

A repository-based framework for evolutionary software development

Martine Devos, Director Information Systems Departement (mdevos@argo.be)

Michel Tilman, System Architect Unisys (mtilman@argo.be)

Abstract

This paper describes a framework developed to support the Argo administration. Argo is a semi-government organization managing several hundred public schools. It uses this framework to develop its applications, which share a common business model, and require database, electronic document, workflow and Internet functionality.

The framework is based on a repository in two ways. First, it consists of a set of tools for managing a repository of documents, data and processes, including their history and status. These tools enable users to select applications, enter and view data, query the repository, access the thesaurus and manage electronic documents, workflow processes, and task assignments. More importantly, the framework behavior is driven by the repository. The repository captures formal and informal knowledge of the business model. None of the business model is hard coded. The tools consult the repository at runtime. End-users can change the behavior of the system by using high-level tools to change the business model. Thus we separate descriptions of an organization's business logic from the application functionality.

The framework helps us develop applications iterative and interactively. It enables us to build increasingly complete specifications of applications that can be executed at once. The resulting system is easy to adapt to changes in the business. Argo can develop new applications through modeling and configuration, rather than through coding.

X.1. Context

Argo is a semi-government organization that manages the public schools (non-denominational) within the Flemish community in Belgium. It consists of a central administration and a few hundred semi-autonomous local sites (boards and schools).

The project started early 1994. At the start, Argo was going through an external audit, which made it likely that there would be major changes of procedures, organization structure, accountability rules and delegations of responsibility. Future requests for changes were a given at the start of the project.

So Argo needed a framework to build applications, sufficiently flexible to cope with changing and emerging needs and opportunities, both functional and technological, imposed or desired. It wanted an environment that not only would survive the expected reorganization and decentralization but that would stimulate users to envision better ways of working.

We started with three pilot applications to provide input for initial framework requirements and design:

1. A documentation center that provided flexible search and retrieval of data and documents such as legislative texts and parliamentary decrees concerning education. Many documents have a complex structure of interrelationships like appendixes, references to subsequent changes and juridical decrees based on multiple laws.
2. A system to support the central board's decision procedures from initial draft version to the final text, and then to monitor the decisions' implementation by the administration or the local boards and schools. These decisions include motions that have been passed or tabled, contracts with teachers and changes to curricula and budget reports. The path through the administration is not clear-cut and it changes depending on parties involved or on the subject of the decision. The procedures are not always limited to a strict hierarchical structure; some involve temporary workgroups or external committees.
3. An application to import, index, store and route large volumes of documents sent in by the local boards, such as attendance records, decisions, meeting notes and large addenda.

Halfway through the project the pending re-organization was effectively carried out. Although the nature and extent of it were largely unknown at the onset, we were able to adapt the existing applications through re-configuration without additional coding.

Since then, several new end-user applications have been delivered. We are now in the process of extending the framework; most notably to give schools and local boards access to the central data and information through the Internet. We are converting large existing relational database applications and we added web functionality and an Internet discussion system.

These applications, together with new technological possibilities, are our major sources of requirements for the evolving framework. Within this process we further refined and re-designed many framework components. The persistency component, for instance, has been completely re-implemented. In addition, we created some applications to support the development process, such as a bug reporting system with follow-up, training-session management and framework documentation management.

Not only do we build end-user applications with data, document, workflow and Internet functionality. We also use the same framework to develop the tools that we use to build and manage these applications. In this way users can use their interface, known from their business objects, to adapt applications.

End-users at Argo are gradually taking over management of applications and configuration from the development team. Key-users teach and write parts of the documentation. User groups organize workshops to discuss how to put the technology to better use. This process was initially driven by small teams of highly motivated pioneer-users, and has gradually started to embrace the end-user community at large. Given the high degree of computer-illiteracy within Argo at the start of the project, we feel this to be a significant achievement on behalf of the end-users.

Project requirements

X.1.1. Functional

The technology has to support applications containing a mixture of database (e.g. school database, patrimonies, personnel, budget, inventory); electronic documents (e.g. educational documentation, research papers, articles, laws and decrees); workflow management (e.g. task assignments and follow-up, decision routing, process management); and has to give local sites access to the central repository of data and documents.

Database

Users have to be able to enter, modify, remove, query, list, browse, report, print, import and export data. They need tools to set up appropriate views, and to store queries for later re-use.

Electronic document management

Users need to create, index, search, edit, annotate, print, export and process documents. These can be either electronic (such as files or E-mail) or scanned paper documents. Structured information, thesaurus keywords and full-text-indexing are used for indexing and searching. Optical character recognition enables users to extract text from scanned documents. Users need tools to manage document versions. They must be able to choose the most appropriate alternative representations of the same document, depending on environment or job at hand. A document created and adapted with a word processor must be frozen once it has passed the decision process and can be presented in PDF format to schools.

Workflow management

Users must be able to plan and handle incoming tasks, to keep track of outgoing task assignments and to manage workflow processes, which can be more or less well defined or completely ad-hoc. Argo wants workflow functionality to suggest opportunities to its users, guiding rather than restricting them.

Alternative access modes

In addition to the applications running in client / server mode at the central site, Argo needs alternative ways of accessing its applications, data, documents and processes in the central repository:

- using an off-line version of the system so that employees can occasionally take work home for specific tasks like preparing a proposal for the central board;
- using the system on-line by means of a direct dial-up connection, e.g. modem or ISDN for board members and employees working remote;
- using E-mail and Web browsers to access the repository through the Internet for schools and teachers.

Infrastructure

The applications have to operate in a Windows - Novell environment. The database is an Oracle server. Internet mail is used for external communication, cc:Mail for internal communication. The framework has to be developed in VisualWorks\Smalltalk.

X.1.2. Support for change

To cater for changing needs, the system has to forego hard coding as much as possible: the technology has to support iterative and incremental development through the use of prototypes. Even more, it must act as a catalyst for a learning process when introducing this new technology, thus effectively complementing a Business Process Re-engineering project at Argo.

X.1.3. Integrated environment

The applications must be integrated in one environment. All applications must use a common business model. This model provides the cohesion expected in a large administration, providing formal procedures and policies (e.g. hierarchy, responsibilities, indexing vocabulary of documentation, and global validation- and authorization rules), but giving ample scope and freedom to adapt the technology to the most appropriate local or personal working practices. For some legal activities all departments have to abide by strict rules requiring a prescribed number of signatures, while less formal processes are in use for more everyday tasks.

X.1.4. Configuration and administration tools

There must be tools to define the object model, constraints and behavior, to configure and manage applications, views and stored queries, to define and manage workflow processes, to set up access control, to manage the database, and to configure caching and back-up of documents.

At the end of the project, Argo's users need to be able to apply these tools. Hence the deliveries do not only include end-user applications, but also the necessary tools to empower users to model and configure applications.

X.2. Using the framework

We give a short overview of a documentation center application, a decision procedure application and an incoming mail registration application. We will illustrate how the user applies the various tools to build and configure applications.

X.2.1. Examples

X.2.1.1. Documentation center

The documentation center applications let users manage and search authors, publishers, and several types of books, periodicals, editions and other kinds of documentation, including the framework documentation. Librarians use one application to import, index and administrate the documents. Regular end-users use a simpler application to search for the documents they want to consult.

Some documentation exists only on paper. In this case we maintain references to the physical location. Most documentation is made accessible on-line, in original format or as scanned documents. Part of the documentation is made available on the Internet, in an alternative representation, such as PDF.

For indexing and querying we rely extensively on thesaurus and full-text-indexing. Librarians manage distribution lists to notify users automatically about new documentation. They also keep track of loans. In addition to the standard printing facilities, librarians print documentation in ISBD-specific format.

The object model is complex. To reduce complexity and make the application more dynamic, we use some explicitly typed objects and associations. For instance, several types of documents share the same attributes and associations, but require different query and forms screens. We use the Type Object pattern for increased flexibility. Similarly, we model several types of associations by one, attributed association. The attribute dynamically types the association, representing for instance “replaces” and “is an addendum to” associations between documentation.

X.2.1.2. Workflow processes

Argo does not use workflow as a means to single-mindedly automate existing processes. Too many exceptions occur in the office to prescribe the flow in most cases. Although some processes follow a strict plan, in general, users need the freedom to deviate from this plan. And in some cases there is no fixed plan.

The central board decision procedure formally describes the process for submitting a dossier for approval by the central board, starting from a department’s manager. The preparation phase prior to this process depends on each department’s culture, which in some cases is very team-oriented and in other cases is inspired by a hierarchical structure. The general manager has the right to advise on the dossier, but he is not part of the actual process flow. To keep him informed of upcoming dossiers, notifications are sent when the process enters a particular phase. When decisions result in a concrete action plan, management needs to check when and if the goals are actually met, by monitoring the state of specific data, documents and processes.

A typical ad-hoc process is the mail registration application for handling incoming (paper) mail. The users register these documents in the system, setting properties such as date, reference of sender, subject, summary and keywords using Argo’s thesaurus. They indicate if the information is confidential. The actual (paper or electronic) documents may be stored in the system using native Windows documents, e.g. HTML, PDF and scanned TIFF. The documents are routed to the persons who have the authority to assign tasks and set priorities and deadlines. Depending on the nature of the documents, they are routed through the organization for further handling. The actual process flows depend on the contents of the documents. The recipients of the tasks are notified by E-mail. They change the status of a task when they complete it or delegate it to someone else. The employees responsible for the dossiers keep track of the status of the tasks and report to superiors if needed.

X.2.2. Building applications

Starting from the model in the repository at a certain moment, we build applications by going through a set of steps. These steps can be performed in any order, even interactively. The development process unfolds incrementally. We can try out and correct ideas dynamically, together with our users.

An end-user application usually starts by extending an existing model. Developing an application consists of:

- extending (and possibly re-factoring) the object model;
- setting up application environments and defining window lay-outs;

- extending or refining authorization rules;
- defining action rules;
- defining workflow process templates.

This process is illustrated in the following picture. Solid arrows denote the products of the tools; other arrows denote dependencies of the tools.

Figure X.1 Building applications

X.2.2.1. Extending the object model

With the object model editor we define the object model: object and association types and basic constraints, such as uniqueness of value and totality of associations.

Object behavior and user-defined constraints may be added now or later on, e.g. to ensure correct initialization. Object behavior and constraints are common to all applications.

The object model editor defines the super-types and subtypes of a new object type. It also defines the properties of a type, which can be either attributes, containing e.g. strings or dates, or references to associated objects.

When we have defined the model, we generate the necessary database structures. We select the mapping strategies for representing objects in the database. These strategies may be mixed, and even changed later on, to optimize the database mapping and to ease integration of existing databases.

Figure X.2 Object model editor

X.2.2.2. Setting up the application environment

This step starts with defining a filter on the shared business model: objects and properties to be accessed, created and queried within the context of an application. Once this basic environment has been stored in the repository and access privileges have been set, the user gets a default application that is available for immediate use. The user can log on, choose this environment, enter data, query the repository, list or print results. He can keep track of task assignments, manage electronic documents and access the thesaurus. Functionality can be added or removed. The environment can be further refined for individual or shared use through views, queries and sub-environments.

In the example below we define the environment for the mail registration (“Indicateur”) application. We select the object types to be used in the application environment. For each object type we indicate accessible properties and how they are used and created.

Figure X.3 Application editor – definition of the mail registration application environment

X.2.2.3. Using the default application

Having defined the application environment, the user has access to a fully functional default application by logging on in the main application window. What the user perceives as an application is essentially determined by the objects he can manage and the functionality of the tools we provide, filtered by the application environment. Thus we always present the user with a similar set of tools, but these tools adapt themselves to the specifications of the application.

The main window tool (here on the default mail registration application) presents the user with the query editor and an overview list. It also provides access to other functionality, such as additional query and list windows, forms, document management, the thesaurus browser, workflow processes, in- / out-baskets, preference settings and on-line help.

Queries can be saved for individual or shared re-use within the application. Each query can have its own list lay-out. This list lay-out is selected automatically when the query is executed.

Figure X.4 Default mail registration application

Users create, view, edit, print, export and browse objects through the form tools. Each form is generated on the fly, according to a default strategy, subject to the authorization rules. The form generator selects appropriate editors for the individual properties and the object type at hand. In this figure we opened a form on the two elements selected in the overview list.

Figure X.5 Forms and document viewer

X.2.2.4. Defining window lay-outs

Lay-out editors enable users to define new forms and overview lists and change old ones. The form lay-out editor presents the available editors for each type of property. For instance, the default association editors may be replaced with embedded list views and forms, enabling re-use of existing lay-outs. In the list lay-out editor, hierarchical relationships can be used to define threaded list views. Furthermore, association chains can be followed when defining the list of properties to be displayed, allowing users to de-normalize views. Lay-outs can be defined as the default lay-out for a particular object type within a given application. In that case they will be selected automatically in overview lists and forms for the corresponding object type.

Figure X.6 Form lay-out editor

X.2.2.5. Extending or refining authorization rules

Access control is based on a set of authorization rules. These are not tied to any specific business model and can be used to set up both very fine- and course-grained access. Authorization rules can be context- and contents-sensitive. Some rules apply to everyone; other rules apply to groups of people whose members are determined at run-time. This way, authorization rules support cohesion, while giving considerable flexibility in supporting different team cultures.

Authorization rules contain four parts. The agent (= whom we are granting or denying access) and object (= what objects the rule is about) conditions are defined using the query editor component.

The aspect part specifies what aspects (e.g. properties) the rule applies to. The rights part denotes the privileges (e.g. read and export).

Figure X.7 Authorization editor

X.2.2.6. Defining action rules

These rules capture business-semantics, set defaults, perform extra validation, impose constraints and add functionality. Action rules can cover an entire organization, particular functions or they can be limited to a specific application environment. Hence action rules give extensive support to autonomous teams while preserving overall consistency. Action rules come in two flavors: high-level template rules for defining workflow processes and script rules for handling more specific cases.

Figure X.8 Script rules editor

X.2.2.7. Defining workflow process templates

Workflow process templates specify default scenarios that the user typically follows in the context of a business process. Users can (usually) deviate from these scenarios and snap back into the pre-defined flow later on. Processes can be configured incrementally, e.g. between departments first, and within each department later on. Depending on team culture, these individual sub-processes are more or less strictly defined. In addition, there are automated tasks and dedicated private or shared in- / out-baskets to manage incoming and outgoing work.

Figure X.9 Workflow editor

X.3. Our framework approach

We aim to combine the high-level modeling power of CASE-tools with the open-ended nature of object-oriented frameworks. Our approach is based on the following observations:

- Users tend to describe their requirements in terms of how they currently work. It is difficult for them to picture how new technology can or will influence their future working practices. Hence they need to try out new software, to learn from it and to correct it. This requires extensive prototyping. Too often, however, prototypes are used to validate the design, rather than help users and developers capture and explore the real requirements. Prototypes can be expensive, as they usually have to be thrown away. This is particularly true when developing many applications, each one requiring several iterations to get the requirements right.
- A business model is particular to each organization at a particular point in time [Sto94]. However, functional requirements, such as data-entry, querying, reporting and document management, are generic.
- Extensions or enhancements to functionality can often be made into re-usable assets, independent of the actual business model, and vice versa.

Architecture

The core of the architecture is a three-level repository that is consulted dynamically by two sets of tools.

X.3.1. Business objects repository and end-user tools

The framework uses the first repository to capture knowledge of a particular business model, including documents, data, state of processes, organization structure and the roles each person can play.

We provide the user with tools to access this repository. These tools enable him to select applications, to enter and view data, to query the repository, to display, print and export the results of a query, to access the thesaurus, to manage electronic documents, workflow processes and task assignments.

This way we deliver the functionality demanded by the end-user: applications for database, document management and workflow and access to all these over the Internet.

Figure X.10 Repository – Business objects

X.3.2. Configuration tools and meta-information

We want the end-user to be able to tailor, configure and manage his applications. Therefore we provide a second set of high-level tools. These tools define and manage the specifications of end-user applications, such as: the structure and behavior of documents, data, processes, organization and roles; private and shared views on data; pre-defined queries for later re-use; authorization and business rules; additional end-user functionality.

We store this meta-information in a second (meta-)repository.

Figure X.11 Meta-repository

We do not generate code for the end-user applications. Both sets of tools consult the meta-repository at run time. Changes made to the meta-information, e.g. object structure and views, are immediately available to clients.

Example: We use the meta-information to dynamically generate query-screens and forms. This process is based on the object-model. Because we store the meta-information about the objects in the repository, the screens in the end-user tools will adapt to changes in the object-model, e.g. to addition or removal of properties.

All tools are fully operational and provide out of the box both end-user functionality together with the means to configure and fine-tune applications. They only need the meta-information to get the system up and running.

X.3.3. Secondary goals

To make these tools easily accessible to users, we give them the same interface and functionality as the regular end-user applications. Ideal would be to have one single set of tools, thus enhancing reuse and ease of maintenance.

Support of a bootstrapping process became an important secondary design goal: administration and configuration tools are gradually expressed in terms of the framework itself.

We replace most hard-wired administration and configuration tools with configured applications. This leads to a small but powerful kernel of generic and orthogonal tools.

The meta-model used to describe the object model, is expressed in terms of itself and is stored in the repository too. This is a third-level repository, a meta-meta-repository. This started as an experiment to push the limit of flexibility, but proved to be useful. It enables us to extend the semantics of an object model.

Example: In the first version of the framework the semantics of cascaded deletes were hard-wired. Now we define and store rules in the repository that overwrite built-in default behavior. More intelligent delete-rules can now be adapted to cope for particular needs in an application.

In practice the three repositories are one and the same, managed with the same tools.

Figure X.12 Repository

X.4. Framework components

In this section we look into more detail at some key elements stored in the repository: object model, meta-model, system objects, application environments, electronic documents, in / out baskets, workflow process templates, authorization rules, action rules, background processing rules.

We borrow some terminology from the GOF design pattern system [Gam95].

X.4.1. Example - the mail registration application

The mail registration application is used to register information about incoming documents and dossiers, to assign tasks for handling them and to keep track of their status.

When the user receives a document he creates an instance of the `IndicateurDocument` object type and sets the appropriate attribute and association properties, including thesaurus keywords. Some properties, e.g. the identity of the employee who treats the document and the creation and modification date, are given default values. The actual document may be stored in the system as an electronic document linked to the `IndicateurDocument` object.

Handling mail requires several types of objects: `IndicateurDocument`, `ElectronicDocument`, `Employee`, `OrganizationalUnit` (department), `ThesaurusDescriptor` (taxonomy of keywords and synonyms). We do not write code to represent these objects but model them in the repository. Nor do we code the relationships between these objects. Association types like “`IndicateurDocument` created by `Employee`” and “`Employee` member of `OrganizationalUnit`” are stored in the repository along with their cardinality constraints e.g. “every `IndicateurDocument` is created by exactly one `Employee`” and constraints on attributes e.g. “the name of the `Employee` is mandatory”.

We re-use objects from existing applications, e.g. Employee, ThesaurusDescriptor (managed within the documentation application), ElectronicDocument. The IndicateurDocument inherits common attributes from an abstract type (WorkflowDocument).

Figure X.13 Simplified diagram of object model

X.4.2. Repository

Hotspots should be easy to change. The framework stores hotspot objects in the central repository so that they can be changed without any programming. They can be changed by end-users, configurators or administrators, rather than by developers.

X.4.2.1. Object model

The repository contains three kinds of objects. The first kind are end-user objects, which are the objects that make up the business model. These are the objects that the end-user talks about. The second kind are system objects, which are objects that the developers or configurators talk about, but that are not part of the business model. For example, authorization rules and the internal structure of electronic documents, such as versions and various representations, are system objects. The third kind are the objects that describe the business model and system objects (object model).

The object model describes the elements stored in the repository: object structures, associations and constraints. The meta-model describes the structure of the object model. It is expressed in terms of itself and is stored in the repository too.

End-user objects

End-user objects model organization structure, processes, tasks, data and relationships. The tools make no a priori assumptions about these end-user objects. This enables us to construct different business models. In fact, several models of the organization can be made to co-exist, which is useful for simulating future scenarios.

A IndicateurDocument keeps track of its creator and the organizational units its creator works in. There are different kinds of organizational units, for example departments and temporary project teams. Instead of modeling these different kinds of units through different object types, we use the Type Object pattern to classify the units dynamically. This enables us to replace the organization model without having to change the object model. A mere change of repository population suffices.

Meta-model

Each object has a specific type. An object type has a name and defines a set of properties. Each property has a name, a type, and also has other attributes such as whether it is mandatory, unique and represents a symmetric relationship.

If an object is of a particular type, then it is allowed to have properties that the type specifies. The values of a property of an object of type T must match the type of the property specified in T. Attribute property values have primitive types, such as dates, strings or numbers, and enumerated values over these types. Association properties reference other objects.

Object types can have subtypes. The property type may be overridden in subtypes. This is particularly useful for specializing abstract associations between two abstract types.

Figure X.14 Simplified diagram of meta-model

The instance diagram views a small portion of the meta-model population, part of the mail registration object model.

Figure X.15 Simplified instance diagram of meta-model

Our model supports N-ary and attributed associations. We use an example of attributed associations in the documentation center application. A document may be related to other documents. For example, it may be an addendum to, or a modification of other documents. We model this relationship as a symmetric association. The nature of the relationship is represented by means of an attribute of the association. The attribute takes its values from an enumeration set.

Individual properties can be constrained to take values in a subset of the type value set. Object and inter-property constraints describe conditions to be satisfied prior to committing changes to (groups of) objects. These constraints are specified by means of script expressions (Smalltalk code). Property types and constraints may be overridden in subtypes.

As an abstraction mechanism to hide details of the object model or to de-normalize the user's perception of the object model (e.g. when using the query editor or defining lay-outs in the forms tools) we use derived properties. These are read-only properties representing values or aggregates of properties. They are defined by means of query expressions (virtual properties) or script expressions (calculated properties). Derived properties can be used as regular properties in most cases (calculated properties cannot be used in queries).

Our framework treats objects that make up the object model in the same way as objects managed by end-user applications. Thus we can re-use the forms tool to edit the object model. Authorization and action rules apply to this object model editor application. For instance, we can define action rules that are triggered whenever the model changes, e.g. to update database tables and columns, or to invalidate lay-outs. And expressing the meta-model in terms of itself enables us to go one step further: the definition of our kernel meta-model can be extended to specify new types of constraints; the semantics can then be implemented by means of action rules.

System objects

The end-user and configuration tools require several types of 'system' objects.

We use many of these system objects to specify end-user applications, such as application environments, stored queries, action rules and object behavior. These types of objects are often associated with meta-level objects, e.g. object types and properties.

Several types of system objects, such as action rules and stored queries, contain parts that would require complex models and need not be queried or updated separately. These parts are modeled as encoded objects that provide an interface to encode (decode) Smalltalk objects into (from) a string representation when updating or accessing the object.

Example: The Indicateur example shows how initialization rules can set the default values of new IndicateurDocuments (e.g. IndicateurDocument is not confidential and the

CreationDate is today). The CreationDate rule is defined in the super-type WorkflowDocument and inherited by all its subtypes.

All these rules are stored in the repository. Rules have an association with object types. Condition and action scripts of script rules are modeled as strings containing the (Smalltalk) source code. The compiled versions of these scripts are modeled as encoded objects, associated with the rules. When a script rule is changed, another script rule is triggered that validates the script's expressions and generates the compiled code.

Figure X.16 Simplified rule object model

Other system objects represent special business objects, such as login-users, electronic documents and thesaurus keywords, that need some dedicated framework components. These objects are usually presented to the end-user through special tools, such as the query editor and viewers for scanned documents. But since system objects are explicitly modeled and stored in the repository, for administration purposes as querying and reporting we don't need additional tools. The generic end-user tools can be used instead to set up management applications.

Object representation

By default, repository objects are represented at run-time by instances of a Smalltalk class (FwEntity). These instances keep track of the identification (object and type references) of the repository objects they represent. In addition the class provides an interface for accessing and updating the properties. Instances typically contain partial information, i.e. we load only those properties from the database we actually need. Accessor methods are created on the fly when using the scripting language, and object behavior is loaded when needed.

This approach usually suffices for end-user objects. In the case of system and meta-model objects, the framework components often require extra behavior or state. We provide the means to add subclasses to FwEntity and to link these classes to specific types of repository objects.

X.4.2.2. Application environments

Application environments restrict the view on the repository by specifying what objects can be accessed, queried and created. Environments can be structured hierarchically: nested environments further specialize the view on the shared object model. Action rules, views and queries are inherited.

For every object type, lay-outs can be configured for use in list-, print-, export- and forms tools. These views are either private or shared among all users of the environment. Default views can be defined. Individual views take precedence over shared ones. Views may be re-used in other views.

In a similar way, queries may be saved for private or shared re-use. Defaults can be specified. Queries can be re-used in other queries. List views can be coupled to specific queries.

Action rules can be defined within the context of the environment, e.g. to set appropriate defaults, trigger a cascaded delete or to add specific functionality, for instance a batch update of a list of objects.

The use of nested environments typically leads to the following strategy: definition of an abstract application with re-usable assets, followed by dedicated sub-applications, such as full management and retrieval applications, and even simplified applications with many pre-defined queries and lay-outs for novice users.

X.4.2.3. Electronic documents

Electronic documents represent unstructured information and require dedicated tools for creating, viewing, editing and printing. Making electronic documents available in end-user applications is straightforward: we only need to associate documents with the relevant object types in the application. Thus electronic documents are not perceived as stand-alone objects by the end-user, rather they are enclosed in objects having more domain specific meaning.

The apparently simple document structure hides a more complex object model that contains the necessary information for the various tools, such as versions, different representations, storage information and instructions to delete or copy documents.

Versions

Documents are perceived by the end-user as logical documents, which have a name and can exist in one or more versions. One version is the preferred one. We do not maintain a version tree, as it would needlessly complicate use of the tools, neither are new versions created automatically by default. If necessary, we define rules to override this default behavior.

Document versions come in two flavors:

- basic documents, corresponding to physical documents, which have a well defined format and type
- complex documents, containing different representations (basic documents) for the same logical document version.

Basic documents are stored on physical media in one ore more copies, which are managed by the document caching process manager.

Representations

Representations enable users to use different views on the same logical document, e.g. to create a text representation of a scanned document, which can subsequently be stored in a full-text-indexing database. Other examples are alternative representations using a more widely available format, e.g. PDF for Internet, or a TIFF representation of a frozen Word document that may only be annotated.

One of these representations is the master document. Invalidating the master invalidates the other (slave) representations as well.

Document caching process manager

A document can exist in one or more copies on several caches and the optical storage medium. A background process (implemented by means of background processing rules) keeps track of the status of these copies to make sure that the up-to-date copy is stored and retrieved.

Instructions to cache or delete documents are not necessarily carried out immediately. Instead, they are listed as instructions to be executed by the background process. This condition takes into account model and contents of the objects in which the document is enclosed (the document owner). This enables us to set up very flexible caching strategies.

Whether an electronic document is persistent is ultimately derived from the repository information: an electronic document must be owned by a repository object. Hence, in case of failure, our recovery strategy makes the document base consistent by starting from the information in the database.

X.4.2.4. In- / outbaskets

In- / out-baskets maintain lists of incoming and outgoing work. What appears in a basket is defined by means of a query expression, hence it is not restricted to task objects only. Each user has a generic in-basket containing tasks sent to him or her. In addition, dedicated baskets may be set up, e.g. to follow up particular dossiers.

The in- / out-baskets tool uses a special application environment that re-uses several system objects from other environments, such as lay-outs.

X.4.2.5. Workflow

Workflow processes represent cause-connected activities to be performed in a (not necessarily unique) particular order to achieve a specific goal, e.g. the central board decision process, ad-hoc processes to register and handle incoming documents, or the process to create and distribute new norms and regulations towards the school.

Activities are modeled by means of task types. Tasks may involve several participants, e.g. the person who issued the task assignment and the recipient of the task.

Pre-defined process templates provide options to guide the user through default process scenarios. Processes can also be triggered automatically by specific events, for instance at the start of each day.

Tasks

A task specifies the interaction between participants necessary to perform a specific activity. Participants can be any object, since object types are independent of the actual rule engine, process editor or in / out basket model.

Some workflow tools, such as the graphical process editor, assume the existence of a minimal set of common properties defined for all task objects:

- the status property defines the various relevant states for a task with regards to commitment of the participants involved in the task;
- the previous task / next task association maintains task dependencies;
- timing information includes date sent, reminder date and deadline.

Other properties can be added as needed in the various subtypes of the Task object type. These properties can be viewed or edited using regular list views and forms, or listed in the graphical process history overview.

Processes

Processes represent the running history of an ad-hoc or pre-defined workflow process (or combination of both), i.e. the tasks performed so far. Tasks are always associated with exactly one process. A process will automatically be created when the user activates a process template from scratch. Within the context of a process, sub-processes may be started by activating process templates.

Process templates

A process template defines default scenarios for a workflow process, by listing the default tasks (more precisely types and default values) to be performed, and, at each step, the options to move from one task to another. These default scenarios are modeled by means of a collection of action rules.

Whether users may deviate from this default process is ultimately determined by the authorization mechanism. Instead of relying on programmed exceptions, the process template definition gives users the means to let non-default processes snap back into the default flow if some rule condition is satisfied. Alternatively, users can synchronize explicitly with one or more key steps in the template by activating a rule specifically configured for these purposes.

The mixture of ad-hoc / pre-defined processes and tasks / sub-processes, allows us to define process templates incrementally. For instance, a process template may be initially defined at the level of individual departments, and refined later within each individual department. These sub-processes may be more or less strict, depending on each department's working practices.

Implementing other process models

Other workflow process models can be implemented too. We illustrate two existing approaches here:

Workflow processes based on intelligent work objects [Kar90] can be modeled using script rules. The work object contains the necessary data and state information. A script rule describing the routing algorithm is activated whenever the user opens a form on the work object.

Process models based on Speech Acts [Flo93] can be modeled using high-level rules. Conversation objects contain the necessary data and state information. The condition expression refers to a matrix describing permissible state transitions. Default acts on conversations, such as promise, decline or counter-offer can be configured by means of action rules.

X.4.2.6. Authorization rules

The authorization mechanism makes no a priori assumptions about a particular model. It allows simple user / user group access controls, as well as more sophisticated, fine-grained context- and contents-sensitive access privileges. It relies on a knowledge-base of authorization rules. Each rule either explicitly grants or denies rights to users for certain aspects (mostly properties) of objects, or tells us nothing at all (i.e. the rule may be irrelevant

or undecided). Thus the authorization mechanism revolves around a strategy to find the most appropriate rule that conclusively answers the question ‘does this user have these rights for these aspects of this object’ with yes or no. The main idea is to capture the default cases in rules with a large scope, and to add rules catering for exceptions in a piecemeal fashion. The actual solution depends on the problem domain (e.g. do we restrict access in the default cases and explicitly grant access in the exceptions, or vice versa), but often reflects global policies within the organization regarding access control and privacy of information.

Example:

- All users may read all documents (default case for all users, aspects and types of documents, granting).
- No user, except administrators, may update documents (default case, denying in the former case, granting in the latter case).
- The colleagues of an employee may update a document he or she created (specific case, granting).
- In the mail registration application we use documents that may be flagged as private. Only members of the same department as the document creator or employees higher up in the hierarchy have access to private documents (specific case, denying).

Since rules may conflict, it is necessary to provide conflict-resolution strategies.

Definitions

Each authorization rule consists of the following 4-tuple:

- the agent set (Ag) specifies the users to whom access is granted or denied
- the object set (O) specifies the subjects of the rule
- the object aspects (As), e.g. a selection of properties
- the rights (R), such as read, write, add, remove and export.

These sets (called dynamic sets) can be dynamic, based on a condition that specifies the elements belonging to the set. Conditions make use of query expressions defined with the end-user query editor, or, in some cases, Smalltalk expressions. In addition, static sets can be defined by explicit enumeration. In practice, agent and object set are usually dynamic, the aspects and rights static. Dynamic sets can refer to context variables, such as the login-user.

Each of the sets can be made *restrictive* (denoted as [S], S a dynamic set). A rule is restrictive if at least one of its sets is restrictive. A rule can also be *granting*. Rules can be both granting and restrictive. Using these rules, we explicitly grant or deny access control.

Rule semantics

The following example illustrates the semantics of authorization rules.

Applying a rule $\{Ag, [O], [As], R\}$ to an agent-object-aspect-right tuple

$\{ag, o, as, r\}$ will answer true, false or undecided according to the following strategy:

- true if $(ag \in Ag)$ and $(o \in O)$ and $(as \in As)$ and $(r \in R)$ and (the rule is granting)
- false if $(ag \in Ag)$ and $(o \notin O)$ and $(as \notin As)$ and $(r \in R)$ and (the type of o is a subtype of the type of the object set O , i.e. the rule must be *relevant*)
- undecided otherwise.

If we substitute $Ag = \{\text{Department heads}\}$, $O = \{\text{Proposals}\}$, $As = \{\text{Visa property}\}$ and $R = \{\text{Edit}\}$, then this rule translates into: “Department heads may only edit the visa property of proposals”. This rule is not relevant for objects of types other than proposals.

Combining rules

Rules can be inherited according to the object set type hierarchy. Rules defined for a particular object type have precedence over inherited rules. Rules defined within the same object type are explicitly prioritized to resolve conflicts.

As rule evaluation may be undecided for a particular tuple $\{ag, o, as, r\}$, at least one default rule is needed that answers true or false for all tuples. This can be a rule defined at the common supertype of all object types.

Using authorization rules

Authorization rules can capture complex requirements succinctly.

Example: “A user has access to all documents created within a workflow process if at least one of these documents has been created by a member of his organizational unit or a unit further down in the hierarchy”. This rule can be defined in the configuration tool without any scripting. Rules like these will typically use context variables, for instance to refer to the login-user.

Rules can be applied to all elements in the repository, including system objects such as document annotations, applications and even authorization rules.

The authorization mechanism is very expressive, but managing the rule-base may become less evident if number and complexity of the rules grow too large. Two techniques can be applied to simplify maintenance of access control:

- Reification of hidden elements in the rule definitions.

Example: to specify the applications a user has access to, a simple end-user application has been configured. This application maintains explicit associations between users and environments. Several authorization rules have been reduced to one meta-rule checking this information. Note that this rule applies to the application itself.

Similar techniques can be used to manage e.g. large numbers of users.

- Modeling additional, simpler authorization rule types to be managed through end-user tools. Transformation methods are needed to convert these rules into system-level rules.

Thus we can configure specific authorization applications, targeted towards a particular use with a minimal effort.

Optimizing performance

By default, authorization rules are executed locally. In some cases, this generates considerable overhead when many objects are retrieved from the repository, only to find out that the user has no access. And even if the user has access to a particular object, the authorization rules may require additional associated objects to be downloaded in order to verify the access privileges. To optimize performance, rules are combined with database queries if appropriate.

Design

The authorization mechanism presents one aspect of the framework [Xer96]. At first sight, its functionality should be woven together with the object store's to ensure consistency. However this behavior is not always required nor even wanted. Performance issues take precedence in some cases. The Facade-pattern provides a clean mechanism to have different framework components access the object store in different ways, for instance, with or without access control. Ideally, however, this requires context-sensitive access control for objects [Ric92], whereby only trusted components may access the object store directly.

X.4.2.7. Action rules

Objects obey constraints defined in the model and exhibit behavior that is used across all applications. End-user applications sharing these objects may however require additional, often context- and contents-sensitive functionality and semantics (business rules) [Gra94]. Action rules capture these requirements. While action rules allow semantics to be different across applications or workflow processes, they do not violate the global constraints defined in the model.

The main elements of action rules are events, conditions, actions and rule scope. The type of event determines whether the rules implement some business rule or provide some additional functionality to the user.

Examples:

- When creating documents we set the creation date and creator. When committing a modified document, we update modification date and remember the last author. Documents in the mail registration application need extra initialization to generate a unique id based on the context: is the document the start of a new workflow process or is it created within the scope of an existing process?
- When removing objects, we delete attached electronic documents (if any) by default.
- We add extra functionality to perform operations not directly provided by the generic tools (e.g. a particular report), or to support repetitive work. In the context

of workflow processes we often need to synchronize the state of different related objects and guarantee the overall consistency; instead of complicating the tasks of both users and configurators, we provide extra functionality to perform the necessary operations in a controlled way. All these rules are typically triggered by menu events corresponding to menus that only appear in the relevant context (application environment and object type).

- Several tasks must be performed automatically, often at regular times. Rules can be triggered by time events generated by the clock.

Events

Action rules are triggered by events. These may be generated by the system or by explicit user actions.

Rules triggered by system events generally affect the semantics and include: creating, duplicating, saving, deleting or modifying an object. The transaction state specifies when the event is actually generated. For instance, rules triggered by changes of individual property values will typically be executed at once, whereas rules triggered by more global object updates are usually delayed until pre-commit time.

User action events either correspond to extra menu options providing additional functionality or to generic events, such as switching applications and selecting an object or layout.

A single event can trigger several rules, which enables us to add behavior in a modular way. All enabled rules (i.e. rules for which the condition is satisfied) will be activated.

Activating rules

Action rules have access to context-variables, such as the login-user or the interface component that triggered the rule. To activate a rule we apply it to a particular object (the receiver) in a given context. If the condition is satisfied, the action will be executed.

A rule is scoped: it can only be applied to objects of a given type (including subtypes). The scope can be further limited to a particular application or workflow process.

Types of rules

We provide essentially two types of rules: script rules and high-level rules

- **Script rules**

Script rule condition and action expressions require our scripting language (basically Smalltalk with dynamic object accessor methods, implicit access control and high-level query access to the repository), hence they are not targeted towards the average end-user. The rule scripting language has access to several context-variables, and can re-use object behavior (the latter uses the same scripting language, but has no access to the context-variables, making it context-independent).

- **High-level rules**

To allow end-users to add their own rules, we provide high-level rule types. High-level rules explicitly model common script practices in the repository, such as:

- initializing objects to default values (constants or results of query expressions)
- updating an object at commit time or when a property has been modified
- creating a follow-up task in a workflow process, based on certain conditions, and updating the state of the current task.

High-level rules

Conditions are specified by means of query expressions. These can be created by the query editor. The action part uses one or more *update expressions* to specify new values for object properties. The first expression refers to the receiver. Executing the expression updates the receiver. Additional expressions specify type and initial values for new objects to be created when the rule is activated.

Examples:

- Rules with one expression are typically used to set default values for a newly created or duplicated object, or to update an object's properties as the result of a user action. The latter is often used in batch procedures to update a series of objects.
- Rules with two expressions are used to support definition of workflow processes. Given a task (the receiver), a rule can be defined to create a follow-up task and update the current task state. We define a workflow process consisting of a set of tasks by describing for each task the options (and conditions) to start other tasks. Each option then corresponds to a rule.

High-level rules can model recursive processes, e.g. to implement an approval procedure that mirrors the hierarchy.

Update expressions

Update expressions can be edited by means of extended forms. No scripting is needed. These expressions behave similarly to regular objects, i.e. they have a type and properties, but the latter accept expressions in addition to values.

The following expressions can be used:

- constant expressions, such as a date value (attribute property) or a reference to another object (association property);
- query expressions;
- property expressions (chains of associations ending in an attribute or association property);

- scripts (using typed methods).

For multi-valued properties, individual elements may be flagged for addition or removal.

Example: The following approval rule (in pseudo-language) illustrates the idea:

Figure X.17 High-level rule example

This can be read as: “When the user has approved a request (condition) and has selected the ‘forward’ menu item (event), the current task will be flagged as finished (current task expression). A new request for approval will be sent to the Finances manager. The documents enclosed in the current task will be attached and the date will be set to today’s date (new task expression).”.

Stepwise structuring

We will formalize other common practices in the future. Operations on objects such as saving, deleting or locking often require that the same operation be performed on related objects (cascaded saves, deletes and locks). A variant of the high-level rules can be used to model these practices. Extended forms can equally well be used to edit these new types of rules. This incremental process of formalizing practices in the repository is an example of stepwise structuring [Häg89].

Authorization

When executing rules, it may be necessary to override the current user’s privileges. Hence, each rule can be given an owner, whose privileges will be used instead. In addition, the user must have the necessary privileges to execute the rule.

X.4.2.8. Background processing rules

Background processing rules specify activities to be performed automatically on elements in the repository. Typical uses are process managers handling requests to migrate document copies to a particular storage medium or to convert notification requests generated by the system into E-mail messages, and automated tasks within the context of a particular application or workflow process.

Background processing rules used to be dedicated objects and framework components. Now they are implemented by means of action rules triggered by time events. Time events enable us to set up elaborate scheduling strategies, including recurrent events.

The action rule manager automatically schedules any background rules defined in a specific application whenever the user selects the application. Whether the background rules run within the context of an end-user application on his or her client, or on a dedicated client PC is solely a matter of configuration. Thus we can easily set up appropriate applications on separate clients in order to achieve the right degree of load balancing.

X.5. Bootstrapping process

The bootstrapping process is driven by our observations that most configuration tools essentially manage specific types of objects in the repository, to which we need to add some extra functionality, validation, consistency rules and specific caching. Often, an alternative

plug-in property editor is all that is required to access these objects easily through end-user tools.

Examples

To bootstrap the system, we initially developed a hard-wired object model editor. About one year into the project we were able to configure end-user tools for accessing the object model. Now we are enhancing this application to support the functionality offered by the new persistency component.

While access to in- /out-baskets is still managed through a dedicated application, baskets are defined through regular forms. We applied a bit of white-box re-use to make the query editor fit in the property editors scheme, and we modeled the basket definition explicitly. Its properties include a query expression specifying the basket's contents.

Most of our latest tools are configured applications, or are developed with configuration in mind. Examples of the former are the object behavior and action rules applications. The graphical workflow on the other hand is designed as an alternative property editor for use in forms.

Reflection

We are increasingly using existing tools to build new (typically configuration) tools. Reflection [Foo96] is a useful technique to support this kind of approach. Many people still view reflection as rather obscure and difficult. Yet, given our framework approach, we consider this idea of bootstrapping as a very natural thing to do: the system already contains so many tools to build applications, that we would take a step back by not trying to re-use this functionality. The locality of change that can be achieved with reflection is also a very important asset. In fact, reflection makes our design both cleaner and simpler.

X.6. Framework evolution

We started the framework initially building three applications. At this moment we have built and maintain many more. We tuned, adapted and extended our framework and business model, based on what we learned from the needs and shortcomings of our applications and from feed-back from configurators and end-users.

Since the business objects, object model, scripts and application specifications are all stored in the same repository, we use the repository as a sort of centralized knowledge base. Some recurring business practices or functional requirements may be less evident at the start, but may come to surface later on, e.g. in the form of similar custom scripts. When we identify these patterns, we try to make them more explicit either in the repository or in the tools.

Example: The form layout editor initially provided the means to include Smalltalk scripts describing how to represent associated objects on the screen. After a few years, these scripts have been replaced by parameterized strings. The parameters refer to properties or property chains of the associated objects, and may include conditional expressions. These expressions capture the existing scripting patterns more explicitly. They are also much simpler and more robust than the original scripts.

When developing applications with the framework we identify:

- **Working practices specific to the organization.** We create re-usable assets in the repository, re-factor the object model, or we develop more specific end-user and configuration applications using the existing tools.
- **Practices which transcend applications.** We change the framework functionality, and re-factor the framework [Foo95,Rob96].
- **Functionality specific to a particular application.** Depending on its nature, we change the functionality of the repository, the framework or both. For instance, some applications require a simple procedure to automatically generate standard reply letters based on repository information. In this case, we store extra behavior in the repository. If a new external tool is required, we usually add a component to the framework.
- **Technological needs and opportunities.** We use sub-frameworks for components that are strongly coupled to technology, such as the persistency layer and the document storage. These sub-frameworks can easily be replaced.
- **Additional functionality,** to support new types of applications, e.g. Internet applications.

This way, the framework explicitly supports evolution and re-use at different levels [Til96].

X.7. Conclusion

The main goal of the framework design was a small kernel of generic components acting dynamically upon the repository. Hard-coding was to be avoided as much as possible. Initially, we focused on achieving this goal for end-user applications. We developed hard-wired administration and configuration tools to help us bootstrap the system. As the framework evolved, and the tools became more flexible, we re-used components of end-user tools in some of our administration tools. In a third phase, we started to replace some of the hard-wired tools with applications configured in the system.

The framework configuration functionality will be enhanced, in order to further increase the expressiveness of the tools, and to allow end-users to adapt the tools to their own needs even better, thus obviating the need for developers to a larger degree.

Some additional components are being developed, most notably to make the repository Internet-aware. The main component will act both as a client of the repository and as a generic Internet application server. This will empower Argo to develop Internet applications through modeling and configuration too.

At first sight, one might wonder whether this approach does not make the design more complex, or affects performance negatively. We do not feel that the design is more complex. In fact, the framework goals help us focus more clearly on the responsibilities of the various components, in particular with regards to re-use and evolution. As is often the case, reifying implicit responsibilities actually makes the design simpler. And the bootstrapping principle is an additional asset in validating the design.

Although we initially encountered some performance problems, these were not essentially related to the approach as such, but to the use of a persistency component which was not really suited for our purposes. We must be aware, however, that the flexibility of the system allows end-users to build e.g. equally good and bad queries. Proper training is of primary importance, as is the need to hide more complex elements of the tools from novice users.

After three years, we are convinced that the approach is a valid one. End-user applications can be developed iterative and incrementally, even interactively. We do not code (apart from some custom scripting) and we do not generate end-user applications. Neither do we use or need throw-away prototypes. Instead, we build increasingly complete specifications of end-user applications. These specifications are available for immediate execution. Thus we help close the gap between specification, development and use of the applications. This way we can afford to stimulate the users to think and rethink their processes and redesign them.

Acknowledgments

We would like to thank the people in our team: Rudy Breedenraedt, Hilde Deleu, Jan Geysen, Els Goossens, Wouter Roose, Fred Spiessens, Du Thanh-Son, Danny Ureel and Leen Van Riel. We would also like to thank the Programming Technology Lab at the Vrije Universiteit Brussel, for the many good ideas, and the pioneer-users at Argo who helped us make this framework happen. Finally, we would like to thank Ralph E. Johnson for the many helpful comments during the review process.

References

- [Foo95] Brian Foote and William F. Opdyke, Lifecycle and Refactoring Patterns that Support Evolution and Reuse, Pattern Languages of Program Design, Addison-Wesley, 1995
- [Foo96] Brian Foote and Joseph Yoder, Evolution, Architecture, and Metamorphosis, Pattern Languages of Program Design, Addison-Wesley, 1996
- [Flo93] Fernando Flores, Michael Graves, Brad Hartfield and Terry Winograd, Computer Systems and the Design of Organizational Interaction, Readings in Groupware and Computer-Supported Cooperative Work, Morgan Kaufmann, 1993
- [Gam95] Erich Gamma, Richard Helm, Ralph Johnson, John Vlissides, Design Patterns, Elements of Reusable Object-Oriented Software, Addison Wesley, 1995
- [Gra94] Ian Graham, Object Oriented Methods, Addison-Wesley, 1994
- [Häg89] S. Hägglund, Iterative Design and Adaptive Maintenance of Knowledge-Based Office Systems, Proceedings of the IFIP WG 8.4 Working Conference on Office Information Systems: The Design Process, 1988, North-Holland
- [Kar90] B. Karbe, N. Ramsperger and P. Weiss, Support for Cooperative Work by Electronic Circulation Folders, Proceedings of the ACM OIS'90 Conference, 1990
- [Ric92] J. Richardson, P. Schwarz, L.-F. Cabrera, CACL: Efficient Fine-Grained Protection for Objects, Proceedings of the ACM OOPSLA'92 Conference, 1992
- [Rob96] Don Roberts and Ralph Johnson, Evolving Frameworks: A Pattern Language for Developing Object-Oriented Frameworks <http://st-www.cs.uiuc.edu/users/droberts/evolve.html>

[Sto94] Frank Stowell and Duane West, Client-Led Design, A systemic approach to Information System Definition, McGraw-Hill, 1994

[Til96] Michel Tilman and Martine Devos, Object-Orientation and Evolutionary Software Engineering, Position paper for the OOPSLA'96 Workshop on Object-Oriented Software Evolution and Reengineering, 1996

[Xer96] Xerox Parc Aspect-Oriented Programming Project, A position paper on Aspect-Oriented Programming, position paper for the ACM Workshop on Strategic Directions in Computing Research, Working Group Object-Oriented Programming, MIT, June 14-15 1996.

Glossary

Action rules: Rules triggered by events that conditionally perform some action. Action rules are used to add extra semantics (business rules) or to add extra functionality.

Authorization rules: Access control relies on a rule-base of authorization rules.

End-user objects: Represent the specific business model.

Indexcards: Forms.

In- / out baskets: Baskets maintain lists of incoming and outgoing work.

Meta-model: Describes structure and constraints of an object model.

Object model: Describes objects, associations and constraints.

Object store: The persistency component.

Properties: Describe attributes and associations (references to other objects) of an object.

Property expressions: Describe chains of properties, by following associations.

Query expressions: Represent in an abstract way queries to be executed. A query (expression) can be stored and re-used.

Re-factoring: The process of redesigning e.g. a class library without changing the overall behavior, for instance to remove redundant code.

System objects: Represent objects in the repository required for proper functioning of the framework tools.

Thesaurus: Indexing vocabulary having its own semantically meaningful structure, such as a generic / specific taxonomy of descriptors (indexing keywords) and synonyms (providing additional access paths during searches).