

ICODES: A Multi-Agent System in Practice

Jens Pohl, Ph.D.

Executive Director, Collaborative Agent Design Research Center (CADRC)
California Polytechnic State University (Cal Poly)
San Luis Obispo, California, USA

Abstract

This paper describes the Integrated Computerized Deployment System (ICODES) from both an architectural and evolutionary vantage point.

First, ICODES is a logistic software application of ship load-planning tools that utilizes intelligent software agents in a human-computer collaborative mode. As an example of a new generation of intelligent military decision-support systems, ICODES includes expert agents with automatic reasoning and analysis capabilities. This is made possible by an internal virtual representation of the load-planning environment, in terms of conveyance and cargo characteristics and the complex relationships that constitute the context within which load-planning operations are performed. ICODES agents monitor the principal determinants of cargo stowage, including: the placement and segregation requirements for hazardous cargo items; the trim, list, stress, and bending moments of the conveyance structure; the accessibility of stow areas through ramps, cranes, elevators, hatches, and doors; the correct placement of cargo items in respect to fire lanes, no-stow areas, reserved stow areas, and inter-cargo spacing tolerances; and, the accuracy of cargo characteristics (e.g., dimensions, weight, type, and identification codes) relative to standard cargo libraries and associated reference tables.

Second, ICODES is a system that has evolved over the past 10 years and is continuing to evolve from a stand-alone application focused on the load-planning of ships to a distributed environment capable of addressing the assembly and planning for any kind of surface or air conveyance. This transition from single domain to multiple domains and from stand-alone to distributed has been made possible by a scalable service-oriented architecture that emphasizes a multi-layered, multi-tiered design approach.

Background

In 1996, ICODES was selected as the *migration* system for ship load-planning by the United States (US) Department of Defense (DoD). It has been deployed by the US Transportation Command (TRANSCOM) through the Military Traffic Management Command (MTMC)¹ to the US Army since 1999, and the US Marine Corps since 2002. Other users include the US Navy and the British Army. ICODES currently interfaces with several external sources that provide it with cargo data, including the World-Wide Port System (WPS) for the US Army, the Transportation Coordinators' Automated Information for Movement System (TC-AIMS II) for several military services, the MAGTF Deployment Support System (MDSS II) for the US Marine Corps, and the Integrated Booking System.

¹ MTMC was renamed in 2004 as the Surface Deployment and Distribution Command (SDDC).

In 2007, ICODES was designated by the Distribution Steering Group (DSG), co-chaired by TRANSCOM and the US Joint Forces Command (JFCOM), as the Single Load Planning Capability (SLPC) for all conveyances with a planned release date of 2010. In this new role ICODES Global Services (GS) will be required to integrate and provide seamless access to communities of planners that have previously operated in separate and largely autonomous domains. The fragmented nature of these logistic enclaves has promoted data quality and exchange problems that have manifested themselves in multiple failure points, leading at times to severe supply chain inefficiencies.

The multiple objectives of the SLPC initiative include the following: to improve the flow of data in a multi-modal transportation and distribution environment; to ensure in-transit visibility from origin to destination; to support collaborative planning efforts; to accelerate the staging, loading and unloading of supplies; to reduce labor requirements through automation; and, to increase throughput without sacrificing in-transit visibility during surge periods. These are indeed ambitious objectives considering the number and diversity of personnel involved and the enormous quantity of supplies and equipment involved. For example during 2005, in its assigned role of DoD's Distribution Process Owner, TRANSCOM and its component commands moved over 2.34 million short tons of cargo and more than 1.1 million passengers.

Load-Planning as a Complex Problem

The rapid deployment of military assets from the US to overseas locations is a complex undertaking. It involves the movement of large numbers of tracked and wheeled vehicles, weapon systems, ammunition, power generating and communication facilities, fuel, food supplies, and other equipment and goods, from military bases to the area(s) of operation. Several modes of transportation are typically involved. Depending on the location of the military base the assets are preferably moved by road to the nearest railhead, from where they are loaded onto railcars for transportation to the appropriate air or ocean port of embarkation.

Alternatively, if rail transportation is not an option, all of the cargo must be shepherded through the public road corridor from the base to the port. At the port of embarkation the assets are briefly assembled in staging areas and then loaded onto aircraft or vessels for shipment. Points of debarkation may vary widely from a commercial air or ocean port with fairly good facilities to a secure airfield in the theater or an amphibious landing on a hostile shoreline under fire. Once the cargo has been disembarked in or near the theater it must be transported to its final destination by road, rail, air, or barge. In many cases this becomes an inter-modal affair with the need for frequent re-planning due to changes in priority or as routes in the theater become temporarily unavailable due to inclement weather conditions or enemy activities.

Speed and in-transit visibility are of the essence (Figure 1). The total time required for the loading and unloading of the conveyance is a critical factor and largely determined by the quality of the load-plan. Ship load-planning, for example, has many of the characteristics of a complex problem situation (Figure 2). First, there are continuous information changes. The vessel that arrives at the port may not be the vessel that was expected and that has been planned for. This means that the existing load-plan is no longer applicable and a new plan has to be developed. Similarly, last minute cargo changes or inoperative lifting equipment may require the existing plan to be modified or completely revised. Second, there are several complex interrelationships. The cargo on any one ship may be destined for several ports of debarkation, requiring careful

consideration of loading and unloading sequences. However, these sequences must take into account unloading priorities that may be dictated largely by tactical mission plans. In addition, the placement of individual cargo items on board the ship is subject to hazardous material regulations and practices. These regulations are voluminous, and complex in themselves. At times they are subject to interpretation, based on past experience and detailed knowledge of maritime risks and practices. Finally, the trim and stability characteristics of the ship must be observed throughout the planning process. This includes listing, draft and deck stress limitations.

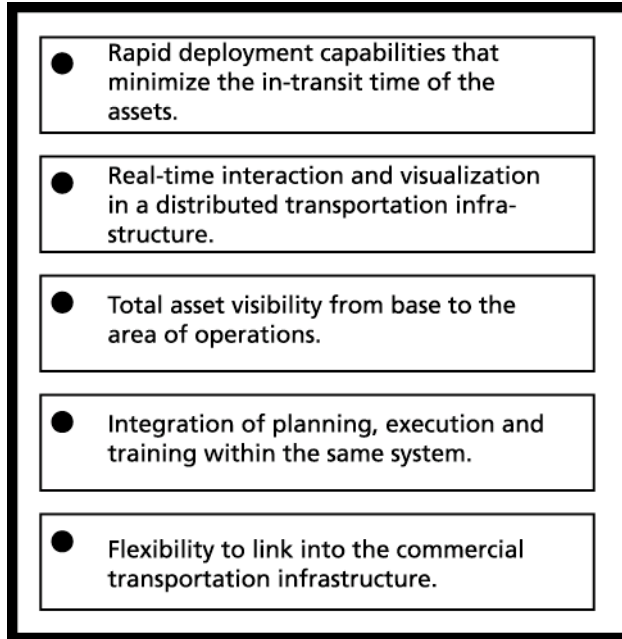


Figure 1: Military deployment objectives

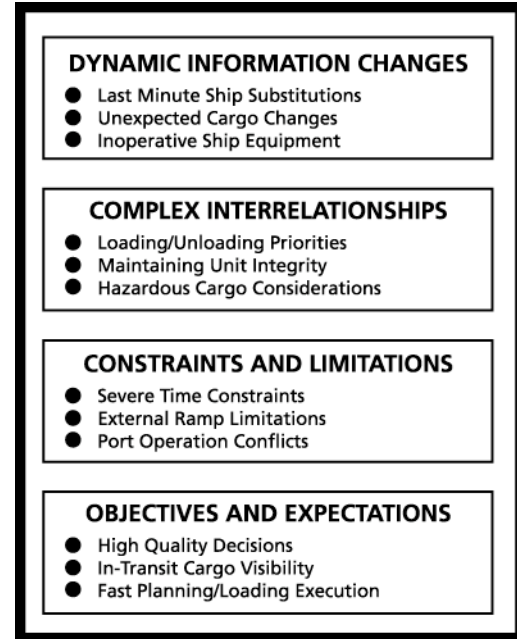


Figure 2: Complexity of ship load-planning

Third, there are many loading and unloading constraints. Some of these constraints are static and others are dynamic in nature. For example, depending on the regional location of an ocean port external ship ramps may not be operable under certain tide conditions, or an airfield may be able to accommodate only a small number of aircraft concurrently on the ground for loading purposes. Local traffic conditions, such as peak hour commuter traffic and rail crossings, may seriously impact the movement of cargo into staging areas or from staging areas to the pier or aircraft loading area. While these constraints are compounded whenever loading operations occur concurrently, the general complexity of the load-planning problem is exacerbated by the number of parties involved. Each of these parties plays an important role in the success of the operation, but may have quite different objectives. Certainly, the objectives of the commercial stevedore crews that may be under contract to carry out the actual loading tasks are likely to differ markedly from the prevailing military objectives that include rapid loading and unloading operations, safety, unit integrity, load density, documentation accuracy, and security.

Initial Functional Requirements

Several general and specific operational and technical objectives were specified by the military sponsor (MTMC) at the beginning of the project in 1994, when ICODES was conceived solely as a stand-alone ship load-planning application. Foremost, it was the vision of the sponsor that

ICODES should present itself to the user as a set of collaborative and expert tools, rather than a conglomeration of predefined solution templates. Experience had shown that the problems encountered in the real world of ship load-planning are driven by dynamically changing factors that are often unpredictable. Accordingly, any predetermined solutions based on preconceived requirements were unlikely to adequately address the nuances of the cargo stowage problem encountered under actual operational conditions.

From a general operational viewpoint the ICODES application was required to be magnitudes faster than the existing DOS-based ship load-planning application. It should allow the concurrent planning of four ships, provide the user with continuous assistance in the form of alerts and warnings throughout the load-planning process, incorporate an automatic cargo placement capability, link to several external systems but be capable of operating in a stand-alone mode, and offer a friendly and flexible, graphical user-interface that could be customized by the user to suit individual needs.

Specifically, the ICODES application was required to automatically alert the user of cargo placements within stow areas that are in violation of hazardous material mandates, the trim and stability requirements of the ship, deck strength limitations, or a host of cargo stowage rules such as adjacency tolerances, fire lanes, boom clearances, and movement restrictions (e.g., door and hatch dimensions, crane lifting capacities and reach, ramp and elevator constraints, and stow area heights). For example, in the hazardous material domain these specific objectives required ICODES to be capable of differentiating among the internationally recognized nine classes of hazardous materials, and the sub-groupings or divisions that exist in five of these classes. In addition, ICODES was required to interpret and apply the regulations prescribed in the following four principal reference sources:

- The 49 Code of Federal Regulations (49 CFR) that specifies segregation requirements for hazardous cargo shipments in the Continental United States (CONUS).

- The International Maritime Dangerous Goods (IMDG) library that applies to all international shipments of hazardous materials.

- The Department of Defense Identification Code (DoDIC) library that applies specifically to Class 1 hazardous items (i.e., explosives), namely munitions.

- The Dangerous Cargo Manifest National Stock Number (DCMNSN) library that is used primarily by the Marine Corps for identifying and load-planning hazardous cargo items.

Technical Objectives and the Development Environment

The general technical objectives established for ICODES in 1994 included the requirement of an open architecture, the ability to add new and enhance existing user-assistance capabilities over the lifetime of the application, the ability to add future modules to support related functional areas such as inter-modal transportation (i.e., air, rail, and truck convoys) and the management of staging areas), as well as the ability for the user to create cargo lists and vessels within the application if these were not available within ICODES and could not be imported from existing external sources.

Like most of the planning and decision-support systems developed by the CADRC Center over the past decade ICODES was designed as a suite of Knowledge Management Enterprise

Services (KMES[®]) and implemented within the Integrated Cooperative Decision Making (ICDM) software environment². ICDM is an application development framework for distributed decision-support systems incorporating software agents that collaborate with each other and human users to monitor changes (i.e., events) in the state of problem situations, generate and evaluate alternative plans, and alert human users to immediate and developing resource shortages, failures, threats, and similar adverse conditions. A core component of any ICDM-based application is a virtual representation of the real world problem (i.e., decision-making) domain. This virtual representation takes the form of an internal information model, commonly referred to as an ontology. By providing context (i.e., data plus relationships) the ontology is able to support the automated reasoning capabilities of rule-based software agents.

Knowledge Management Enterprise Services (KMES[®]) are self-contained software modules with clearly defined functional capabilities and interface specifications. They are designed to be platform independent and to be reusable. Some of these services may have quite narrow capabilities such as the mapping of cargo data imported from an external source to an internal information model or ontology, while others will incorporate larger functional domains such as the optimum placement of cargo of multiple dimensions in a given space, such as a stow area on board a ship or in a marshalling yard.

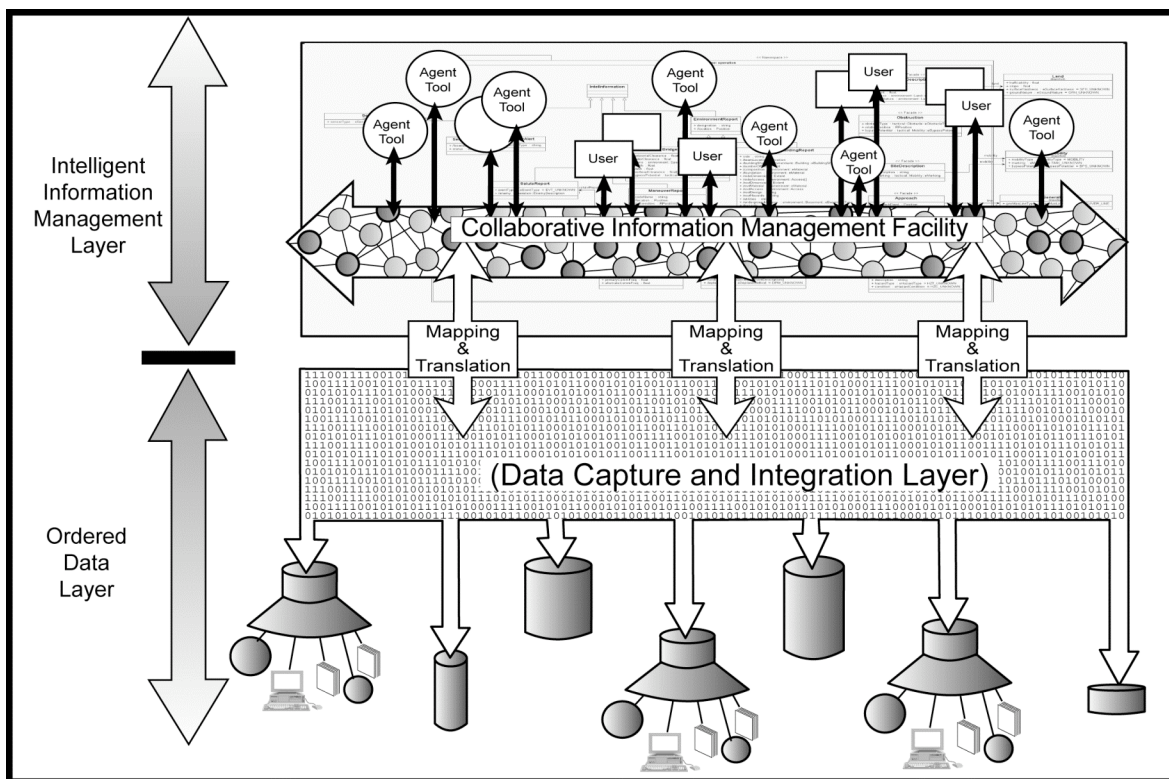


Figure 3: Conceptual KMES-based net-centric architecture

From a general point of view the KMES[®] approach to software systems incorporates intelligent agent technology to provide an internal staff of software agents. These agents analyze and categorize incoming signals and data, and then issue warnings and alerts as appropriate. The

² ICDM is a software development toolkit that is proprietary to CDM Technologies, Inc. and available to third parties under licensing agreements (Pohl et al. 2004).

agents manipulate the incoming data within an internal information-centric representation framework to publish statements of implication, and if so empowered, proceed to develop plans for appropriate action. Legacy data-centric systems can become clients of such an agent-based KMES[®] software environment through the use of interoperability bridges that map the data model in one system to the information model of the other and allow a two-way exchange of data. Conceptually, a KMES[®]-based application consists essentially of two component layers (Figure 3): a data-centric Data Capture and Integration Layer that incorporates linkages to existing data sources; and, an Intelligent Information Management Layer that resides on top of the data layer and utilizes software agents with automatic reasoning capabilities, serving as decision-support tools.

A multi-tier architecture is used to logically separate the necessary components of the data layer into levels. The first tier is the data repository, which ensures the persistence of the data level and provides the necessary search capabilities. The second tier is the service level, which provides the interface to the data level and at the same time supports the data access requests that pass through the mapping interface from the Intelligent Information Management Layer to the Data Capture and Integration Layer. It is designed to support request, response, subscribe, and publish functionality. The third tier is the control level, which routes information layer and user requests to the service level for the update, storage and retrieval of data. Finally, a view layer representing the fourth tier serves as a user-interface for the Data Capture and Integration Layer.

The Intelligent Information Management Layer consists of KMES[®] components in the form of a group of loosely coupled and seamlessly integrated decision-support tools. The core element of each KMES[®] component is an ontology that provides a relationship-rich model of the particular decision-support domain. Typically, KMES[®] components are based on a three-tiered architecture incorporating technologies, such as distributed-object servers and inference engines, to provide a framework for collaborative, agent-based decision-support that offers developmental efficiency and architectural extensibility. The three-tiered architecture clearly distinguishes between information, logic, and presentation. Most commonly an information tier consists of a collection of information management servers (i.e., information server, subscription server, etc.), while a logic tier incorporates an agent engine, and a presentation tier is responsible for providing interfaces to human operators and external systems.

The notion of service-oriented is represented as much in the elements of each of these tiers as it is in the functional capabilities of each KMES[®]. Therefore, even the internal elements of a KMES[®] communicate through standard interfaces as they provide services to each other. They are, in essence, decoupled software modules that can be replaced with improved modules as the technology advances. Each of these modules functions in an integrated fashion to form a comprehensive agent-based decision-support execution framework (Pohl 2007).

The ICDM Software Development Framework: For the past two decades the CADRC Center and more recently CDM Technologies have pursued the design and development of agent-based decision-support systems³ utilizing the ICDM software development toolkit. Not only does ICDM function as an accelerator (i.e., rapid development) and stabilizer (i.e., built-in robustness and fault tolerance) in the development of decision-support systems, but it also provides a concrete vehicle for representing the key concepts and philosophies that the CADRC Center and

³ CDM Technologies, Inc. is the commercial arm of the Collaborative Agent Design Research Center (CADRC) at California Polytechnic State University (Cal Poly), San Luis Obispo.

ICDM have found to be useful for the success of KMES[®]-based systems (Pohl et. al 2004, Pohl 1997). The key design principles on which ICDM is founded are collaboration-intensive, context-based representation, flexibility and adaptability, multi-tiered and multi-layered, within the framework of a service-oriented, distributable architecture. An ICDM-based application is based on an information-centric premise, in the sense that it incorporates an internal information model of objects, their characteristics, and the relationships that associate these objects to each other and the functional capabilities of the application (Myers and Pohl, 1994; Pohl et al. 1992; Pohl K. 2002).

The term information-centric refers to the representation of information, as it is available to software modules, not to the way it is actually stored in a digital machine. This distinction between representation and storage is important, and relevant far beyond the realm of computers. When we write a note with a pencil on a sheet of paper, the content (i.e., meaning) of the note is unrelated to the storage device. A sheet of paper is designed to be a very efficient storage medium that can be easily stacked in sets of hundreds, filed in folders, folded, bound into volumes, and so on. As such, representation can exist at varying levels of abstraction. The lowest level of representation is wrapped data. Wrapped data consists of low-level data, for example a textual e-mail message that is placed inside some sort of an e-mail message object. While it could be argued that the e-mail message is thereby objectified it is clear that the only objectification resides in the shell that contains the data and not the e-mail content. The message is still in a data form offering a limited opportunity for interpretation by software components.

A higher level of representation, commonly referred to as an ontology, endeavors to describe aspects of a domain as collections of inter-related, constrained objects. This allows context to be captured and represented in a manner supportive of software-based reasoning. Apart from their role as services, however, distributed behavioral objects can also be employed as a mechanism for supporting the notion of facades. As one of the fundamental patterns employed in object-oriented design, facades provide a level of derivation attained from the particular representation or ontology on which they are based (Pohl K. 2001). In the case of ICDM and the kinds of ontologies it manages, facades offer a method of supporting and managing an alternative perspective from that modeled in the ontology from which they are derived. In other words, ICDM-based facades allow the perspective inherent in a particular model of a domain to be augmented, or in some way altered to support a more appropriate (i.e., to the façade user) representation of the concepts, notions, and entities over which that user is operating.

Facades can also be utilized to support real-time calculations. In this sense, the façade derivation would involve a calculation or algorithm perhaps based on one or more attributes of the base object(s). For example, consider a stow area on board a ship with its length, width, and height dimensions described in American pound/foot units that is to be accessed by a planning service that understands only Metric kilogram/meter units and also requires space volumes. Utilizing ontology-based facades a model can be easily developed in which, not only the length, width, and height, but also the volume of the space are calculated and presented to the planning service in terms of Metric units. Although there are a number of approaches to supporting calculated attributes in the case where an alternative perspective is to be supported, the façade approach permits an extensible (i.e., one perspective extended from another) and encapsulated (i.e., easily maintainable) solution.

One of ICDM's primary goals is to support a high degree of flexibility in respect to the configuration of its components both at the development and execution levels. ICDM supports

the addition, replacement, and reuse of software components in the context of agent-based, decision-support systems, and achieves this goal by reducing inter-component coupling to an absolute minimum. Two key ICDM properties permit this flexibility. First, all collaboration between clients takes place via, and in terms of the informational ontology (i.e., distributed objects). No direct communication exists between collaborators. The result is a collaborative environment in which client identities are essentially irrelevant in respect to this process.

The second property deals with the manner in which clients access and interact with the ontology. ICDM offers a standard interface component known as the Object Management Layer (OML) that both shields accessing clients from the complexity of ontology management as well as providing an abstracted view of the ontology. Clients of OML interact with the ontology via object wrappers based on a set of corresponding ontology-specific templates. Promoting the notion of adaptability, these templates are discovered by OML as a runtime activity. The resulting support for dynamic definition permits elements of the ontology to be extended, eliminated, or even redefined during the course of a runtime session.

From an architectural organizational point of view ICDM strictly adheres to the principle of separation between areas of functionality at both the conceptual (i.e., tier) level and the more concrete (i.e., layer) level. Conceptually, the architecture of an ICDM-based decision-support system is divided into three distinct tiers namely, information, logic, and presentation. To manage its particular domain each tier contains a number of logical layers that work in sequence. As the name suggests the information tier houses both the information and knowledge (i.e., ontology) being operated on, in addition to all of the mechanisms needed to support management, transport, and access. The information is further delineated into layers. The first of these is the OML described above. Below the OML resides the Object Access Layer (OAL) responsible for managing access to the information tier. The OAL exists as a level of abstraction below OML and interfaces directly with the Object Transport Layer (OTL). Based on the CORBA specification (Mowbray and Zahavi 1995) the OTL is responsible for communicating the various requests and subsequent replies for distributed information and behavior issued through the OAL throughout the system. The OTL is the only layer that forms a dependency on an underlying communication protocol. As such, support for alternative communication facilities can be implemented with minimal impact on either the OAL or the OML. This exemplifies the benefits of a layered architecture in supporting component reuse and replacement.

The Logic Tier contains the business rules (i.e., agents) and analysis facilities by which these rules are managed. Although extensible to include other forms of reasoning the current version of ICDM focuses on opportunistic rule-based analysis. Regardless of which form of reasoning is employed this capability is supported by two layers namely, the Business Rule Layer (BRL) and the Business Engine Layer (BEL). The BRL is primarily system-specific and contains the agent-based analysis facilities resident in the system. Execution of agents is in turn managed by the BEL. To integrate the Logic Tier with the Information Tier the BEL interfaces with OML permitting the agents to both access and contribute to the ontology.

The final tier is the Presentation Tier. This tier is responsible for interfacing with the various users of the system. In this sense a user may be a human operator or an external system. In the case of a human operator support is provided through a Graphical User Interface Layer (GUIL) that presents and promotes interaction with the contents of the Information Tier. In the case of an external system, support takes the form of a Translation Layer that manages the mapping of

representations between systems. Like the GUIL, access to and from the Information Tier is supported by OML.

User-Interface and Functionality

Implemented in a typical Windows 2000 operating system environment the main screen of ICODES Version 5.2 is shown in Figure 4, as consisting of six components or sections.

- (1) The **Main Menu Bar** provides access to the nine principal ICODES option groups in the form of pull-down menus.

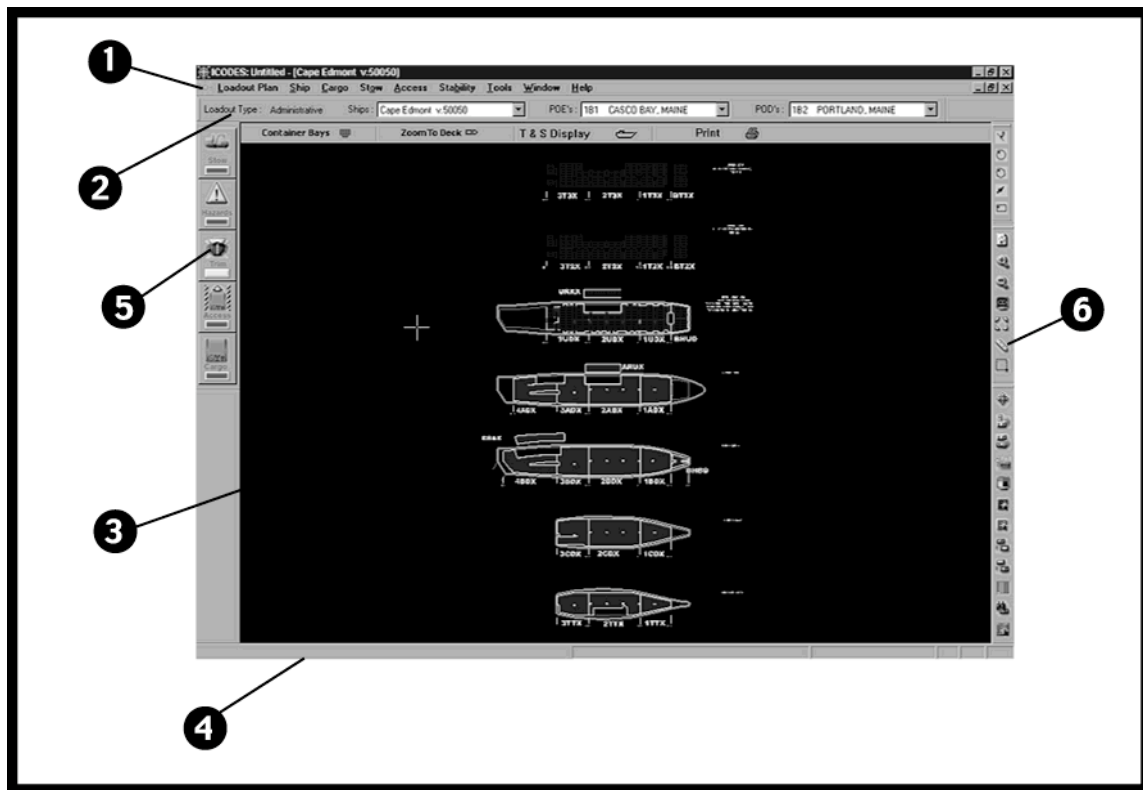


Figure 4: The Main Screen of ICODES Version 5.2 (Released August 2002)

- (2) The **Loadout Banner** provides information about each of the currently displayed load plans such as plan type(s), ship name(s), ports of embarkation and debarkation, and the measurement units used in each plan.
- (3) The **Graphics Window** displays the ship drawing(s). It can accommodate multiple ships, with the number of ships that are concurrently displayed limited only by the constraints of the screen size and the memory capacity of the computer.
- (4) The **Message Window**, found at the bottom of the main screen, provides the user with messages relating to the current status of ICODES (e.g., the status of an option selected by the user, or instructions relating to the use of a particular tool).
- (5) The **Agent Status Bar** on the left side of the main screen provides access to agent reports and explanations of warnings and alerts.

- (6) The **Tool Bars** on the right side of the main screen contain three groups of tools: stow tools (e.g., rotate, flip, unstow individual cargo items); view manipulation tools (e.g., zoom, pan); and, drawing tools that allow the user to superimpose lines, circles, polygons, and rectangles, on a displayed ship drawing.

ICODES offers a very comprehensive set of editing, saving, restoring, reporting, and special operations options (MTMC 2002). In addition, ICODES recognizes the differences among tactical (emphasizing mission accomplishment), pre-positioning (accommodating the maintenance requirements of pre-loaded regionally positioned ships) and administrative (focusing on the maximum utilization of troop and cargo space) load-plans.

The development of a load-plan can be undertaken in either of two modes. In the *User Stow* mode the user selects a cargo item from a textual cargo lists, ICODES automatically converts the selected item into the appropriate graphic cargo symbol, and once the user has placed the cargo symbol in a stow area the agents assess the impact of the cargo item in that position on both the validity of the load-plan and the condition of the ship. The agents take into account: the path of the cargo item from the dock to its final location on the ship (e.g., availability of ramps, cranes and elevators, and the dimensions of doors, hatches and openings); the segregation and other special requirements related to hazardous materials; and, the trim and stability conditions of the ship.

In the *Assisted Stow* mode the user is able to define specific parameters at the cargo and ship levels and then request ICODES to automatically stow the cargo on one or more ships. Parameters include the establishment of preferences for individual stow areas, the exclusion of stow areas, the specification of spacing distances between cargo items, the orientation of cargo items, and the selection of subs

ets of the cargo list. Once the parameters have been specified (either by default or user selection) ICODES will automatically prepare a load-plan that does not violate any of the rules and regulations known by the agents.

Expert Agent Capabilities

There are many definitions of software agents in the literature (Wooldridge and Jennings 1995; Bradshaw 1997). To the author, a software agent in its simplest form is a software module (i.e., service) that is capable of communicating with other software modules or human agents to facilitate some action. However, at this level of definition an agent is not necessarily intelligent. An intelligent agent would need to communicate using a common language (such as the ontology represented by the Semantic Network in ICODES) to support reasoning capabilities. In addition, an agent may have deep information and expert skills within a narrow domain and would then be referred to as a knowledge-based agent that has the ability to act on its own initiative. Such agents typically collaborate with other software and human agents to accomplish goals, and use local information to manage local resources.

The expert agents in ICODES are designed to assist the load-planner in the knowledge domains of hazardous material, trim and stability of the ship, cargo access paths, cargo attribute verification, and the actual placement of cargo in stow areas. The agents do not communicate directly with each other, but are totally decoupled. In fact, they do not know about each others

existence. They collaborate indirectly as clients through a subscription service that allows them to post interests to data changes within the context provided by the ontology.

When the user is developing a load-plan while operating in User Stow mode, the agents will alert the user to any violations by turning the surround of the appropriate agent status window red. The user can then click on the status window to display a window with an explanation of the violation. In fact ICODES provides several different types of agent warnings:

- A **yellow** surround of an agent status window provides warning of a situation that could lead to a potential violation.
- An **orange** surround of an agent status window indicates that a warning has been acknowledged but still exists.
- A **red** surround of an agent status window indicates a violation (i.e., the existence of a serious problem).
- A **purple** surround of an agent status window indicates that an alert (i.e., serious violation) has been acknowledged but still exists.

If the user operates in Assisted Stow mode the agents will collaborate to place the cargo in such a manner that there are no violations. Cargo items that could not be placed in any stow area without causing a violation are simply not stowed. Brief summaries of the functional capabilities of each ICODES agent are provided below.

The **Stow Agent** supports both manual and automatic load-planning operations. Using default settings in the automatic mode (i.e., Assisted Stow), the Stow Agent attempts to place the heaviest cargo items as low as possible on the ship without causing a violation. This results in a low center of gravity for the ship, which is desirable in most cases. The Assisted-Stow mode provides a comprehensive set of settings. This allows the user to define exclusive and inclusive constraints and preferences in respect to both the cargo that is required to be stowed and the stow areas that have been designated as being available. The Stow Agent checks to see that the placement of a cargo item does not overlap another cargo item, a fixture of the ship such as a stanchion or fire lane, or if the item is not entirely within a stow area. In Assisted-Stow mode, the user can also set the front/back and side to side spacing requirements of a cargo item (e.g., 18 inches front and back and 6 inches side to side) and the Stow Agent will abide by these settings so as not to stow within that imagery boundary around each cargo item.

Other parameters checked by the Stow Agent include the ports of embarkation and debarkation to ensure that they match the ports indicated in the voyage documents, and the height of each cargo item to ensure that the latter can reach their final stow positions. The Stow Agent automatically adds a safety cushion (specified by the user) to the actual height, which is set by the end-user, to make sure that height plus the cushion does not exceed the maximum allowable height for cargo in that stow area and the access path to the stow area..

While in the Assisted Stow mode ICODES will ensure that the automatically generated load-plan has no violations, in manual mode (i.e., User Stow) ICODES will allow the user to stow cargo items that are in violation. However, the Stow Agent will alert the user of the violations and provide an explanation on request.

The **Trim and Stability Agent** checks the placement of cargo items on the ship to see if they violate any desired (i.e., user specified) or mandated maximum draft settings, strengths (i.e., bending of the ship) or deck stress limitations. The Stow Agent in automatic mode will rearrange the placement of cargo during the Assisted Stow process if the placement of cargo causes the upper limits of the strengths properties of the ship to be exceeded. For example, if the predefined stow order requires the middle two stow areas of a deck to be stowed first and second, this would result in a sagging condition of the deck. Under these conditions the Stow Agent will automatically redefine the stow order used by the Assisted-Stow process, so that the placement sequence of the cargo will begin with the forward and aft areas of the deck (thereby preventing the occurrence of a sagging condition).

ICODES calculates the effects of the exact placement of every cargo item stowed on the ship in three different planes. These planes are: forward to aft often referred to as the Longitudinally Center of Gravity; side to side or Transverse Center of Gravity; and, up and down or Vertical Center of Gravity. The Trim and Stability Agent takes into account the combined effects of all of the cargo items, the ballast, and the original condition of the ship to provide the user with fairly accurate estimates of the center of gravity in each of the three planes, as well as an overall assessment of the stability of the ship.

The **Access Agent** checks all paths to ensure that a cargo item can be stowed in a particular stow area. This includes openings, doors and hatches, differentiating between cargo that is loaded with cranes through hatches (i.e., LOLO: Lift On Lift Off) and cargo that is driven or pulled into stow areas (i.e., RORO: Roll On Roll Off). Under Assisted Stow conditions, if there is a violation in the stow path of a particular cargo item the Stow Agent will not place this cargo item in that stow area but will attempt to place it in another stow area. In this situation the violation is transmitted indirectly from the Access Agent to the Stow Agent without notification of the user.

In manual mode (i.e., User Stow), on the other hand, if a cargo item is placed in a particular stow area for which all of the possible stow paths register an access violation then the Access Agents will inform the user that the cargo item has a violation for every path to the stowed location. In addition, the Stow Agent will identify for the user the shortest stow path and the nature of the violation that is associated with that path.

ICODES allows the user to edit the ship characteristics, including the usability properties of the cranes and the dimensions of doors, openings and hatches. Since the Access Agent utilizes the current ship characteristics as the existing constraint conditions, these changes will be reflected in the actions of the Stow Agent in automatic mode and the alerts provided by the Access Agent in manual mode.

The **Cargo Agent** checks the characteristics of each cargo item against the expected characteristics for that cargo item recorded in the Marine Equipment Characteristics File (MECF) or Tech Data cargo libraries. Not all cargo characteristics can be verified in this manner. These cargo libraries currently contain more than 20,000 items, but are restricted in terms of the attributes that are provided for each cargo item. Typically, this verification process is complete and reliable only for dimensional (i.e., length, width and height) and weight attributes. If discrepancies are detected the Cargo Agent generates warnings.

The **Hazard Agent** verifies the proper placement of hazardous cargo items in reference to the various hazardous material codes and regulations discussed previously. It considers issues such as: Is the cargo item stowed in an acceptable deck location according to its stowage requirements? What are the segregation requirements for the cargo item, taking into account both the type of cargo item (e.g., break-bulk, container, vehicle) and the proximity of any other hazardous cargo items? In the case of containers, the Hazard Agent considers the hazard category of each item in the container in assessing the hazard condition of the container and its location relative to any other hazardous cargo item on the ship.

Operational Performance Assessment

During the more than eight years of ICODES releases and wide-spread military use no quantitative metrics have been collected to compare military ship load-planning from the period prior to the availability of ICODES (i.e., prior to 1997) and after ICODES became the system of record for Army, Marine Corps, and Navy surface load-planning. However, it is generally accepted within the military load-planning community that ICODES has been responsible for a dramatic improvement in decreasing the loading time of ships and berthing costs. In addition, ICODES further proved its utility in unanticipated areas, such as ship selection for the movement of supplies, cargo in-transit visibility, historical analysis of cargo movements, and ship design. The following selected areas of military load-planning operations may serve as indicators of the improvements in operating efficiency and cost savings that have been achieved through the deployment of the ICODES suite of adaptive tools over the past several years.

Load-planning efficiency: Previous to the fielding of ICODES, the creation of a pre-stow plan would often take one load-planner using the DOS-based CODES software at least two days. Once the cargo list had been cleansed, through the laborious manual process of comparing the data pertaining to each cargo item with the official equipment library, often a day long process, the load-planner would copy-and-paste the cargo symbols on the ship deck drawings. Then other planners with expertise in hazardous cargo stowage, trim and stability, and cargo flow would check the plan, which often took another day. This time consuming cycle would begin again for each time the cargo list was updated, often up to 30 times during the development of a pre-stow plan.

With ICODES, and in particular through its agents (i.e., Cargo, Access, Trim and Stability, Hazard, and Stow Agents), a load-planner is able to create a similar pre-stow plan in about three hours. When updated cargo lists arrive the ICODES *merge* function allows the same plan to be updated within minutes without re-starting the planning process.

Marine Corps cargo specialists have indicated that prior to the availability of ICODES the planning of the equipment for a Marine Expeditionary Unit (MEU) involving 10 to 14 ships would take an Operation Planning Team five to seven days. With ICODES this task has been reduced to about 14 hours.

In-transit visibility: An area of support that did not exist prior to ICODES is the electronic submission of cargo manifests and cargo ship placement reports to the ship personnel and to the Port of Debarkation (POD) staff. This capability has provided visibility of cargo to the ship to assist with in-transit issues, to the POD for off-load-planning and/or load-planning of new loads, and to military administrative personnel for tracking and historically reporting on cargo movements.

At a POD, prior to ICODES, immediately after the arrival of a vessel a cargo survey and meeting would be held to discuss cargo placement and off-loading strategies. With the availability of ICODES documentation this half-day delay is no longer necessary resulting in a significant saving of berthing costs. In addition, the off-load-planning that can now be accomplished with ICODES prior to ship arrival results in substantial labor and off-load space assignment savings.

For ships with multiple ports of loading and discharge, ICODES load-plans are now passed electronically from port to port to determine the effects of the loads and off-loads on the ship and provide a common operating picture. Beyond the port, the Army Logistic Operations Center uses a database of ICODES-generated load-plans to estimate off-load times. In the past this has been a labor intensive operation, often resulting in missed deadlines.

Trim and stability analysis: Since the ICODES Trim and Stability Agent utilizes certified formulas for ship trim and stability calculations, the results are not only used by load-planners but also by the ship's crew to confirm ship loading conditions. Because of the trusted quality of the validated ICODES trim and stability analysis, ships are much less prone to unsafe stow configurations and further, sail up to a day earlier than in the pre-ICODES era. The earlier departure of ships leads to fuel savings since ships are able to proceed at reduced speed and still stay on schedule. In addition, ships stowed using the precision and operational knowledge offered by the ICODES system experience decreased port costs associated with berthing and service fees.

Prior to the availability of ICODES ships were often loaded with little concern for the distribution of weight along the ship's perpendicular axis, eventually causing several classes of ships to develop stress fractures. The continuous monitoring of the condition of the ship during load-planning has led to better load distributions and the resultant reduction in costly ship repairs.

Reconciliation of planned cargo placement: Using the ICODES Automatic Information Technology (AIT) capabilities, the staging area cargo placement and the ship as-loaded plan is confirmed with hand-held Personal Digital Assistants (PDA), as opposed to manually drawn sketches and tally sheets. Using the ICODES AIT functionality, personnel costs have been reduced to about 20% of the cost of the manual process and the number of port cargo administrative personnel have been reduced by about 50%. With the increasing availability of AIT wireless communications at ports cargo locations are updated automatically to an ICODES computer in the port command center, allowing near real-time visibility of cargo to port administrative personnel and preventing the misplacement of hazardous materials.

Since its first release as a system of record in 1997, the granularity of the cargo data has increased greatly as ICODES moved from Level 4 to Level 6 detail. A typical Army cargo list in 1997 seldom included more than 2,000 individual cargo items. From 2004 onward ICODES has been required to process Marine Corps cargo lists with more than 30,000 individual cargo items. Despite this increase in the volume of data the performance of ICODES, in terms of response time, has continued to reduce as well. The typical performance results shown in Table 1 are based on periodic metrics collected by CDM's ICODES Test Group over the past eight years.

Table 1: Historical ICODES performance metrics

Tested Procedure	V 3.0 (1998)	V 5.0(2001)	V 5.4 (2005)
------------------	--------------	--------------	--------------

Create two-ship load-plan with 2,400 normal cargo items	20 min	8 min	1.5 min
Create two-ship load-plan with 1,200 hazardous cargo items	25 min	11 min	2.5 min
Unstow inventory of 2,400 items from two ships	10 min	5 min	1.0 min

ICODES System Architecture

As a KMES[®]-based system, designed according to service-oriented architecture principles and implemented within the ICDM development environment, ICODES incorporates a three-tier architecture that draws a clear distinction between representation, logic and presentation. It is a multi-agent system based on a knowledge management premise, in the sense that it incorporates an expressive information model consisting of context-oriented objects, their detailed characteristics, and the relationships that associate these objects to each other and the functional capabilities of the application (Diaz et al. 2006).

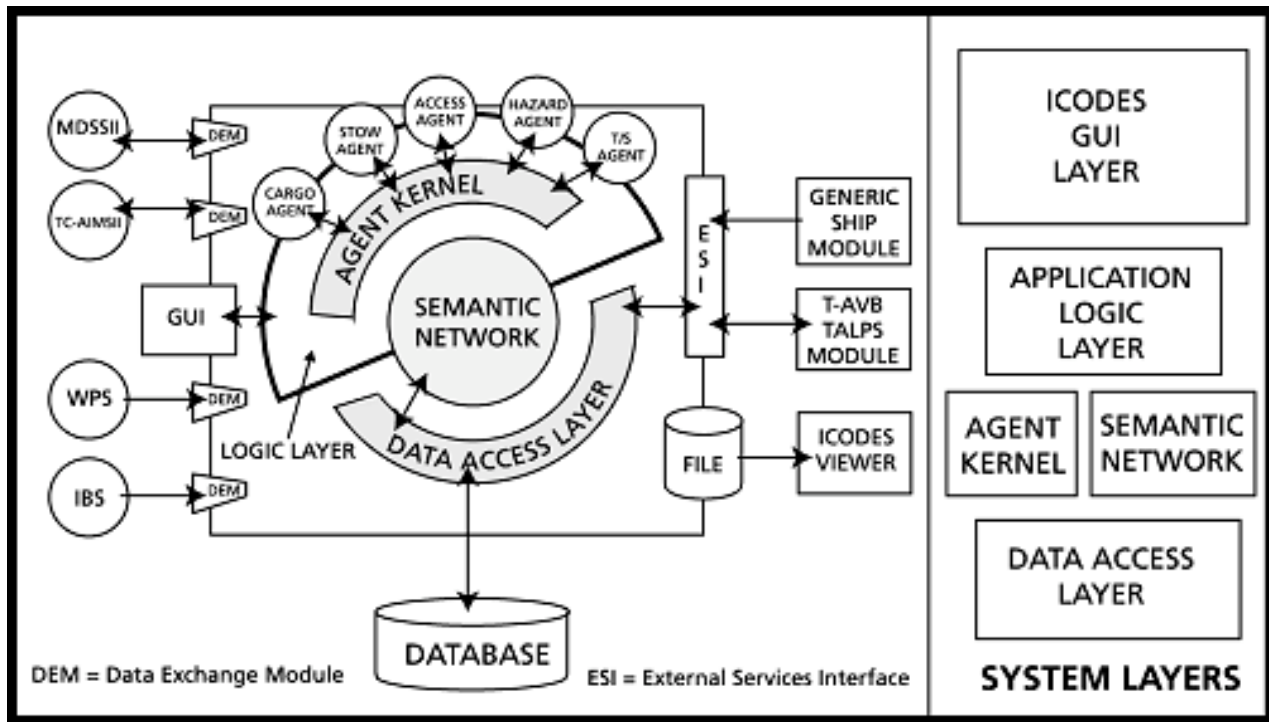


Figure 5: ICODES conceptual architecture

This internal information model provides the necessary context to enable reasoning agents to collaborate with each other and the human user to collectively evaluate events and generate warnings and alerts. Figure 5 provides a view of the ICODES conceptual architecture. Among other aspects, this diagram clearly illustrates the attention to the separation of concerns (i.e., representation, logic, and presentation) that is embedded in the design. Figure 6 provides an additional view of ICODES, presented in the Unified Modeling Language (UML) symbology (Fowler and Scott 1997), describing the significant components along with their inter-dependencies. This alternative view also emphasizes the attention given to inter-component visibility and the objective of a decoupled architecture.

Domain Tier: This tier (also referred to as the semantic network) of the ICODES multi-tiered architecture is essentially comprised of a set of services that govern access to and general lifecycle management of the contextual objects representing the space-planning domain. The scope of this contextual description goes beyond representation of the physical elements that comprise the load-planning environment (i.e., vessels, cargo, etc.) but also include the more intangible, and sometimes much more subtle, concepts and notions that are vital to a complete description of the problem space and evolving solution(s). For example, included in the domain model and housed within this tier are expressive descriptions of the notions of restrictions, violations, recommendations, and accessibility. These are the intangible concepts that empower the ICODES agents with the ability to grasp a deeper understanding of the space-planning activity than would be possible with a data-centric representational paradigm.

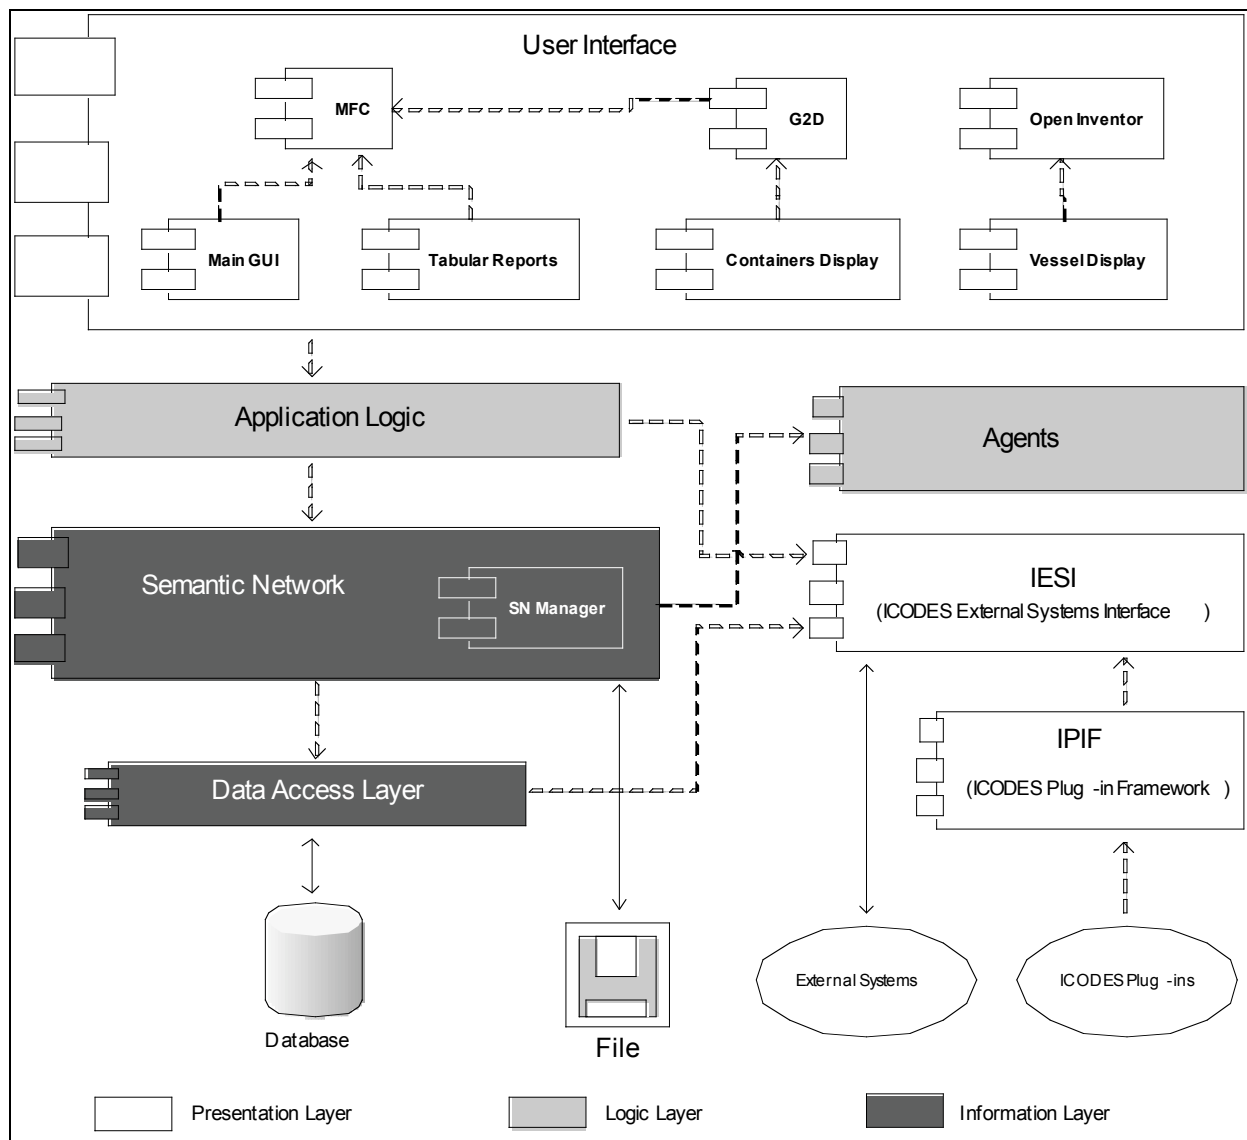


Figure 6: Key component dependencies

The primary service within this tier is the *Semantic Network Service*, which has five related roles. First, it acts as a repository for objects and certain metadata in the semantic network. Second, it provides transaction management for any actions operating over this network of objects. Third, it arranges for the maintenance of the persistent storage of objects. Fourth, it stores dependencies of objects or tagged data on objects housed within the semantic network. Finally, it arranges for notification of interested components (i.e., agents, plug-ins, etc.) when objects in the semantic network change.

The Semantic Network Service engages the *Subscription Service*. This service provides a robust mechanism permitting semantic network clients to register interests in events occurring within the semantic network. The Subscription Service allows subscribers to be notified of events including object creation (i.e., objects being added to the semantic network), object destruction, and object modification. Employing a call-back design pattern, the subscriber provides a method or function to be called when the particular event of interest occurs. When the Subscription Service determines that the specific event has occurred, appropriate notification is sent to all interested parties. As an added means of efficiency, such notification may include state information regarding the event, thus obviating the need for the receiver to issue a set of follow-on queries to examine the details of the event.

The Domain Tier also houses a *Transaction Service* that maintains the overall integrity and consistency of the ICODES information environment. At the heart of this service is the notion of a transaction, as taken from the world of database systems. Any change to the semantic network occurs within the context of a transaction, although that transaction is sometimes provided implicitly. Multiple transactions may be active simultaneously, although no object may take part in more than one transaction at any given time.

Upon completion of a transaction, the Transaction Service employs the *Persistence Service*, which is responsible for reflecting the affects of this transaction within the persistent form of the semantic network. This involves the addition of new objects to storage, the updating of persistent storage to reflect changes made in modified objects, and the removal of objects that have been removed from the semantic network. Although not restricted to any particular vendor, ICODES currently uses the Microsoft SQL Desktop Engine database for its persistence needs. Further, ICODES is flexible enough to actually use more than one form of storage concurrently.

The primary interface offered to Persistence Service clients takes the form of the Independent Data Access Layer (IDAL). This interface is used for both saving an object's state to a data store as well as sharing objects across processes. IDAL provides a *handler* class that defines the interface for persistence of data throughout the ICODES system.

The Semantic Network Manager is capable of attaching tagged data and/or objects to any existing object in the semantic network in the form of *Attachments*. The attached data benefit from the normal lifecycle management offered by the semantic network and are therefore available to any internal or external component (e.g., agent, plug-in, etc.). The first of these two types of attachable elements, tagged data, is textual in nature and although it imparts no contextual information, has presumably a degree of meaning to the contributing component. The tagged data are given a name by the contributor requesting the housing of the attachment. Once named, any component may then retrieve the attached element by providing the name and the identifier of the semantic network object to which the element is attached.

However, far more powerful than attachment of meaningless text is the ability to attach a contextual object. Since an attached object becomes dependent on the object it is attached to (i.e., is owned by), the Semantic Network Manager will automatically remove the attached object if the object it is attached to is removed from the semantic network. This allows a component to create objects in the semantic network that are to be integrated with a load-plan without being concerned about cleaning up objects that are no longer relevant. This ability to add external, non-native types of content to the semantic network is yet another mechanism empowering ICODES with the ability to support the evolving needs of the user with minimal to no impact on the existing code base.

Logic Tier: This portion of the ICODES three-tier architecture not only houses the standard application-level logic typically found in such a tier but also contains the community of decision-support agents that provide the analytical depth empowering the ICODES space planning environment. The agents described in a previous section of this paper collaborate in an indirect fashion to assist in the development of efficient and correct load-plans. A brief discussion of how this community of agents works together to formulate violation-free load-plans in both the assisted stow and the manual stow modes, follows.

The main function of the agent community operating within the *Assisted Stow* mode is to automatically find a valid stow location (i.e., free of any issues relating to accessibility, hazardous materials, trim and stability, etc.) for cargo items within some space. This is achieved through a round-robin style of agent collaboration. Once configured with user preferences and restrictions, the Assisted Stow capability takes advantage of the Cargo Agent by filtering out cargo items that are not deemed valid for stow. In most cases cargo items are filtered out because of missing information relating to dimension or weight. The next step is to locate an empty space on the vessel where the cargo item can be placed without overlapping other cargo items. However, the search for an empty space is constrained by factors such as the weight of the cargo item. For example, the heaviest items in the cargo list should be placed near the bottom of the ship for reasons of overall stability.

This candidate location is then presented for evaluation to both the Access Agent and the Hazardous Material Agent. In accordance with their particular domain expertise, if either agent finds an issue with the candidate placement the Assisted Stow capability nominates an alternative location and the evaluation repeats. This process is repeated until a violation-free location is found or the expanse of possible stow locations is exhausted. In either case, the Assisted Stow capability continues onto the next cargo item to be stowed.

In contrast to the collaborative assessment model applied in Assisted Stow, the *Manual Stow* (i.e., User Stow) mode of operation instructs the agents to function independently of one another. In other words, as the user places, or templates, a cargo item within a stow space each agent reacts concurrently identifying any outstanding issues incurred by such placement in accordance with their individual domain of expertise. Any such violations or warning are presented to the user in the form of agent reports comprised of a concise depiction of the issue along with any possible resolutions the agent may be able to offer. However, regardless of the severity of the issue, the user makes the final determination of how, and even if, the issue is to be resolved.

Presentation Tier: Like most industry-standard applications, ICODES offers a graphical user-interface (GUI) as the primary means of exchanging information with the user. Because

ICODES is designed to run on a Windows operating system platform its user-interface adheres to the standard Windows logo compliance design pattern. ICODES offers its presentation capabilities in two forms, standalone and web-based.

The design of the ICODES standalone GUI is comprised of several internally developed as well as off-the-shelf components. Each of these components operates in unison to provide a robust means for interacting with the user. Following is a brief discussion of several of these components.

The *Vessel Display* portion of the ICODES stand-alone user-interface is based on the Open Inventor low level graphics library (Wernecke 1994). Open Inventor is a commercial object-oriented three-dimensional (3D) graphical toolkit built on top of OpenGL (Shreiner 2000). This toolkit uses a programming model based on a 3D Scene database that dramatically simplifies graphics programming. Open Inventor offers the software developer a rich set of objects including cubes, polygons, text, materials, cameras, lights, trackballs, handle boxes, 3D viewers, and editors. Together these elements allow for the development of robust, interactive graphical software applications.

The *Graphics 2D (G2D) Viewer* is a cross-platform implementation of a set of two-dimensional graphics tools. G2D displays its graphical elements in a layered manner. The viewer tool supports drag-and-drop operations as well as the drawing of graphic primitives, including multi-polygons and Bezier curves⁴. The G2D Viewer also supports the vector and matrix mathematics functionality that is required for the use of transforms.

The fundamental G2D viewer architecture essentially consists of three components: the G2D surface that forms the conceptual container in which all displayed information resides; multiple graphical layers displaying the graphical elements comprising the scene; and, the user-interface layer where items are drawn during drag or drawing operations.

In the *Web-Based User-Interface* (or Thin Client) ICODES load-plans are described as documents that represent all of the information available to the ICODES application during the preparation of a load-plan. When a stow-planner uploads a load-plan to the ICODES File Share, the ICODES Thin-Client uses standard XML parsers along with other tools to extract information from the file and place it in a database, thereby making that information available for later use in the user-interface.

The Thin-Client uses a standard SAX⁵ parser to process the load-plan file. Instead of building a tree representation of an entire XML file in memory as a DOM⁶ parser would do, a SAX parser identifies the individual parts of an XML document as it reads the file and immediately passes those parts to an object that implements the *org.xml.sax.ContentHandler* interface. When the parser identifies, for example, the start of an XML element and processes its attribute list, the parser will call the *ContentHandler.startElement* method, passing the element's name and universal identification (if applicable), and the list of attributes as *name/value* pairs. SAX parsers eliminate the need to parse the entire document before

⁴ Developed by [Pierre Bézier](#) in the 1970s for CAD/CAM operations, a Bezier curve is a cubic equation that can be used in computer graphics for the construction of non-linear shapes.

⁵ SAX is an acronym for Simple API for XML (see: <http://www.saxproject.org/>).

⁶ DOM is an acronym for Document Object Model (see: <http://www.w3.org/DOM/>)

processing can begin, which is important when dealing with notably large XML documents such as ICODES load-plans.

Extended Functional Requirements

After an extensive evaluation of several existing military planning systems in 2007 it was determined by a joint TRANSCOM and JFCOM assessment team that both the functionality and scope of ICODES should be extended to support all military conveyance load-planning and staging requirements. With a release date scheduled for October 2010, it is expected that ICODES GS will incorporate a Collaborative Information Workspace (CIW) that will provide global user-access to a seamless environment of intelligent load-planning and cargo staging tools. These tools must function equally well in net-centric and stand-alone mode, through either a web-based Thin-Client or a Thick-Client user-interface.

Initial explorations have indicated that the technology is now available to implement a Thin-Client with full interactive graphics functionality at an acceptable performance level. For example, in ICODES the user must have the ability to drag-and-drop the graphical symbol representing a cargo item located in the graphical display of a ship (in plan view) from one deck at the top of the screen to another deck at the bottom of the screen without any noticeable response delay. Until recently such an action would have required the redisplay of the entire graphics screen with the attendant severe performance penalty due to communication bandwidth limitations. Open source and commercial tools are now available that allow such a web-based user-interface to be designed so that only those portions of the graphic display that have been affected by the user's action need to be refreshed.

The existing ICODES Master Vessel Library that currently includes over 300 objectified ships (i.e., mostly commercial vessels leased by the military for the transportation of supplies) will need to be greatly extended to also include aircraft, train cars, trucks, and marshalling yards. The resulting Shared Object Library is being designed as a set of generic services that can objectify, store and retrieve the graphical representation of any kind of two-dimensional storage space together with its characteristics. It is expected that the size of the new Shared Object Library will be at least one order of magnitude larger than the current Master Vessel Library.

The CIW, designed as an Information Management Framework (IMF), will form the core of the ICODES GS service-oriented architecture implementation. In this role the IMF must provide the base services responsible for the discovery, integration, persistence, and exchange of data, as well as associated mediation and security functions. While these services are largely hidden from the users, they enable the users to perform tasks that require access to multiple data sources in a seamless fashion. Typical examples of the functions performed by these services include: *distribution* (subscription, prioritization and synchronization); *validation* (data cleansing and mapping); *discovery* (metadata management); *persistence* (archiving, indexing and aggregation); *security* (authentication, intrusion detection and virus prevention); *monitoring* (performance measurement, evaluation and optimization); and so on.

Accessible to the users through the IMF will be the tools that they require to perform their tasks. These include the various agents that assist in the load-planning process, the Thin-Client and Thick-Client user-interfaces, user-controlled data access and cleansing tools, and several utilities. Based on service-oriented architecture (SOA) principles these tools function as services and

depending on individual requirements incorporate internal representations of context and agents capable of automatically reasoning about data changes within that context. The degree of intelligence and automation embedded in each of these services depends on the functional requirements. For example, the Assisted-Stow Agent will have a fairly high degree of intelligence because of the complexity of the load-planning problem and the dynamic nature of the operational environment.

It is in the nature of an information-centric system environment that the operators are shielded not only from the internal system services of the IMF, but also from the physical location, configuration requirements and data access requirements of the functional services that the operators utilize to perform their tasks. After signing on through a single login entry point, it should be transparent to the operator whether the required capability (i.e., tool) is a single service or several services, and whether it is an external legacy application linked to the IMF by an intelligent interoperability bridge service or an internal service incorporating artificial intelligence methodologies.

Conclusion

The ICODES application currently provides a comprehensive tool-set of software agents to assist the cargo specialist in the development of ship load-plans for military deployments. It is one of the earliest military examples of information-centric software that incorporates an internal, relationship-rich information model to provide context for the reasoning functions of collaborative software agents. Over the next two years ICODES will be extended in functionality to serve as a toolset for the load-planning of all types of conveyances (i.e., aircraft, trains, and trucks) and assembly areas. At the same time the ability of its underlying SOA-based design will be severely tested as ICODES scales from a standalone application to a global environment of integrated intelligent services.

As an ICDM-based application, ICODES adheres to three notions that are fundamental to its decision-assistance capabilities.

1. ICODES processes information (i.e., data with relationships) as opposed to legacy systems that normally process data only (even though the data may be in the form of objects with characteristics). The key to the assistance capabilities of ICODES is that the system has some understanding of the information that it is processing. In the internal Semantic Network cargo items are described in terms of characteristics that relate each item to hazard, trim and stability, accessibility, and ship configuration, constraints. This internal information model provides context for the automatic reasoning capabilities of software agents.
2. ICODES is a collection of powerful collaborative tools, not a library of predefined solutions. This overcomes the deficiencies of legacy systems in which built-in solutions to predetermined problems often differ significantly from the complex operational situations encountered in the real world. In this respect ICODES is a collaborative decision-support system in which the operator interacts with computer-based agents (i.e., decision making tools) to solve problems that cannot be precisely or easily predetermined.

3. ICODES incorporates agents that are able to reason about the characteristics and the relationships of cargo items, the internal configurations of conveyances and the constraints that must be considered during the development of load-plans. Although these agents are decoupled (i.e., do not know about each others existence) they are able to indirectly collaborate through a data blackboard and subscription services, as they assist the user throughout the load-planning process.

The advantages of an information-centric software system have been evidenced in three areas by the performance of ICODES in the field over the past three years. First, if all necessary data are available ICODES is capable of automatically generating the load-plans of four medium-size ships in around two hours. This is a significant improvement in load-planning speed over the legacy application that it replaced. The predecessor application typically required two person-days for the development of a single load-plan. Second, the assistance capabilities of the ICODES agents elevate the performance of a novice load-planner to at least an acceptable level. This is an important consideration in view of the attrition rate of military cargo specialists during the past decade. The performance of an expert load-planner, on the other hand, is raised to an exceptionally high productivity level. Third, the ability of ICODES to continuously evaluate the evolving load-plan in respect to accessibility, hazardous material, and trim and stability conditions, has greatly increased the quality and accuracy of the resulting load-plan.

References

- Bradshaw J. M. (ed) (1997); 'Software Agents'; AAAI Press / The MIT Press, Massachusetts, (pp. 3-11).
- Diaz C., W. Waiters, J. Pickard, J. Naylor, S. Gollery, P. McGraw, M. Huffman, J. Fanshier, M. Parrott, S. O'Driscoll-Packer, Boone Pendergrast and Evan Sylvester (2006); 'ICODES: Technical and Operational Description'; Technical Report CDM-20-06, CDM Technologies Inc., San Luis Obispo, California, November.
- Fowler M and K Scott (1997); 'UML Distilled: Applying the Standard Object Modeling Language'; Addison-Wesley, Reading, Massachusetts.
- Mowbray T. and R. Zahavi (1995); 'The Essential CORBA: Systems Integration Using Distributed Objects'; Wiley, New York, New York.
- MTMC/ICODES (2002); 'ICODES Version 5.2: USMC Basic Training Manual' and 'ICODES Version 5.2: Advanced Training Manual'; Military Traffic Management Command (MTMC), US Army.
- Myers M. and J. Pohl (1994); 'ICDM: Integrated Cooperative Decision Making – in Practice'; 6th IEEE International Conference on Tools with Artificial Intelligence, New Orleans, Louisiana, November 6-9.
- Pohl J., A. Chapman, K. Pohl, J. Primrose and A. Wozniak (1992); 'Decision-Support Systems: Notions, Prototypes, and In-Use Applications'; Technical Report CADRU-11-97, Collaborative Agent Design Research Center, Design and Construction Institute, College of Architecture and Environmental Design, Cal Poly, San Luis Obispo, California, January.

Pohl K. (2001); 'Perspective Filters as a Means for Interoperability Among Information-Centric Decision-Support Systems'; Office of Naval Research (ONR) Workshop hosted by the Collaborative Agent Design Research Center (CADRC), Cal Poly (San Luis Obispo, CA) in Quantico, Virginia, June 5-7.

Pohl J., K. Pohl, R. Leighton, M. Zang, S. Gollery and M. Porczak (2004); 'The ICDM Development Toolkit: Purpose and Overview'; Technical Report CDM-16-04, CDM Technologies, Inc., San Luis Obispo, California, May.

Pohl J. (2007); 'Knowledge Management Enterprise Services (KMES): Concepts and Implementation Principles'; Pre-Conference Proceedings, Focus Symposium on Representation of Context in Software, InterSymp-2007, Baden-Baden, Germany, 31 July.

Pohl K. (2002); 'The Underlying Design Principles of the ICDM Development Toolkit'; Preconference Proceedings, InterSymp-2002, Focus Symposium on Collaborative Decision-Support Systems, Baden-Baden, Germany, July 29-August 30.

Shreiner D. (ed.) (2000); 'Open GL Reference Manual'; Addison-Wesley, Menlo Park, California.

Wernecke J. (1994); 'The Inventor Mentor: Programming Object-Oriented 3D Graphics with Open Inventor'; Addison-Wesley, Menlo Park, California.

Wooldridge M. and N. Jennings (1995); 'Intelligent Agents: Theory and Practice'; The Knowledge Engineering Review, 10(2), (pp. 115-152).

