IRM: the Integrated Reference Model for Open Systems Environments

Mieczyslaw M. Kokar
Department of Electrical and Computer Engineering
Northeastern University
360 Huntington Avenue
Boston, MA 02115

Nancy Daly
The MITRE Corporation
Bedford, MA 01730

September 24, 1996

Submitted to:
Journal of Systems Integration

©1996

Abstract

This paper reports results of a study of reference models for representing open system environments (infrastructures). One of the issues considered in this paper is how to compare particular reference models known in the literature. We propose a set of features of reference models on which to base such comparisons. We have found that in some cases, due to the ambiguity of definitions, it is difficult to assess whether a particular model has some features or not, and that none of the five analyzed models possesses all the features listed in this paper. As a result, we developed the Integrated Reference Model (IRM) which capitalizes on the concepts included in the previously developed models. We provide a description of the basic components of the IRM and suggest directions for refining the model if necessary. We analyze the IRM from the point of view of the desirable features of reference models. We also outline ways in which the IRM could be used in the processes of the design of software infrastructures and in developing software using infrastructures.

1 Introduction

Software systems become increasingly large and complex. Typically, they are required to be able to run on different distributed target systems. They require such services as communications, data base management, sophisticated user interfaces, security services, and others. Additionally, they are expected to be easily adaptable to new updated requirements. For systems of this complexity traditional software development paradigms are not suitable. In order to control development risks, costs and development time the paradigms called *the spiral model* [1] and *software reuse* [9] have been proposed. The main point of the reuse paradigm is the reusability of existing components. Collections of components are kept in *libraries*. When the components include operating systems, communications services, data base services, user interface services, and such, they are called *infrastructures* or *environments*. When interfaces to such environments are defined through nonproprietary specifications they are called *Open Systems Environments (OSEs)*.

Since OSEs are part of software systems, they need to be considered in both software development and software acquisition. Various OSEs need to be specified, characterized, analyzed, compared, evaluated, communicated, and learned. For all of these activities, first of all, there needs to be a language in terms of which such environments can be represented. *Reference models* are used to represent some aspects of OSEs.

Many reference models have appeared in the literature during last five years, each with its own special features, advantages and, unfortunately, limitations. Each model was developed with some special applications and features in mind, and therefore they are limited in the scope of applicability. In this paper we present an attempt to develop a reference model that covers the advantages of several existing reference models in one consistent framework and is still flexible enough to be used as a representation language for open systems environments. We call it the *Integrated Reference Model (IRM) for Open Systems Environments*.

The IRM is based upon the following existing reference models:

- National Institute of Standards and Technology (NIST) Open System Environment (OSE) Reference Model [6]
- Portable Operating System Interface for Computing Environments (POSIX) OSE Reference Model [8]
- Department of Defense (DOD) Technical Reference Model [3]
- Department of Defense Intelligence Information Systems (DODIIS) Reference Model [4]
- NIST Reference Model for Software Engineering Environments Framework [7].

In our study, we first identified the goals for developing open systems environments. These goals are listed in Section 2.1. In general, the goals are related to software development using the environments. Clearly, the existence of a representation of an OSE in terms of a reference model is a very significant, if not necessary, component of both the software development and the software acquisition processes that leads to the achievement of the goals of OSEs.

To rationalize the comparison of various reference models we first identified a set of *features* of open systems environments that need to be represented in a reference model; these features are related to the goals of OSEs. Therefore, one way to compare reference models is by their ability to express those features. Additionally, since reference models are representational languages, they must possess some generic features, common to all representational languages. We present all these features in Section 2.3.

In this study, we analyzed the five reference models [6, 7, 4, 8, 3] with respect to the features that a reference model should possess. Brief descriptions of these five models are given in Appendix A. Our first conclusion is that, in some cases, it is difficult to assess, without reasonable doubt, whether a particular model has some features or not. This is due to the ambiguity of definitions of reference models. Our second conclusion was that none of the five analyzed models possesses all the features listed in Section 2.3. Based on these conclusions we developed the IRM; we present the description of this model in Section 3.

Section 4 contains the analysis of the IRM with respect to the reference model features described in Section 2.3. Section 5 describes how the IRM can be used in the process to support the evaluation of open system infrastructures. Section 6 outlines the use of the IRM in the software development process. And finally Section 7 presents conclusions and suggestions for future research on the IRM.

2 Goals of OSEs and Features of Reference Models

2.1 Goals of Open Systems Environments

Through this study of reference models we identified the goals for OSEs that are listed below. More details can be found in the descriptions of particular models [6, 7, 4, 8, 3].

Improve development efficiency. According to [3] the development efficiency can be improved by the following: common development (applications common to multiple mission areas developed only once), common operating environment (sharing hardware, software, and data), the use of commercial-off-the-shelf (COTS) products (reduce dependence on custom development and thus reduce time of development and ease maintenance process), and reuse of software components in custom development applications.

Enable application portability and scalability. Software developed on a particular infrastructure can be transferred from one platform to another with minimal or

no modifications. Also, hardware upgrades should not affect operations. Scalability means the ability to configure applications for platforms of different size, from micros to mainframes.

Increase application flexibility. A flexible infrastructure allows for ease of incorporating changes and enhancements into the existing infrastructure at relatively low cost.

Improve system interoperability. An application and infrastructure can exchange and mutually use information with other applications and systems. This may be facilitated through a common infrastructure (operating system, database management, data interchange, network services, user interfaces), standardization (common set of services to all applications), or other means.

Improve user productivity. This can be achieved through consistent user interfaces, integrated applications, and data sharing. Consistent user interfaces simplify training as well as facilitate development of future applications and ease of use across applications. Integrated applications, such as office automation and electronic mail applications, integrated with mission specific applications, significantly reduce problem-solving time. Data sharing gives a similar advantage. Database sharing across software systems gives a similar advantage by reducing the time for database creation and maintenance. One database instead of multiples is created and maintained.

Improve user portability. With a minimum of retraining, user's skills, mental models, working strategies, and motor tendencies can successfully be transferred to different applications and systems.

Promote vendor independence. Interchangeable components, both hardware and software, supplied by different vendors can be used in the same system. One factor that promotes this goal is the availability and the use of nonproprietary specifications and open standards during infrastructure development.

Improve security. This goal is affected by such factors as uniform security policy. A uniform accreditation procedure can diminish security problems related to interoperability. Consistent security interface can result in fewer human errors.

Reduce life cycle cost. All the principles listed above support this goal which is, specifically, achieved by reducing duplication such as overlapping functionality and redundancy. This results in reduced software maintenance cost (since there is less software to maintain), and reduced training cost. Also, users moving to other organizations don't need be retrained, since newly developed software has the same user interface principles.

2.2 The Need for Reference Models

According to [6] an *Open System Environment (OSE) is an extensible framework that allows interfaces, services, protocols, and supporting data formats to be defined in terms of non-proprietary specifications that evolve through open (public) consensus-based forums.*

An OSE is intended to provide a set of information system building blocks with associated interfaces, services, protocols, and data formats.

Following the spirit of the above definition, software development that is based on OSEs is an evolutionary process that begins with an OSE as a starting state and continues until the software specifications are met. An OSE thus is a significant part of the product. As such, OSEs need to be specified, designed, compared, evaluated, characterized, learned and communicated. All of these activities require a representation language, and for OSEs, reference models can play such a role. According to [6], a *reference model is a generally accepted representation of a particular application domain that allows people interested in that domain to agree on definitions and build a fundamental understanding within the scope of the model.* In other words, a reference model serves the purpose of representing the application domain and the open system environment. More specifically, a reference model should be able to represent all the elements of a computer system and the interfaces among the elements. This is especially important for Open Systems Environments, since particular elements can be developed by different vendors and a proper specification of these elements is necessary for both their correct implementation and for evaluation.

2.3 Desirable Features of Reference Models

As was mentioned earlier, we decided to compare reference models as representation languages for OSEs. First of all, starting with the goals listed in the previous section, we identified a set of features that reference models should possess in order to be able to adequately represent the important aspects of Open Systems Environments. Some of these features are generic, common to all representational languages, while others are specific to the representation of the Open Systems Environments. In the following we briefly describe the desirable features of reference models.

Interpretability. The meaning of particular representational terms should be unique and easy to interpret. For instance, it is important to clearly define the meaning of graphical primitives, such as boxes, circles, lines, and their interrelations. The topology of the graphical representation of the model should carry some clearly defined meaning. It should be stated explicitly whether adjacent horizontal/vertical lines indicate a relationship between the two services, and what kind of relationship it is. It should be clear whether proximity of two blocks means anything, etc.

Scope and flexibility. A reference model should be able to represent relationships imposed by system designers, for instance a specific reference architecture that needs to be followed due to the regulations imposed by the customer or the government [5]. A reference model is not a reference architecture, but it should be able to represent various reference architectures. A reference model should be flexible enough and have sufficient expressive power to describe many different designs of open system architectures. At the same time, a reference model should not be biased for or against any design solution. If a graphical representation is insufficient to represent all the relevant aspects of an OSE,

formats for supporting documents should be part of a reference model.

Conciseness. A reference model should have a relatively small number of basic representational concepts. The fewer the number of concepts employed, the easier the reference model will be to understand. Also, a smaller number of the basic concepts means a more abstract representation of the system.

Representation of major system components. It should represent all major components of the OSE and of the domain. Specifically, it should include: the software platform, the hardware platform, software-to-software interfaces, software-to-hardware interfaces, protocols and supporting data formats, users of the computer system, external data stores, and communication links.

Representation of relationships among system components. It should represent all important relationships among the system components that are always present in an open system. Two relations that are normally represented in system analysis are *uses* and *part-of*. The former relation means that one of the elements needs and uses services of the other element to deliver its functionality. The latter relation means that one element is part of another element. Specifically, a reference model should fulfill the following requirements expressed in terms of these two relations.

- Any software element *uses* hardware to implement its function and therefore both software-to-software and hardware-to-hardware interfaces are needed for any software element.
- A software element possibly *uses* a number of other software elements, and this fact should be representable in a reference model.
- Users interact with the computer system through hardware (through keyboard, mouse, display) and therefore users are *part-of* hardware. Consequently, user interaction specifications are part of software-to-hardware specifications.
- External platform elements interact with the computer system through (communications) hardware and therefore communications links are *part-of* hardware. Consequently, specifications of external interaction are *part-of* software-to-hardware specifications.
- External data stores are connected with the computer system through hardware and therefore external data stores are *part-of* hardware. Consequently, specifications of interaction with external data stores are *part-of* software-to-hardware specifications.
- In open systems, applications, once incorporated into an OSE, should be accessible to other software elements in the same way as other elements of the OSE. This fact should be representable in the model as a *uses* relation between the infrastructure and the application.

Software standards. It should be able to explicitly represent standards that are followed in the Software Interface Definitions (SIDs) through which particular services

can be accessed within a reference model. Also, the model should facilitate the assessment of conformance to standards.

Hardware standards. It should be able to explicitly represent standards for the Hardware Interface Definitions (HIDs) through which particular services can be accessed. Also, the model should facilitate the assessment of conformance to standards.

Processing model. In a time-critical multiprocessing system, many complex control flow issues must be resolved including shared access to resources, such as memory, CPU utilization, and mass storage. Issues that arise include: deadlock avoidance and/or recovery, process starvation, race conditions, serialization, and synchronization (possibly in real time). To represent all of these issues, the overall processing model for the architecture needs to be specified. It is not necessary to specify the lower level details of the behavior of the underlying operating systems. However, it is important that the scheduling and prioritization policies (e.g., round-robin, pre-emptive scheduling) and time management be specified, since these will directly affect how the applications software architecture of multiple processes will behave.

Whether the processing model is object-oriented needs to be specified, since this will impact the way application software components should be designed. If the architecture allows processes to be specified independent of the data they operate on, application processes will need to be designed as data managers. On the other hand, if the architecture assumes that all processing is performed via concurrent objects, which package processing and data together, then the applications software components will either need to be designed that way, or will need to be encapsulated or shadowed by proxy objects which make the components appear to work that way to their clients.

COTS and Nondevelopmental Items (NDI). It should be useful to identify COTS and NDIs included in the infrastructure and assessment of the degree of integration of COTS with the infrastructure, both in terms of COTS/NDIs being able to access other services, and other services being dependent on the COTS/NDIs.

Software development tools, methods and process. It should be useful for representing software development tools, tool interfaces, tool variety (coverage), degree of integration, and conformance to standards. These tools may include software process development and management services. This representation should be unbiased, i.e., independent of any specific software development process and methods.

Security services. It should allow for the identification and representation of the services requiring special security handling. In addition, it should represent which of these services will invoke security services.

Compatibility with other reference models. Many reference models are already known in the literature, so any newly proposed reference model should be as close as possible to existing reference models. This will allow for comparison of various reference models in case a choice needs to be made.

3 Description of the IRM

3.1 IRM Physical Representation

Figure 1 is a top level view of the IRM. It shows the system elements (represented as boxes) and relations among the system elements (represented as arrows and lines).

- The hardware (HW) layer, which includes: CPUs, memory, keyboards, monitors, external data stores, communication channels, and any other hardware devices. Since the user interacts with the system through hardware, the user is also considered part of the hardware layer.
- The software (SW) layer, which includes the applications programs and all the system services (including software development tools). All software items are commonly termed *services*. Descriptions of the services are provided in the documents describing other reference models [6, 4, 8, 3].
- The interface definitions, which include protocols and supporting data formats, consist of two kinds: software interface definitions (SIDs) and hardware interface definitions (HIDs).
- The arrows represent the *uses* relation among the software elements (services). The IRM shows that these services interact through interfaces defined in the SID layer. The lines represent the *uses* relation between the software services and the hardware. The IRM shows that the software interacts with the hardware through interfaces defined in the HID layer. The *part-of* relation is represented by showing a particular service or a subsystem inside of a box.

3.2 SIDs

SIDs are either textual descriptions or formal specifications of interfaces among software components (software-to-software interfaces). They specify the *uses* relations among particular services. Additionally, they specify the ways in which particular services (or subservices) can be invoked by other services and/or applications, as well as what their functional and nonfunctional characteristics should be. Unlike services, these definitions do not contain any procedures or executable code.

3.3 HIDs

HIDs are either textual descriptions and/or formal specifications of interfaces between software components and the hardware platform. They specify the *uses* relations among

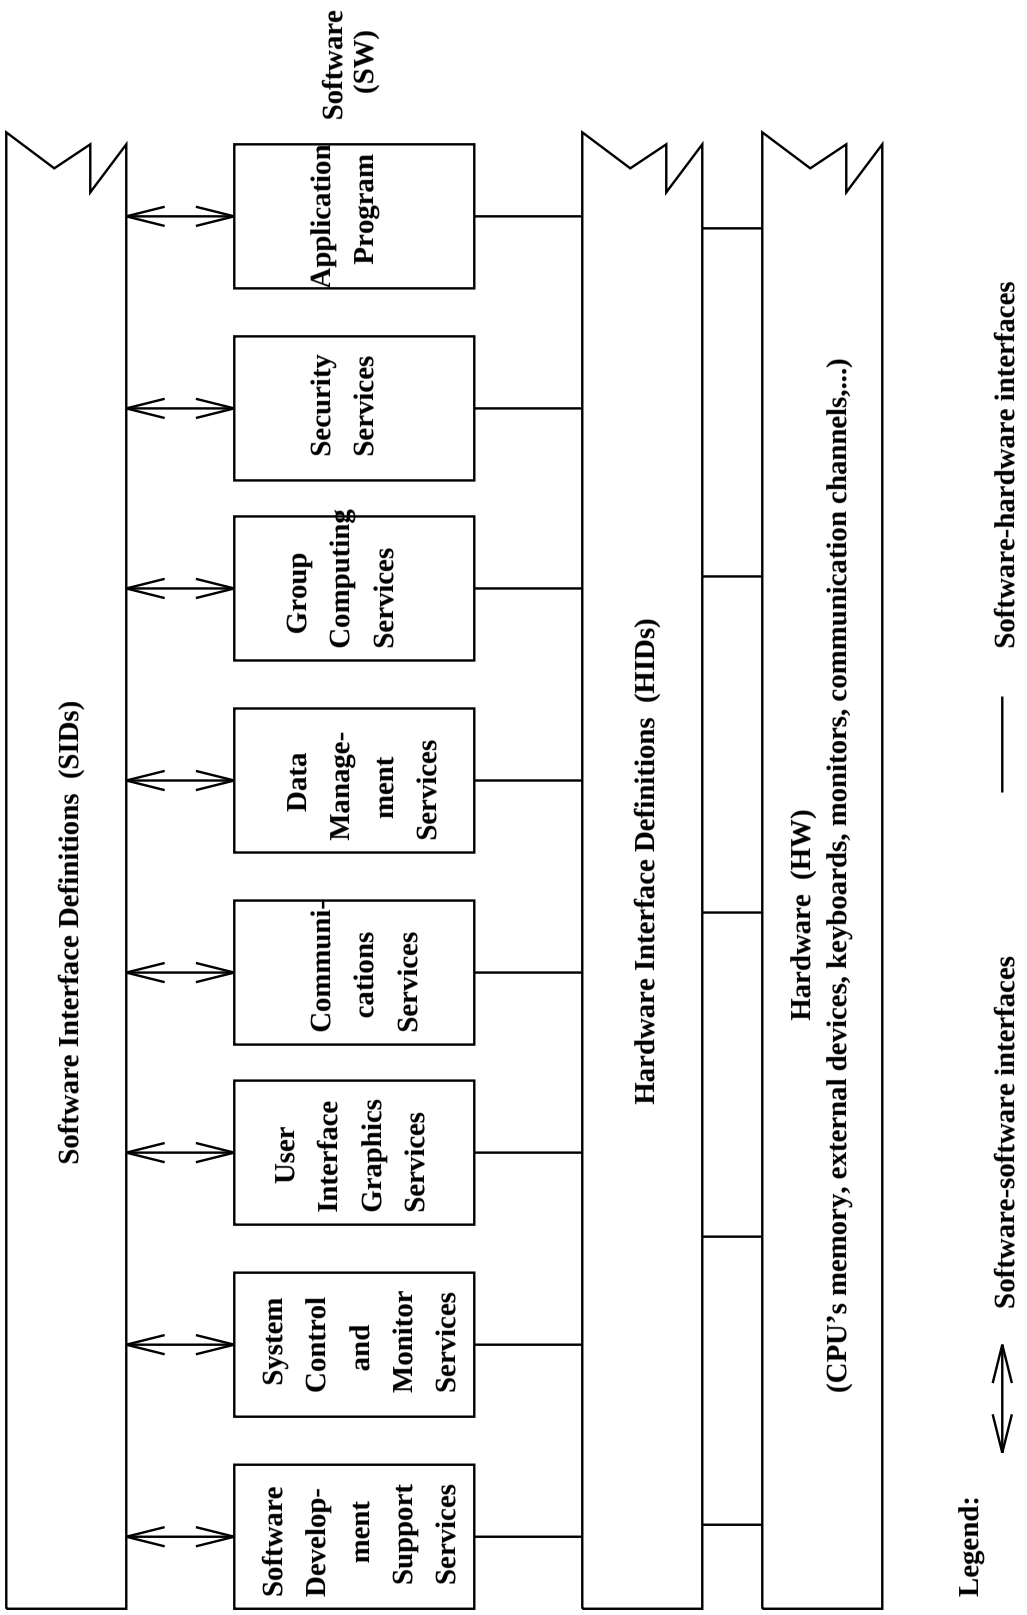

Figure 1: IRM: Integrated Reference Model

the software services and the hardware. This includes interfaces to communication channels, CPUs, memory, external data stores and input/output (I/O) devices, including keyboards and monitors. Since users interact with the computer system through hardware, these specifications also cover user interaction. As a general rule, HID has information needed by system programmers to develop compilers, interpreters, and drivers. In particular, the HID layer contains the following definitions:

- Programming language specifications
- Event (interrupt) interfaces and latency
- Bus interfaces and latency
- Multiple CPU handling
- I/O interfaces and performance
- Data transport protocols (connectivity, interoperation, transfer)
- Time services
- Memory management services
- File services
- User interface protocols
  - Command interface (getting system's attention, interrupting)
  - Look and feel of character-based, block-mode, and format-based interfaces (display characteristics)
  - Protocols for window interaction
  - Window look-and-feel (display characteristics)
  - Graphics protocols and definitions.

3.4 Supporting Documents

In addition to the graphical representation of the IRM, we propose to use, if appropriate, two additional graphical representations: an "integration matrix", showing which of the subservices are accessible to subservices in other service groups, and a layered internal structure of particular services ("hierarchy of subservices"), showing the internal hierarchical structure of some or all the services. The concept of the integration matrix is presented in Figure 2. This matrix defines the *uses* relations among the services shown as boxes in Figure 1. The general idea of the layered representation is presented in Figure 3. The purpose of this representation is to show the *uses* relation among particular subservices within a given group of services.

Services	A1	A2	B1	B2
A1		X		
A2			X	
B1	X			X
B2		X		

Figure 2: An Integration Matrix

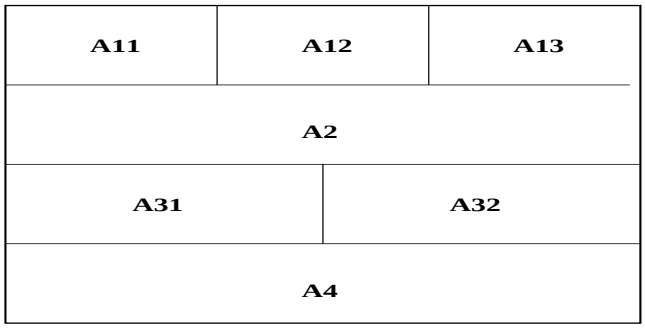

Figure 3: Hierarchy of Subservices

4 The IRM and the Desirable Features

In this section we discuss the IRM in terms of the desirable features of reference models listed in Section 2.3. We also relate the IRM to other reference models.

Interpretability and conciseness. The difficulty with the interpretability of the five analyzed models was one of the reasons we developed the new IRM reference model. We identified the following major issues:

- How to interpret the topology of the graphical representation?
- How to distinguish between internal and external interfaces?
- How to interpret "interfaces"; are they software elements or just specifications of formats and protocols?
- How to assign particular services to particular groups of interfaces, given a list of groups of services and a list of groups of interfaces?

The IRM, on the other hand, has a limited number of representational concepts (thus it is concise) with clearly defined meaning. The graphical representation of the IRM (Figure 1) shows that software runs on hardware, that interaction with the external world is carried out by means of hardware, and that there are definitions of interfaces for

both software-to-software and software-to-hardware interactions. The software services (tools) represented in Figure 1 can at this moment be treated as examples. The reference model does not imply any specific connections (instances of the *uses* relation) between particular services and applications. But it allows to fix such relations at the time of defining a *reference architecture*. The fact that a particular service is next to another service does not imply that the two services interact. Instead, the *uses* relations among services for a specific reference architecture should be represented in a separate document, for instance the integration matrix (Figure 2) and the hierarchy of subservices (Figure 3).

In the IRM, the definitions of the interfaces can be clearly associated with particular system components. The programming interface definitions are associated with particular services and the hardware platform interface definitions are associated with particular hardware components. This should eliminate the vagueness present in other reference models related to the assignment of particular specifications to groups of services and hardware elements.

Scope and flexibility. Because of its simple one-layer software structure the IRM provides an unbiased and flexible representation that can accommodate various reference architectures. Through the accompanying mechanisms for imposing constraints on the interfaces and the interconnections among services, it can support and enforce deliberate constraints imposed by the designers to comply with some standards or design philosophies.

The two kinds of representational constructs used in the IRM (the main graphical representation and supporting documents) allow for representing a reference architecture at two levels of detail. At the higher (more abstract) graphical representation level, only the blocks for particular services are shown. Their interconnections and/or accessibility can be shown in the supporting documents. Within the IRM we are able to explicitly represent the dependency of any services on any other services which is a consequence of the fact that all services are structured in one layer.

In other reference models the integration (the interservice dependencies) is not clearly represented. If we interpret the graphical representations of other models as collections of layers, where only the layer on top of a given layer *uses* that layer, then we are strongly constrained in interconnecting the layers. If on the other hand, as is stated in most of the reference models descriptions, we do not consider a reference model as a layered structure, then this implies that the feature of integration is not representable in those models.

Among the five reference models analyzed in this project, only POSIX Reference Model does not impose any layered structure. All other models do not present this kind of flexibility. The DOD TRM and the DODIIS Reference Model in particular represent the OSE in a highly structured layered way. The NIST Software Engineering Environment (SEE) Reference Model represents the infrastructure as a collection of groups of services, each of the groups represented by a block. The graphical representation indicates some relationships among the groups through the locations of the blocks. For

instance, communication service is positioned at the bottom with all other services laid on top of this service; application software (tools) is connected to Object Management Services and Process Management Services, which in turn are connected to User Interface Services. Although the document describing the NIST SEE states explicitly, "It is very important to realize that the diagram must not be interpreted as a set of layers," the very fact that the diagram contains meaningful topological distinctions either inclines the reader to interpret it as a set of real interconnection constraints, or at least inhibits the nonstructured interpretation. The NIST OSE reference model has little structure – it shows security services as a layer that separates all the services from the External Platform Environment.

Representation of major system components. The elements of a reference model are software, hardware, and definitions, which include protocols and data formats. All of these elements are represented in the IRM, although the specifics of the interface definitions, i.e., protocols and data formats, are not represented as separate graphical objects. (These elements are not represented in any of the other models considered in this study.)

The IRM represents some of the components of an infrastructure that other models do not cover. In particular, the IRM explicitly represents the hardware, and makes a distinction between the hardware interfaces and the software interfaces. Explicit representation of all system elements, including the hardware, is essential because a complete system specification should include all protocols and formats for interaction with these devices.

As was mentioned earlier, the IRM contains two types of interface definitions: SIDs and HIDs. Although at the first glance they look similar to the API and EEI layers in the POSIX, NIST, and DOD Technical Reference Model (TRM), the interpretation is quite different. In the other models, interfaces to external platforms were distinguished as separate entities. In the IRM, on the other hand, the distinction goes along the software/hardware lines.

For instance, the interface to the operating system services, e.g., as defined by POSIX, is defined at SID; only the interactions with the hardware (e.g., language bindings, interrupts) are defined at HID. From the application software point of view, the distinction between external platform and internal platform is not important. Whether a particular piece of data or software is residing on the same or on a different machine should be transparent to the application, and may even change with time. The application itself may also be distributed over the whole system; and, therefore, the internal/external distinction may be meaningless. For these reasons, we do not introduce such a distinction in the IRM model.

However, an explicit representation of the definition of the interface between the software and the hardware platform is essential. This is a true picture of the real configuration in which the application runs directly on the CPU, without interaction with any of the services, until a trap takes place, which causes the transfer of control to the operating system. This side of the specification is especially important for time-critical systems;

the way in which the operating system is invoked, the priorities of the application in relation to other services, and the list of circumstances under which the context switch takes place, should be a part of a complete specification of a system. These specifications are contained in the HID layer of the IRM.

Representation of relationships among system components. The IRM unambiguously represents the *uses* and the *part-of* relations among system components. This is not the case with the other five reference models.

The IRM also represents the instances of these relations that are always present in any architecture. It explicitly shows that software *uses* hardware. This relation is important for the sake of completeness of a system specification. Since users interact with the computer system through hardware (through keyboard, mouse, display), they are *part-of* hardware. Consequently, user interaction specifications are part of software-to-hardware specifications. The User Interface (UI) definition contained in the hardware interface definition (HID) layer is a part of the realization of the concept of "uniform look and feel." This is the place to define these aspects, since the look and feel is perceived through displays, which are part of the hardware platform.

Since external platform elements interact with the computer system through (communications) hardware, communications links are *part-of* hardware. Consequently, specifications of external interaction are *part-of* software-to-hardware specifications. Similarly, since external data stores are connected with the computer system through hardware, they are *part-of* hardware, specifications of interaction with external data stores are *part-of* software-to-hardware specifications.

One of the important aspects of this reference model is that it shows applications at the same level as services (see Figure 1). This means that they can be integrated with other services, which allows us to talk about interdependencies among services and applications. Here infrastructure services can be implemented in an integrated fashion, and functions of one service can call functions of other services. Since some functions can simply be reused instead of re-implemented, cost of software development is minimized. Among the four reference models we studied, only the NIST SEE [7] reference model takes this aspect into account. The other three models deal with integrating an application with the infrastructure but do not cover either the integration across different applications or the integration across services within the infrastructure.

Software and hardware standards. The IRM clearly defines that standards for software-to-software interaction should be located in the SDI layer and that standards for software-to-hardware interaction should be located in the HDI layer. Other reference models do not make such a clear distinction between these two classes of standards. The IRM, similarly as the other reference models, does not impose any set of standards.

Processing model. Since an integrated software environment is an extension of the operating system, the model should support two different views of the software environment, the software view and the hardware view, and similarly the concept of the "user's view" and the "system's view" of the operating system. This distinction is im-

portant for the understanding of the processing model – how does software interact with the hardware. It was stated earlier that the IRM model allows for definition of both software-software and software-hardware interfaces and thus addresses this requirement. In the other reference models this issue is not clearly defined. Similarly as the other reference models, the IRM does not provide means for distinguishing object oriented systems from procedure based systems.

COTS and Nondevelopmental Items (NDI). As it was stated earlier, the IRM shows applications at the same level as services (see Figure 1). This means that COTS and NDIs can be integrated with other services. Among the reference models we studied, only the NIST SEE [7] reference model allocates slots for integrating additional software elements.

Software development tools, methods and process. Both software being developed and software to support the software development process and management (CASE tools) can be explicitly represented in the IRM model. Therefore, this model can be used to discuss the software engineering-related issues. The top level representation of the IRM can be used as a top level representation of either a reference architecture or a specific architectural design. At this level specific groups of services and subservices, as well as elements of computer hardware, can be shown. More architectural information can be added to this representation by showing how particular services are structured into layers and/or parallel blocks. Interfaces and/or standards can be defined in the SDI and HDI layers.

The IRM can also be used for discussing the software development process. For instance, in the development phase of an application, the application is not the program that is being executed. During the development time the developer executes some software development services such as editors, compilers, debuggers, linkers, simulators, testers, etc., as well as some software engineering process services like software metrics calculation, cost and effort estimation, project tracking, etc. Only when development is finished will the application become the final product that runs on the hardware platform and calls on other services. This situation may be reversed again should a need for further enhancements arise.

Other reference models include software development tools and services, but the use of these models for analyzing software engineering issues suffers from all the problems discussed above.

Security services. The IRM allows for the identification and representation of the services requiring special security handling (see Figure 1), but it does not enforce the use of security services by all other services. Through the use of the integration matrix and/or the subservice hierarchy (Figures 2 and 3) the designer of a reference architecture can impose such a requirement selectively.

The NIST reference model [6] shows security services at the bottom of the layered hierarchy, which suggests (although it is not clearly defined) that all other services must use security services. The TRM model [3] shows security services as one half of the

bottom layer, with the other half occupied by the system management services. It is not clear how to interpret such a representation. Some other reference models do not mention security services.

Compatibility with other reference models. The IRM was developed as a result of an analysis of five reference models [6, 7, 4, 8, 3]. We have incorporated as many features of these five models as possible into our model.

5 Using the IRM for Evaluating Open System Infrastructures

One of our objectives was to use the IRM as a mechanism to evaluate and compare infrastructures during the source selection process and during the infrastructure development process while on contract. For source selection infrastructure evaluations, the offerror would be required to complete the IRM, which describes their infrastructure components and architecture, in order to allow to understand the infrastructure to verify that it satisfies the objectives of the system development. One of the most important reasons to use the IRM is that it provides a uniform and unbiased representation of the offeror's or contractor's infrastructure. In the following section we present the description of the offerror's or contractor's part in this process. Then we briefly describe examples of features that can be analyzed using the developed IRM-based representation of the infrastructure.

5.1 Developing a Representation of the Infrastructure in the IRM framework

The first activity in the infrastructure evaluation process is to represent the infrastructure in terms of the IRM. This activity consists of the following steps:

Develop the pictorial representation of the infrastructure. The offerror or contractor shall develop a pictorial representation (Figure 1) of the infrastructure that is being considered during the evaluation. Included in this picture should be all the major groups of services that constitute the infrastructure.

Provide a textual description of the infrastructure. The textual description should include a discussion of which of the general goals and features of open systems environments are included in the offerror's or contractor's infrastructure (see Sections 2.1 and 2.3.) The goals and features of the infrastructure should be discussed in the context of the major groups of services and subservices within these groups.

Develop an integration matrix. Using Figure 2 as the sample format show, which of the subservices in particular groups are accessible to other services or subservices from other groups, and which of the services or subservices from other groups are allowed to

access particular subservices. For instance, referring to Figure 2, service A1 can call only service A2 within the same group, A2 can call B1, B1 can call A1 and B2, and B2 can call only A2.

Show the hierarchy (layers) within each of the groups of services. Using Figure 3 as the sample format, this representation should show the hierarchical calling dependency within particular groups of services.

Provide the interface specifications for each of the services. The interface specifications relate to the software interface definitions (SIDs) shown in Figure 1. The interface specifications should be provided on a per service basis and should include and distinguish between inter-service interfaces and interfaces available to the application software. The interface specification should identify the specific standards followed, if applicable.

Provide the list of standards with which hardware interfaces comply. The interface specifications relate to the hardware interface definitions (HIDs) shown in Figure 1. The hardware interface specifications should be provided on both a per service basis (including application software as one of the services) and on a per group basis, as listed in Section 3.3. The interface specification would identify the specific standards followed, if applicable.

In order to gain a better understanding of the issues involved into developing such a representation we used an abbreviated example of the user interface service and developed its representation using the format of the IRM and the instructions listed above. This exercise is described in [2]. It exemplifies the level of detail that an offerror or contractor should include in the response to the infrastructure evaluation using the IRM.

5.2 Evaluating the Infrastructure

Once an IRM-based representation of the open system infrastructure exists, it can be analyzed from various points of view, depending on the importance and priorities associated with particular features. Ideally, this evaluation should be performed by a team of experts from both the offerror (contractor) and the requestor (sponsor). Below we present an example of features that can be analyzed using this representation.

Consistency of SIDs. Using the integration matrix (Figure 2), identify all possible interactions among services. Check the consistency of all of the software interfaces.

Coverage. Check if all of the generic services and application-specific services (needed for the particular application) are included in the proposed infrastructure and represented in the IRM.

Degree of integration. Using the integration matrix (Figure 2) representing dependencies of particular services on other services, assess the degree of this interdependency.

Conformance to standards of the SIDs. Using both the hierarchical representation (Figure 3) and the integration matrix (Figure 2), identify all the software interface definitions (SIDs) through which particular services need to be accessed. Using the textual description provided by the offerror, identify standards used in their specification. Analyze the conformance of the interfaces to these standards.

Conformance to standards of the HIDs. Using the IRM representation (Figure 1), identify the hardware interfaces. Using the textual description provided by the offerror, identify the standards used in their specification. Analyze the conformance of the interfaces to these standards.

Variety and quality of software development tools. Among services represented in the IRM (Figure 1), distinguish software development tools. Analyze the functionality of tools (variety) and the degree of the integration.

Compatibility of COTS. Analyze SIDs and HIDs of the COTS that will be integrated into the infrastructure. Analyze the compatibility of the COTS with the infrastructure.

Security services. Identify which of the services require special security handling. Analyze the security services provided by the infrastructure. Analyze the integration matrix and check if all of those services invoke security services.

Flexibility and adaptability of the infrastructure. Consider a number of scenarios of change in the requirements and functionality of the application. Assess what and how much would need to be changed in either the application or the infrastructure or both to accommodate the required changes.

6 Using the IRM for Developing Open System Infrastructures

The problem that the software developer is facing is how to develop applications using open system infrastructures. Anybody who had to develop software that incorporates X Windows or Windows had to overcome such a problem. First of all, one needs to understand the functionality, the structure, the scope and the limitations of such an environment. One needs to develop an internal model of the whole environment. The documentation that comes with such environments is either too general or too specific. The documentation usually gives a very general description of the whole environment, relatively good description of the functionality of particular modules, and the source code. A reference model, like the one presented in this paper, provides the structure that serves as a glue for all of the pieces necessary for the understanding of all of the aspects of the environment. Only such an in-depth understanding allows one to begin to develop software for a particular environment.

Leveraging on an existing infrastructure will reduce the initial development time and

let the contractor focus on developing the mission-critical software rather than the support services. An existing software infrastructure that possesses a stable software development environment and the necessary support services will lend itself to supporting an incremental build and integration system development strategy. An existing infrastructure will also allow the contractor to use the actual software to develop rapid prototypes and demonstrate the system performance as it is incrementally developed.

While the software infrastructure offers many benefits, significant challenges exist. Two of the most significant challenges not addressed via the software infrastructure approach are:

- Maintaining interoperability between different infrastructures
- Ensuring that the selected infrastructure does not preclude the use of future standards and technologies.

The software infrastructure approach discussed in this paper leaves the interoperability between infrastructures to be addressed and resolved by the application software, except as facilitated by common communication standards and/or infrastructure services.

It may not be possible to ensure that the infrastructure does not preclude the use of future standards. However, there are a few things that can be done to mitigate this risk.

- Evaluate each infrastructure on a case-by-case basis for the software dependencies on standards and technologies.
- Evaluate the APIs for their ability to support extensibility and flexibility of the infrastructure services.
- Perform a series of flexibility exercises while the system is under development, to ensure that the design characteristics are being considered during the development process.
- Evaluate the risk of using the selected infrastructure for a prespecified minimum duration life cycle.

7 Conclusions

This paper reports results of a study of reference models for representing open system environments. One of the main issues in this study was the decision on how to compare particular reference models known in the literature. Our conclusion was that it can be based on goals of open system environments and features of reference models. Through our study, we first identified a set of goals for open system environments. Secondly, we proposed to consider reference models as representational structures, and identified

features of reference models that are appropriate for representing open system environments in a reference model. As the next step, we studied several reference models. The initial objective of the reference model investigation was not to develop a new one, but to select a reference model that satisfied our view of the C^2 processing domain and open systems. However, upon review of the reference models, we have found that (1) in some cases, due to the ambiguity of definitions of reference models, it is difficult to assess, without reasonable doubt, whether a particular model has some features or not, and (2) that none of the five analyzed models possesses all the features that are important to the development of a C^2 system. As a result, we developed the IRM which capitalizes on the concepts included in the previously developed models and includes the benefits of open systems provided in the models. Some of the advantages of the IRM are listed below.

- It does not impose any model-related biases on the system design and representation, but it allows one to impose constraints, if necessary.
- It explicitly spells out the meaning of all representational elements and, thus, it is clear how to interpret open system environments represented in the IRM.
- It allows us to express both the hardware and the software constraints that need to be imposed on the interfaces.
- It allows us to show not only how a given application accesses particular services, but also how services can access the application and other services.
- We can include software process management services in the definitions of reference models in exactly the same way as other services.

One of our conclusions in this study was that we can more effectively communicate the goals of open systems and the software infrastructure services required to satisfy a specific system via a tailored reference model (IRM). Also, we found that reference models can be used to gain insight into particular infrastructure designs and implementations during the acquisition process. Therefore, we suggest that reference models be used in a different manner than they are often presently being used (i.e., a vehicle in which to identify and levy standards on system development contractors). In this paper, we recommend that the software requestor should provide the IRM (or a tailored version of another reference model) to an offerror with the Request for Proposal (RFP) as a way of describing in detail the requirements and characteristics of the software infrastructure required to implement the system being acquired. We also recommend that the offerrors describe their infrastructure in the context of the reference model. This will allow the requestor to gain an understanding of an offerror's software infrastructure design and implementation in a common and more structured format.

It was not our intention to imply that the reference models analyzed as part of developing this paper are not adequate or that other existing reference models are not adequate for describing the infrastructure services. However, it is strongly recommended

that the service descriptions provided in the reference models be all inclusive of the capabilities and attributes required in the target system. Further, the service description in the selected reference model may require tailoring to accommodate the requirements of the target system under development.

References

[1] B. Boehm. A spiral model for software development and enhancement. *Computer*, 21, no. 5:61–72, 1988.

[2] R. G. Couture, N. Daly, and M. M. Kokar. The command and control (c2) infrastructure reference model (irm). Technical report, MITRE, Bedford, MA, 1994.

[3] DISA. Technical reference model for information management, version 1.3. Technical report, Defense Information Systems Agency, Center for Information Management, Washington, D.C., 1992.

[4] DODIIS. Dodiis reference model for the 1990's. Technical report, Department of Defense Intelligence Information Systems, Arlington, VA, 1991.

[5] R. Hayes-Roth. Architecture-based acquisition and development of software: Guidelines and recommendations from the arpa dssa program. Technical report, Teknowledge Federal Systems, Palo Alto, CA 94303, 1994.

[6] NIST. Application portability profile (app) the u.s. government's open system environment profile, ose/1 version 1.0. Technical Report 500-187, National Institute of Standards, Gaithersburg, MD, 1991.

[7] NIST. Reference model for frameworks of software engineering environments. Technical Report 500-201, National Institute of Standards, Gaithersburg, MD, 1991.

[8] POSIX. Standards project: Draft guide to the posix open system environment. Technical Report P1003.0/D15, IEEE, New York, 1992.

[9] W. Tracz. *Software Reuse: Emerging Technology.* IEEE Computer Society Press, 1988.

Appendix A: Overview of Reference Models

In this appendix we provide a brief overview of five reference models. More details can be found in the documents describing particular models [6, 7, 4, 8, 3]. For each of the reference models, we provide a short discussion in which we identify some of the shortcomings of the models and the accompanying documents.

NIST's Open System Environment (OSE) Reference Model

The NIST OSE Reference Model (see Figure 4) consists of two kinds of elements: *entities* and *interfaces*. The model entities are: Application Software, Application Platform (both software and hardware), and External Environment.

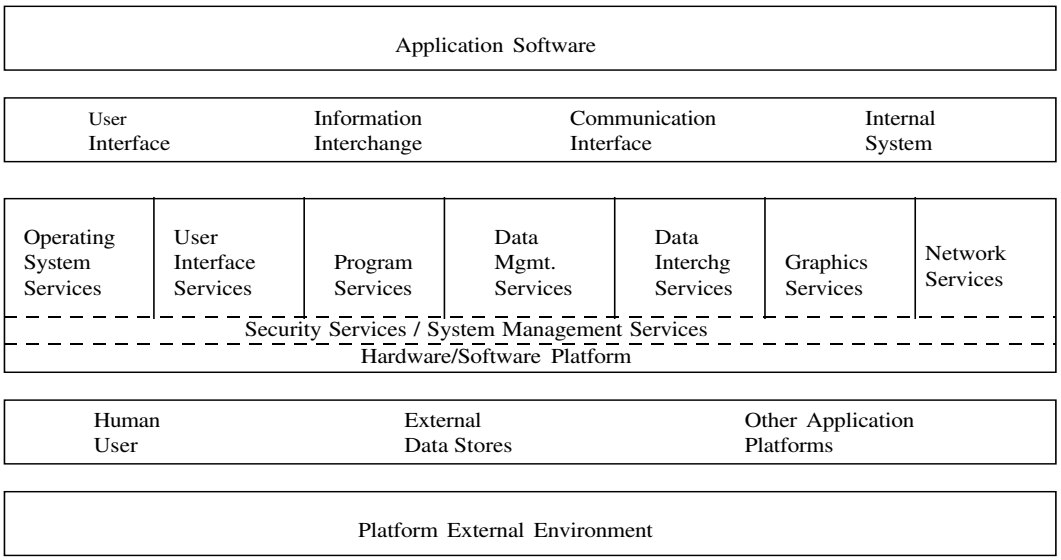

Figure 4: NIST OSE Reference Model

The Application Software is the software that is specific to a customer's application that supports particular business needs. This element includes data, documentation, training and programs. The Application Platform is a collection of hardware and software components that provide services used by the application programs. The Application Platform entity consists of seven groups of services, as shown in Figure 4. The platform External Environment entity consists of system elements that are connected to the application platform, including Human Users, External Data Stores and systems executing on Other Application Platforms.

There are two kinds of interfaces: *application program interfaces (API)* and *external environment interfaces (EEI)*. According to [6], the API is a set of functions between the Application Software and the Application Platform. They are grouped into four groups:

User Interface, Information Interchange interface, Communication Interface and Internal System interface. The EEI consists of three *transfer services* from/to: Human User, External Data Stores and Other Application Platforms.

NIST OSE RM: Discussion

Some of the definitions provided in [6] are not clear. For instance, the definitions of reference model and OSE reference model describe the reasons for having a reference model rather than what a reference model is. The definition of the interfaces are ambiguous. In [6], they are defined as "functions", which suggests that they are executable software. Other documents that use the NIST model interpret them as "specifications". Additionally, it is not clear what is the relationship between the APIs and the application platform services. Also, the NIST OSE model does not explicitly represent hardware. Instead, hardware is treated simply as part of the application platform. The consequence of this is that not only one of the important computer system elements is not explicitly represented, but also the relations between hardware and software (software-to-hardware interfaces) are not representable. The application in the NIST model is represented as an element that uses other services. The model cannot represent the fact that the application can also be used by other applications and/or services. The description of the model [6] does not provide information on the external environment interfaces (EEI).

The POSIX Open System Environment Reference Model

The POSIX Reference Model has been developed to describe, among others, the POSIX Open System Environment (OSE), and in particular to compare and classify standards. The POSIX OSE is a collection of concepts that provide for user requirements and standards specification. The POSIX OSE, in addition to the reference model, includes service definitions, standards and profiles [8].

The top level view of the POSIX OSE Reference Model is represented in Figure 5. Similarly as the NIST model, it consists of three *entities* (Application Software, Application Platform, and Platform External Entity) and two *interfaces* (Application Program Interface (API) and External Environment Interface (EEI)).

The Application Software entity is further decomposed into *elements*: Program Element (source code, command/script files), Data Element (user data, application parameters, screen definitions), and Documentation Element (on-line documentation only).

Although the Application Platform includes both hardware and software, the finer view of this entity shows *services*, which represent software. Apparently, the intent here is to make the configuration related details of the platform transparent to the application software. There are four groups of services: System Services (Language Services, System Services), Communication Services, Information Services (Data Services, Data Interchange Services, Transaction Processing Services), Human Computer Interaction

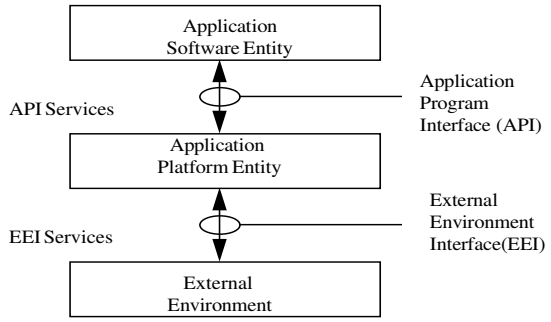

Figure 5: POSIX OSE Reference Model (outline view)

Services (User Command Interface Services, Character-Based User Interface Services, Windowing System Services, Graphics Services, Application Software Development Services). For each of the services a more specific reference model is provided in [8].

The External Environment includes (see Figure 6) People, Information Interchange Entities (removable disk packs, floppy disks, security badges), and Communication Entities (phone lines, local area networks, packet switching equipment).

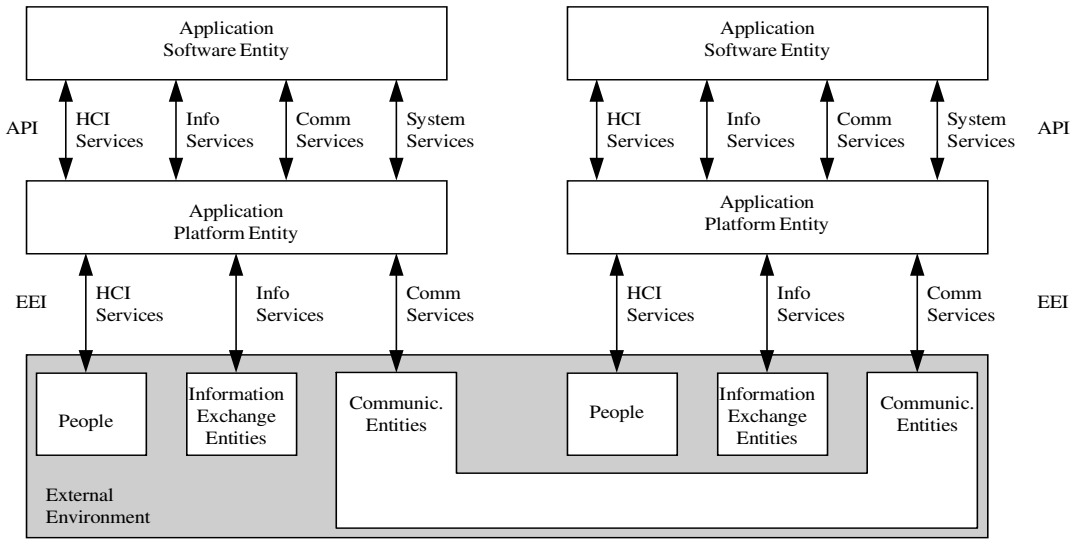

Figure 6: POSIX OSE Reference Model

The API is represented by arrows. The API is subdivided in four groups according to the partition of services. Therefore the interfaces are associated with particular services as described above. Only the System Services API is subdivided further into programming language API (language primitives), language-independent service specifications (platform's services such as interprocess communication, inter-object messages, access to user interface and to data storage), and language binding API (services for translating language-independent specifications into language-specific specifications).

The EEI consists of three groups: Human-Computer Interaction Services (standards related to CRT displays, keyboards, mice, audio input/output), Information Services (format and syntax for interacting with data storage devices), and Communication Services (protocol, syntax and format for communication).

POSIX OSE Model: Discussion

Although in the POSIX OSE model the explicit association of the API with particular services provides a clear interpretation of particular groups within the API, the specification of the EEI is still not precise. Especially, the separation between the API and the EEI (see the desired feature of *interpretability* in Section 2.3) and the impact of these two specifications on the goals of open system environments, i.e., portability and interoperability, are not clear. For instance, although the primary EEI's responsibility is to provide interoperability, it also supports user portability. Similarly, user and data portability are provided by both types of interfaces. E.g., for windows, the program interface standards for accessing windows are associated with the API, while the user standards for windows are associated with the EEI. Requests for services issued through the API are internally translated into requests at EEI.

Additionally, the POSIX OSE has the same drawbacks as the NIST model due to the fact that it does not explicitly represent hardware and that the application is represented at a higher level (on top of) than the application platform (see Section 7).

DoD/DISA Technical Reference Model (TRM)

The DoD TRM reference model [3] is a combination of the NIST model [6] and the POSIX model [8]. We do not provide a separate figure for the TRM, since the graphical representation of the TRM is the same as the POSIX representation: three entities and two interfaces. The entities are represented as boxes, while the interfaces are represented as arrows. But the components of the Application Platform entity – the services – include, in addition to the groups of services specified in the NIST model [6], the Kernel Operations, Commands and Utilities service. The six groups of services are further decomposed into a number of subservices. This decomposition was helpful in matching standards for particular services. Standardization was the main goal for developing the TRM.

TRM: Discussion

Since the TRM is a combination of the POSIX and NIST models, all of the discussion of the two previous sections apply to this model. One feature that is specific to the TRM is that since it is a refinement of the other two models, with more detailed information about subservices, it includes a graphical representation of the subservices. The DoD document

[3] states that the representation does not imply a layered structure. Yet for some reason, some of the service groups have a special place within the box representing the application platform. For instance, Security Services and System Management Services occupy the bottom layer of the box, with Kernel Operations Commands and Utilities, Operating System Service and System Management and Security located on top of them. The remaining six groups of services are on top of the previous three. The more complex topology of the application platform does not seem to be justifiable. Either it should be represented simpler, like in the NIST model, or the topology should carry a meaning, in which case the layout would have to be redesigned.

DODIIS Reference Model

The *generic* DODIIS reference model consists of eleven major components, as shown in Figure 7. The services in this model are classified into five groups; the classification is different from any other reference model. Although three of these groups use the term "object", it is not meant to be an object as in object-oriented software development. Objects in this model refer to "data objects". Each of the services is defined further by a *layered* set of functions in the *detailed model* (not shown in this paper). The Operating System Services are shown as a layer underneath the five groups. Although the model shows the Hardware Platforms, the document [4] indicates that the goal is to have so

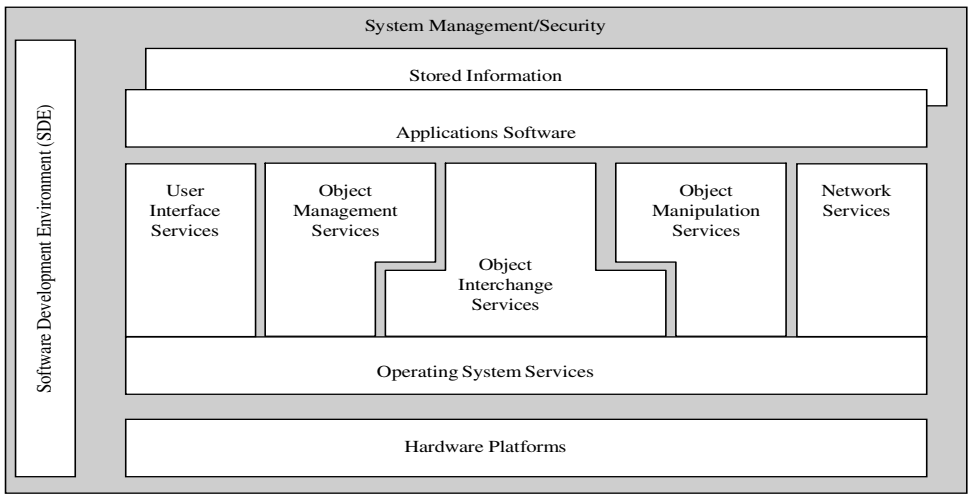

Figure 7: DODIIS Reference Model

Since the main focus of the DODIIS model is information, it includes a separate entity, Stored Information, in which the information architecture is represented. System Management and Security services play a special role; they are represented as a base for the rest of the system. The DODIIS document [4] lists five system management functions: Configuration Management, Fault Management, Performance Management, Accounting Management and Security Management. This document states that all five system management functions should be included in each of the service areas.

DODIIS Reference Model: Discussion

The main difference between the DODIIS model and all other models is that it does not explicitly represent interfaces, even though API and EEI are mentioned in [4]. Consequently, even though it represents some of the system elements, including hardware and data, it still is ambiguous about representing relations among these elements. This ambiguity can be resolved by selecting specific standards for the functionality contained within each rectangle in the model representation. The semantics of the graphical representation is not clear. Although the DODIIS document [4] states that the five services are composed of layers, it is not clear whether the position of a rectangle in the model above another rectangle always means the hierarchy in the system, or not. Many questions are difficult to answer. For instance, does each of the five service groups have interfaces to the operating system? How is Applications Software related to the five groups of services? Can Application Software have interfaces to the operating system? How is Stored Information related to other components? How is the Software Development Environment related to the other components, like Operating System Services and Applications Software? These and other questions are related to the desired features of reference models: interpretability, scope and flexibility, and representation of system components.

NIST's Reference Model for Software Engineering Environments (SEE) Frameworks

The main aim of the SEE reference model [7] is to provide means for describing and comparing software engineering environments, i.e., of systems which provide automated support of the engineering of software systems and the management of the software process. Therefore the goal of the SEE model is somewhat different than the goals of the other models discussed in this paper. The main focus of this model is integration.

The SEE model groups functional capabilities of software engineering frameworks into groups of services. It includes the five groups shown in Figure 8: Object Management Services, Process Management Services, User Interface Services, Communication Service, Tool Slots (place holders for extending a framework), and two groups of services that are not included in the figure: Policy Enforcement Services, and Framework Administration and Configuration Services. Each of the groups is further decomposed into subservices and extensive description is provided for each of the subservices in [7].

The SEE reference model requires that each of the subservices be described in a structured way using the following eight "dimensions": *conceptual* (only functionality described at a conceptual level, no implementation detail), *operations* (list of operations and their functionality), rules (pre and post conditions and constraints), types (data model and meta-data information), *external* (how a particular service can be made available to other services), *internal* (implementation issues), *related services* (which services interact and how), and *examples*. The NIST document [7] includes descriptions of all the basic services in all these dimensions.

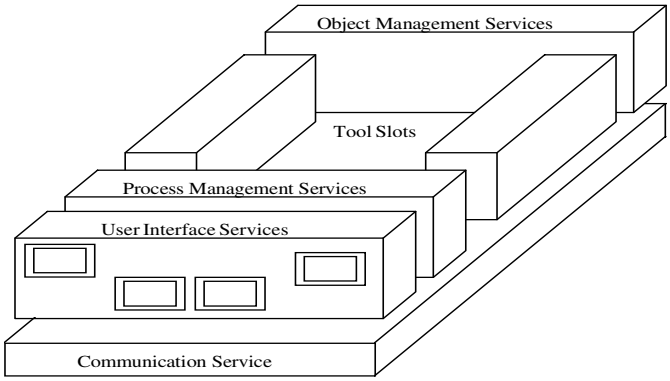

Figure 8: Software Engineering Environments (SEE) Frameworks Reference Model

SEE Reference Model: Discussion

The goal of the SEE reference model can be considered as a subgoal of the other reference models, since all of them include software development tools. The SEE model, however, is much more specific in describing software engineering environments than any of the other models.

The information content of the graphical representation of this model is rather limited. The NIST document [7] states that the grouping of the services is just for the convenience of representation and for the purpose of discussing various aspects of software engineering frameworks. The graphical representation should not be interpreted as a set of layers. Consequently, from the point of view of desired features of reference model, it is limited in the areas of scope and flexibility and in the ability to represent major computer system components and relations among the components. However, it should be noted that the associated description in terms of the eight model dimensions provide much more information than other models.