# MAINTAINING THE CONFIDENTIALITY OF INTEROPERABLE DATABASES WITH A MULTILEVEL FEDERATED SECURITY SYSTEM

Marta Oliva[1] and Fèlix Saltor[2]
[1]*Dept. Informàtica i Enginyeria Industrial, Universitat de Lleida, C. Jaume II, 69, E-25001 Lleida (Catalonia).*

[2]*Dept. Llenguatges i Sistemes Informàtics, Universitat Politècnica de Catalunya, Campus Nord-Mòdul C5, Jordi Girona Salgado, 1-3, E-08034 Barcelona (Catalonia).*

Abstract: When several databases with multilevel security policies are federated to form a tightly coupled federated database management system, heterogeneities such as different accreditation ranges must be overcome. This paper describes an extended methodology to integrate policies that use different lattices as accreditation ranges. A semi-automatic process obtains the federated accreditation range and needed translation functions among accreditation ranges in order to be validated by the security administrator.

## 1. INTRODUCTION

There is a growing need to interoperate several information sources, such as *DataBases* (DBs), to satisfy the demanding requirements of users. Technically, this is accomplished by superimposing a new system, a *Federated DB Management System* (FDBMS), as a layer of software upon the *Data Base Management Systems* (DBMSs) of each one of the pre-existing DBs, then called *Component DBs* (CDBs).

*Security* is a big concern in such an environment. Each CDB has its own security policy, enforced thru its own security mechanisms, to protect its data from unauthorized access. It joins the Federated system if its *autonomy* can be preserved, and wants to ensure that user queries coming from the FDBMS do not compromise the security of its data. These differences in security policies, models and mechanisms among the CDBs form the

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35587-0_24
M. S. Olivier et al. (eds.), *Database and Application Security XV*

*security heterogeneity*, which comes in addition to the *system, syntactic* and *semantic heterogeneities* among them (see [SL90]), which are out of the scope of this paper.

A FDBMS can federate a number of CDBs using a *Discretionary Access Control* (DAC) security model –common in business data processing-, while other CDBs are required to use a more strict *Mandatory Access Control* (MAC) model, particularly in the form of *MultiLevel Security* (MLS, [BL75]).

In this environment, the CDBs that apply MLS systems would not easily accept to share information under the control of a DAC system, because those CDBs would want information flow under control. In addition, DAC systems are not powerful enough to avoid unauthorized logical inference and aggregation, and are subject to Trojan Horse attacks ([Per93]). Consequently, it is up to the FDBMS to have mechanisms to protect interoperation among several CDBs, using or not using MLS, at the highest level of security, i.e. in terms of MLS.

Our work is related to *tightly coupled* FDBMS, in the sense of [SL90]. In such systems, there is a federated administrator, and one or more federated schemas; they are of help in enforcing security. Among the subsystems forming the FDBMS, the one in charge of security will use, as explained above, MLS, and is called the *MultiLevel Federated Security System* (MLFSS).

Even for those CDBs based on the MAC and MLS mechanisms, there are lots of heterogeneities among them, particularly in their *accreditation ranges* ([MLTS92], [DJ94]). The number of classification levels will be different, as will be their names; the *dominance relation* among levels may be a total order or a partial order; frequently a lattice.

The objective of this work consists of extending our previous methodology to integrate multilevel security policies (using total ordered sets), to integrate multilevel security policies that use lattices as accreditation ranges. As a result of this work we are able to integrate any set of MLS DBs.

This paper is organized as follows. Our schema integration process is summarized in section 2, and an overview of our integration methodology, among multilevel security policies that use total ordered set as accreditation ranges, in section 3. Section 4, the core of the paper, presents the extended methodology to integrate multilevel security policies that use lattices as accreditation ranges. Section 5 includes related work and conclusions.

## 2. SCHEMA INTEGRATION PROCESS SUMMARY

When several DBs decide to share their data they can use several interoperation ways. Since we are interested in tightly coupled federations, we follow a building methodology like that presented by [SL90] by means of their 5-level schema architecture. In this case the integration process is based on our 7-level schema architecture ([ROSC97]), where one of the additional schemas (*Authorization Schema*) is related to security aspects. Another important aspect to note is that we use *BarceLona Object-Oriented Model* (BLOOM) as Canonical Data Model (CDM) of the federation.

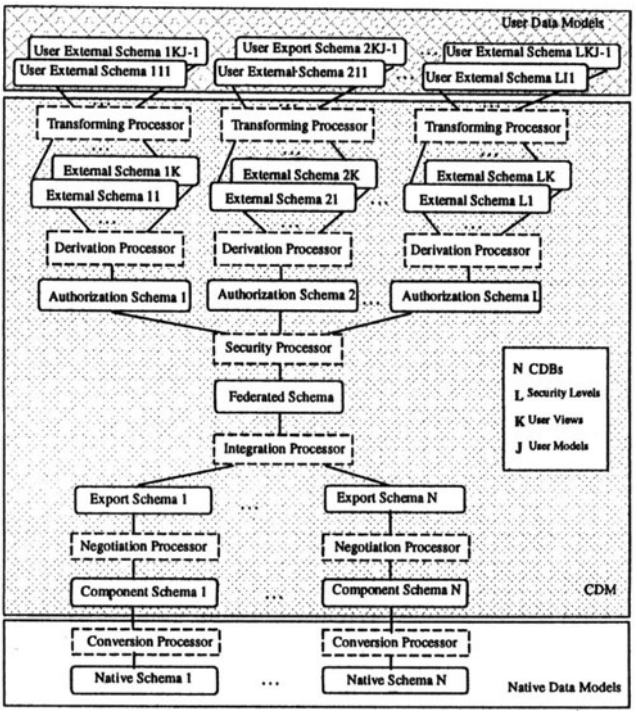


*Figure 1.* 7-level schema architecture

From the point of view of maintaining data confidentiality, it is necessary to extend several processors between schemas in our 7-level schema architecture (see figure 1). First of all, when the *Conversion Processors* convert schemas expressed in data models different from BLOOM into BLOOM schemas to obtain the *Component Schemas*, it is essential to maintain the classification of data, which depends on native accreditation ranges of security levels. Thus, it is also necessary to deduce the classification of the components obtained by the semantic enrichment

process, to maintain the confidentiality of data and to prevent possible inferences. Through the negotiation process administrators decide which data of which confidentiality will be offered to the federation in the form of *Export Schemas*.

After that, the *Integration Processor* integrates the *Export Schemas* into a *Federated Schema*. The integration is performed in two phases: identifying semantically related objects in the detection phase, and integrating the CDBs BLOOM schemas after conforming them, by means of a discriminated form of generalization, in the resolution phase). Thus, the *Integration Processor* also has to solve heterogeneities (set of security levels and their semantics) among the *Component Accreditation Ranges* (CARs), obtaining the *Federated Accreditation Range* (FAR). Then, the translation functions between the FAR and each CAR could be deduced, and later the process will be able to classify each component of the *Federated Schema* in terms of the FAR (at the same time, solving differences of initial classification of equivalent components). Sections 3 and 4 of this paper focus on an integration methodology to be applied by the *Integration Processor* to integrate different CARs.

After that, the *Security Processor* creates as many *Authorization Schemas* as security levels the FAR has. Each *Authorization Schema* contains the subset of the *Federated Schema* that a federated user, with certain clearance level, can access (read-only). At the same time, *Authorization Schemas* allow changing the form of the *Federated Schema* in order to hide confidential data that could be inferred because of the semantics of BLOOM abstractions ([AOSS00]).

## 3. OUR METHODOLOGY TO INTEGRATE MULTILEVEL SECURITY POLICIES

Our proposed *semi-automatic* methodology, presented in [OS01], to integrate multilevel security policies and to solve their heterogeneities collaborates with the schema integration process summarized in the previous section. We use the term semi-automatic because this process needs to be validated by the *Security Administrator* of the federation.

The methodology was developed, at first stage, to integrate MLS CDBs that apply total ordered sets as accreditation ranges, for instance $(S_1, \leq_1)$ and $(S_2, \leq_2)$ (see figure 2.a).

To detect the relationships among security levels belonging to different ordered sets it uses the following information:

a) The confidentiality of data stored in the Export Schemas (see the classification of the components of each CDB schema in fig. 2.a).

b) The similarities between components of different Export Schemas, obtained by the schema integration process, these are more precisely, Equivalence Semantic Relationships (E-SRs) and Specialization Semantic Relationships (S-SRs) (see fig. 2.a).

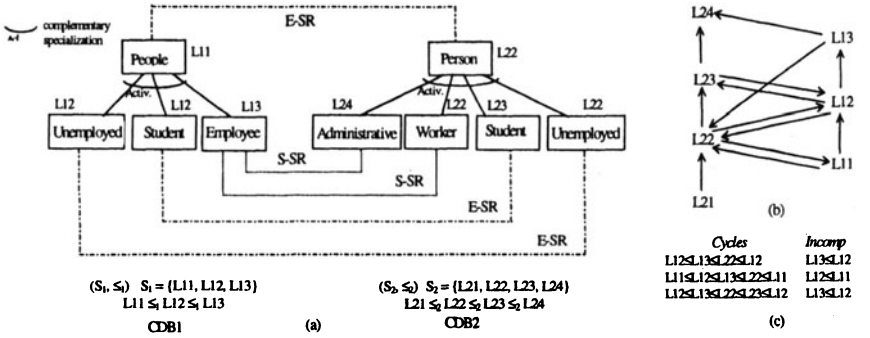


*Figure 2.* Multidigraph building and analysis example

The methodology is based on a multidigraph representation of the information described above (see an example in figure 2.b). The multidigraph is built from Hasse diagrams of ordered sets of the CDBs and a set of arcs representing possible dominance relations between security levels of different CARs. The first proposal of the set of dominance relations is obtained from the set of similarities between components of different Export Schemas (it is easy to think that similar concepts have similar confidentiality) without any analysis, so the greatest quantity of information is considered. For each E-SR (see fig. 2.a and 2.b) an arc is added to the multidigraph from the classification level of a class to the classification level of the class that is assumed the equivalent concept, and also the inverse arc. For S-SR (see fig. 2.a and 2.b) an arc is added to the multidigraph from the classification level of a class to the classification level of the class that is assumed to be the subclass concept.

Then, the multidigraph is analyzed to detect incompatible arcs. Two arcs are incompatible if their existence breaks the dominance relation between security levels of the initial ordered sets. A penalization approach is used to automatically solve incompatibilities losing the least number of arcs. The penalization approach is based on the detection of cycles with more than two participant vertices. An arc accumulates as many penalizations as the number of cycles it is involved in. For instance, in figure 2.b, the arc (L13, L22) is involved in 3 cycles and produces 3 incompatibilities (see fig. 2.c).

After that, the most penalized arc is removed and the penalization of the arcs are calculated again, repeating this process until the remaining arcs do not have any penalization. When there is no incompatibility, the resulting

multidigraph is converted into an ordered set to be used as FAR, and at the same time translation functions are obtained. Note that in the final multidigraph, each cycle of only two participant vertices indicates that both security levels are the same. Finally, the components of the Federated Schema are classified taking the FAR and translation functions into account. The classification of classes that initially appears in Export Schemas are translated using appropriate translation functions, and new classes, built by means of the discriminant generalization, are classified at the least level of the set of classification levels belonging to their subclasses.

# 4. LATTICES INTEGRATION METHODOLOGY

As pointed out in the introduction of this paper, in MLS systems sometimes the dominance relation among security levels forms a partial ordered set or a lattice. The main characteristics of a lattice are that a greatest and a least element, among all elements forming the lattice, exist. Because of all ordered sets also have greatest and least elements they are lattices as well. Besides, as we can always include a fictitious greatest element as well as a fictitious least element into the set of elements forming a partial ordered set, then we can also work with partial ordered sets as they were lattices. So, a methodology that allows integrating lattices will also be able to integrate total and partial ordered sets.

Taking into account the integration methodology of multilevel security policies summarized in the previous section, the main difference between ordered set integration and lattice integration is the set of possible incompatibilities that can be obtained after building the corresponding multidigraph. In subsection 4.1 these possible incompatibilities are analyzed. After that, subsection 4.2 covers the penalization approach and the set of criteria used to solve the incompatibilities. Finally, three special cases that need a different automatic solution are treated in subsection 4.3.

## 4.1 Incompatibilities

From the point of view of possible incompatibilities, the difference between an ordered set and a lattice is the presence, in the latter, of no comparable elements. As a result, apart from incompatibilities originated by cycles produced by more than two vertices, another incompatibility is produced when the multidigraph contains a path between two incomparable security levels, belonging to the same lattice, which implies a dominance relation between them.

Figure 3.a shows an example where the set of dominance relations between lattices $L_1$ and $L_2$ produces a path. Because of this path relates two no comparable levels, belonging to the same lattice, it causes an incompatibility (see fig. 3.b).

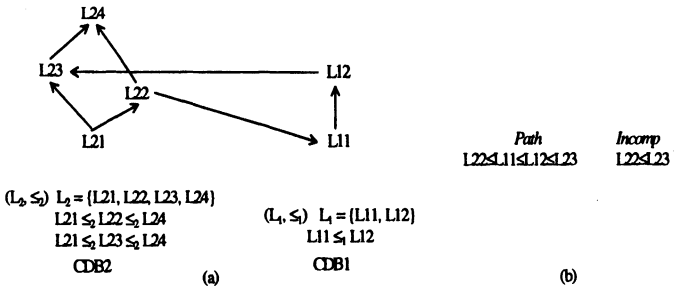


*Figure 3* Example of a path between no comparable security levels

## 4.2 Penalization approach

Following the same penalization approach as in the methodology that integrates total ordered sets, each arc of the multidigraph (originated by the representation of the possible dominance relations between security levels of lattices) accumulates as many penalizations as the number of cycles and paths it is involved in.

After penalizing all arcs involved in cycles or in paths, the process uses the following criteria to decide which arc is the best to remove:

1. Look for the most penalized arc.
2. If there is more than one arc with the same greatest penalization, then the one that appears the least number of times is removed. Remember that we use a multidigraph because the same arc can be added several times, depending on the times that the schema integration process finds similarities that produce the same dominance relation between two security levels.
3. If it is possible to apply neither the first nor the second criterion, then the process chooses the arc that relates more confidential levels. Generally, it is easier to share less confidential than more confidential data.

Once an arc is removed, the penalization of the arcs are calculated again, and the previous process is repeated until all penalized arcs disappear.

Figures 4, 5 and 6 show an example of the application of the first, second and third criterion respectively using the lattices $L_1$ and $L_2$ previously depicted in figure 3. The part (a) of each figure contains the initial multidigraph, in the part (b) is depicted the resultant multidigraph after

removing the arc. Finally, the part (c) shows the resultant FAR after integrating the initial lattices.

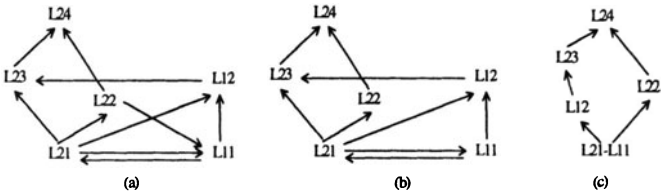


*Figure 4.* First criterion applicability example

Multidigraph in figure 4.a has the cycle L21-L22-L11-L21 and the path L22-L11-L12-L23, so the arc (L22, L11) is chosen for removing because it has 2 penalizations in front of the arcs (L11, L21) and (L12, L23) that only have a penalization. After removing the arc (L22, L11) no incompatibility remains (see fig. 4.b). Figure 4.c shows the resultant FAR where, as the cycle in figure 4.b points out, L21 and L11 are the same level, so the arcs (L11, L12) and (L21, L12) imply the same dominance relation, and the arc (L21, L23) disappears because is a redundant arc.

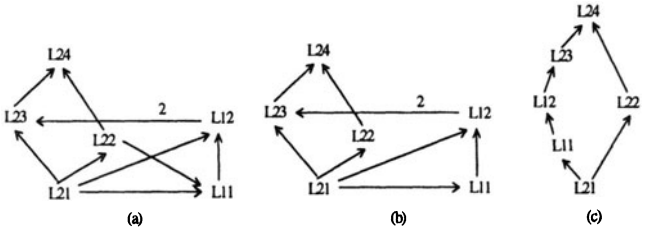


*Figure 5.* Second criterion applicability example

Multidigraph in figure 5.a only has the path L22-L11-L12-L23, so both arcs (L22, L11) and (L12, L23) has just a penalization. But as the arc (L12, L23) appears twice (see the number with which the arc (L12, L23) is tagged) then the arc (L22, L11) is chosen to be removed (see fig. 5.b). The resultant FAR depicted in figure 5.c is obtained in the same way that the FAR in figure 4.c.

Multidigraph in figure 6.a only has the path L22-L11-L12-L23, but this time the first and the second criterion are not applicable because both arcs have the same penalization and appear just the same number of times. Consequently, the arc (L12, L23), that relates more confidential levels than the arc (L22, L11), is removed (see fig. 6.b). The FAR in figure 6.c is obtained just as the previous ones, but this time we use a fictitious *Federated Top* (FT) security level in order to obtain a lattice.

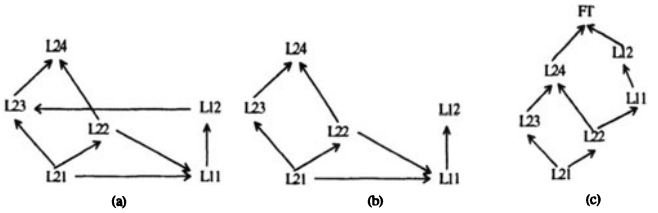


*Figure 6.* Third criterion applicability example

## 4.3 Special cases

In spite of the criteria described in previous section, there are three special cases that need a specific treatment. In all of these special cases there is at least a path between two no comparable levels, but it only passes through a unique level belonging to the other lattice. Although it is impossible to apply any of those criteria to solve the incompatibility, there is a different semi-automatic solution. Note that, in a MLS DB, the set of elements classified at one of the no comparable levels has nothing in common with the set of elements classified at the other one. So, it is possible to consider the set of elements classified at the level through which the path passes just as two different subsets. These subsets can be identified by means of the information about similar objects (E-SRs and S-SRs), obtained thru the schema integration process. Figures 7, 8 and 9 depict an example of each of these special cases. In the part (a) of each figure is depicted the multidigraph obtained thru the methodology summarized at section 3, while the part (b) of each figure shows the resultant FAR after applying the proper semi-automatic solution. In each figure, the part (c) includes a simple integration example of two CDBs. Each integration example depicts the CDBs Export Schemas with their classification and the similarities between objects. The classification, in terms of the FAR, of the resultant Federated Schema is shown as well. Consequently, you can see how applying our methodology to integrate security policies a secure interoperation without security breaches is obtained.

Multidigraph in figure 7.a shows the special case where there are two no comparable levels of the first lattice that are related with the same level of the second lattice. Thus, because the relation is an equivalent relation in both cases, it would imply that the L23 and L22 levels would be the same level (in this case L12). We avoid this by splitting L12 into L12a and L12b (because of the reason explained above), and including them, together with L12, in the resultant FAR (see figure 7.b). This is because each federated user only has to be assigned one clearance level, and at the same time, interoperation access using this clearance level has to comply with the

autonomy and security principles of [GQ96]. In addition, FAR in figure 7.b includes a fictitious *Federated Bottom* (FB) security level, like the inclusion of FT in figure 6.c.

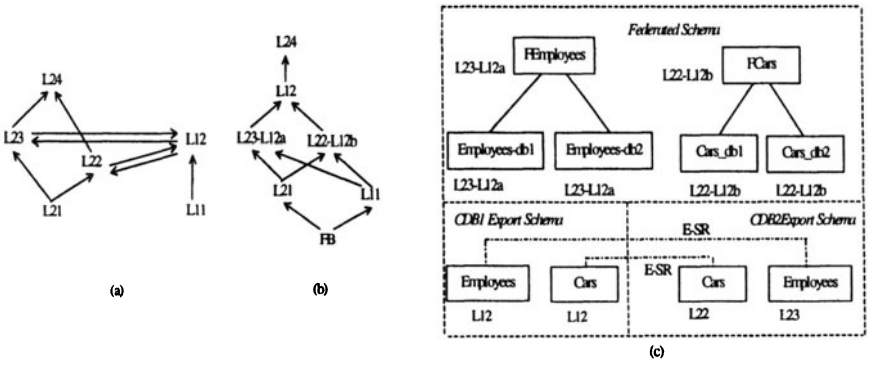


*Figure 7.* First special case example

Remember that the principle of autonomy and of security ([GQ96]) say, respectively, that any access permitted within an individual system must be also permitted under secure interoperation; and that any access not permitted within an individual system must be also denied under secure interoperation.

*Table 1.* Translation functions

| | | |
|---|---|---|
| $f_2^F$(L24)=L24 | $f_F^2$(L24)=L24 | $f_F^1$(L24)=L12 |
| $f_2^F$(L23)=L23-L12a | $f_F^2$(L12)={L23, L22} | $f_F^1$(L12)=L12 |
| $f_2^F$(L22)=L22-L12b | $f_F^2$(L23-L12a)=L23 | $f_F^1$(L23-L12a)=part{L12} |
| $f_2^F$(L21)=L21 | $f_F^2$(L22-L12b)=L22 | $f_F^1$(L22-L12b)=part{L12} |
| | $f_F^2$(L21)=L21 | $f_F^1$(L21)=∅ |
| $f_1^F$(L12)=L12 | $f_F^2$(L11)=∅ | $f_F^1$(L11)=L11 |
| $f_1^F$(L11)=L11 | $f_F^2$(FB)=∅ | $f_F^1$(FB)=∅ |

For better understanding, table 1 describes translation functions between each component lattice and the federated lattice, and translation functions between the federated lattice and each component lattice as well. The first functions are used to translate the component user clearance into the federated user clearance ($f_i^F$, where *i* is $CDB_i$). The last ones are used to translate a federated user clearance assigned to a query to the component user clearance ($f_F^i$), assigned to the corresponding component subqueries.

*$f_F^2$(L12)={L23, L22}* means that a query, belonging to a federated user with L12 clearance level, has to translate into two subqueries when the MLFSS passes the query to the $CDB_2$: first subquery tagged with the L23 clearance level, and second tagged with L22. Finally, the FDBMS has to consolidate the result of both subqueries. On the other hand, *$f_F^1$(L23-L12a)=part{L12}* means that a query belonging to a federated user with the

L23-L12a clearance level, has to be translated into a query tagged with L12 when it was passed to the $CDB_1$. After that the MLFSS will need to filter the result to return only part of elements classified at L12 security level.

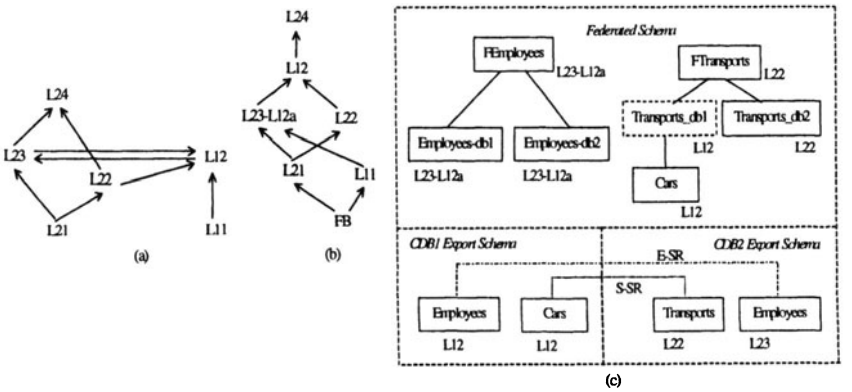


*Figure 8.* Second special case example

Figure 8 shows an example of the second special case. In this case, one of the two no comparable levels also has an equivalence relation with the security level of the other lattice, but the other one is dominated by the level of the other lattice (see fig. 8.a). As in the first special case, the solution consists of splitting the set of elements classified at level L12, but in this case it is enough to only take a subset into account, just the subset L12a (see fig. 8.b). Seeing the example of figure 8, note that the Federated Schema depicted in figure 8.c includes the class `Transports_db1` (not appearing in the CDB1 Export Schema) because it is originated thru the conformation of the CDB1 Export Schema during the schema integration process.

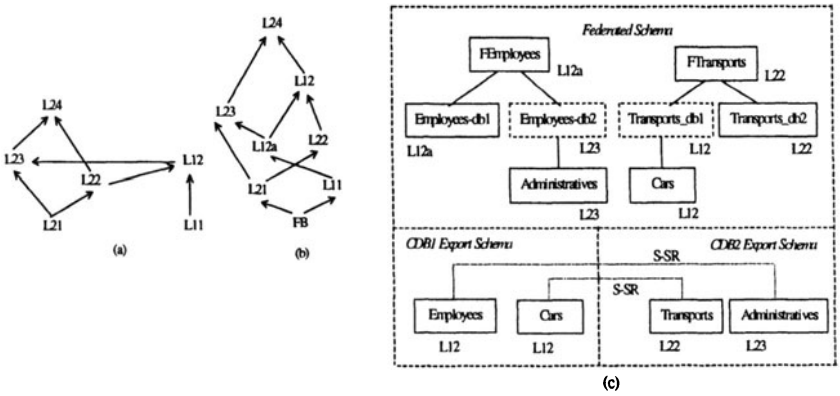


*Figure 9. Third special case example*

You can see an example of the third special case in figure 9. Both incomparable levels are not equivalent to the security level of the other lattice. Whereas L22 is dominated by L12, L23 dominates L12, so these dominance relations form a path (see fig. 9.a). As in the second special case, only considering the subset L12a is enough to solve the integration (see fig. 9.b). Translation functions between lattices and the FAR in examples in figures 8 and 9 can be obtained in the same way as those in example 7.

## 5. RELATED WORK AND CONCLUSIONS

Several papers in the literature are related to secure authorized accesses during interoperation among heterogeneous DBs. Some of them work in the context of FDBMSs and in particular in tightly coupled FDBMSs. [JD94], [JD96] and [DCdVS97] base their security systems on the DAC mechanisms; [Per93], [TF97] use canonical security systems, that support both the DAC and MAC mechanisms, and [Osb01] protects interoperable accesses by means of a *Role-Based Access Control* (RBAC) system. As far as we know, only [IGQ94] and [BSS96], apart from [DQS00] (that is related to loosely coupled using a mediator/wrapper approach), base the security system of the federation on a MLS system.
Particularly, [IGQ94] presents an integration methodology to solve confidentiality conflicts produced when a federation is built. The methodology is also based on a technique to create a global data schema, but the security systems of the CDBs use MLS systems with identical accreditation ranges. The problem is reduced to the resolution of the classification of each object belonging to the global data schema.

On the other hand, [BSS96] describes techniques (a logic programming approach and a graph-based approach) to obtain a unified accreditation range through a way of combining different accreditation ranges that preserves a set of interoperation security constraints. In this case, these techniques use a predefined set of interoperation security constraints, but the paper does not indicate how they are obtained.

Although [DQS00] is related to loosely coupled systems, there are some relevant aspects related to our work. The paper includes a way to check the correctness (consistency, nonambiguity and nonredundancy) of the specifications of mappings between levels of an external application and levels of the sources. They also analyze inconsistencies by means of cycles and paths among security levels of different accreditation ranges. But in this case, the analysis is not performed between CARs; it is

performed among the accreditation range of an external application, defined manually by a human integrator/administrator, and the accreditation ranges of the sources. Mappings among accreditation ranges also are determined by the same human integrator/administrator and in case they are incorrect, the human integrator/administrator has to propose some changes in order to obtain a correct solution.

Finally, [GQ96] shows a study of the problem of secure interoperation among heterogeneous systems. This study insists on the detection and resolution of the security breaches caused by the interoperation but without specifying which security policy is applied. It starts from the assumption that mappings among security attributes and among distinct heterogeneous security systems have been previously solved (for example by the administrator of the FDBMS).

In this paper we extend our methodology to integrate multilevel security policies, presented at [OS01], that takes into account MLS CDBs that use total ordered sets as accreditation ranges, in order to integrate MLS CDBs that use lattices as accreditation ranges. The described methodology can also be used to integrate multilevel security policies using total or partial ordered sets. Using this methodology a semi-automatic process is obtained, and it only has to be supervised by the security administrator in order to validate it. As result, the accreditation range of the federation and translation functions between each CAR and the FAR are obtained.

## REFERENCES

[AOSS00] A. Abelló, M. Oliva, J. Samos and F. Saltor. Information System Architecture for Data Warehousing from a Federation. In M. Roantree, W. Hasselbring and S. Conrad, editors, Engineering Federated Information System, (Proceedings of the 3rd Workshop EFIS 2000, June 19-21, 2000, Dublin (Ireland)), pages 33-40, infix, 2000.

[BSS96] P.A. Bonatti, M.L. Sapino and V.S. Subrahmanian. Merging Heterogeneous Security Orderings. In E. Bertino, G. Kurth, H. Martella and E. Montolivo, editors, Computer Security - ESORICS 96 (4th European Symposium on Research in Computer Security, Rome, Italy, September 25-27, 1996, Proceedings), volume 1146 of LNCS, pages 183-197, Springer-Verlag, 1996.

[BL75] D.E. Bell and L.J. LaPadula. Secure computer systems: Unified exposition and multics interpretation. Technical Report MTR-2997, (AY/W 020 445), The MITRE Corporation, Bedford, MA, Jul 1975.

[DJ94] K.R. Dittrich and D. Jonscher. Current Trends in Database Technology and Their Impact on Security Concepts. In J. Biskup, M. Mongersten and C.E. Landwehr (eds), Database Security VIII (A-60),. Elsevier Science B.V. (North Holland), pages 11-33, 1994.

[DCdVS97] S. De Capitani di Vimercanti and P. Samarati. Authorization Specification and Enforcement in Federated Database Systems. Journal of Computer Security, 5(2):155-188, 1997.

[DQS00] S. Dawson, S. Qian and P.Samarati. Providing Security and Interoperation of Heterogeneous Systems. Distributed and Parallel Databases, 8(1): 119-145, Jan 2000.

[GQ96] L. Gong and X. Qian. Computational Issues in Secure Interoperation. IEEE Transactions on Software Engineering, 22(1):43-51, January 1996.

[IGQ94] N.B. Idris, W.A. Gray and M.A. Qutaishat. Integration of Secrecy Features in a Federated Database Environment. In T.F. Keefe and C.E. Landwehr, editors, Database Security VII (A-47), pages 89-109. Elsevier Science B.V. (North-Holland), 1994.

[JD94] D. Jonscher and K.R. Dittrich. An Approach for Building Secure Database Federations. In Proccedings of the 20th VLDB Conference, pages 24-35, 1994.

[JD96] D. Jonscher and K.R. Dittrich. Argos – A Configurable Access Control System for Interoperable Environment. In T.C. Ting and D. Spooner, editors, Database Security IX: Status and Prospects. Chapman and Hall, 1996.

[MLTS92] M. Morgenstern, T. Lunt, B. Thuraisingham and D. Spooner. Security issues in federated database systems: panel contributions. In C.E. Landwehr and S. Jajodia, editors, Database Security V (A-6): Status and Prospects, pages 131-148. Elsevier Science B.V. (North Holland), 1992.

[OS01] M. Oliva and F. Saltor. Integrating Multilevel Security Policies in Multilevel Federated Database Systems. In B. Thuraisingham, R. van de Riet, K.R. Dittrich, Z. Tari, editors, Data and Applications Security: Developments and Directions. Kluwer Academic Publishers, 2001.

[Osb01] S. Osborn. Database Security Integration using Role-Based Access Control. In B. Thuraisingham, R. van de Riet, K.R. Dittrich, Z. Tari, editors, Data and Applications Security: Developments and Directions. Kluwer Academic Publishers, 2001.

[Per93] G. Pernul. Canonical Security Modeling for Federated Databases. In D.K. Hsiao, E.J. Neuhold, and R. Sacks-Davis, editors, Interoperable Database Systems (DS-5) (A-25), pages 207-222. Elsevier Science Publishers B.V. (North-holland), 1993

[ROSC97] M.E. Rodríguez, M. Oliva, F. Saltor and B. Campderrich. On Schema and Functional Architectures for Multilevel Secure and Multiuser Model Federated DB Systems. In S. Conrad, W. Hasselbring, A. Heuer, G. Saake, editors, Proceedings of the International CAiSE'97 Workshop on Engineering Federated Database Systems (EFDBS'97, Barcelona), Otto-von-Guericke-Universität Magdeburg, Fakultät für Informatik, preprint Nr. 6, pages 93-104, 1997.

[SL90] A.P. Sheth and J.A. Larson. Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases. ACM Computing Surveys, 22(3):183-236, September 1990.

[TF97] Z. Tari and G. Fernandez. Security Enforcement in the DOK Federated Database System. In P. Samarati and R.S. Sandhu, editors, Database Security X: Status and Prospects, pages 23-42. Chapman and Hall, 1997