10th International Command and Control Research and Technology Symposium

June 13 – 16, 2005
McLean, Virginia

Topic: Coalition Interoperability

Title: **Integration of the MIP Command and Control Information Exchange Data Model into National Systems**

Author: Dr. Michael Schmitt
Organization: FGAN/FKIE
Address: Neuenahrer Straße 20
53343 Wachtberg-Werthhoven
Germany
Phone: +49 228 9435 414
Fax: +49 228 9435 685
E-Mail: m.schmitt@fgan.de

Integration of the MIP Command and Control Information Exchange Data Model into National Systems

Michael Schmitt

FGAN/FKIE
Neuenahrer Straße 20, 53343 Wachtberg-Werthhoven, Germany
m.schmitt@fgan.de

Abstract

Interoperability of command and control information systems gains an ever-increasing importance. The aim of the Multilateral Interoperability Programme (MIP) is to achieve international interoperability in order to support land component commanders in joint and combined operations. For that purpose, MIP defines the *Command and Control Information Exchange Data Model* (*C2IEDM*) and the *Data Exchange Mechanism* (*DEM*). However, when implementing the MIP Solution, it is not sufficient to simply add new interfaces to existing systems. Instead, far-reaching modifications to the core of national C2ISs have to be made to ensure true semantic interoperability.

In this paper, we address several interoperability and implementation issues of the MIP C2IEDM. We point out that a shared tactical picture only becomes reality if the commanders are fully aware of the extent of interoperability that is given by their national C2ISs. Moreover, the subtle problems of coupling a geographic information system (GIS) with the MIP solution are discussed. On the data base level, we show that the co-existence of a proprietary national data model and the C2IEDM results in systems that are extremely hard to maintain. To hide away the complexity of the C2IEDM from C2 applications, we propose a data access stack that provides a canonical, business objects view on the data model.

Key words: MIP, Multilateral Interoperability Programme, C2IEDM, Integration, Implementation, Interoperability, Data Replication

---

[1] The author would like to thank Henriette Schüller and Gerhard Bühler for proofreading and providing valuable comments.

1 Introduction

In the context of combined and joint missions, interoperability of command and control information systems (C2ISs) plays a critical role. Thus back in 1998, the Multilateral Interoperability Programme (MIP) has been established. What began as a voluntary initiative of six nations has turned into one of the most important interoperability programs with participation of 26 nations and organizations.

According to the MIP Tactical C2IS Interoperability Requirements (MIP MTIR, 2004b, p. 7 and 8),

> "The aim of the Multilateral Interoperability Programme (MIP) is to achieve international interoperability of Command and Control Information Systems (C2IS) at all levels from corps to the lowest appropriate level, in order to support combined and joint operations; [...] MIP meets the requirements of the Land Component Commander of Allied Joint and Combined Operations (including Article 5 and Crisis Response Operations)."

In order to fulfill these requirements, the *MIP Solution* is defined. Essentially, it covers two technical aspects:

- A common data model, called *Command and Control Information Exchange Data Model* (*C2IEDM*; MIP, 2004c)
- A set of procedures and protocols that allow replicating data among different C2ISs, called *MIP Data Exchange Mechanism* (*DEM*).[2]

The architecture of the MIP Solution is given in figure 1. Although (MTIR, 2004b, p. 9) states explicitly that the "function, implementation and the display of the host C2 applications is not the concern of MIP", the MIP Solution is not a "plug-and-play" technology for existing national C2ISs. In order to ensure true semantic interoperability, far-reaching modifications to the core of national C2ISs are necessary rather than just the addition of mapping adapters as new interfaces to the existing systems.

In this paper, we discuss the impacts of the MIP Solution on the national C2ISs with focus on the integration of the MIP C2IEDM. In section 2, we discuss the relationship between national information requirements and information exchange requirements. We show that the concept of a shared tactical picture can only become reality if the commanders are fully aware of the extent of interoperability that is given by their national C2ISs.

Due to the fact that the C2IEDM is an information exchange data model,

[2] Actually, MIP defines a second exchange mechanism, called *Message Exchange Mechanism* (*MEM*). It is used for transmitting NBC reports, plans & orders, and some MIP gateway management information. However, C2IEDM data are exchanged exclusively with the DEM.

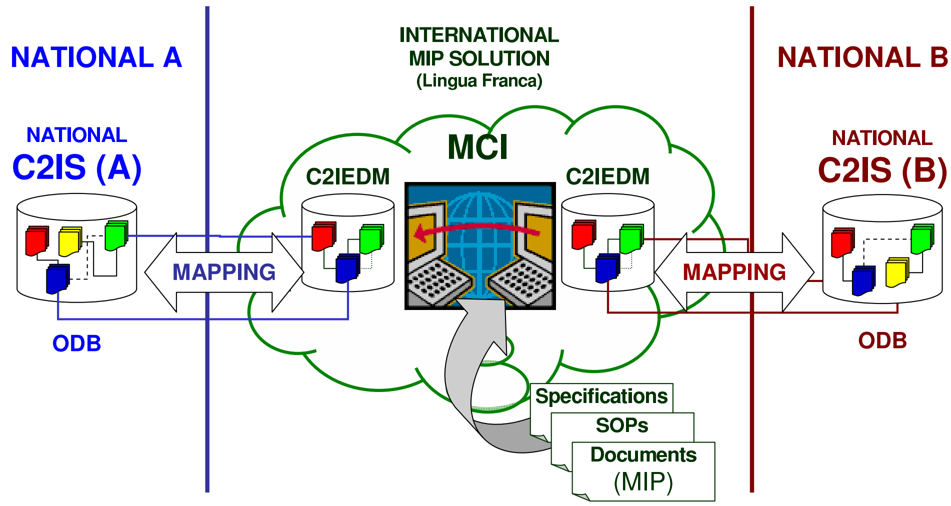

Figure 1. The MIP Solution (MIP, 2005)

there is an evident need for importing data from other data sources into the C2IEDM such that these data can be replicated. By describing the task of coupling a geographic information system (GIS) with a C2IEDM data base in section 3, we illustrate the subtle problems that application developers have to consider in order to ensure usability and interoperability among heterogeneous systems.

In section 4, we demonstrate that the co-existence of a proprietary national data model and the C2IEDM also results in non-trivial problems on the data base level. In particular, the role of synthetic keys in the C2IEDM and information dissemination over multiple echelons are examined.

The C2IEDM or, more precisely, its technical realization as an RDBMS, is not meant to be accessed directly by any C2 application. To hide away the complexity of the C2IEDM from C2 applications, we propose a multi-layered data access stack. It abstracts from the relational model and provides a canonical, object-oriented view on the data model as well as business functions. In section 5, we motivate and illustrate each layer of the proposed data access stack.

Finally, a summary and conclusion is given in section 6.

2 Ensuring the Shared Tactical Picture

The MIP C2IEDM models the information that combined joint component commanders need to exchange (MIP, 2004c, page xx). In recent years, the data model has been extended continually and the forthcoming *Joint Consultation Command and Control Information Exchange Data Model* (*JC3IEDM*) will cover even more information exchange requirements.

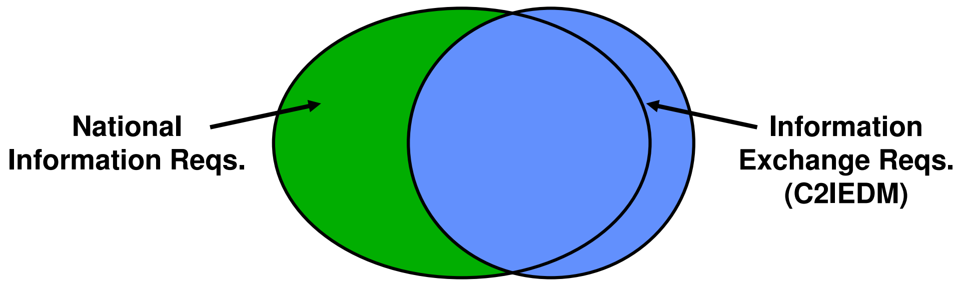

Figure 2. National Information Requirements vs. Information Exchange Requirements

However, the increasing expressiveness of the data model does not imply that the national C2ISs automatically keep pace with the standardization efforts. For MIP block 1, the former *Land C2IEDM* (*LC2IEDM*) was divided into entities that were supported by all national C2ISs and entities that were implemented optionally. The diverse support for the MIP data model is expected to continue in the future.

On the other hand, by definition the MIP data model does not cover the full set of national information requirements. For instance, information of specific functional areas which is not directly related to command and control is out of the scope of the C2IEDM.

The relationship between national information requirements (covered by a national data model) and the information exchange requirements (covered by the C2IEDM) is shown in figure 2.

From an interoperability perspective, the given situation is problematic. To ensure a common operational picture (shared tactical picture), it is essential that each commander is aware of what information is available to all other commanders. While in general no assumptions can be made on the C2ISs of other nations, the national C2IS must clearly answer the following questions to the commander:

- Sending of information:
  - What kind of information is sent through the MIP common interface?
  - To which commanders is a specific information sent?
  - In what form is this information sent, i.e., did it have to be transformed in order to fit into the C2IEDM?
- Reception of information:
  - What kind of information is received through the MIP common interface?
  - From which commander has a specific information been received?
  - Did information have to be transformed to fit into the national data model?
  - Which kind of information cannot be handled by the national C2IS?

The fact that the national C2IS may not be able to handle some information is critical, since the sender of the information may rely on that the receiver is

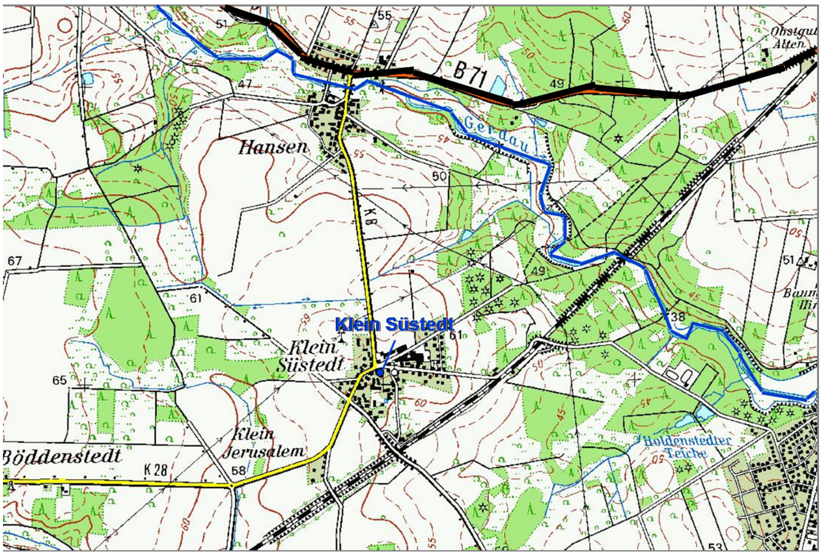

Figure 3. Missing Synchronization: Roads, Rivers, and Towns are Drawn Twice

able to get and evaluate the information. As a consequence, even if the C2IS does not support some entities, attributes, or domain values of the C2IEDM natively, it should provide a fall-back solution to display corresponding information in a generic way.

3 Synchronization with National Data Sources

The C2IEDM is designed for exchange of C2-relevant information; information which has no definite military relevance is not considered for replication. However, in order to get a complete operational/tactical picture, the commander needs further information which military relevance cannot be decided in advance.

Geographic information systems (GIS) provide a lot of such information. Modern GIS are based on vectorized maps and provide various information on features (e.g., roads and rivers) and facilities (bridges, police stations, governmental offices, etc.). This information is typically stored in data base management systems that are highly optimized for memory usage and efficient data access.

Due to the tremendous amount of (potentially relevant) information, it makes no sense to map all these data to the C2IEDM and replicate them; in particular, as the structures of the C2IEDM are not optimized for representing large volumes of homogeneous GIS information. The C2IEDM provides a reference mechanism that allows to point to sources outside the model. However, since nations may use different GISs, it does not help either.

As a consequence, a distinction has to be made between GIS data that are kept nationally and GIS data that are subject to exchange. Since the commander will want to enrich GIS data during an operation (e.g., he may want to declare a bridge as destroyed), redundant data storage is inevitable.

The technical constraints require substantial development efforts to create user-friendly applications. Otherwise, the situation arises where objects are displayed twice (see figure 3). Ideally the following procedures should be supported:

- On the sender side:
  When the commander edits a feature or facility on the map, the C2IS must decide whether this information is already subject to MIP data exchange or was retrieved from the underlying national GIS data base. If the information stems from the GIS, the application must be able to extract the information from the GIS data base and store it in a way that allows adding C2IEDM-specific information and replicating it by the MIP gateway. In addition, the C2 application must ensure that the same information is not output twice on screen.
- On the receiver side:
  When a C2IS receives new data from the MIP gateway, it must be able to synchronize these data with the data of its associated GIS. Since even identical data may not fully match, e.g., the coordinates of a road may differ slightly, the comparison must be made within a range of tolerance.

The problem outlined above is not specific to GISs. It applies to all data sources that are only loosely coupled with the core operational data base. The problem may also arise if only one kind of C2IS is used. However, the problem gets worse in multinational operations, since no assumption can be made on the national data sources.

4 Integration of the C2IEDM on the Data Base Level

The MIP solution as presented in figure 1 on page 3 conceptionally differs between an operational data base (ODB) that covers the national information requirements and the C2IEDM that is used for information exchange. From the MIP point of view, the ODB may be based on any proprietary data model as long as it can be mapped to the relational schema of the C2IEDM (see figure 4(a)). The C2IEDM itself does not necessarily have to be implemented physically by means of an RDBMS; replication can also operate on transformed data taken directly from the ODB.

The differences between the underlying data model of the ODB and the C2IEDM may be technical or logical. For instance, unlike the C2IEDM, the operational data base may be technically based on an object-oriented model.

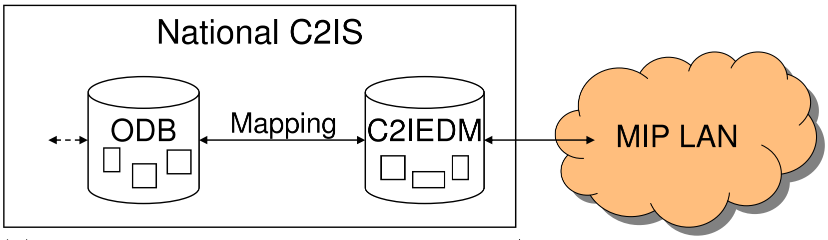

(a) ODB data model differs technically/conceptionally from the C2IEDM

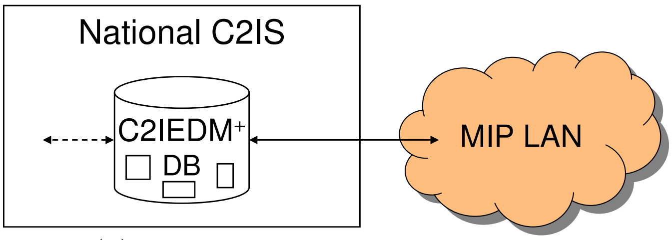

(b) ODB is an extended C2IEDM DB

Figure 4. Operational DB vs. C2IEDM

Logical modifications comprise insertions, deletions, and semantic modifications on the level of domain values, attributes, entities, and relationships.

The required mapping rules can be very complex in practice. In particular, this holds in cases in which there is no clear 1:1 mapping of concepts. For instance, n attributes of the ODB might have to be mapped onto m attributes in the C2IEDM where the attributes may be distributed over several entities.

If the expressive power of the proprietary data model and the IEDM is not identical, no bijective mapping is possible, i.e., data exchange will inevitably result in information loss. In this case, the commanders must be warned (cp. section 2).

If the data model of the ODB is a conceptual extension of the C2IEDM, then the relational schema of the C2IEDM should be used as the basis for a national implementation (see figure 4(b)). No mapping is needed as long as the national extensions are not subject to information exchange with other MIP-compliant C2ISs.

4.1 Maintenance of C2IEDM Keys

One important issue during mapping is the maintenance of C2IEDM keys. To uniquely identify data, the entities of the C2IEDM have identifier and index attributes. These attributes contain automatically generated, synthetic keys. Synthetic means that the keys are 15- or 18-digit numbers which have no operational meaning. They are composed of a party prefix (e.g., *180* for Germany), a national node prefix, and an entity-specific sequence number.

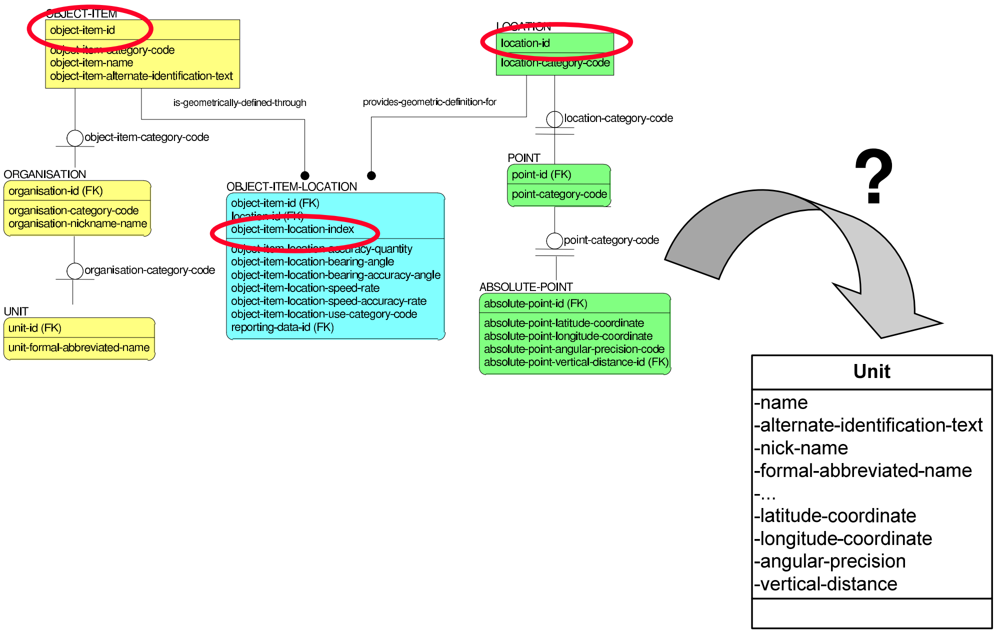

Figure 5. Key Management

This policy ensures that each data record gets a unique key within the network of MIP-compliant systems.

In order to support continuous communications, the recipient of data must preserve these keys. If the ODB is not based on the C2IEDM, there are two ways to achieve this: The first solution is to extend the mapping adapter between the ODB and the C2IEDM by a proxy table that keeps track of all keys and associates MIP objects with internal objects.

The second alternative is to extend the ODB by new attributes. Depending on the structure of the ODB, multiple MIP key attributes might have to be added to a single entity of the ODB. For instance, in the MIP data model, units and their locations are stored in separate entities. The relationship between both is established by entity *OBJECT-ITEM-LOCATION* (see figure 5). In contrast, in a proprietary ODB a unit and its most recent location might be stored in a single object or table. In order to keep track of all MIP keys, one would have to extend it by the three attributes that are highlighted in the figure. As a rule of thumb the more the ODB is denormalized, the more key attributes have to be inserted into each ODB entity.

4.2 Information Forwarding

Theoretically, a national C2IS only has to implement the MIP common interface in order to achieve interoperability – the core of the C2IS remains untouched at the cost of complicated mappings. However, in order to work correctly, the MIP solution and the C2IEDM – although the latter is an information *exchange* data model – impose some requirements on the internals and the behavior of the C2IS.

In block 2, MIP has introduced the Operational Information Group (OIG) concept (see MIP, 2004b, p. 102 and MIP, 2004a, p. 60). The purpose of OIGs is to structure information according to logical aspects and to disseminate information differently depending on the affiliation to a particular OIG. For instance, information in OIGs of category *Composed Plan* is only distributed to the superior unit, all subordinate units, and to the flanking units, whereas information in OIGs of category *Globally Significant* is distributed to all units.

As the latter example indicates, information may have to be disseminated via several echelons. Since the C2ISs involved form a (locally) distributed system, information might have to be forwarded through the national network. According to the MIP System Requirement Specification (MIP, 2004a, p. 33), "the National implementation of MIP Gateways shall allow data that is received by one Gateway on a MIP LAN to be available at other gateways on other MIP LANs. The internally forwarded data must be identical at all gateways."

The requirement that data must be passed unchanged, in particular implies that the integrity of keys must be preserved. What are the consequences of this requirement on the implementation of the ODB?

Figure 6(a) shows a scenario in which the national C2ISs each have a proprietary non-C2IEDM operational data base and use a mapping adapter at the MIP gateway. Internally, data exchange between the distributed C2IS is performed by means of a national, proprietary exchange mechanism. A proxy table is used at each MIP gateway to store the received MIP keys. However, due to the fact that the proxy tables work independently, all of MIP's synthetic keys that are received at the upper MIP gateway will get lost on their way through the national network. The mapping adapter at the lower MIP Gateway may create new keys but from the point of MIP this means the creation of completely new data. So in essence, a purely "interface-based" solution will not work correctly in the context of information forwarding.

In the second scenario, the MIP Data Exchange Mechanism is used in addition to the national exchange mechanism (see figure 6(b)). This scenario reflects the need to do MIP replication on the one hand and to exchange ODB-specific data on the other hand. Unfortunately, this solution does not work either. The problem here is that the lower gateway receives data from two different sources: via DEM replication from the upper MIP gateway directly and via proprietary data exchange through the lower C2IS. Since the synthetic keys are the only criterion to uniquely identify C2IEDM data, it is impossible for the lower gateway to determine identical data received from both sources.

As a consequence, the only solution that works correctly is to fully rely on a single exchange mechanism that preserves keys within a network of MIP-compliant systems (figure 6(c)). This means that the ODB must be extended

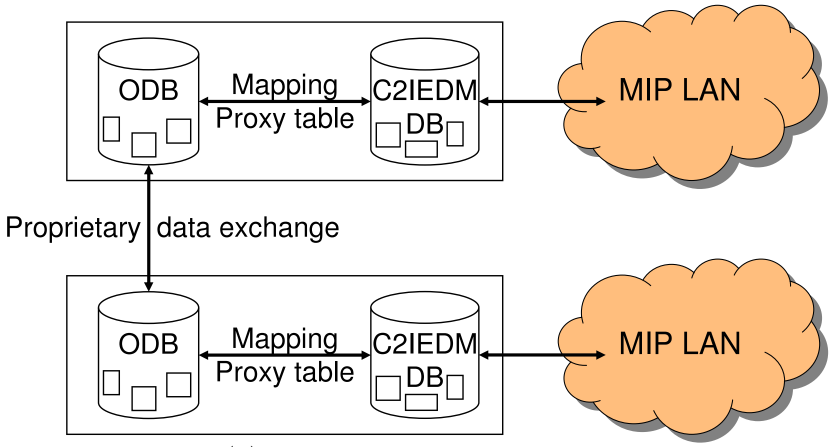

(a) Proprietary data exchange

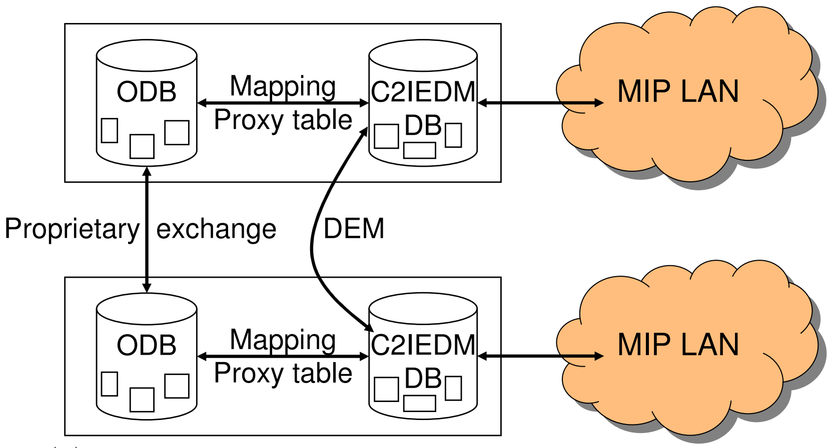

(b) Combined use of proprietary data exchange and DEM

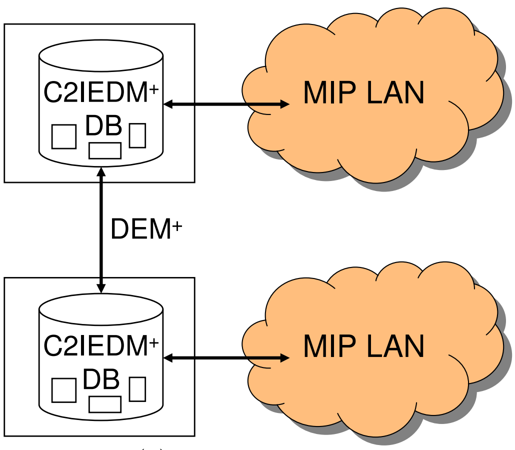

(c) Extended DEM

Figure 6. Information Dissemination

for attributes that store all MIP keys. Since the effort to modify and maintain the ODB can be enormous if its data model differs significantly from the C2IEDM and since the DEM imposes specific requirements on the integrity of transmitted data, it seems reasonable to use the C2IEDM as the basis for the national data model and use an extended version of the DEM for information exchange.

5 A Model for C2IEDM Data Access

The C2IEDM is specified in terms of an entity-relationship (E-R) model. Since the relationships between the entities are described explicitly by foreign and primary key attributes, it is possible to automatically derive a relational data base schema from the MIP data. This data base schema may serve as the basis for a national implementation.

However, a C2IS application should never operate on this schema directly for several reasons: First, the MIP data model undergoes regular updates. According to the MIP schedule, a new version of the MIP data model is approved about every two years. In addition, bug fixes and interim releases are provided if necessary. The different versions of the MIP are neither backward nor forward compatible. A C2IS that relies on the relational schema has to be adapted with every model update resulting in poor maintainability.

Second, the C2IEDM is primarily designed for information exchange and is not optimized for fast data access. As illustrated in section 5.4, complex transformations and algorithms have to be applied in order to get data structures that can be processed and evaluated easily. To improve efficiency, high-level data objects should be defined that abstract from the complexity of the MIP data model.

Third, the MIP data model has some potential pitfalls. If the model is used improperly, this may lead to an erroneous shared tactical picture. In order to ensure the correctness of all existing C2IS applications, common operations (data updates and queries) should be handled centrally.

Taking into account the above-mentioned quality criteria — maintainability, efficiency, and correctness —, a data access stack is suggested that hides away the complexity and subtleties of the C2IEDM from the applications.

The data access stack and its abstraction layers are shown in figure 7. Each layer abstracts from some of the technically and operationally motivated properties of the layer below. Two services have to be realized in each layer, namely *access control* and *notification*. Access control serves the authentication and authorization of users and applications.

The notification service informs objects in higher layers about data updates

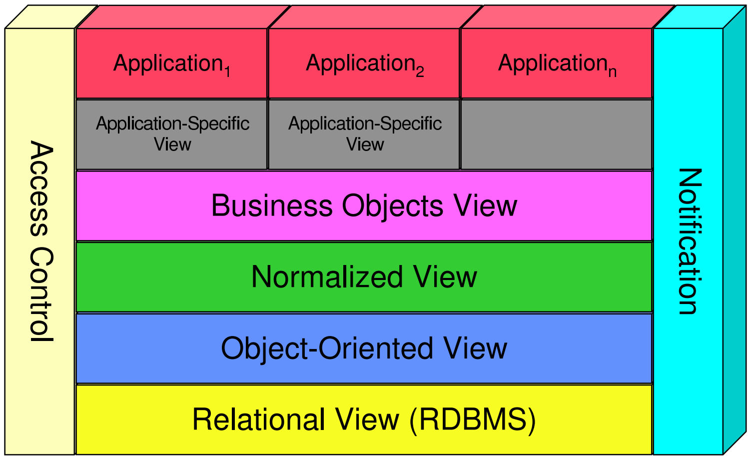

Figure 7. Data Access Stack

in lower layers. This service is essential to provide the commander with an up-to-date shared tactical picture without having to poll for updates repetitively. Notification is an asynchronous operation, because an upper layer object cannot predict the exact time and state in which it receives information updates. Technically, notification can be realized by an event listener in the upper layer that registers to the notification service of the lower layer. The event listener is called by the service whenever a specific kind of information is updated.

In the following, the major properties of each layer are described.

5.1 Relational View

The first layer of the data access stack provides a relational view on the MIP data model. This view reflects the table structure as given by the underlying RDBMS. As mentioned before, a data base schema can be derived from the C2IEDM. This schema also defines the exchange format for the DEM.

In addition to the entities and attributes given by the C2IEDM, it may be reasonable to add new elements to the relational model. Such extensions might serve two purposes: covering national information requirements and improving data access efficiency.

For instance, it can be useful to introduce new attributes that, e.g., keep track of whether a status information is outdated. If existing domains are extended, it must be ensured that the new values are mapped onto existing C2IEDM values before they are replicated. The introduction of new domain values mainly affects codes. E.g., a category code must be extended if a new subtype is introduced. Complementing the data base schema on the table and attribute level does not affect the technical functionality of the DEM. If data are replicated with another MIP-compliant system, the additional elements are simply ignored.

5.2 *Object-Oriented View*

The MIP data model suggests the use of a relational data base management system for storing data. While RDBMS are still the preferred way for making data objects persistent today, modern software applications are developed according the object-oriented paradigm. Therefore, the entities of the relational model have to be mapped to Java, C++, or UML classes in a first step.

The semantical gap between the relational and the object-oriented world, also called O-R impedance, is a well-known problem in computer science. In a relational schema, one-to-many relationships are expressed by foreign keys. In an OO model, a one-to-many relationship is represented by either an object reference or a list of object references in either one of the two classes involved or both. In a relational model, many-to-many associations require additional entities that allow establishing the link. In an OO model, these association entities are not necessarily needed. Subtyping can be described in different ways in a relational model. For the MIP data model, it was decided to have a hierarchy of entities where the subtype is specified explicitly by means of a discriminator code in the super entity. For the OO model, incomplete subtyping should be transformed to complete subtyping. This allows declaring entities such as *ObjectItem* to be abstract, i.e., only the leaf classes of a class hierarchy can be instantiated.

Example: In order to define a new type for UN mandated buffer zones that are affiliated to the NATO, records have to be inserted into six tables of a C2IEDM data base. The part of the C2IEDM that covers the corresponding entities and records is given in figure 8.

A UML class diagram that provides an object-oriented view on the same entities is given in figure 9. Instead of having to fill 6 data base tables, an application that operates on the OO view can describe the same information with three simple statements that are also listed in figure 9.

There are many open source and commercial tools, frameworks, and application programming interfaces (APIs) that support object persistence by means of an RDBMS. Solutions for the Java programming language include Hibernate (Hibernate, 2005), Java Data Objects (JDO, 2005), J2EE Container Managed Persistence (CMP) and ObJectRelationalBridge (OJB, 2005). In all cases, mapping rules are defined – implicitly or explicitly – from objects to data base structures. These rules adhere general-purpose patterns. Some tools are complemented by reengineering tools that create OO classes from a relational schema in a generic way.

When using the MIP data model, the relational schema is fixed and based on specific design rules. For instance, subtyping requires a discriminator code. The above-mentioned O-R mapping tools do not allow to specify detailed mapping rules that realize these specific design rules. Unless the object-relational

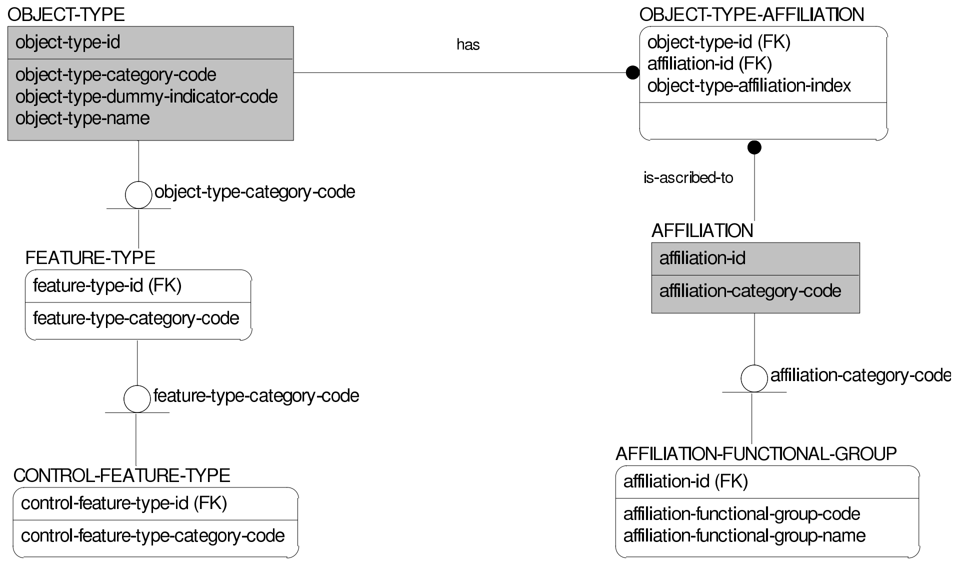

OBJECT-TYPE			
*-id	*-category-code	*-dummy-indicator-code	*-name
260110000000201	**FE**	**NO**	**UN mandated buffer zone**

FEATURE-TYPE	
*-id	*-category-code
260110000000201	**CF**

CONTROL-FEATURE-TYPE	
*-id	*-category-code
260110000000201	**NOS**

AFFILIATION	
*-id	*-category-code
260110000000001	**AFLFNC**

AFFILIATION-FUNCTIONAL-GROUP		
*-id	*-code	*-name
260110000000201	**MULTIN**	**NATO**

OBJECT-TYPE-AFFILIATION		
object-type-id	affiliation-id	*-index
260110000000201	**260110000000001**	**260110000000001**

Figure 8. Example "UN mandated zone" – Relational View

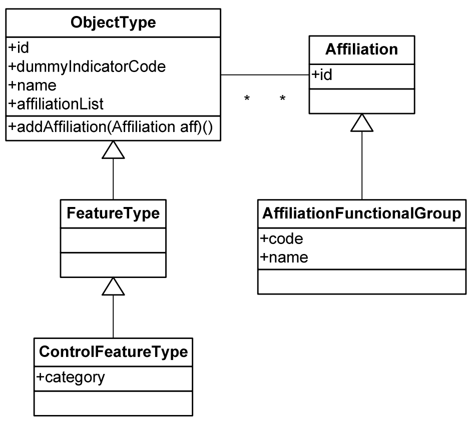


```
bufZone = new ControlFeatureType(No, UN mandated buffer zone, NOS);
affNato = new AffiliationFunctionalGroup(Multinational, NATO);
bufZone.addAffiliation(affNato);
```

Figure 9. Example "UN mandated zone" – Object-Oriented View

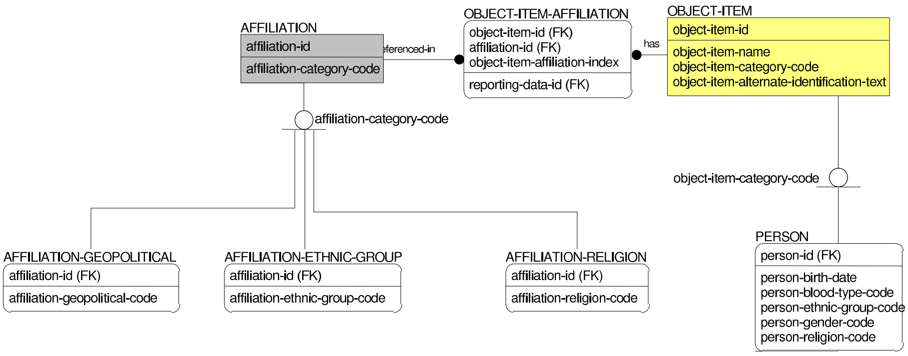

Figure 10. Ambiguity in the C2IEDM

mapping is realized manually, this leaves two options: OO classes are created by reengineering tools; the classes may contain some attributes that are only relevant in the relational model. OO classes are tailored manually; the O-R mapping tool maps onto a relational schema that is not fully identical to the MIP schema.

5.3 Normalized View

In the C2IEDM, semantically identical information can be modeled in different ways. Thus, the task of the third layer of the data access stack is to provide a normalized view. That means that all data have to be transformed in a canonical form.

There are two reasons why the representation of information can vary:

- Ambiguities in the model
  - · Violation of orthogonality
  - · Missing business rules
- Duplicated (type) information

An ambiguity in the model is shown in figure 10. The religious and ethnic affiliation of a person can be stated in two ways in the model: Either by attributes *person-ethnic-group-code* and *person-religion-code* in entity *PERSON*, or by associating records of entities *AFFILIATION-ETHNIC-GROUP* and *AFFILIATION-RELIGION* with the respective person.

In order to prevent different representations of the same information – which may also result in conflicting information – the C2IEDM has to be orthogonalized. There are also some high-level ambiguities (how to represent the location of a bridge?) that must be resolved by business rules. The MIP community is continually improving the model but there will always be some unresolved problems.

The second motivation for data normalization is duplicated type information. In multinational operations, two or more nations may create records for the same object type. These records will have identical values except that their synthetic keys differ. The same situation may arise if a data base is reused for several operations. For instance, two operators may – independently from each other – need to introduce an object type for churches. Similarly, there can be multiple definitions of the same geopolitical affiliation.

The situation where two distinct object types are essentially semantically equivalent is explicitly anticipated by MIP and must be handled by the national C2ISs. Determining duplicated data is essential for handling data queries from the national C2IS correctly. In particular if the queried data are compared or used for statistical analysis, the comparison must not rely on the equivalence of the synthetic keys only.

5.4 Business Functions

The C2IEDM is primarily designed for information exchange rather than for efficient data access. The structural complexity of the C2IEDM is caused by the fact that a multi-dimensional data space has to be mapped onto a flat model. The dimensions of the C2IEDM are:

- Time
  To satisfy loggability and traceability requirements, the C2IEDM does not allow to delete old or faulty data. Instead, any reportable data is associated with a *REPORTING-DATA* record that provides meta information such as the effective start date and time. If an operator accidentally enters erroneous data, these data must not be removed afterwards but a new meta data record has to be added that marks them as erroneous.
- Operational Information Groups (OIGs)
  OIGs are used to group related pieces of information according to logical aspects and to disseminate them depending on OIG-specific rules (see also section 4.2).

Taking both dimensions into account, the computation of the current shared tactical picture turns out to be a difficult undertaking. In order to determine the most recent status of a unit with regard to a specific OIG, the following factors have to be considered:

- The unit may or may not be an element of the given OIG and the membership may change over time.
- The most recent status of the unit may be different in each OIG. E.g., the status may be different in OIGs *Friendly and Neutral (Org)* and *Composed Plan*.
- Each status is associated with an effective start date/time that may be some point in the past, present, or future.
- Each status may also be associated with an effective end date/time at which the status information becomes invalid.

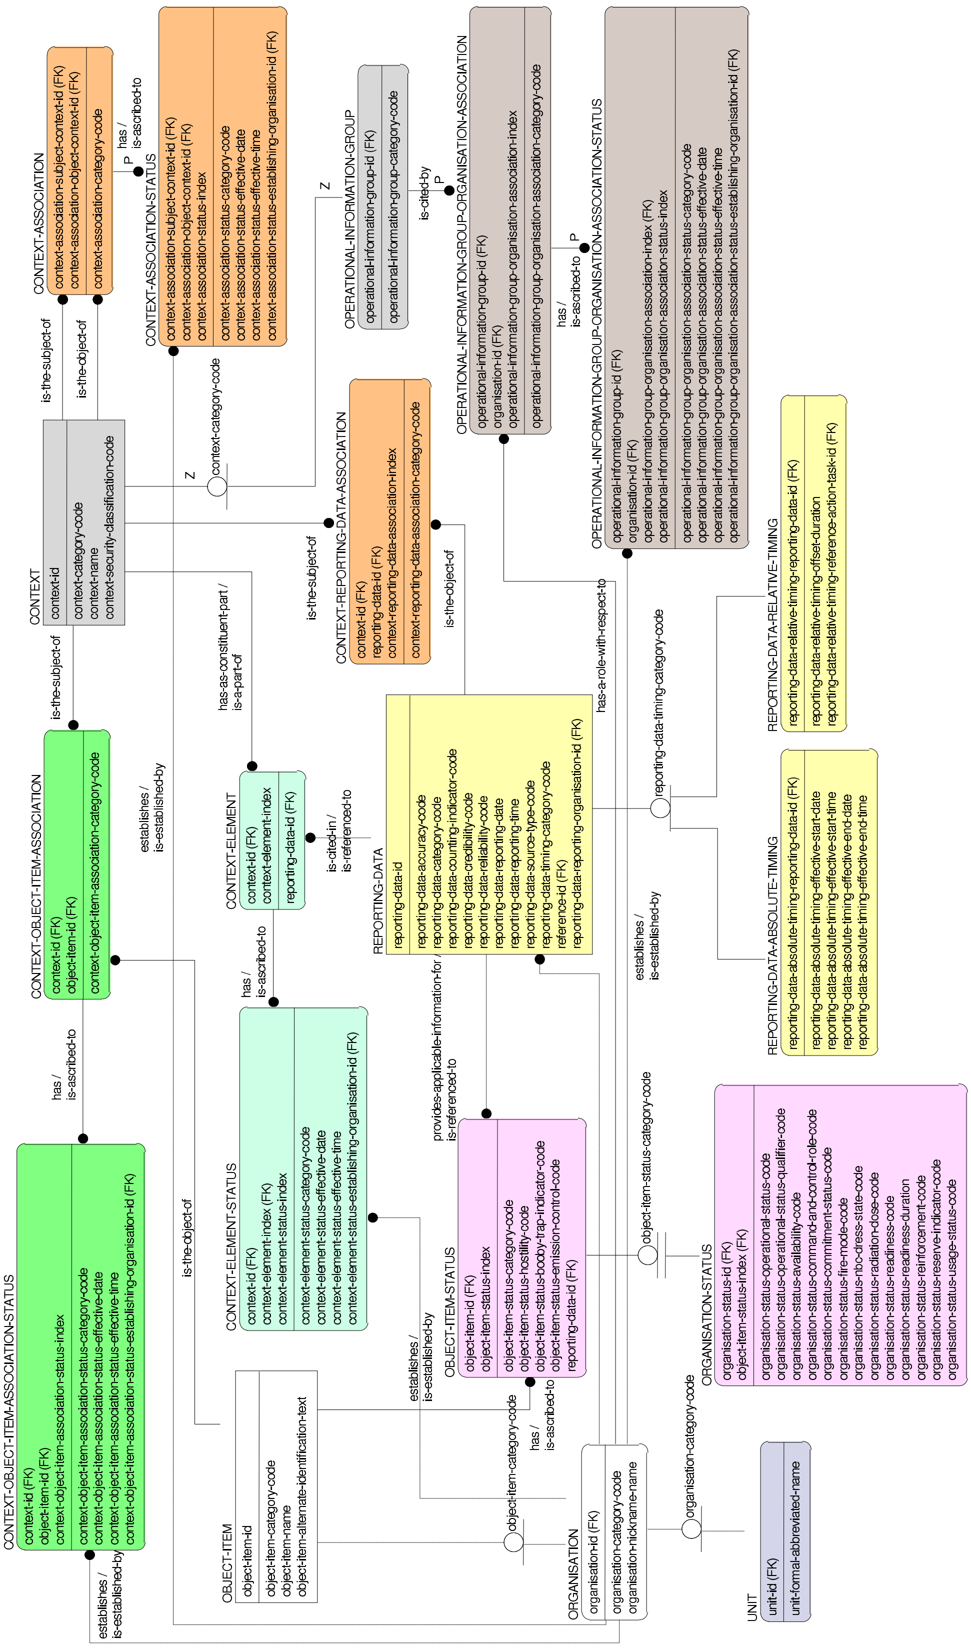

Figure 11. Computation of Most Recent Organization Status

- Status information can be negated.
- An organization can have multiple instances of the same OIG category with only one OIG instance being valid at each point in time.

Considering all factors, the records of no less than 19 (!) entities of the C2IEDM have to be evaluated. For illustration, all entities involved are shown in figure 11.

The objective of the forth layer of the data access stack is to provide C2IS applications with services for recurring and general-purpose tasks and queries (business functions). The services will return simplified data structures (business objects) for efficient processing. By means of caching of objects, e.g., caching of the current task organization, data access can be sped up significantly.

6 Summary and Conclusion

In this paper, we have shown that the MIP Solution – despite satisfying information exchange requirements only – has substantial influence on national C2ISs. For a seamless and correct implementation, a number of technical and organizational aspects have be considered that even impact the design of end user applications.

Although we focused on MIP and its C2IEDM, many of the issues described should apply to other data models and interoperability efforts in a similar way. It turns out that multinational interoperability cannot be achieved at the interfaces – it needs to be established in the core of national systems!

References

Hibernate (2005). `www.hibernate.org`.

JDO (2005). Java Data Objects. `http://java.sun.com/products/jdo/index.jsp`.

Multilateral Interoperability Programme (2004a). MIP System Requirement Specification (SRS), Edition 2.1. `http://www.mip-site.org`.

Multilateral Interoperability Programme (2004b). MIP Tactical C2IS Interoperability Requirement (MTIR), Version 2.2. `http://www.mip-site.org`.

Multilateral Interoperability Programme (2004c). The C2 Information Exchange Data Model (C2IEDM Main), Edition 6.15. `http://www.mip-site.org`.

Multilateral Interoperability Programme (2005). MIP Standard Briefing. `http://www.mip-site.org/Public_documents/MSB.ppt`.

OJB (2005). ObJectRelationalBridge. `http://db.apache.org/ojb`.