payment, and uses the central CRM system and call center. The Car department has its own monolithic system, but uses the central CRM system and call center. The Legal Aid department has its own back- and front office systems (Figure 2).

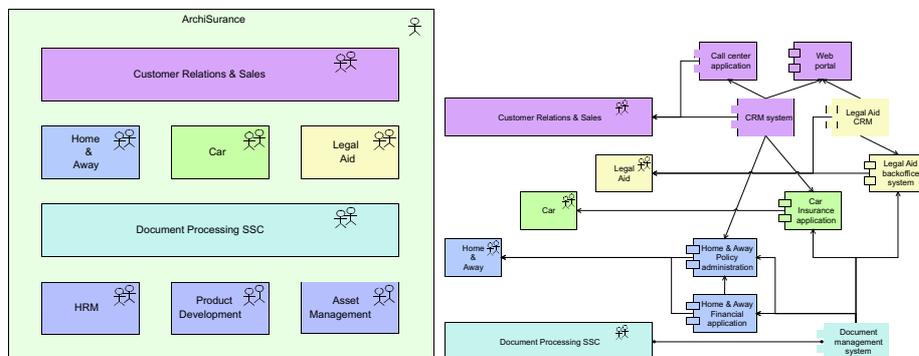


Fig. 2. Actor diagram showing ArchiSurance departments; Applications used by departments

An important prerequisite for the changes in ArchiSurance's IT is that the IT integration should be 'invisible' to ArchiSurance's clients: products and services remain the same. However, this is not a straightforward requirement. To illustrate the complex relationships among organization, products, business processes and IT support, Figure 3

shows the relations among the Damage Claiming process, its IT support, and the organization. Note that this figure shows these relations for only a single business process. In general, many business processes within the back office link the external products and services with the internal systems. E.g., Figure 3 shows the ‘travel insurance’ product, comprised of a number of services realized by different business processes.

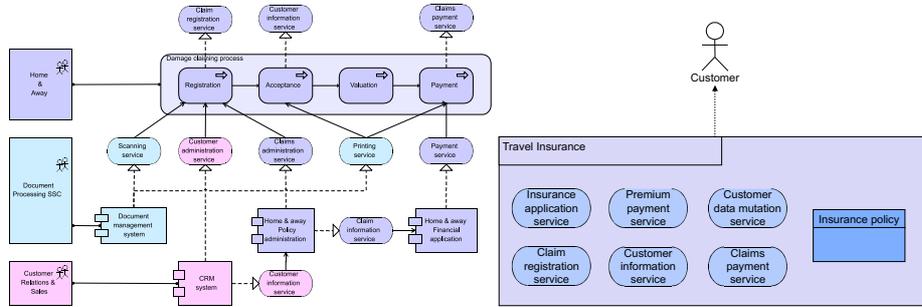


Fig. 3. Relations among Damage Claiming process, its IT support, and the organization; The ‘travel insurance’ product

This web of relations creates a major problem if we want to create insight in the IT support of ArchiSurance. Many systems used by many processes realizing various products and services comprise too much detail to display in a single figure. This is a typical example of where landscape maps can help. As shown in Figure 4, a landscape map of ArchiSurance’s IT applications in relation to its business functions and products provides a high-level overview of the entire IT landscape of the company.

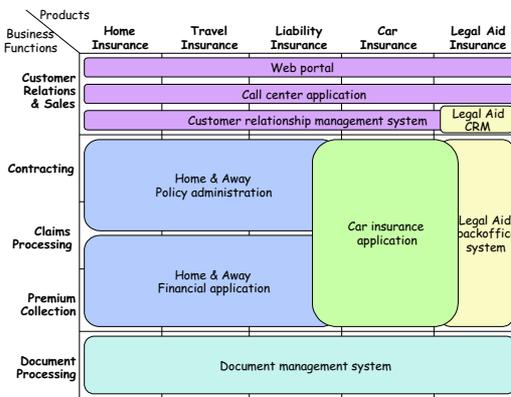


Fig. 4. Landscape map of ArchiSurance

From this figure, it is apparent that there is an overlap in ArchiSurance’s IT support, both in the Car insurance application and in the Legal Aid CRM system. This insight is difficult to obtain from the previous figures. It requires the composition of relations such as ‘product contains business service’, ‘business service is realised by business process’, ‘business process is part of business function’, ‘business process uses application service’, and ‘application service is realised by application’.

3 Definition of landscape maps

In this section we give a formal definition of landscape maps, which is used to facilitate the construction of landscape maps from underlying architectural models, the interaction between the visual part of a landscape map and its underlying model, and the analysis of landscape maps. So far landscape maps have been used as a notation without formal underpinnings. Though visualization and interaction techniques can also be developed without a formal model in the background, we believe that the uniform and abstract view on landscape maps given by our formal definition facilitates their design and realization.

The formal definition of a landscape map is based on a well-defined relation to an underlying architectural model and on our notion of the signature of a landscape map, which in turn is loosely based on the notions of architectural description informally defined in the IEEE 1471-2000 standard document [11]. In the subsection below we explain what we mean by a signature of a landscape map. Next, we distinguish symbolic and semantic models, based on this signature. Finally, we discuss landscape map viewpoints, and the distinction between views and visualizations. This terminology is borrowed from formal methods [15], and the use of this terminology in enterprise architecture is discussed in more detail in [3].

3.1 The signature of a landscape map

Intuitively, architectural descriptions such as landscape maps visualize a set of generic concepts and relations. A concept is interpreted as a set, and elements of this set are called concept instances. Of course, in many cases concepts like function and application in our running example are interpreted as objects, not as sets. In such cases we have to add the constraint that the set is a singleton set. Representing concepts by sets is the most general approach, and applied in most modelling languages.

There is a set of concepts on the X-axis (categories, e.g., products), a set of concepts on the Y-axis (functions, e.g., business functions), and one or more sets of concepts displayed on the matrix (e.g., applications). Moreover, there is a ternary relation that represents a landscape map. Thus, the ArchiSurance example contains the following concepts and one relation. Note that the set of concepts does not make explicit which concepts occur on the X-axis, the Y-axis or on the plane.

$$\begin{aligned}
 C &= C_A \cup C_X \cup C_Y \cup C_Z & R &= \{\text{use}\} \\
 C_A &= \{\text{product, function, application}\} & C_X &= \{\text{home_ins, travel_ins, \dots}\} \\
 C_Y &= \{\text{customer_rel, claim_processing, \dots}\} & C_Z &= \{\text{web_portal, call_center, \dots}\}
 \end{aligned}$$

Furthermore, we assume an implicit ‘is-a’ or subset relation on the concepts, since each element of C_X is a product, each element of C_Y is a function, and each element of C_Z is an application.

$$\text{home_ins} \subseteq \text{product, travel_ins} \subseteq \text{product, \dots}$$

Finally, the relevant information in the landscape map in Figure 4, e.g., that web portal is used by customer relations & sales, but not contracting, can be represented in two ways. First, we may say that the relation is defined on the set of concepts, i.e., use $\subseteq C \times C \times C$. We have product on the X-axis and function on the Y-axis, i.e. use $\subseteq C_X \times C_Y \times C_Z$. Under this interpretation of the relation, the landscape map in Figure 4 can be represented by the following relation.

$$\text{use} = \{ \langle \text{home_ins}, \text{customer_rel}, \text{web_portal} \rangle, \langle \text{home_ins}, \text{customer_rel}, \text{call_center} \rangle, \\ \langle \text{home_ins}, \text{customer_rel}, \text{crm} \rangle, \langle \text{travel_ins}, \text{customer_rel}, \text{web_portal} \rangle, \dots$$

Alternatively, inspired by the notion of signature in formal methods [15, 3], we can interpret relations on concept instances use $\subseteq \text{product} \times \text{function} \times \text{application}$. More precisely, the relation can be interpreted as follows.

$$\text{use} \subseteq \text{home_ins} \times \text{customer_rel} \times \text{web_portal} \cup \text{home_ins} \times \text{customer_rel} \times \text{call_center} \\ \cup \text{home_ins} \times \text{customer_rel} \times \text{crms} \cup \text{travel_ins} \times \text{customer_rel} \times \text{web_portal} \dots$$

In this paper we use the latter option, which is analogous to the notion of typing information in signatures of formal languages [15]. For the techniques developed in this paper both options could be used, but for other techniques such as analysis and simulation of enterprise architectures, the latter option is preferred [3]. The difference is that in the second case, the relation is not defined on the concepts, but on the concept instances. Consequently, the relation between concepts may be called typing information of the landscape map. This is explained in more detail in the following section, when we discuss the interpretation of a signature, and semantic models. A further discussion can also be found in [3].

The combination of a set of concepts together with the pre-defined is-a relation, and the relation together with typing information, is called the signature of a landscape map.

3.2 Semantic and symbolic models

The notion of an architectural model is notoriously ambiguous. The ambiguity becomes clear when we use our formal machinery. First, we have to distinguish between semantic and symbolic (syntactic) models of an architecture. The former are an abstract description of the structure and ‘meaning’ of the architecture itself; the latter are its denotation, i.e., part of the architectural description. This distinction between the architecture and its description is also made in the IEEE 1471 standard.

A semantic model consists of a domain and an interpretation function. In such a model, each concept is interpreted as a set from a domain, which represents that concepts are generic. The ternary landscape map relation is defined on the concepts instances, not on the concepts themselves. For example, the following describes a simple model. Assume that all concepts that are not mentioned are empty sets. In this example, use is defined on instances of concepts like h_1 and c_1 .

$$\text{home_ins} = \{h_1, h_2, h_3\}, \\ \text{customer_rel} = \{c_1, c_2, c_3\},$$

$$\begin{aligned} \text{web_portal} &= \{w_1, w_2, w_3\} \\ \text{use} &= \{\langle h_1, c_1, w_1 \rangle, \langle h_1, c_2, w_3 \rangle\} \end{aligned}$$

Second, the notion of architectural model in the IEEE 1471 standard corresponds to what we call symbolic models, which are logical theories based on a signature. Symbolic models are part of the architectural description, and thus they can describe elements of an architecture. In our formal model, an architectural description is more than just a signature that can be visualized in a view: it also contains constraints and actions which play an important role in analysis and interaction.

Landscape map constraints are logical expressions expressed in terms of the landscape map signature. They further describe the architectures (the semantic models) that fit the signature of the landscape map. For example, there may be constraints that each concept is non-empty, or that concepts are singleton sets. Examples of such constraint languages are OCL for UML [7] and description logics for first-order models [1].

Landscape map actions are descriptions of how a view can be modified, for example due to interaction with the user or as triggered through another view. An action specifies both the interaction dialogue with the user (which kind of information must be obtained from the user when he clicks a button), as well as the consequence of the interaction (e.g., whether and how the underlying model must be modified after interaction with the user).

3.3 Viewpoints, views and visualizations

In the IEEE-1471 standard [11] a view is a representation of a whole system from the perspective of a related set of concerns, that may consist of one or more architectural (i.e., symbolic) models. A viewpoint is a specification of the conventions for constructing and using a view.

In our formal model of landscape maps we abstract away from stakeholders and their concerns, because they are notoriously hard to formalize. Moreover, we distinguish between a view and its visualization. A landscape map viewpoint contains a partial mapping from the signature of the architecture to the landscape map signature, and a landscape map view is the result of applying this mapping to an underlying model. Moreover, the landscape map viewpoint contains a mapping from the view to visual structures.

The visual structure can be formalized in many ways, for example by a signature that expresses in a mathematical way that there is an X-axis, a Y-axis, a plane and rectangles in this plane. Under this formalization, the landscape map viewpoint contains a partial bijective mapping between the signature of the landscape map and the signature of the visualization. Such a mapping is partial because some elements of the architecture will not be visualized and some elements of the visualization (e.g., colors) may have a meaning outside of the model. However, visual structures can also be formalized in other ways, and we do not constrain ourselves to this particular formalization.

Summarizing, in our model we see a landscape map as a composition of a symbolic model and a visual structure, together with a partial bijective mapping of the model on the visual structure. This is less ambiguous than the informal IEEE standard.

4 Visual aspects of landscape maps

The goal of a landscape map is to give an overview of and insight into some architectural relations. In general, a landscape map represents two relations in one map: on the one hand the relation between the entities along the vertical axis and the entities in the plane, and on the other hand the relation between the entities along the horizontal axis and again the entities in the plane. Through the entities in the plane, an indirect relation is established between the two dimensions of the axes. For instance, if an application supports some business function in relation to a certain product, then the business function can be said to support that product.

4.1 The axes

An intuitive and easy to understand choice for the axes is essential for landscape maps to be useful. In the infrastructural approach of Ordina [13], the axes are chosen as follows. The vertical axis represents business functions, i.e., business behavior categorized with regard to results and independent of resource or deployment choices. The horizontal axis represents cases, which still can be specialized to different types of entities. For instance, products or services can be considered as cases business functions add value.

To be useful for managing and designing for change, it is important that the choice of axes is stable, i.e., that the same axes can be used for different usages of the map, and for representing different situations over time. Another requirement is that the choice of axes results in a useful decomposition of the domain. The map is useless if all entries are assigned the same value, or if the matrix becomes sparse in all situations.

The axes themselves allow some freedom in how columns (or rows) are ordered. Sometimes this freedom can be used for arranging columns such that the plane consists of nice rectangular regions like in Figure 4. In other cases, there are semantic constraints, like ordering of business functions in time or an ordering according to priority.

Finally, it is possible to add a hierarchical structure to an axis. For business functions or processes this is an obvious approach to allow more detail in a landscape map.

4.2 The cells

The cells of a landscape map, which are the third dimension of the landscape map, depend on the purpose of the landscape map. If the map is to be used for enterprise application integration, the cells will represent applications or systems; see, for instance, the landscape map in Figure 4. If the purpose is to give insight into the use of data elements, the cells will hold references to data types.

The landscape map in Figure 4 is a view on an underlying model, with its own signature. For example, it may be based on the same sets of concepts, but with two other relations:

$$\begin{aligned}
 C &= C_A \cup C_X \cup C_Y \cup C_Z & R &= \{\text{support, realize}\} \\
 C_A &= \{\text{product, function, application}\} & C_X &= \{\text{home_ins, travel_ins, \dots}\} \\
 C_Y &= \{\text{customer_rel, claim_processing, \dots}\} & C_Z &= \{\text{web_portal, call_center, \dots}\}
 \end{aligned}$$

Figure 2 illustrates how an application supports a business process that spans a number of business functions, and assumes, moreover, that the underlying model specifies how an application realizes a product. Now we have: $\text{support} \subseteq C_Y \times C_Z$ and $\text{realize} \subseteq C_Z \times C_X$. In this particular case, we can directly find the use relation as the product of support and realize. To be precise, combining support and realize leads to a relation $C_Y \times C_Z \times C_X$, so we still have to reshuffle the order of the parameters to find the use relation. Moreover, as explained in Section 3.1, support, realize and use are defined on concept instances, not on concepts themselves. $\text{use} = \text{support} \times \text{realize}$

Note that in this case, because the ternary use relation of the landscape map is constructed as a cross product of two binary relations, its components can always be visualized as a rectangle. That is, if there is an application say z , for (x_1, z, y_1) , and (x_2, z, y_2) , then there are also components for (x_1, z, y_2) and (x_2, z, y_1) .

At the first sight, it may seem that the visualization in Figure 4 of the three dimensional use relation is straightforward. However, a closer inspection reveals that several choices must still be made. First, the three dimensional relation does not specify in which order the items on the axes are presented.

Second, the three dimensional relation does not specify how applications are ordered within a cell of the matrix. Consider for example the top right cell, which visualizes the applications used for Customer Relations & Sales, and Legal Aid Insurance. There are four application components in this cell: Web portal, Call center application, Customer relationship management system, and Legal Aid CRM. However, the relation does not specify that the web portal must be on top, that it is bigger than the two others below, etc. This is what we call visual information, which must be deduced and/or produced by a layout algorithm.

4.3 Automatic layout of landscape maps

An important condition for landscape maps to be effective for problem identification is that the visualization must be intuitive and easy to understand. To a large extent, the choice of the axes and the ordering of the rows and columns determine the layout of a landscape map. If adjacent cells in the plane have the same value assigned, they can be merged to form a single shape. If there are no other criteria for ordering the axes such as time or priority, the ordering can be applied to optimize the layout of shapes the plane, and also to limit their number.

Creating the layout of a landscape map can be seen as a search process. We must define the search space, what it means to have a ‘good’ or ‘nice’ layout, and we must find smart ways to search.

For the search space, a cell that has multiple values assigned can be visualised by multiple combinations of overlaps and ways to split the cell. The input for an automatic layout algorithm is an empty matrix with a per-cell list of values. A cell that has multiple values assigned may be visualised by using overlaps and/or cell splitting. For instance, the top right cell of Figure 4 has four values assigned and is visualised by splitting the cell in three rows of which the bottom one is overlapped by the fourth value. A layout algorithm should be able to derive such a visualisation (semi-)automatically. Examples of rules that can be used to evaluate possible layouts are to minimize the number of objects on the plane, minimize the number of corners on the objects. maximize the

convexity of the objects. make the smallest object as large as possible. We must define also the precedence of the rules in cases of conflict.

For the order of searching the search space, we must find some good heuristics because the search space is huge, already for a simple example as in Figure 4. For instance, assume that a cell has four values assigned, i.e., four applications are used by a particular combination of a product and a business function. The basic layout of this cell can be chosen from 168 possible options:

- 4 overlaps, no cell splitting $\rightarrow 4*3*2$ (top bottom order!)
- 4 subcells, no overlaps $\rightarrow 4*3*2 = 24$
- 2 subcells with 2x2 overlap $\rightarrow 4*3*2 = 24$
- 2 subcells, one with 3 overlaps $\rightarrow 4*3*2 = 24$
- 3 subcells, one with 2 overlaps $\rightarrow 4*3*2 = 24$
- 3 subcells, one overlapping two $\rightarrow 24$
- 3 subcells, one overlapping three $\rightarrow 24$

Clearly, the total number of possible visualizations for a landscape map grows exponentially. Therefore, the search space must be constrained by rules. Some rules to guide the search process are to consider only applications that are allowed by the model, for subcells at the border of a cell, choose applications that also occur in neighboring cells, start with borders, then go for corners, and finally choose centers of cells.

A particular kind of search process works as follows. We first try to find a good initial layout and thereafter we try to improve this initial layout. In this case, it is important to find a good initial layout, since improvement is slow. Here we can use variants of the search heuristics mentioned above. The landscape map of Figure 4 has been generated using the following set of rules:

- Choose applications for borders.
- Choose applications for corners.
- For applications that occur only in one cell, put the application in the center of the cell (e.g., top right corner).
- Fill rows and columns (e.g., top three rows of applications).
- Fill neighbors in a ‘smart’ way.

In the particular case of Figure 4, these construction rules directly yield the presented landscape map without the need for any improvement rules. However, in general there will be room for improvement. Examples of improvement rules are to enlarge one application in a cell, as long as it does not exclude another application from the cell, or to swap two subcells in a cell. The improvements are again measured by the same kind of rules as in the previous case. Now, the additional possibility is not to search the whole search space in an exhaustive way, but to randomly apply the improvement rules, as in evolutionary learning techniques.

5 Interaction with landscape maps

So far, landscape maps have been used as a static one-way presentation format, and landscape map tools contain only editors that allow architects to create landscape maps, with no provisions to relate them to more formal underlying architectural models. We use landscape maps as an interactive medium. Landscape maps are used as a starting point for more detailed models and specifications and they can be used for entering relations between the chosen dimensions. Changes in the landscape map can also be analyzed for impact on other elements of the map.

We have developed new techniques to define interactive landscape maps. In this section we discuss the notion of landscape map action, which has already been defined in Section 3.2 as a description of how a view can be modified, for example due to interaction with the user or as triggered through another view.

We say that landscape map actions create new views and visualizations from existing ones, and can therefore be formally described as mappings between views and their visualizations. At this abstract level, they have something in common with our notion of viewpoint. However, intuitively they are clearly different in important ways, and this is reflected also in our formal definition.

First, when we change the landscape map view we may also have to change the underlying model, and vice versa. For example, consider a stakeholder that works with multiple views at the same time, or multiple stakeholders with multiple viewpoints. In such cases, we visualize the changes directly in all views (with the problem of calculating new visual attributes). On the contrary, when we create a new view from a viewpoint, then existing views do not change. Actions that change the underlying model necessarily have a strict semantics, whereas actions that change only the visualization of a model in a view can be used to make a landscape map more ‘suggestive’, e.g., by using colors and sizes of objects to signify their relative importance. Although important in practice, we do not discuss visual and psychological aspects of landscape maps here.

Second, actions may be interpreted in different ways, depending on the stakeholder and its viewpoint. For example, some stakeholders may change the underlying model, while others may not. For this reason, we represent actions explicitly in views, in the sense that viewpoints or landscape map actions can also modify the landscape map actions. In this, we use an extended notion of ‘view’ compared to the IEEE 1471 standard, in which views only relate to the architecture itself.

Third, actions typically require some interaction with the user, before they can be executed. We therefore extend the notion of action by associating an interaction protocol with it. Thus, an action specifies both the interaction dialogue with the user (which kind of information must be obtained from the user when he clicks a button), as well as the consequence of the interaction (e.g., whether and how the underlying model must be modified after interaction with the user).

In this section we discuss two kinds of interactions. In Section 5.1 we describe the creating and navigating a landscape map. In Section 5.2 we discuss the more complex case in which the underlying model can be changed by editing the landscape map.

5.1 Creating and navigating a landscape map

The first contours of a landscape map are usually drawn on a white-board, flip-over or piece of paper. Together with the stakeholders the architect tries to address their concerns. The map should be such that it concentrates on the choices that must be made. The drawing must also be such that consequences are visible. In this interaction the architect chooses the concepts on the axes and on the plane, the level of detail, leaving out the facts that are less important. For the sake of readability, understandability, and acceptance the architect juggles a little bit with the (unwritten) rules of the landscape map. With pen and paper this can obviously be done.

Back at the desk and using the tool we envisage, the landscape map must be constructed in a more formal way. First, the architect needs to select the type of concepts used on the X-axis, on the Y-axis, and on the plane (see Figure 5). In our ArchiSurance example, the X-axis contains products, the Y-axis signifies business functions, and the plane holds applications.

Next, the objects on these axes must be chosen (the X_1, \dots, X_m and Y_1, \dots, Y_n in the figure). If a landscape map is used to define a new architecture, these objects can be freely chosen (of course conforming to the type of the axes). Alternatively, if an existing architecture model is visualized they may be selected from this model. By choosing the concepts for the axes the field of play is defined.

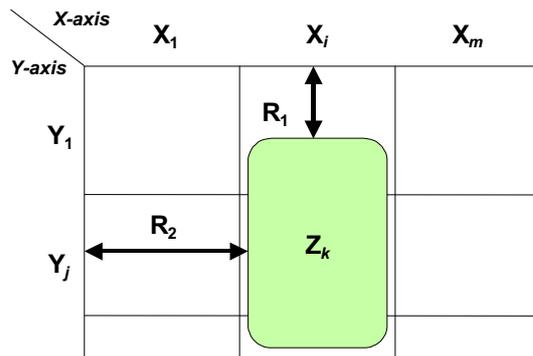


Fig. 5. Elements of a landscape map

After this, the architect must choose the type of assertions that are made by putting an object Z_k somewhere on the plane, i.e., the relations R_1 and R_2 . In our example, he chooses business functions on the vertical axis, products on the horizontal axis and applications on the plane. The most obvious, intuitive assertion is that an application is used by activities required within the business function in realizing the product, giving us $R_1 = \text{support}$ and $R_2 = \text{realize}$. For every object the architect places on the plane, these relations between X_i , Y_j , and Z_k are instantiated.

Furthermore, if the rectangle of the application Z_k is not exactly aligned within a row and/or column, then the relations with the X and Y elements are in a sense

‘incomplete’. For example, an application may deliver only some of the functionality needed to support a business function.

In a similar fashion, an existing landscape map can be used as a starting point for navigation. In this case, relevant interactions include:

- Open a rectangle: detailed specifications or detailed models are shown in a separate window.
- Close detailed specification or detailed models.
- Change granularity of an axis; for instance, business processes can be changed to business activities.
- Link two rectangles by a relation supported by the underlying concept. For instance, if rectangles represent systems, use or composition can be used.

5.2 Changing a landscape map

If an architect or stakeholder wishes to change an existing landscape map, the effects of this change on the underlying architecture model need to be assessed. Some changes may be purely ‘cosmetic’ in nature, e.g., changing the color of an object. Other changes need to be propagated to the underlying model, e.g., if an object is added or deleted.

Mapping a seemingly simple change to the map onto the necessary modifications of the model may become quite complicated. Since a landscape map abstracts from many aspects of the underlying model, such a mapping might be ambiguous: many different modifications to the model might correspond to the same change of the landscape map. Human intervention is required to solve this, but a landscape map tool might suggest where the impact of the change is located.

In the example of Figure 4, the architect may, for instance, want to remove the seemingly redundant Legal aid CRM system by invoking a ‘remove overlap’ operation on this object. This operation influences both the visualization and the architectural model. Figure 6 illustrates the effects of the operation on the underlying model.

First, the architect selects the object to be removed, in this case the Legal Aid CRM system. The envisaged tool colors this object and maps it back onto the underlying object in the architecture model (an element of the set CZ as defined in Section 3.1).

Next, the relations connecting this object to its environment are computed (the second part of Figure 6). Here, this concerns the relations of Legal Aid CRM with the Web portal and the Legal Aid backoffice system. These relations will have to be connected to one or more objects that replace the objects that are to be removed. Since we have chosen a ‘remove overlap’ operation, the landscape tool computes with which other objects Legal Aid CRM overlaps, in this case the CRM system. The relations formerly connecting Legal Aid CRM are then moved to the other CRM system, unless these already exist (e.g., the relation with the Web portal).

Naturally, this scenario presents an ideal situation with minimal user intervention. In reality, a tool cannot always decide how a proposed change is to be mapped back onto the model, and may only present the user with a number of options. For example, if the functionality of the Legal Aid CRM system would overlap with more than one other system, remapping its relations requires knowledge about the correspondence between these relations and the functions realized by these other systems.

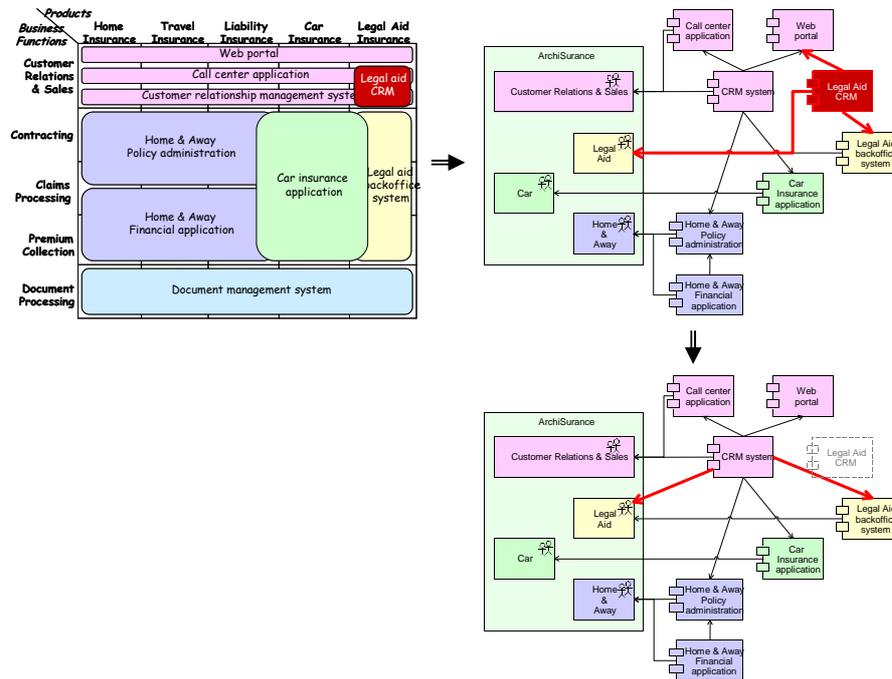


Fig. 6. Editing a landscape map

6 Conclusions and Future Work

In this paper we introduce landscape maps for enterprise architectures, which is an instrument that has proven its value in the architecture and consultancy practice of Ordina. The landscape map is an easy-to-read informative format that provides overview of and insight into architecture relations between different domains. We define a formal semantics for landscape maps based on a mapping to and from the ArchiMate language. We show how landscape maps can be automatically derived from enterprise architectural models and explain how they can be automatically visualised. Finally, we explain how interaction with landscape maps is realised by propagating changes back to the underlying model by means of actions.

The development of landscape maps in this paper carefully balances formalization and informal discussion. Enterprise architectures are often informal, because they must not constrain the architect. However, to provide the architect with useful tools, some formal definitions are necessary. In this paper we only adopt a minimal formalization in terms of the signature of a landscape map, but we do not constrain for example the visual structures, or the landscape map actions. The minimal formalization has been sufficient to define the visualization and interaction techniques.

Based on the formalization of landscape maps presented in this paper, we have developed a prototype that illustrates the added value of the formal semantics of landscape maps and the benefits of interactive landscape maps (based on actions). The prototype

serves as a proof of concept of the formalization of, automatic layout of, and interaction with landscape maps. Furthermore, it illustrates to members of the ArchiMate Forum the tool support that we envisage for enterprise architecture. Some of the ideas have been incorporated by BiZZdesign in their ArchiMate-compliant Architect tool.

The use of viewpoints for architecture-level impact analysis (see, e.g., [13, 4, 2, 12]) is a subject for further research. The type of editing illustrated in this paper tends toward this type of analysis. By propagating changes to the landscape map through the architectural model, a high-level impression of the effects of a change can be obtained. Landscape maps can also be used to visualise the results of other types of analysis. For example, a cost or performance analysis may yield quantitative results. Several visual techniques such as colors, line styles, and fonts can be used to highlight these effects within the landscape map.

References

- [1] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and editors P. F. Patel-Schneider. *Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2002.
- [2] F. de Boer, M. Bonsangue, L. Groenewegen, A. Stam, S. Stevens, and L. van der Torre. Change impact analysis of enterprise architecture. In *Proceedings of IEEE International Conference on Information Reuse and Integration (IRI'05)*, 2005.
- [3] F. de Boer, M. Bonsangue, J. Jacob, A. Stam, and L. van der Torre. A logical viewpoint on architectures. In *Proceedings of EDOC'04*, pages 73–83. IEEE, 2004.
- [4] F. de Boer, M. Bonsangue, J. Jacob, A. Stam, and L. van der Torre. Enterprise architecture analysis with xml. In *Proceedings of 38th Hawaii international conference on system sciences (HICSS'05)*, 2005.
- [5] A. Finkelstein, J. Kramer, B. Nuseibeh, L. Finkelstein, and M. Goedicke. Viewpoints: A framework for multiple perspectives in system development. *International Journal of Software Engineering and Knowledge Engineering*, 2(1):31–57, 1992.
- [6] H. Jonkers, M.M. Lankhorst, R. van Buuren, S. Hoppenbrouwers, M. Bonsangue, and L. van der Torre. Concepts for modelling enterprise architectures. *International Journal of Cooperative Information Systems*, 13(3):257–287, 2004.
- [7] A. Kleppe and J. Warmer. The object constraint language and its application in the UML metamodel. In *Proceedings UML'98 Beyond the Notation*, Mullhouse, France, 1998.
- [8] M. Lankhorst *et al.* *Enterprise Architecture At Work*. Springer, 2005.
- [9] N. Lassing, D. Rijsenbrij, and H. van Vliet. Viewpoints on modifiability. *International Journal of Software Engineering and Knowledge Engineering*, 11(4):453–478, 2001.
- [10] B. Nuseibeh, S. Easterbrook, and A. Russo. Making inconsistency respectable in software development. *Journal of Systems and Software*, 56(11), 2001.
- [11] IEEE Computer Society. *IEEE Std 1472-2000: IEEE Recommended Practice for Architectural Description of Software-Intensive Systems*. 2000.
- [12] A. Stam, J. Jacob, F. de Boer, M. Bonsangue, and L. van der Torre. Using xml transformations for enterprise architectures. In *Proceedings of ISOLA'04*, 2004.
- [13] W.A.M. van der Sanden, P. Bergman, J.T.P. Campschroer, and H.R. de Reus. Realisatie van flexibele informatievoorziening (in dutch). *Informatie*, 41:58–65, 1999.
- [14] W.A.M. van der Sanden and B.J.A.M. Sturm. Informatiearchitectuur, de infrastructurale benadering (in dutch). 2000.
- [15] J. van Leeuwen, Ed. *Handbook of Theoretical Computer Science, vol. B: Formal Methods and Semantics*. Amsterdam, 1994.