Component Configuration Management

Magnus Larsson
Development and Research
ABB Automation Products AB
721 59 Västerås, Sweden
+46 21 342666
Magnus.Larsson@mdh.se
www.idt.mdh.se/personal/mlo

Ivica Crnkovic
Department of Computer Engineering
Mälardalen University
Box 883, 721 23 Västerås, Sweden
+46 21 103183
Ivica.Crnkovic@mdh.se
www.idt.mdh.se/personal/icc

Abstract: Component-based programming is now a widely recognized approach in software development. There remain many open problems related to both technical and non-technical aspects of the components. In this paper, we point out the problem of component identification. Since the components are usually binary units deployed in the system at run-time, we do not have the same insight into their characteristics as into those of the software units which we manage at development time. This problem could be solved if the components had this information integrated together with the binary code, which can be achieve by defining a standardized identification interface. As such interfaces do not exist in standard component models today, this concept can only be used with components built in-house. For external components, extensive tests can, to some extent, compensate for the lack of information. To perform a successful testing efficiently we must limit the number of test cases. Which parts of our system can be affected by the introduction of a component, or by its updating? We can answer this if we can keep track of changes introduced in the system and their impact on the system. These problems are similar to the problems arising at development-time solved by Software Configuration Management (SCM) disciplines. In this paper we point out these problems and make proposals for their solutions at run-time using SCM principles.

1 Introduction

When developing a component independently of system development, we meet a number of problems due to the fact that information we usually have during the component development process is not available. One type of problem is related to the components themselves – the component interface, pre- and post-conditions and the nonfunctional component characteristics such as reliability, resource requirements, timing requirements, etc. Another type of problem is associated with the relationship between the component and the rest of the system. In this paper we address this second type of problem.

When integrated in a system, the new component has an impact on a part of the system. The new component may refer to certain components, and it can also be used by other components. In addition to these explicitly defined dependencies, we also have indirect dependencies, derived from the components which are used by the new

component. Finally, we have implicit dependencies, which are related to the system environment (for example timing or other resource constraints). In general, we can expect that some parts of the system are not affected by a change when introducing a new component or a new component version. This situation is shown in Figure 1.
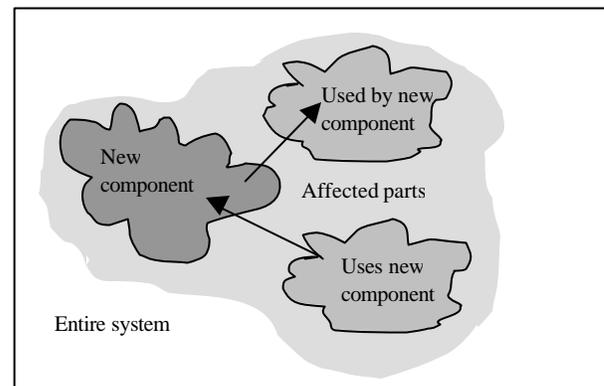

Figure 1. Dependencies between components

The dependencies are not directly visible in component models available today, such as COM [1][2] or EJB [3].

To limit the uncertainty of the system behavior, we must identify those parts of the system which might be affected by the introduction of a component.

If we could identify the component versions explicitly, we could specify the entire system as a set of component versions. Two systems, or two versions of one system can be compared and differences on the component level can be identified.

If we could automatically identify the dependencies between components and their versions we could avoid the well-known problem with different versions of shared libraries. The problem is illustrated in Figure 2: We have two programs *Foo.exe* and *Bar.exe,* which share a *Common.dll* library, version v_1. We then upgrade Bar.exe, obtaining a new version of *Common.dll*, v_2. The replacement could be successful if version v_2 of *Common.dll* is compatible with version v_1, but if this is not the case the *Foo.exe* program can fail. Even if the new version is interface-compatible, *Common.dll* may contain undetected errors, which appear in a combination with

Foo.exe. *Foo.exe* may then import some erroneous code and crash, even if the library was interface-compatible.
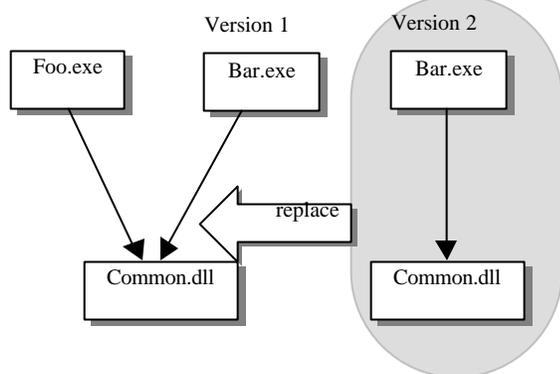

Figure 2.   Uncontrolled update of a component

One way to handle multiple versions of libraries is to insert version information into the actual library name as Microsoft does in MFC [1]. For example, names such as *MFC40.dll* and *MFC42.dll* can be used for version 4.0 and 4.2. This prevents name collision problems but can introduce a vast number of versions over which we have no control.

To identify the parts of the system which can be affected by the change we must:

- Identify components including their versions.
- Identify direct and indirect dependencies.
- Obtain sufficient information to localize the implicit dependencies.

Identification and dependency management is a typical subject of SCM. The SCM disciplines and their possible implementations for managing components are discussed in section 2. In section 3 we discuss the problem with dependency information which is missing from component models available today. In its absence, we discuss a possibility of finding dependencies directly from the code. A Dependency Browser, an application which displays dependencies between binary assets, is depicted in section 4. Finally, section 5 outlines further investigations.

2   Component Management and SCM

As a component is a unit of composition, its management is natural related to Software Configuration Management, the main objective of which is to manage composite entities. However, most of the SCM functions are used at development-time, and are not utilized sufficiently at run-time[5]. The major disciplines of SCM are *Version Management*, *Configuration Management* and *Change Management* [7] [8], and we discuss their use for managing components.

Version management performs the identification of entities and recognizes different versions of enities. We can apply this principle to the components at run-time: Every component in the system should be identified by a name, version number and other version attributes such as creation date, historical information, etc. We need the component version identification for two reasons: Firstly, when we update a component with a new version, we want to be able to identify that change. Secondly, in some cases, we wish to keep several component versions integrated in the system. Managing different versions of components is important for middle-size or large systems. A component might not have been originally designed to cover all the system requirements which evolve. In general, it is better to release a component containing currently required features and to upgrade it later, instead of releasing a fully-fledged component too late. Later, when new features are added to the component, it may happen that the new component version is not compatible with the previous, or that is not fully tested. In that case we want to keep both versions - the new one exploiting new features, and the old one, to be used by those parts of the system we had not yet changed or tested. When the system must support this type of environment, and when several versions of components are used at the same time, the development time and maintenance increase. Experience shows however that this type of evolution is appropriate for large systems [4].

Configuration and build management methods are used to select and identify specific versions of entities (i.e. to generate a *baseline* or a *configuration*) and to integrate them into a new version of the composite entity. It also includes build procedures. The building procedures use information about the dependencies between the entities. These principles can be applied in the run-time system: A system configuration is defined as a set of component versions. By adding a new component or a new version of a component, a new configuration of the system is identified. Similar to *Make* dependencies, which describe the dependencies at build-time, a component should include the specification of the components used (the references to the components used actually exist in the component, but they are hidden in the binary code).

 Change management provides information about the changes introduced in the system on an abstract level, the logical changes which have been introduced in the process, rather then a physical. Change management becomes important when a new entity version is created. In a similar way, every component version can include information about the differences between it and the previous version. This information cannot be automatically generated (which is possible for other type of information, such as version identification and version attributes), and it must be explicitly defined by the component developers. A new component version might be added to introduce new functions in a system, or only to change its behavior, (better performance, better stability), without changing the interface. When replacing a component or a component version we must consider which type of change we permit and which system characteristics we want to preserve, in order to guarantee the system behavior.

To describe this possible impact on the system, we have defined three levels of compatibility:

Input and Output compatibility. A component requires input in a specific format (or perhaps no input at all) and produces results in a defined format. The internal characteristics of the component are of no interest.

Interface compatibility. The interface remains the same, but the implementation can be different

Behavior compatibility. Internal characteristics of the components, such as performance, resource requirements, must be preserved.

The compatibility criteria can be used to decide if a component can be replaced or not. This decision can be especially important in case of a replacement "on the fly" in a run-time environment.

3   Managing Component Dependencies

Binary components are delivered as shared libraries and executables, which usually have no additional information about any dependency between components. To be able to predict what will happen in a system when a component is installed we need to have information about which part of the system will be affected by the component.

As components can be loosely coupled there is no information connecting different versions of components with each other. In COM for example, a component finds components it refers to through the Windows registry. In the Window registry all installed components store their activation data, such as Interface id, class id, library locations and where to find their stubs and proxies. Connections between components are set up first at run-time. A client uses a unique key to find the server component in the registry and then the COM run-time will load the corresponding component or stub into the client memory [2].

To be able to get full dependency graphs over the system with coherent information about all the components, and the type of change introduced in a component, we need meta-data. By meta-data we mean additional information which is not crucial for the component to run but is valuable for the entire system. Meta-data can be provided as a new interface on the component [5] or stored in a repository where it have been placed during the component registration process. Facts about version, name, creation date, compatibility change, interfaces provided and components used, as mentioned in the previous section, are examples of meta-data which can be of help when building a system with consistent configuration management.

The World Wide Web Consortium has defined a standard to describe components and their dependencies. This language is XML-based and is called Open Software Description (OSD) [6]. However, OSD is mainly designed for web components and does not solve the problems with component dependencies. It is important that meta-data is accessible to third part users. A common standard making it possible to describe components in all component models is probably a utopia. We can expect that different types of components will be described in different ways, which is vastly better than their not being described at all.

As we do not have meta-data incorporated in the standard component models, the only information about the components we can get through binary libraries and executables. The information about which shared libraries are linked to other libraries or programs can be gathered fairly easy. In general this information is linked into the binary code and can be extracted. This information can be used to list the dependencies between different programs and libraries.

The following is the formal procedure: A component version c is implemented as a library or an executable. A component version has a set of attributes (name, size, creation-date, and others [11] used in different component models), by which it is identified.

The set of all components installed in the system is designated S. We define a relation $\rightarrow$ called "depends on", where $c_i \rightarrow c_j$ if the correct operation of c_i requires the correct operation of c_j. This relation is transitive which means that we can derive all indirect dependencies from the direct dependencies.

The set of all dependencies is defined as

$$D = \{(c_i,c_j) : c_i,c_j \in S \wedge c_i \rightarrow c_j\}.$$

The dependency set D is stored as a baseline before new components are installed. A snapshot of the current configuration is the set of all components and their dependencies:

$$C = (S,D).$$

When new versions of existing components or new components are installed they will affect the configuration

$$C' = (S',D')$$

We identify all component versions which are placed in only one configuration

$$c_{new} \in S_\Delta \ : \ S_\Delta = \overline{S \cap S'},$$

and the dependencies D_Δ of components c_{new}

$$D_\Delta = \{c : c \rightarrow c_{new}\}.$$

All components in S_Δ and dependencies in D_Δ can change the behavior of the system and are subjects for futher investigation.

For the dependencies where new components use other components

$$D'_\Delta = \{ c_{new} : c_{new} \rightarrow c \}$$

we test if the input-output domain (i.e. expected outputs from c_{new} for inputs to c) have been preserved or not. If a new range of input to the component c occurs, this dependency should be tested in this new domain range [9].

If a system configuration can contain several component versions, specified ranges in input/output domain can be

compared with the current values and used as criteria for selecting a component version to be executed [10][11].

4   Dependency Browser

To show how dependencies can be traced, we have designed an application for Windows NT 4.0, Dependency Browser which parses through the system, finds all shared libraries and generates the dependency graph. A snapshot of the current configuration can be shown and saved in a repository. Different versions of snapshots are placed under version control and treated as configuration items. The current configuration, or an earlier snapshot, can be compared with other configuration snapshots, and the differences between the configurations can be displayed. Typically, before installation of a component, a snapshot can be saved. The component is then installed, and a new snapshot can be taken. The difference graph shows which components have been changed and their relations to other parts of the system. The browser can show the entire system, or a specific component and its dependencies, which makes it possible to see a potential consequence of a component update. System integrators can use the dependency browser to view dependencies in the test system, when a new component has been integrated in the system.

All components which depend on the changed component are highlighted and the user can decide and take action upon this information as shown in Figure 3. The dependency browser helps the integrator of the system to verify that nothing unexpected occurs when the system starts. With this tool, it is possible to see all the files affected when a component has been updated or installed.
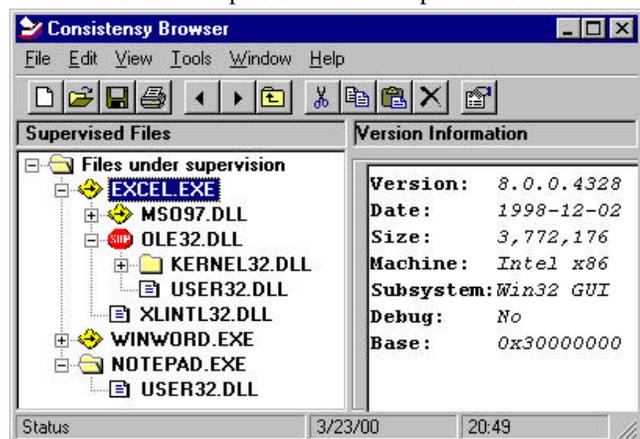

Figure 3.   Affected components are highlighted in the browser to alert the user.

The changed or updated components have the stop sign icon while affected components are marked with an arrow icon. Version information of the component is presented in the right pane view. The browser can be used to browse through the information and to get an understanding of the effects of the introduction of new and updated components in the system. The tool can browse through different configurations and label components as changeable or not changeable. This kind of knowledge is useful if the cause of malfunction in the system is to be traced. An incorrect

version of a library may have been installed by mistake and without dependency information it is difficult to find the real cause of the problem.

5   Conclusion

In this paper we have pointed out the problems encountered during the dynamic configuration of systems. Our contribution is a proposal for component configuration management in which components can be placed under version control. We tie together software configuration management (SCM) and component-oriented programming (COP) with ideas from both disciplines. A simple dependency model is presented and we have shown how to solve the dependency problem for this model when new components are installed. We plan to do more work on a formal description and management of dependencies.

Future work will include the realization of the Dependency Browser, its implementation for different component models and platforms. In this paper we have treated components as binary entities, i.e. executables or shared libraries. How dependencies between loosely coupled components can be recorded, will be studied in a more thorough investigation. The goal is to have ability to predict the behavior of a system before a system update.

6   References

[1]   D. Rogerson, Inside COM, Microsoft Press, ISBN 1-57231-349-8

[2]   D. Box, Essential COM, Addison-Wesley, ISBN 0-201-63446-5

[3]   E. Roman, Mastering EJB, Wiley, ISBN 0-471-33229-1

[4]   M. Larsson, I. Crnkovic, Development Experiences of a Component-Based System, 7th IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2000)

[5]   M. Larsson, I. Crnkovic, New Challenges for Configuration Management, System Configuration Management, SCM-9, Springer 1999, ISBN 3-540-66484-X

[6]   W3C, Open Software Description Format, http://www.w3.org/TR/NOTE-OSD.html

[7]   R. Conradi and B. Westfechtel, Version Models for Software Configuration Management, Software Configuration Management Symposium, SCM-7, 1977, Springer, ISBN 3-540-63014-7, ACM Computing Surveys, Vol. 30, No.2, June 1998

[8]   J. Estublier, S. Dami, M. Amiour, Hifg Level Process Modeling for SCM Systems,

[9]   H. Thane, A. Wall, Formal and Probabilistic Arguments for component Reuse in Safety-Critical Real-Time Systems, Technical report CBSE – State of the Art, Mälardalen University, 2000

[10]   J. E. Cook, J. A.Dage, Highly Reliable Upgrading of Components, 21st ICSE, 1999, ACM ISBN 1-58113-074-0

[11]   Henrik Lykke Nielsen, René Elmström, Proposal for Tools Supporting Component-based programming, Workshop on Component-based Programming, 1999