

The Leordo Computation System

Erik Sandewall

Department of Computer and Information Science,
Linköping University, Linköping, Sweden

and

Unit for Scientific Information and Learning,
Royal Institute of Technology, Stockholm, Sweden

March 12, 2008

Abstract

The purpose of the research reported here is to explore an alternative way of organizing the general software structure in computers, eliminating the traditional distinctions between operating system, programming language, database system, and several other kinds of software. We observe that there is a lot of costly duplication of concepts and of facilities in the conventional architecture, and believe that most of that duplication can be eliminated if the software is organized differently.

This article describes Leonardo, an experimental software system that has been built in order to explore an alternative design and to try to verify the hypothesis that a much more compact design is possible and that concept duplication can be eliminated or at least greatly reduced. Definite conclusions in those respects can not yet be made, but the indications are positive and the design that has been implemented so far has a number of interesting and unusual features.

1 Introduction

Project Goal and Design Goals

Leordo is a software project and an experimental software system that integrates capabilities that are usually found in several different software systems:

- in the operating system
- in the programming language and programming environment
- in an intelligent agent system
- in a text formatting system

and others more. I believe that it shall be possible to make a much more concise, efficient, and user-friendly design of the total software system in the conventional (PC-type) computer by integrating capabilities and organizing them in a new way.

The purpose of the Leordo project ⁽¹⁾ is *to verify or falsify this hypothesis*. This is done by designing and implementing an experimental system, by iterating on its design until it satisfies a number of well defined criteria, and by implementing a number of characteristic applications using the Leordo system as a platform.

The implementation of the experimental system has passed several such iterations, and a reasonably well-working system is in daily use since several years. The following are the requirements that were specified for that system and that are satisfied by the present implementation. We expect to retain them in future system generations.

¹The project and the system were previously called Leonardo, but the name was changed to Leordo in order to avoid a name conflict.

The system is of course organized in a modular fashion, where the modules are called *knowledge blocks* and contain both algorithms, data, and intermediate information such as ontologies and rules. There shall be a designated *kernel* consisting of one or a few knowledge blocks that is used as a basis on which other blocks can be built, for the purpose of additional services and for applications. The following were and are *the requirements on the kernel*:

- It shall contain self-describing information and corresponding procedural capabilities whereby it is able to administrate itself, its own structure, and its own updates.
- It shall provide the extension capabilities that make it possible to attach additional knowledge blocks to it and to administrate them in the same way as the kernel administrates itself.
- It shall provide adequate representations for the persistent storage of all contents of the blocks in the kernel, as well as the representations and the computational services for performing computations on the same contents.
- It shall provide capabilities for adaptation, in particular to facilitate moving a system between hosts, and for defining alternative configurations based on different sets of knowledge blocks.
- Although the experimental system will be based on an existing, conventional operating system and ditto programming language, it shall be designed in such a way that it can be ported to the weakest possible, underlying software base.

The last item in these requirements is included because in principle we believe that the services of the operating system and the programming language and system, should just be parts of one integrated computation system. The longer-term goal is therefore that the Leordo system itself should contain the programming-language and operating-system services.

Furthermore, the facilities in the kernel have been, and will continue to be designed in such a way that they do not merely serve the above-mentioned requirements on the kernel itself; they shall also be general enough to provide a range of applications with similar services.

Above the kernel and below the specific application areas and applications, there shall also be an extensible *platform* consisting of knowledgeblocks that are of general use for a number of applications of widely different character.

Illustration materials and annexes of the present article can be found on the article's persistent webpage at <http://www.ida.liu.se/ext/caisor/pm-archive/leonardo/002/>. It is useful to have access to that webpage while reading the present article.

Main Hypothesis for the Leordo Project

The main hypothesis for this project, for which we hope to obtain either strong positive evidence or a clear refutation, is as follows: *It is demonstrably possible to design the kernel and a platform in such a way that (1) repeated implementation of similar tasks is virtually eliminated in the kernel and platform, and (2) the total software structure that is obtained when several applications are built on this platform can also be essentially free from repeated implementations of similar tasks.*

Approach to the Design

The design of the system does not start by defining a programming language, nor by defining a process structure or a virtual instruction set. In Leordo, the first step in the design is to define an object-oriented information structure that has some points in common with RDF⁽²⁾ and OWL⁽³⁾, although also with significant differences. The notation used for this purpose is called KRE, the Knowledge Representation Expression language. It is used for all information in the system, including application data, procedures, ontologies, parameter structures, and whatever corresponds to data declarations in our system. Full KRE is a knowledge representation language but the present article will only describe those parts of KRE that are used for the Leordo system design.

The element in the KRE structure is called an *entity*, and entities can have *attributes* and *properties*. Attribute values can have structure and are not merely links to other entities; they can be constructed by the formation of sets, sequences, and records, even recursively. Moreover, entities can be composite expressions; they are not merely atoms with mnemonic names. Property values are like long strings and can be used for expressing e.g. a function definition, or a descriptive comment. Because of this expressive power, KRE is best viewed as a knowledge representation language.

We use the term ‘entity’ rather than ‘object’ for the elements in KRE since the term ‘object’ has a connotation of message passing and a fairly restrictive view of class hierarchy, which are not applicable in KRE.

Each knowledge block consists of a set of *entity files*; each entity file consists of a sequence of entities; and each entity has its attributes and properties.

The experimental system which is based on conventional operating systems, has in addition the following design. A *Leordo individual* is a section of the file system in a computer hosting the individual, that is, one directory and all its sub-directories, with all the files contained in them (with the exception of auxiliary files such as `.bak`

²<http://www.w3.org/RDF/>

³<http://www.w3.org/TR/owl-features/>

files). Each entityfile in the Leordo sense (i.e., a sequence of entities) is represented by one file in the sense of the file system; this file is a text ('ascii') file adhering to a particular syntax. An *activation* of the individual is obtained by starting a run with a host programming language, where the run is initialized using some of the files contained in the individual. The run usually includes interactions with a human user, but maybe also with robotic equipment, Internet information sources and resources, or other Leordo individuals. Entityfiles in the individual can be read and written during the activation, for example for storing information that has been acquired during the activation, or for updating the software.

In accordance with the specified goals for Leordo, as described above, the individual shall be self-contained and be able to model its own structure, and to update it. In that sense the individual is able to *modify itself* during the activation. The individual shall also contain facilities for moving itself, or allowing itself to be moved from one host to another, in ways that are reminiscent of mobile agents⁽⁴⁾.

The use of directories and files for representing aggregates of Leordo entities is an intermediate solution. In the longer run we wish to port Leordo to a persistent software system that is able to represent entities directly, so that the structures and services that are traditionally offered by an operating system and in particular by its file system, can instead be implemented in the Leordo kernel or platform.

Both the experimental system and the forthcoming persistent system must use a *host programming language*. Functions, procedures, classes, or whatever other building-blocks are used in the host language will be represented by Leordo entities, and the definition of a function (etc) is expressed in a property of that entity. Our main experimental system has been implemented in CommonLisp; a part of the core has also been implemented in Python. We expect that the persistent system will be based on a language similar to Scheme. Interpretation-oriented languages such as these are the best suited for our approach.

Notation vs. System

The language design and the system design in our approach are strongly interdependent. The language design has come first in the present project, but the system design is by far the largest part of the work and it has arguably the largest novelty value. The main purpose of the present report is to describe the system design, but it is necessary to describe the language design first.

⁴<http://en.wikipedia.org/wiki/Mobile-agent>

2 An Example of KRE

By way of introduction we show two examples of how the Leordo Data Expression language, KRE, is used in Leordo. A more detailed specification of KRE can be found in the report “The Leonardo Representation Language ⁽⁵⁾).

KRE in the Common Knowledge Library

The Common Knowledge Library ⁽⁶⁾ (CKL) is an open-source repository for structured information ranging from ‘facts’ to ‘knowledge’. It presently contains more than 60.000 entities each having a number of attributes. The KRE notation is the ‘first language’ used by the Common Knowledge Library. The reader is encouraged to browse its website in order to obtain an impression of how information can be represented using KRE.

One thing that is not immediately seen on the CKL website is, however, the representation of meta-information. We use two kinds of meta-information: type information and catalog information. Each entity has an attribute called `type` where the attribute value is a new entity representing the type of the given entity. Furthermore, for each entityfile there is one entity that serves as the name of the entityfile; one of the attributes of the naming entity is a sequence consisting of the entityfile’s members. The name of an entityfile is itself the first member of that list.

The type system is quite simple. For use in some applications there is also a notion of *classes* which are similar to the ‘concepts’ of description languages, but this is not used in the system kernel.

Notice that entityfiles are used for expressing both programs and data. Each named unit in a program, such as a function or a procedure, is represented as a KRE entity, with the program code in a property of that entity. The entityfiles that are used within the system differ in some minor ways from those shown on the CKL website. For example, the provenance and IPR information occurs only in the published files and not in system-internal files.

The operation of *loading* an entityfile is performed by activations, and consists of reading the text file for the entityfile, such as the ones used in the CKL, and constructing the corresponding data structures in the activation. The operation of *storing* an entityfile is the reverse operation of re-writing its text file by converting data structures to corresponding, textual expressions. Loading and immediately storing an entityfile has a null effect on its text file.

⁵<http://www.ida.liu.se/ext/caisor/pm-archive/leonardo/001/>

⁶<http://piex.publ.kth.se/ckl/>

Cooperating Agents

The next example shows how a distributed computational process is organized in Leordo, and how the KRE language is used for representing the control information. Consider the following method description in Leordo:

```
-----  
-- method6  
  
[: type method]  
[: plan {[intend: t1 t2 (remex: lar-004 (query: makebid))]  
        [intend: t1 t3 (query: makebid)]  
        [intend: t4 t5 (query: propose-compromise)]}]  
[: time-constraints {[afterall: {t2 t3} t4}]]  
-----
```

This is a plan, i.e. a kind of high-level procedure, for a situation where two separate users have to give their respective bids for some purpose, and when both bids have been received, one user has to propose a compromise. This requires performing the action `query`: three times with different arguments. The time when the first two occurrences are to start is called `t1`; the third occurrence starts at a time `t4` which is defined as being when the first two occurrences have ended. The time when the first mentioned occurrence ends is called `t2`, and similarly for `t3`. The method consists of a set of intended actions, and set of time constraints between them.

This plan is supposed to be executed in a particular individual (called `lar-003` in our specific run of the plan) but the first mentioned action is to be remote executed (therefore `remex`:) in another individual called `lar-004`.

The KRE language is used for representing this plan, or script. In this example there is an entity called `method4` with three attributes `type`, `plan`, and `time-constraints`. The value of the `type` attribute determines what other attributes may be present.

This examples uses more of the KRE expressivity than in the first example. It shows how expressions in KRE may be atomic ones (symbols, strings, or numbers), or may be formed recursively using the operators `<...>` for sequences, `{...}` for sets, `[...]` for records, and `(...)` for forming composite entities. In the example, `(query: makebid)` is a composite entity that has a type and attributes, just like the atomic entity `method4`.

The webpage of the present article contains details from an activation using the method shown above, and it illustrates how KRE is used for the control information as the plan or script is executed, and for retaining some of that control information afterwards.

3 Information Structure

The Structure of Knowledgeblocks

The total information in a Leordo system is organized as a set of knowledgeblocks, and each activation of Leordo is initialized by loading one specific knowledgeblock in that set. Some knowledgeblocks *require* others, however, so that to load a knowledgeblock one first loads those other knowledgeblocks that it requires, recursively, and then one loads the entityfiles that are specified for the given knowledgeblock itself.

Each knowledgeblock consists of a set of entityfiles. One of those entityfiles represents the knowledgeblock as a whole and contains overall information about it; it is called the *index* of the knowledgeblock. The first entity in the index has the type `kb-index` which is a subtype of `entityfile`, and this entity is used to designate the knowledgeblock as a whole. This means that it can have both attributes that pertain to its role as describing its own entityfile, and attributes that pertain to the knowledgeblock as a whole.

One important use of the knowledgeblock index is to specify where the textfiles for other entityfiles in the same knowledgeblock are stored. The kb index specifies the mapping from entities as understood by Leordo, to actual file paths in the computer or filestore at hand ⁽⁷⁾. This makes it straightforward to move entityfiles and to redirect references to them, which has a number of uses including that it makes it easy for several individuals to share some of their files.

A few of the entityfiles in a knowledgeblock have special properties or play special roles, besides its index. This applies in particular for ontology files. To the largest extent possible, entities in the core Leordo system and in its applications are organized in terms of an ontology which is subject to modularization like all other aspects of the system. The kernel contains a ‘core ontology’, and every knowledgeblock contributes additional entities and links to the ontology, thereby extending the core. Each activation of an individual contains a working ontology that has been formed by loading and integrating the ontology files of the knowledgeblocks that have been loaded.

Entities in a knowledgeblock can be of three kinds with respect to mobility: *software specific*, *individual specific*, or *host specific*. These are defined as follows. If an individual is moved to another host, then it shall encounter host specific entities of the new host instead of those it had on the old host, whereas software specific and individual specific entities are retained. On the other hand, if a knowledgeblock is exported from one individual to another then only the software specific entities are exported and they will be used with the individ-

⁷Some Leordo individuals are placed on detachable memory devices, such as USB sticks, which means that they can have activations on different hosts without their file structure having been ‘moved’ in a conventional sense.

ual specific entities of the receiving individual. Individual specific information includes the history and experience of the individual; host specific information includes e.g. the locations and properties of databases, printout devices, and other resources that the host can offer to a visiting software individual.

In the present Leordo design, each entityfile is required to have all its members of the same kind in this respect, so that the distinction between software specific, individual specific, and host specific applies to entityfiles as well. The knowledgeblock index specifies only the locations of software specific entityfiles that belong to it. There are separate catalogs for all host specific and for all individual specific entityfiles.

Considerations for the Design of KRE

Knowledge Representation Expressions (KRE) is a textual representation for information structures, and the above examples have given the flavor of this notation. The details of the syntax are described in a separate memo that is available on the Leordo website (⁸). The present subsection shall discuss the design considerations that guided the definition of KRE.

The idea of allowing data structures to be expressed as text, and to define input and output of data structures accordingly, was pioneered by John McCarthy with Lisp 1.5 (⁹). It has been adopted in several interpretive or 'scripting' programming languages that are used extensively, such as Perl (¹⁰) and Python (¹¹). It is also characteristic of high level message formats, such as the KQML (¹²). With partly different goals, this tradition has also been continued in the XML family of information representation languages, including e.g. RDF and OWL besides XML itself.

There are several possible motivations for representing information structures textually, in text files or otherwise, and in particular:

1. For persistent storage of the information, between runs of computer programs.
2. For presentation of the information to the user, and for allowing her or him to edit the information.
3. As a message format, for transmitting chunks of information from one executing process to another one.
4. For representation of internal system information, such as parameter settings for particular programs or services.

⁸<http://www.ida.liu.se/ext/leonardo/>

⁹<http://www.lisp.org/alu/home>

¹⁰<http://www.perl.org/>

¹¹<http://www.python.org/>

¹²<http://www.cs.umbc.edu/kqml/>

These alternatives apply regardless of whether the text files are used for representing pieces of code, application data, or declarations, ontologies, or other metadata.

The choice of representation may depend on which of these are the intended uses. In particular, if the second purpose is intended then it becomes important to have a representation that is convenient to read for the human. The poor lisibility of XML was apparently accepted because it was thought that XML coded information should mostly be seen and edited through graphical interfaces, and not directly by users or developers.

In our case, we wish to use KRE for all four of the above mentioned purposes. We also have some other design requirements:

- The notation should be suitable for use in textbooks, research articles, and manuals. This strongly suggests that it should stay as close to conventional set theory notation as possible.
- Since the notation is going to be the basis for an entire computation system, including its programming-language aspects, it must be expressive enough for the needs of a programming language.
- The notation is used for the ontology structure that is a backbone for the Leordo system. It must therefore be expressive enough for what is needed in ontologies.

These requirements led to the decision of not using an existing language or notation, but to design our own. The following aspects of the KRE language should be emphasized in particular.

1. *The use of multiple bracket types.* Most programming languages use several kinds of parentheses and brackets, such as (\dots) , $[\dots]$, $\{\dots\}$, and possibly others. On the other hand, representations for information structures often use a single kind of brackets, such as (\dots) in Lisp and $\langle \dots \rangle$ in XML and other languages in the SGML tradition. This is sufficient in principle, but it makes it necessary to rewrite sets, sequences, and other “naturally parenthesized” structures along the lines of

```
(set a b c)
(sequence a b c)
```

and so on. LRX uses the multiple brackets approach and allows expressions such as

```
{a b c}
<a b c>
```

for sets and sequences, and in addition a few other kinds of brackets. This difference is trivial from an abstract point of view, but it makes surprisingly much difference for the ease of reading complex expressions. Compare, for example, the KRE representation of the plan entity in example 2, which was as follows:

```

-----
[: type method]
[: plan {[intend: t1 t2 (remex: lar-004 (query: makebid))]
        [intend: t1 t3 (query: makebid)]
        [intend: t4 t5 (query: propose-compromise)]}]
[: time-constraints {[afterall: {t2 t3} t4}]}
-----

```

with a representation of *just the second line* of that information in an XML-style ⁽¹³⁾ single-bracket notation:

```

-----
<plan>
  <planstep-set>
    <planstep>
      <intendstep>
        <fromtime>t1</fromtime>
        <totime>t2</totime>
        <remote-execute>
          <execute-at>
            <indiv-name>lar-004</indiv-name>
          </execute-at>
          <execute-what>
            <query-action>
              <phrase>makebid</phrase>
            </query-action>
          </execute-what>
        </remote-execute>
      <intendstep>
    <planstep>
-----

```

Even the Lisp-style single-bracket representation is much less convenient to read than the multi-bracket representation:

```

-----
(maplet type method)
(maplet plan
  (set (record intend: t1 t2
        (term remex: lar-004 (term query: makebid)))
        (record intend: t1 t3 (term query: makebid))
        (record intend: t4 t5 (term query: propose-compromise))
        ))
(maplet time-constraints
  (set (constraint afterall (set t2 t3) t4)) )
-----

```

In a historical perspective it is interesting to compare the great interest in legibility issues when programming languages are designed,

¹³<http://www.w3.org/XML/>

with the virtually complete disregard for the same issue in the design of so-called markup languages for representing structured information in general.

2. *The use of composite expressions for entities.* KRE is similar to e.g. OWL in that it is based on the use of entities to which attributes and properties are assigned. In the simplest cases, entities are written as identifiers and attributes are expressions that are formed using set, sequence, and record forming operators. However, entities can also be composite expressions that are formed using a *symbolic function* and one or more arguments which are again atomic or composite expressions, for example as in

```
(payment: (membership: (member-number: 1452)
              (year: 2006) ))
```

for “the payment for the membership during year 2006, by member number 1452”, or

```
(b: 356 (remex: lar-004 (query: makebid)))
```

for “the instance of the action (query: makebid) that was initiated at time 356 for execution in the individual lar-004“. Entities formed by composite expressions share the same characteristics as atomic entities, for example that they have a type and can be assigned attributes and properties, and that they can be included in entityfiles with their assignments.

The YAML (Yet Another Markup Language) ⁽¹⁴⁾ allows assigning attributes to composite structures, but does not make it possible to include such a composite structure as a term in a larger expression.

The use of composite entities has turned out to be very useful in the design of ontologies and other knowledge representations. It is included in the kernel of the Leordo system and is used in various ways even for representing “system” information in the kernel, as the second introductory example has showed. Another example is for intermediate structures in the version management subsystem. On higher levels of the design, there is an extension of LRX for representing formulas in first-order predicate calculus, in which case composite entities are identified with terms in the sense of logic.

3. *The use of event-state records.* Records are formed in LRX using the notation of the following examples:

```
[date: 2006 10 24]
[quantity: meter 42195]
[quantity: (per: meter second) 46]
```

for the date of October 24, 2006, for the quantity of 42195 meters, and for the quantity of 46 meters per second, respectively. Records differ from composite entities in that they are passive data objects that can not be assigned attributes or properties.

¹⁴<http://www.yaml.org/>

One kind of records, called event-state records, actually allow some of their components to change, but in restricted ways. They are used for representing the current state of an action instance that the system is performing during a period of time, or the current observation state of an event that the system observes. An event record such as (simple example)

```
[move-robot: r14 pos12 pos14 :start 16:22
      :velocity [quantity: mps 4]
      :current-pos [xy-coordinate: 47 12]]
```

where `mps` is an abbreviation for (`per: meter second`), may represent the current state of an event where the robot `r14` moves from position `pos12` to position `pos14`. The record contains the three *direct arguments* of the operator `move-robot:`, and after them the three *state variables* of the event with the respective labels `:start`, `:velocity`, and `:current-pos`. The values of the state variables, except `:start`, can be changed while the event executes in order to reflect the current state of the robot concerned. When the event ends then the ending time is added as an additional state variable, other state variables are added or removed so that the record becomes a representation of the event as a whole and its final effects, the record freezes, and no further changes are possible in it.

An event-state record such as this may be the value of an attribute of an action-instance entity as formed by, for example, the symbolic function `b:` that was introduced previously.

4 The Leordo Kernel and Platform

The Structure and Constituents

The Leordo Kernel consists of four knowledgeblocks, beginning with the *core* which is called `core-kb`. By convention, the names of knowledgeblocks end with “-kb“. The core satisfies all the requirements on the kernel except version management. In addition there is `chronos-kb` that implements a representation of calendar-level time and of events in the lifecycle of an individual, `config-kb` that is used for creating new copies (“individuals”) of the system and for configuring old and new individuals, and finally `syshist-kb` that implements version management. Both reproduction and version management add entries to the system’s history of its own activities which is maintained by `chronos-kb`. For example, a so-called synchronization in the sense of version management is treated as an event in the representation provided by `chronos-kb`.

The more basic aspects of self-modification in the system are implemented in `core-kb`, however. This includes, for example, facilities for allowing the user to edit attributes and properties of an entity, and to add entities and entityfiles.

The core part of the Leordo ontology, called `coreonto`, is an entityfile within the the initial knowledgeblock `core-kb`. Every other knowledgeblock, including the other three kernel blocks, can have their own ontology files that extend preceding knowledgeblocks.

The Core Knowledgeblock, `core-kb`

The following are the contents of the core block, as organized in a number of entityfiles:

- The initial loading or 'bootstrap' machinery. It consists of a few entityfiles that are the very first ones to be loaded when an activation is started, and it prepares the ground for subsequent loading.
- The index file of the core knowledgeblock. (`core-kb`).
- The ontology file of the core knowledgeblock, which is at the same time the core or "top-level" ontology of Leordo as a whole. (`coreonto`).
- Miscellaneous additions to the core ontology that are needed for historical or other reasons. (`toponto`).
- Definitions of procedures for loading entityfiles and parsing the textual representation of entities. (`leo-preload`, `leoparse`).
- Definitions of elementary operations on the structured objects in the Leordo data representation, such as sequences, sets, and records. (`leoper`).
- Miscellaneous auxiliary functions that are needed in the other entityfiles but which have a general-purpose character. (`misc`).
- Major timepoints in the history of the present instance of the system (`mp-catal`).
- Functions for administrating entities, entityfiles, and knowledgeblocks, for example, for creating them and for editing their attributes. (`leo-admin`).
- Functions for writing the textual representation of entityfiles, and for producing the textual representation of Leordo data-structures from their internal ones. (`leoprint`).
- Definitions for a simple executive for command-line operation of the system. (`lite-exec`).

The entityfile `mp-catal` mostly serves `chronos-kb`, but it is initialized in the core block which is why it is present in this list.

Many of these blocks are straightforward and do not require further comment here; their details are described in the systems documentation. I have already described and discussed the data format for

the textual representation of entityfiles. The files for loading and storing that representation (`leo-preload`, `leoparse`, `leoprint`) are direct implementations of the data format. Furthermore I shall discuss the ontology, the bootstrap machinery, the machinery for cataloguing entityfiles using knowledgebase index files, and the facility for defining multiple configurations within an individual. Final sections will describe the other parts of the kernel, namely, the facility for administrating and ‘remembering’ information about calendar-time-level events in the history of a Leordo individual, and the facility for version management of entityfiles.

The Leordo Startup Machinery

One of the basic requirements on the Leordo Kernel is that it shall be able to administrate itself, and as well it shall provide facilities for self-administration of other knowledgeblocks that are built on top of the four knowledgeblocks in the kernel. This self-administration requirement includes several aspects:

- All program code in an implementation shall be represented as entityfiles, without exceptions. This guarantees that general facilities for administration and analysis of Leordo software can apply even to the initial parts of the bootstrap process.
- Since interactive sessions with the Leordo system typically involve loading information from the textual representation of entityfiles, modifying their contents, and re-storing those entityfiles, it shall be possible to edit all entityfiles for software in that way as well.
- However, it shall also be possible to text-edit the file representation of an entityfile and load it into an activation of Leordo, in order for the edits to take effect there.
- In addition, there shall be a version management system that applies to software entityfiles like for all other entityfiles.

The first three of these aspects is implemented using the `core-kb` knowledgeblock; the fourth one using the separate `syshist-kb` knowledgeblock. Notice, however, that the first aspect is a step towards (i.e., facilitates greatly) the fourth one.

The startup process for Leordo activations is actually a good illustration of how a somewhat complex process can be organized around its symbolic data structures. Appendix 4 describes this in some detail.

Configuration Management

One Leordo individual may contain the software for a number of applications, for example for simulation, for robotics, for document

management, and so on. However it may not be necessary, or even desirable to have all of that software and its associated application data present in a particular activation of the system. The individual should therefore have several *configurations* that specify alternative ways of starting an activation. The startup files that were described above serve to define such configurations. In particular, the `kb-included` attribute specifies which knowledgeblocks are to be loaded when the system starts. Knowledgeblock dependencies whereby one knowledgeblock may require some other knowledgeblocks to be loaded first are supported, and are represented by particular attributes on the knowledgeblocks themselves.

Each configuration may also make some other specifications, for example for extra information that is to be loaded in order to start it. Furthermore, each configuration shall specify its user interface, in the sense of a command-line interpreter, a GUI, and/or a web-accessible service. This is done with the `execdef` attribute on the startup-file that was described in Appendix 4.

The Knowledgebase Index Files

Each Leordo individual is represented as a directory structure, consisting of a top-level directory and its various subdirectories on several levels, with their contents. In a predecessor to Leordo, the Software Individuals Architecture, we used fixed conventions for where the entityfiles would be located within the directory structure, and relative addressing for accessing them. This turned out to be too inflexible, and for Leordo we have a convention where each entity representing an entityfile is associated with the path to where the textual entityfile is to be found.

At first it would seem that this should be one of the attributes of the entity that names and describes the entityfile, and that is the first element in the entityfile. However, it would be pointless to put that attribute within the file itself, since the system needs it in order to find the file so it can load it. One can think of two ways out of this dilemma: either to divide the attributes of an entity into several groups that can be located in different physical files, or to construct a composite entity with the entityfile entity as its argument.

Both approaches have their pros and cons. Leordo does provide a mechanism for *overlays* whereby one can introduce entities and assign some attributes to them in one entityfile, and then add some more attributes in an overlay, which is a separate file. However, that facility is not part of the kernel, and we are reticent of putting too much into the kernel. Also, overlays require the entity as such to have been introduced first, before the overlay is added. The attribute for the location of an entityfile is needed before the entity itself is available.

We have therefore chosen the other alternative. The following is a typical entity in an index file for a knowledgeblock, such as `core-kb`:

```
-- (location: leoadmin)

[: type location]
[: filepath "../../../leo-1/Coreblock/leoadmin"]

@Comment
Loading entityfiles and knowledgeblocks, creating new ones,
etc.
```

It defines the location of the entityfile `leoadmin` by introducing a composite entity (`location: leoadmin`) whose type is `location`, and assigning a `filepath` attribute to it⁽¹⁵⁾. Among the files that occur at the beginning of the startup phase, `self-kb`, `kb-catal` and `core-kb` consist mostly or entirely of such entities.

5 Other Kernel Knowledgeblocks

Until this point we have described the design of the core knowledgeblock, `core-kb`. The Leordo kernel also contains three other knowledgeblocks, beginning with `chronos-kb` that enables the Leordo activation to register events and to have an awareness of the passing of time and a notion of its own history. Based on it there is the reproduction facility, `config-kb`, and the versions management facility, `syshist-kb`.

Both reproduction and version management are essential for the evolution of the Leordo software through concurrent strands of incremental change in several instances of the system, i.e. several Leordo individuals. This is the decisive factor for considering these to be an integral part of the system kernel. In addition, by doing so we also provide a set of tools that can be used in applications of several kinds. – The importance of having software tools for version administration do not need to be explained; it has been proven through the very widespread use of tools such as CVS⁽¹⁶⁾.

The following are brief summaries of the services that are provided by these knowledgeblocks in the kernel:

Awareness of Time in the Leordo Individual

The basic contributions in `chronos-kb` are the following:

- A facility for defining and registering significant *timepoints*. Such a timepoint is registered with its date, hour, minutes, and seconds, and it can be associated with the starting or ending of events.

¹⁵Actually this attribute is called `filename` in the current system, for historical reasons. This is due to be changed.

¹⁶<http://www.nongnu.org/cvs/>

- A facility for introducing *events* in a descriptive sense: the system is told that a particular event starts or ends, and registers that information.
- A facility for defining *sessions* which are composite events corresponding to the duration of one activation of the Leordo system, and for defining individual events within the session.

All of this information is built up within the Leordo system, and is maintained persistently by placing it in entityfiles.

System History and Version Management

The system history is a kind of skeleton on which several kinds of contributions can be attached. The first of these is the version management facility which consists of two parts, one that is local within an individual, and one that requires the use of two individuals.

Local version management works as follows. The individual maintains a sequence of *archive-points* which are effectively a subset of the timepoints that are registered by *chronos-kb*. Archive-points have names of the form *ap-1234*, allowing up to 9999 archivepoints in one individual. Each archive-point is associated with the archiving of a selection of files from one particular knowledgeblock. The archiving *action* takes a knowledgeblock as argument, obtains a new archive-point, and for each entityfile in the knowledgeblock it compares the current contents of the file with those of the latest archived version of the same file. It then allocates a new directory, named after the new archive-point, and places copies there of all entityfiles where a non-trivial difference has been identified. The archive-point is an entity that is provided with attributes specifying its timepoint, its knowledgeblock, the set of names for all entityfiles in the knowledgeblock at the present time, and the set of names for those entityfiles that have been archived.

However, the comparison between current and archived version of the entityfile also has a side-effect on the current file, namely, that each entity in the file is provided with an attribute specifying the most recent archive-point where a change has been observed in that particular entity. This makes it possible to make version management on the level of entities, and not merely on entire files, which is important for resolving concurrent updates of the same entityfile in different individuals.

Local version management is useful for backup if mistaken edits have destroyed existing code, but it does not help if several users make concurrent changes in a set of entityfiles. This is what *two-party version management* is for. In this case, there is one 'server' individual that keeps track of updates by several users, and one 'client' that does its own updates and sometimes 'synchronizes'⁽¹⁷⁾ with the

¹⁷This is the usual term, although it is of course a terrible misuse of the word 'synchronize'.

server. Such synchronization must always be preceded by a local archiving action in the client. Then, downward synchronization allows the client to update its entityfiles with those changes that have been incorporated into the server at a time that succeeds the latest synchronized update in the client. If the current entityfile version in the client is not a direct or indirect predecessor of the version that is presently in the server, then no change is made. After that, an upward synchronization identifies those entityfiles whose contents still differ between the server and the client. If the version in the server precedes, directly or indirectly, the current version in the client, then the current version in the client is imposed on the server.

In the remaining cases, the system attempts to resolve concurrent changes in a particular entityfile by going to the level of the individual entities. If that is not sufficient, the user is asked to resolve the inconsistency.

A particular technical problem arises because these synchronization actions require the Leordo activation to read and compare several versions of the same entityfile. The problem is that normally, reading such a file makes assignments to attributes and properties of the entities in the file, but for synchronization purposes one does not wish the definitions in one file to replace the definitions that were obtained from another file. This problem is solved using composite entities, as follows: The procedure for reading an entityfile in KRE format has an optional parameter whose value, if it is present, should be a symbolic function of one argument. If it is absent then the file is read as usual. If it is present, on the other hand, then that function is applied to each entity that is read from the file, obtaining a 'wrapped' entity, and the attributes and properties in the file are assigned to the wrapped entity. After this, the comparisons and updates can proceed in the obvious way.

We have now seen two examples of how symbolic functions and composite entities have been useful even for internal purposes within the kernel. This illustrates the potential value of reorganizing the overall software architecture so that certain, generally useful facilities are brought into, or closer to the system kernel, instead of treating them as specific to applications.

Configuration and Reproduction of Individuals

One of the important ideas in Leordo is that the system shall be self-aware, so that it is able to represent its own internal state, to analyze it and to modify it, and it shall be able to represent and "understand" its own history. Furthermore, all of this shall occur in persistent ways and over calendar time, and not only within one activation or "run" of the system.

We believe that these properties are important for a number of applications, but in particular for those that belong to, or border on artificial intelligence, for example for "intelligent agents". A sys-

tem that acquires information during its interactions with users and with the physical world, and that is able to learn from experience for example using case-based techniques, will certainly need to have persistence. It does not make sense for the system to start learning again each time a new activation is started. It is then a natural step to also provide the system with a sense of its own history.

One must then define what is “the system” that has that persistence and sense of its own history. What if the software is stored in a server and is used on a number of thin clients that only contain the activations? What if several copies of it are taken and placed on different hosts? What if a copy of the system is placed on a USB stick so that it can be used on several different hosts?

In the case of Leordo, the answer is in principle that each individual is a self-contained structure that contains all of the software that it needs. Different individuals may contain equal copies of that software, but in addition each of them contains its own history and its own “experience”. However, it is also perfectly possible for each individual to modify its software so that it comes to differ from the software of its peers.

What if additional copies (individuals) are needed, for example because additional persons wish to use the system? The simplest solution is to have an archive individual from which one takes copies for distribution, but in any case that archive individual will change over time, so a notion of version or generation of the entire individual will be needed. But more importantly, separate strands of the Leordo species may develop in different directions, and a particular new user may be more interested in obtaining a copy of his friend’s Leordo rather than one from the archive.

In principle, a new individual that is obtained from a Leordo individual by copying its software but erasing its history and other local information, is to be considered as an “offspring” and not as a “copy”. If the copy is perfect and all history is preserved in it, then it shall be called a “clone”. The administration of clones offers additional problems that will not be addressed here.

For offspring, the following conventions are adopted. The making of an offspring from an individual is to be considered as an action of that individual, and is to be recorded in its history. Each individual has a *name*, and the offspring of a particular individual are numbered from 1 and up. No individual is allowed to have more than 999 offspring. The first individual under this scheme was called `1ar`, and its direct offspring are called `1ar-001`, `1ar-002`, etc. The offspring of `1ar-002` are called `1ar-002-001`, `1ar-002-002`, and so forth. The abbreviation `1ar` stands for “Leordo Ancestry Root”.

The overall convention for the population of Leordo individuals is now that new individuals can only be produced as offspring of existing ones, so that the parent is aware of the offspring being produced and so that no name clashes can occur in the population. Additional information about when and where offspring are produced is of course

valuable, but can be considered as add-on information.

Notice in particular that version management information is not inherited by offspring, and they start with an empty backup directory as well as an empty memory of past events.

In principle, each new individual should obtain a copy of all the software of its parent. In practice this is quite inconvenient when several individuals are stored on the same host; one would like them to be able to share some of the software files. This has been implemented as follows: Each individual may identify another individual that is known as its “provider”, and when its index files specify the locations of entityfiles, they may refer both to files in its own structure, and files in its provider. An individual is only allowed to update entityfiles of its own, and is not supposed to update entityfiles in its provider⁽¹⁸⁾. When a new individual is created, then it is first produced with a minimal number of files of its own, and it relies on its parent as its provider for most of the entityfiles. After that, it is up to the offspring to copy whatever software it needs from its provider to itself, until it can cut that umbilical cord. Only then is it in a position to migrate to other hosts. Besides, given adequate software, it may be able to import knowledgeblocks and entityfiles from other individuals and not only from its parent.

What has been said so far applies to Leordo-specific software. In addition, applications in Leordo will often need to access other software that is available in the individual’s host for its current activation, for example text editors and formatters. The kernel contains a systematic framework for administrating this.

Facilities for reproduction of individuals were first developed in the earlier project towards the *Software Individuals Architecture*. In that project we considered reproduction and knowledge transfer between individuals to be very central in the architecture, besides the abilities for self-modelling. In our present approach reproduction has been relegated to a somewhat less central position, due to the experience of the previous project.

Other Facilities in the Kernel

The four knowledgeblocks in the kernel also contain a number of other facilities that have not been described here. In particular, there is a concept of a “process” in a particular sense of that word. Leordo processes are persistent things, so they exist on calendar time and not only within one activation of the system. Each process has its own subdirectories where it maintains its local state between activations, and each activation is an activation *of* one particular process. Each process can only have one activation at a time, but different processes can have activations concurrently.

¹⁸This restriction is not enforced at present, but users violate it at their own risk.

6 Platform Facilities

The next layer in the Leordo software architecture, after the kernel, is called the *platform*. This layer is under construction and is intended to be open-ended, so that new contributions can be added continuously as the need is identified and the implementation is completed. The following are some platform-level knowledgeblocks that exist and are in use at present.

Channels

Leordo channels are a mechanism for sending messages between individuals, for the purpose of requesting actions or transmitting information. Each channel connects two specific individuals for two-way, asynchronous communication and is associated with a number of attributes, including one specifying the data format to be used for the messages. The KRE data format is the default choice.

Communicable Executive

The initial example in this article describing the interactions between two Leordo individuals was executed using our *communicable executive* (CX). The basic command-line executive in the kernel is not sufficient for it. CX performs incessantly a cycle where it does three things:

- Check whether an input line has been received from the user. If so, act on it.
- Check what messages have arrived in the incoming branch of the currently connected channels for this individual. If so, pick up the messages and act on them.
- Visit all the currently executing actions in the present individual, and apply an update procedure that is attached to each of them. This procedure may perform input and output, update the local state of the action, and terminate the action with success or failure, if appropriate.

The communicable executive is a natural basis for several kinds of applications, including for some kinds of robotic systems, dialog systems, and simulation systems.

7 The Implemented Leordo System

History of the Experimental Implementation

The design for Leordo started in early 2005. It was based on the earlier experience with the Software Individuals Architecture (SIA), and with several earlier systems before that. The SIA was used as the

platform the a major part of the Linköping-WITAS Robotic Dialog Environment, RDE ⁽¹⁹⁾, which contributed valuable background for the present system.

During the almost three years of Leordo development we have tried to make ‘laboratory notes’ documenting what steps were taken, what design changes were made, and so on. We shall study the possibility of extracting a more concise account of essential design decisions and design changes from these laboratory notes.

Current Usage

The goal of the Leordo project, as stated in the introduction to this article, is to validate or refute the project’s hypothesis concerning the possibility of a fully integrated software system architecture. In order to test this hypothesis it is necessary both to implement the system kernel, and to use it for a few different applications of widely different character.

Two such applications have been fully implemented and are in regular use. This is a way of checking that the system is always kept operational while it is being revised and extended continuously.

The present author uses a Leordo-based software application as his standard tool for the preparation of articles and other documents and for website pages, including the extensive CAISOR website ⁽²⁰⁾. This application has been in daily use since year 2006.

Secondly, Leordo is used for the management and extension of the Common Knowledge Library and its website. This support system is a fairly large collection of routines for acquisition, revision, and presentation of structured information, including version management, IPR management, and type checking of large information structures.

Plans for the future include the porting to Leordo of the previously written applications for simulation of a robotic environment and for user dialog with such a robot.

8 Discussion: the Need for Software System Consolidation

The main goal of the Leordo project, as we stated initially, is to explore the possibility of obtaining a much simpler design of the overall software system in a computer, in particular by reorganizing and re-aligning its major parts so as to eliminate duplication of concepts and of software facilities. It is not yet possible to evaluate the concrete, experimental Leordo system design against that goal, but it is possible to identify how the new design relates to some of the concrete redundances in conventional systems. They are as follows:

¹⁹<http://www.ida.liu.se/ext/casl/>

²⁰<http://www.ida.liu.se/ext/caisor/>

Duplication of procedural language between operating system (shell scripts) and programming languages. In Leordo there is a host language which may vary between generations of the system, but which shall in any case be a language of the ‘interpretive’ or ‘script’ type, such as Scheme, Python, etc. The Leordo kernel provides the command-script situation, and the language can be extended with more facilities, and restricted using e.g. type system, in order to satisfy the needs of other usage situations.

Duplication of notations and systems for type declarations of data structures, between programming languages, database systems, communication systems e.g. CORBA, etc. The two-layered approach to the type system in Leordo was explained in the beginning of section 2. Exactly because the type system is not built into the system kernel, we foresee that it shall be possible to design it in such a flexible way that it can satisfy the varying needs of several kinds of contemporary type systems. This is of course one aspect of the main design hypothesis that was stated at the beginning of the present report.

Scripting languages in various software tools, for example spreadsheet systems, webpage languages such as Javascript, etc.. The idea is that such tools ought to be implemented based on the Leordo kernel and inherit its facilities, including in particular the use of the host language.

Duplication between the file-directory system and database systems. Although the present, temporary implementation of Leordo is based on a conventional OS and makes fairly extensive use of its file system, the long-term idea is to replace it with an implementation of entities and aggregates of entities that is done on directly on the base software. This new information structure shall then subsume what the file-directory system does today.

In the continued work on Leordo we are going to build a number of applications for the purpose of obtaining additional experience with these and other aspects of duplication. At the same time we shall be vigilant about what new duplications may arise as the system and the applications grow in size and complexity.

References

Due to the character of this material, most of the references are to websites that provide information about a particular language or system. These references have been placed in footnotes on the page where the reference occurs.

References to published articles and released reports from the Leordo project can be found on the project website, ⁽²¹⁾. References to published articles from the preceding Software Individuals Architecture project (SIA) can be found on its past project website ⁽²²⁾.

²¹<http://www.ida.liu.se/ext/leonardo/>

²²<http://www.ida.liu.se/ext/caisor/systems/sia/page.html>