Mobile and Heterogeneous databases

Security

A.R. Hurson

Computer Science

Missouri Science & Technology

Heterogeneous Distributed Databases

Note, this unit will be covered in two lectures. In case you finish it earlier, then you have the following options:

1) Take the early test and start CS6302.module7
2) Study the supplement module (supplement CS6302.module6)
3) Act as a helper to help other students in studying CS6302.module6

Note, options 2 and 3 have extra credits as noted in course outline.

Heterogeneous Distributed Databases

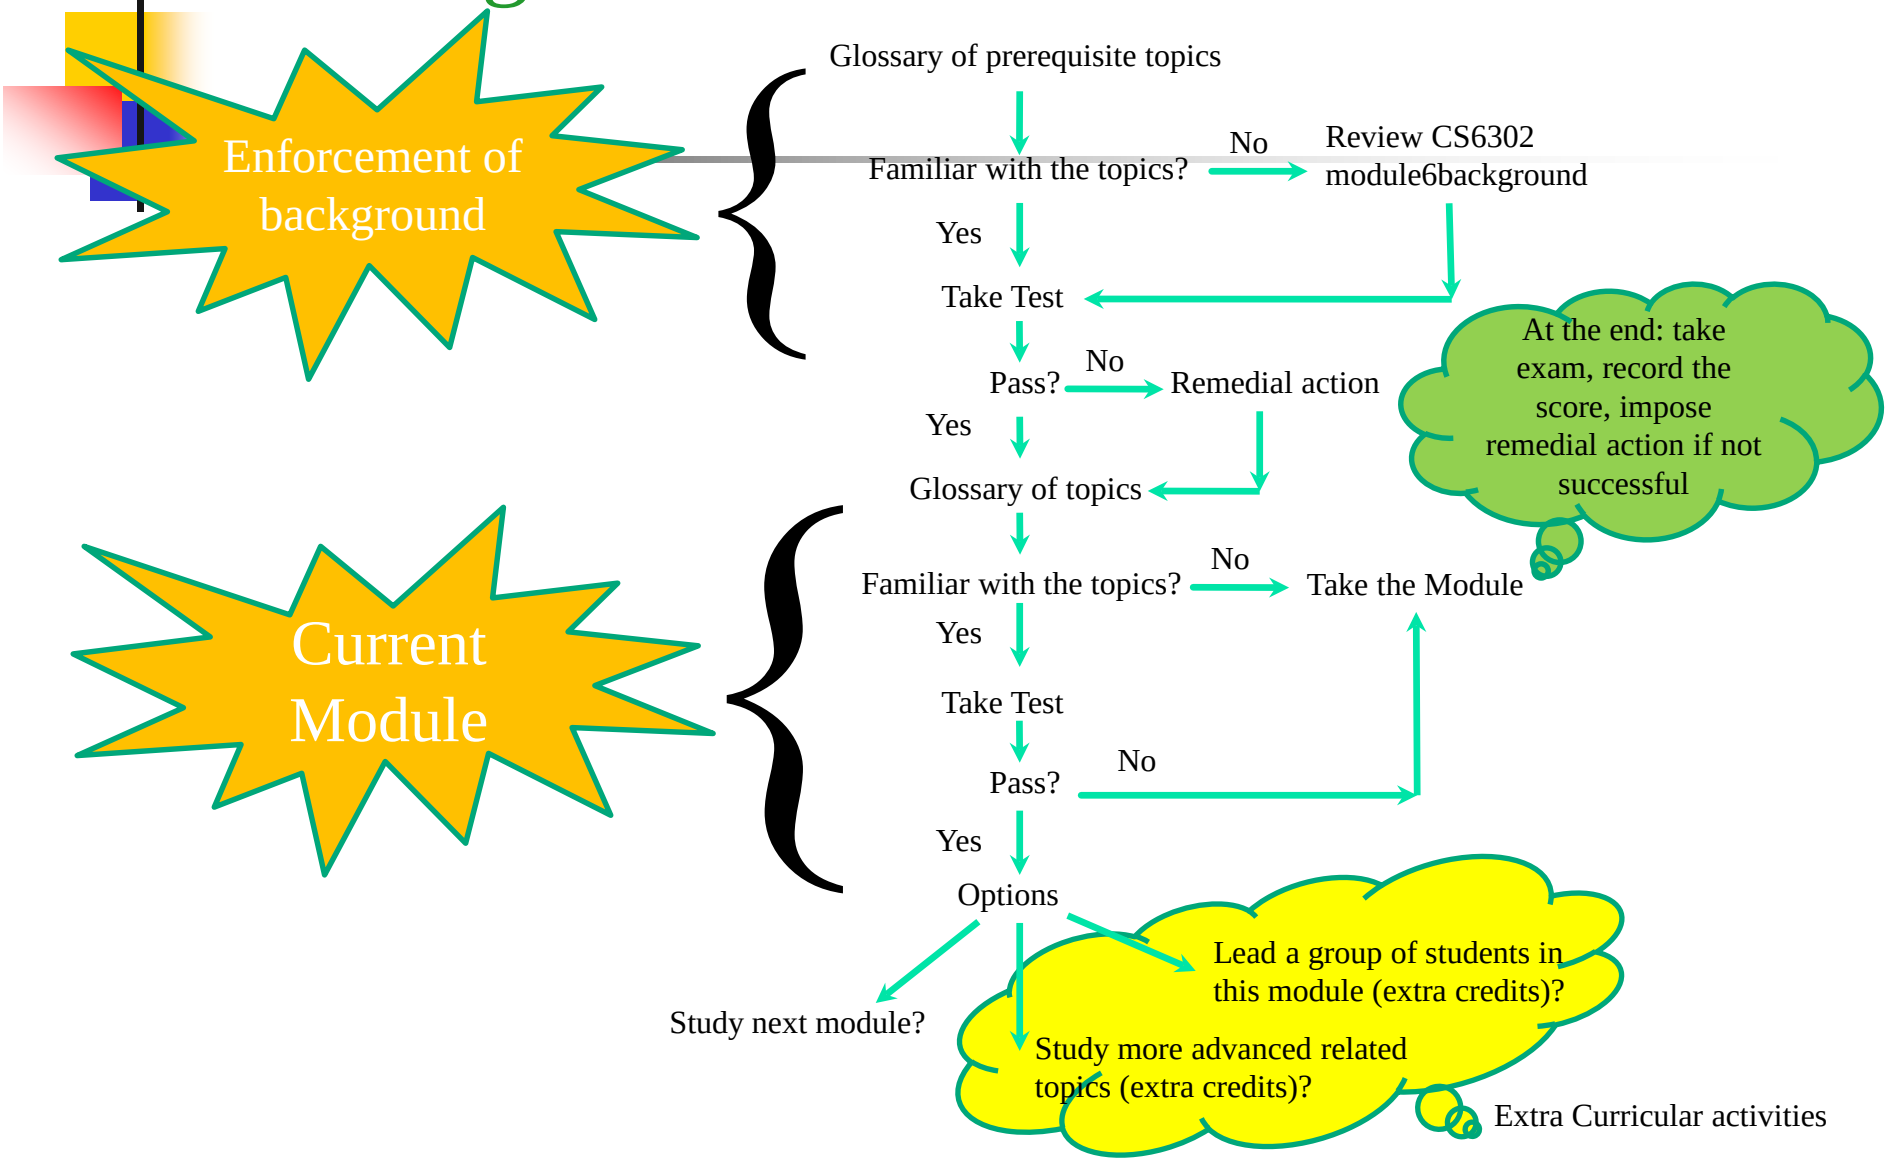

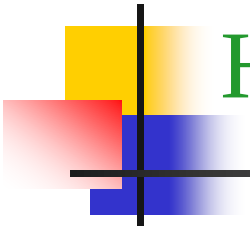

Heterogeneous Distributed Databases

- You are expected to be familiar with:
  - Heterogeneus Distributed Databases,
  - Security issues in centralized database environment
- If not, you need to study CS6302.module4 and module6.background

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security
  - The security issue in a multidatabase system relative to the traditional distributed system is becoming more complicated due to the:
    - Heterogeneity of the local databases, since different sites may use different and incompatible mechanisms for expressing and enforcing the security policies.
    - Autonomy of the local databases, since each site determines what security mechanism can be enforced. In addition, because of the communication autonomy, the local site may decide not to communicate such an information globally.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In the context of multidatabase systems, the goal of authentication is the same as it is in centralized DBMSs – identifying one party to another. However, the problem of authentication is far more complex due to the distribution of the parties that are involved.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Users, workstations, communication channels, and services are the basic components of a distributed system, while users and workstations are the major principals in a centralized system.
  - In a centralized environment, the authentication is between users and workstations. In a multidatabase system, users often log on to local workstations to access services provided by remote workstations.
    - The user must be authenticated at a local workstation first.
    - The local workstation, acting on the user's behalf, is mutually authenticated with the remote service provider.
  - Since, communication links are involved in the authentication process, countermeasures must be taken to handle eavesdropping, replay attacks, and masquerading. Without secure communication channels, the authentication system can be easily compromised.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Authentication in centralized database

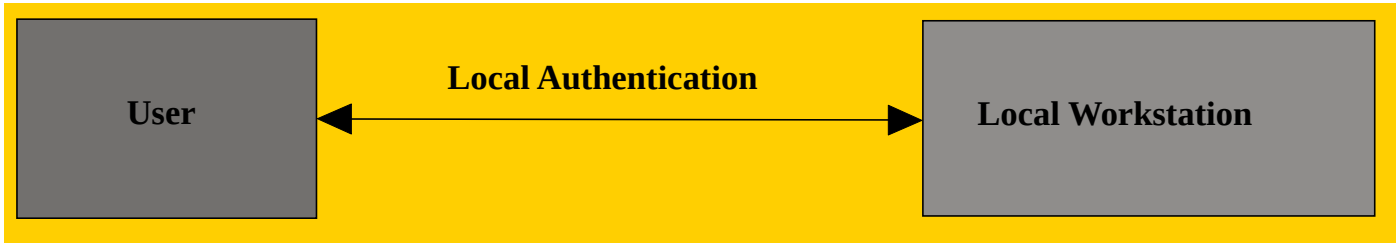

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Authentication in distributed database

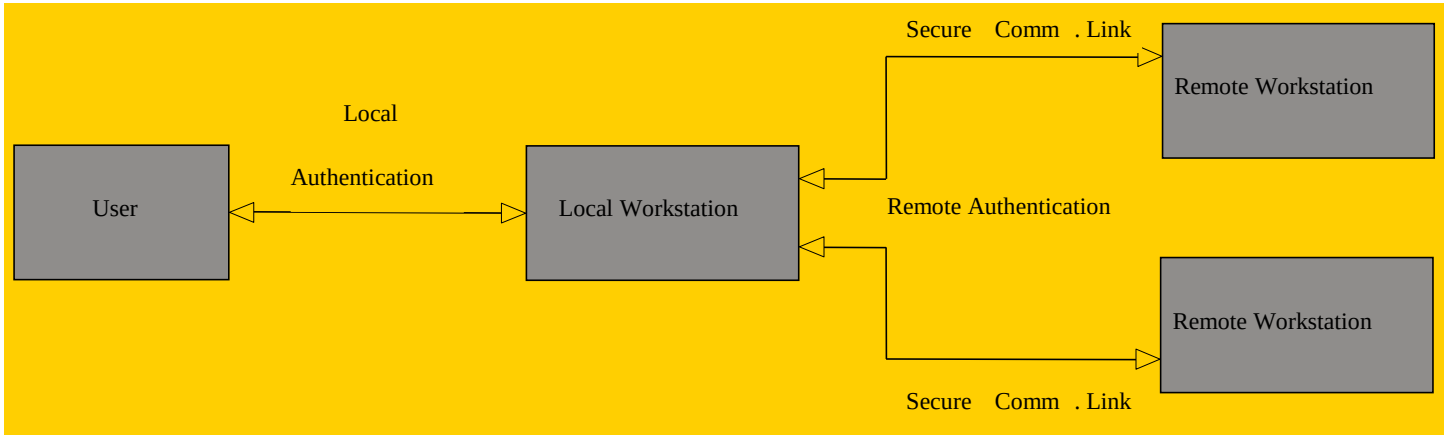

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication

Issue		Solution
Authentication	At Global schema level	• Authenticate users to system only after verification. • Authenticate users to system without verification (This is generally seen in loosely coupled databases)
	At Local database level	• Authenticates users only after verification of the identity of the user (autonomous system) • Automatically authenticate users to the local system if the global schema accepts the user (Delegates the authentication to the multidatabase global schema)

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Authentication protocols for distributed systems can be classified as:
    - Symmetric cryptosystem based challenge-response,
    - Asymmetric cryptosystem based challenge-response,
    - Router-based, and
    - Agent & model based.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Issues increasing the complexity of authentication in multidatabase systems.

Heterogeneity of Local Authentication Components	• Global users have to pass through different procedures to gain access. • The identity of a user can vary from system to system. • Each user should be authenticated once but correct to all relevant participating systems per session.
Local Autonomy	• Component DBSs have the autonomy to decide whether a user is valid.
Uniformity	• The same user may have different identities and identifiers but has to be handled uniformly.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In federated database systems, access to data can be seen at two different levels:
    - The federation level and
    - The local level.
  - At the federation level, users explicitly require access to the federated data, while at the local level, the local requests corresponding to the global requests must be processed.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In federated database systems, with respect to who should enforce authentication, we can distinguish between local and global authentication.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In local authentication users are required to re-authenticate themselves at each local site. Upon reception of a request by the federation, the local site asks the user to identify himself/herself and, after authentication, performs access control and possibly returns the data to the federation.
  - In global authentication, a user's identity is passed to the site by the federation along with the request.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication

Local Authentication	Advantages	• Local access decisions can be taken with respect to identifiers known at the site • Does not require the local site to be informed about remote or federation identities of users
	Disadvantages	• May make access control process very heavy • Each access request on federated data could be split into several access requests on local data (possibly stored at different sites), which would require the user to login to each site involved in the transaction.
Global Authentication	Advantages	• Users are not required to authenticate themselves at each local site
	Disadvantages	• The local system needs to put some trust on the remote or federation identity. • Authorization at local sites needs to be specified with respect to identities not administered by the local site itself.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In federated database systems users can be classified into three groups:
    - Local users with one identity per affiliated system,
    - Global users with one global identity, and
    - Federated users with local and global identities.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In federated database systems three authentication policies have been proposed:
    - Direct Authentication
    - Indirect Authentication
    - Global Authentication

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Direct authentication requires:
    - The user to be authenticated by all participating systems that he/she wishes to access.
  - This approach is suitable under the following situations:
    - High local autonomy and security requirements
    - Low trust between the participating systems
    - Invisible heterogeneity.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - Indirect authentication approach derives the relevant user information for the local authentication indirectly from a special component, not directly from the user.
  - Without a global component, each database stores not only a user's identity and identifier used by that database, but also the user's identities and identifiers used by all other databases.
  - In the presence of a commonly trusted global component, a user can be authenticated using his/her global identity and identifier.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authentication
  - In global authentication approach the FDBMS takes full control of the authentication process.
  - This approach is only suitable for special applications since it sacrifices local autonomy.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Research issues involved in access control in multidatabase system include:
    - Administration of authorization,
    - Authorization specification, and
    - Access control policy heterogeneity

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control

Issues		Solutions
Access Control	At Global Schema level	• Local identity to issue access to user (assigned by local database) • Unique global identity to decide access to user (assigned by multidatabase global schema) • Remote identity to allow access to local component (assigned by third party like trusted key distribution center and agreeable to both multidatabase and local component)
	At Local Database level	• Local identity to issue access to the local component • Unique global identity to decide access to user • Remote identity to allow access to local component

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Administration of Authorization
    - In a multidatabase system, there are objects that belong to only component databases, objects created at the federation level, and objects imported from the component databases to the federation.
    - For objects created directly by the global component or that belong only to component databases, classical administrative policies developed for centralized system can be applied.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Administration of Authorization
    - Managing authorization for imported objects is more complex. In practice, three approaches are often considered:
      - Delegating the administration of the objects to the federation administrator.
      - Leaving the privilege of specifying authorizations to the administrator of the local object.
      - Allowing both the federation administrator and the local administrator to specify authorizations.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control

Type	Description
Full Authorization	•Local access decisions are based on local user identities. •Is cumbersome since the user has to enter a number of passwords to authorize a request made on multiple local systems.
Medium authorization	•Some trust (e.g. global authorization server) is placed on the global system that correctly validates the authority of users. •The global system authorizes a request and associates an identifier with every component request obtained by decomposing the global request. •The identifier is then used by local systems to make access control decisions. •The global and local systems cooperate to provide authorizations of requests.
Low authorization	•The global system alone authorizes requests for data from external users and the component requests are directly executed at local systems. •Is very convenient if the system is to contain only global users. •Requires considerable amount of trust in security mechanisms of the global system.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Authorization Specification
    - When forming a multidatabase system, one important decision to make is how to specify authorization rules at different levels (global level and local level) and how to resolve the conflicts.
    - Three basic approaches Can be used:
      - Independent Approach
      - Top-down Derivation
      - Bottom-up Derivation

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Independent Approach
    - Authorization specified at the global level and the local level are independent from each other.
    - The federation administrator and the local administrator specify their rules independently on federated data objects and local objects, respectively. However, the two administrators must cooperate in order to avoid inconsistent specifications.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Top-down Approach
    - The global administrator specifies the rules for a user to access global objects.
    - Local data objects involved in the authorization are determined and the authorization requests are derived from the global authorization rules.
    - If there is an inconsistency between the authorization specification defined at the local level and the derived authorization request, the authorization is rejected at the global level.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Bottom-up Approach
    - In this approach, when an object is imported into the federation, its global authorizations are derived from its local authorizations.
    - Authorizations defined for an object may differ in different local databases. When conflicts arise, no global authorization can be derived for that object.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Access Control Policy Heterogeneity
    - Access control policy heterogeneity refers to different local sites enforcing different access control policies.
    - A local DBMS may adopt one of the variations of the access control polices: mandatory access control, discretionary access control, or role-based access control policies.
    - Heterogeneity may arise even if all sites enforce the same type of policy.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Access Control Policy Heterogeneity

Policy	Heterogeneity May Occur If Different Sites:
Mandatory	• Use a different granularity of classification • Refer to different classification lattices • Give different meanings to the same security levels
Discretionary	• Allow different types of authorizations to be specified (i.e. one site may enforce a closed policy while another site enforces an open policy)

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Multidatabase Access Control Models
    - Wang and Spooner proposed an approach to enforce content-dependent access control in a heterogeneous federated system:
      - A user must register at every local site being accessed.
      - Authorizations can be specified at both local and global levels through view materialization.
      - Administration of authorization is ownership-based.
      - Local autonomy is preserved by giving local administrator the rights to decide whether a view materialization request from the global level should be granted.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Multidatabase Access Control Models
    - In Mermaid authorizations are specified both at the global level and the local level, but the access control decision is always made locally.
    - Jonscher and Dittrich proposed a model allowing authorization to be specified at both the global and local levels. Global security administrator specifies the local identities corresponding to each global identifier. A global authorization is generated only if all corresponding local authorizations can be granted.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Multidatabase Access Control Models
    - Blaustein et al. introduced the concept of agreement into control access in federated database systems.
    - Agreements are rules regulating the access to the cooperating database systems by users connected from the different sites. Two kinds of agreements are considered:
      - *Action agreements* describe the action to be taken in response to database requests,
      - *Access agreements* allow enforcing exceptions to prohibitions otherwise in effect.
    - The identity of users at the remote site from which they submit the request is used in access control.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control
  - Multidatabase Access Control Models
    - Vimercati and Samarati proposed an access control model where both the federation and local sites are involved.
    - The federation sends each site storing a local object involved in the transaction an access request for the groups to which the user belongs and the remote identity of the user.
    - The user will need to re-authenticate himself/herself at the local site.
    - Each local site will check the local authorizations and grant or deny. In particular, in the case of site-retained or cooperative policy, access will be granted if an authorization exists for the access and no negative authorization exists.
    - At the federation level, access will be granted if no negative authorization exists – Global access is granted if all local sites accept the local requests; it is denied otherwise.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Access Control

Policy	Advantages	Disadvantages
Wang and Spooner	• Allows local systems to preserve authorization autonomy	• Authorizations can be specified only for users • A user must be registered at any local system needed to access
Mermaid	• Preserves authorization autonomy • Supports different degrees of authentication autonomy	• Does not support decentralized authorization at the global level • Users must be registered with Mermaid and local systems • Access control is based on Access Control Lists
Jonscher and Dittrich	• Supports different degrees of authentication autonomy • Decentralized administration • Preserves local autonomy	• Does not allow local systems to share their objects with reference to specific privileges • Authorization must be granted by both the global owner and the local administrator
Blaustein et. al.	• Very flexible • Local autonomy not restricted	• Heavy burden on users responsible for negotiation at each site
Vimercati and Samarati	• Both federation and local sites are involved • Local systems do not need to keep track of identifiers for each single user of the federation	• Mapping a global authorization into a set of local authorizations can be difficult.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Inferential Security
  - Inferential security is defined as the security breach when the user may use logical reasoning to infer a supposedly restricted piece of information.
  - The database management system must take all steps necessary to insure that the user *u* cannot infer any item in a pre-designated set *S(u)* of items that are to be kept secret.
  - Thus, it is possible for external users to infer information in the information repository even when enough access rights are available. This is a security breach and needs to be addressed as early as possible when a query is submitted.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Inferential Security
  - For example, James may/may not have a criminal record. This record has restricted access if it exists and information cannot be publicly available.
  - An unauthorized person wants to determine if James has a criminal record and would be happy with a “Yes/No” answer.
  - The user could query the database for non-availability of James’ information. The query result would return “True/False” instead of failing since information about James is not being accessed. Thus, the unauthorized user can infer that James has/has not a criminal record even though the details were not retrieved and James’ record was not accessed.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Integrity Issues
  - Additional inter-database integrity constraints could be required after integration of component databases in a federated system.
  - However, site autonomy may allow local operations to violate the global inter-database integrity constraints.
  - The global integrity constraints may be violated because local operations are performed outside the control of the federation layer – enforcement of global integrity constraints does not necessarily always guarantee that the database is consistent or that no integrity violations occur.
  - A federated database is consistent if all global integrity constraints as well as all integrated local integrity constraints hold.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Integrity Issues
  - Global integrity constraints are sub-divided into:
    - Global key constraints,
    - Global referential integrity constraints, and
    - Global aggregate constraints.
  - Enforcing these integrity constraints results in a certain reduction of the local autonomy

Heterogeneous Distributed Databases

- MultiDatabase Systems – Integrity Issues

Global Constraint	Description
Global key constraints	• Ensures the uniqueness of an object of the federated schema
Global referential integrity constraints	• Used to describe relationships between two objects from different local databases
Global aggregate constraints	• Used to model constraints on multiple objects in different local databases

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - Local autonomy requirement of multidatabases dictates preservation of authorization models of local databases.
  - Heterogeneity of multidatabases makes the task of enforcing a single global authorization model quite a challenge.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - One approach motivates a bottom up process in deriving a global authorization model from underlying local authorizations of local databases.
  - Authorizations can be derived for integrated or imported objects based on the similarity between subjects. However, subjects among local databases are unlikely to be compatible and may have different and conflicting access authorizations to the same object. As a result, no global authorization can be derived for those subjects.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - As an alternative, one can adopt a top down approach, in which the global authorization is propagated to local databases and enforced when local data accesses are requested. However, due to the local autonomy local databases may accept or reject any dictated global authorization.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - We propose an authorization model for the SSM based on role-based access control (RBAC).
  - The motivation is to define a global authorization model that not only is independent of local authorizations, but also inherits common entities which individual local authorization is mapped onto without changing local authorizations.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - Due to the semantics of the hierarchical structure of the SSM, when more general access terms are formed at a higher level, less degree of authorization is required for accessing those terms.
  - Hence, a global authorization model should be expressed in a hierarchical form.
  - In RBAC, roles can form a role hierarchy and may suite the hierarchical structure of the SSM.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - We consider a role as a common representative for users or subjects in local databases since a role represents a job function defined by an organization that accommodates local databases and the MDBS.
  - Each local database may map some of its local subjects to a global role defined at the MDBS level. In addition, no local subject identification is maintained at the global level.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - The proposed authorization model specifies subjects and objects both at local databases and at MDBS level.
    - At local databases, there are local subjects and local objects. Local objects are objects created and maintained at local databases.
    - Each local subject can access its own local objects according to access control rules defined locally and independently.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - At MDBS level, access terms in the SSM hierarchy are global objects. Since global subjects are allowed to access objects across multiple local databases, it is natural to assume that only a subset of local subjects is allowed to be global subjects. The local databases are responsible for mapping their local subjects to corresponding global roles. A local database may maintain a table that keeps track of which subject is mapped to which global role.
  - If a new role is added or an existing role is deleted, all local databases will be informed and their local subjects can be remapped. When a user logs in at any node, the authentication can be done at a local database where a user has an account. Hence, no global authentication is needed.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - There are two major motivations in enhancing the SSM model by an authorization policy and, hence, limiting accesses to access terms in the SSM hierarchy:
    - Each term has its own degree of sensitivity and should not be accessed by unauthorized subjects. For example, “salary of an employee” in a company should not be publicly accessible while “name of an employee” may be publicly accessible,
    - Any unauthorized access detected as early as possible reduces network traffic and computation and as a result increases the query bandwidth – improving the response time of valid accesses.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - Populating global authorizations:
    - Map an individual local subject to a common role defined at MDBS level.
    - Tag SSM access terms with a set of roles that are allowed to access those terms.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM

Assume that we have two local databases. One database exports an access term *x* accessible to a role r_1 and the other database exports an access term *y* to a role r_2.
An SSM meta-data or access term is formed according to the semantic contents of *x* and *y*, and r_1 and r_2 role relationship:

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - Populating SSM meta-data

(a) If x and y are semantically different, then two SSM access terms are formed as [hypernym of x, r_1] and [hypernym of y, r_2] at the next SSM level.

(b) If x and y are semantically similar and z is the hypernym of x and y, then we consider r_1 and r_2 as follows:

(i) If r_1 and r_2 are partially ordered in the role hierarchy, an SSM term [z, minimum (r_1, r_2)] is formed at the next SSM level.

(ii) If r_1 and r_2 are not related, an SSM term [z, r_1 or r_2] is formed at the next SSM level.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM

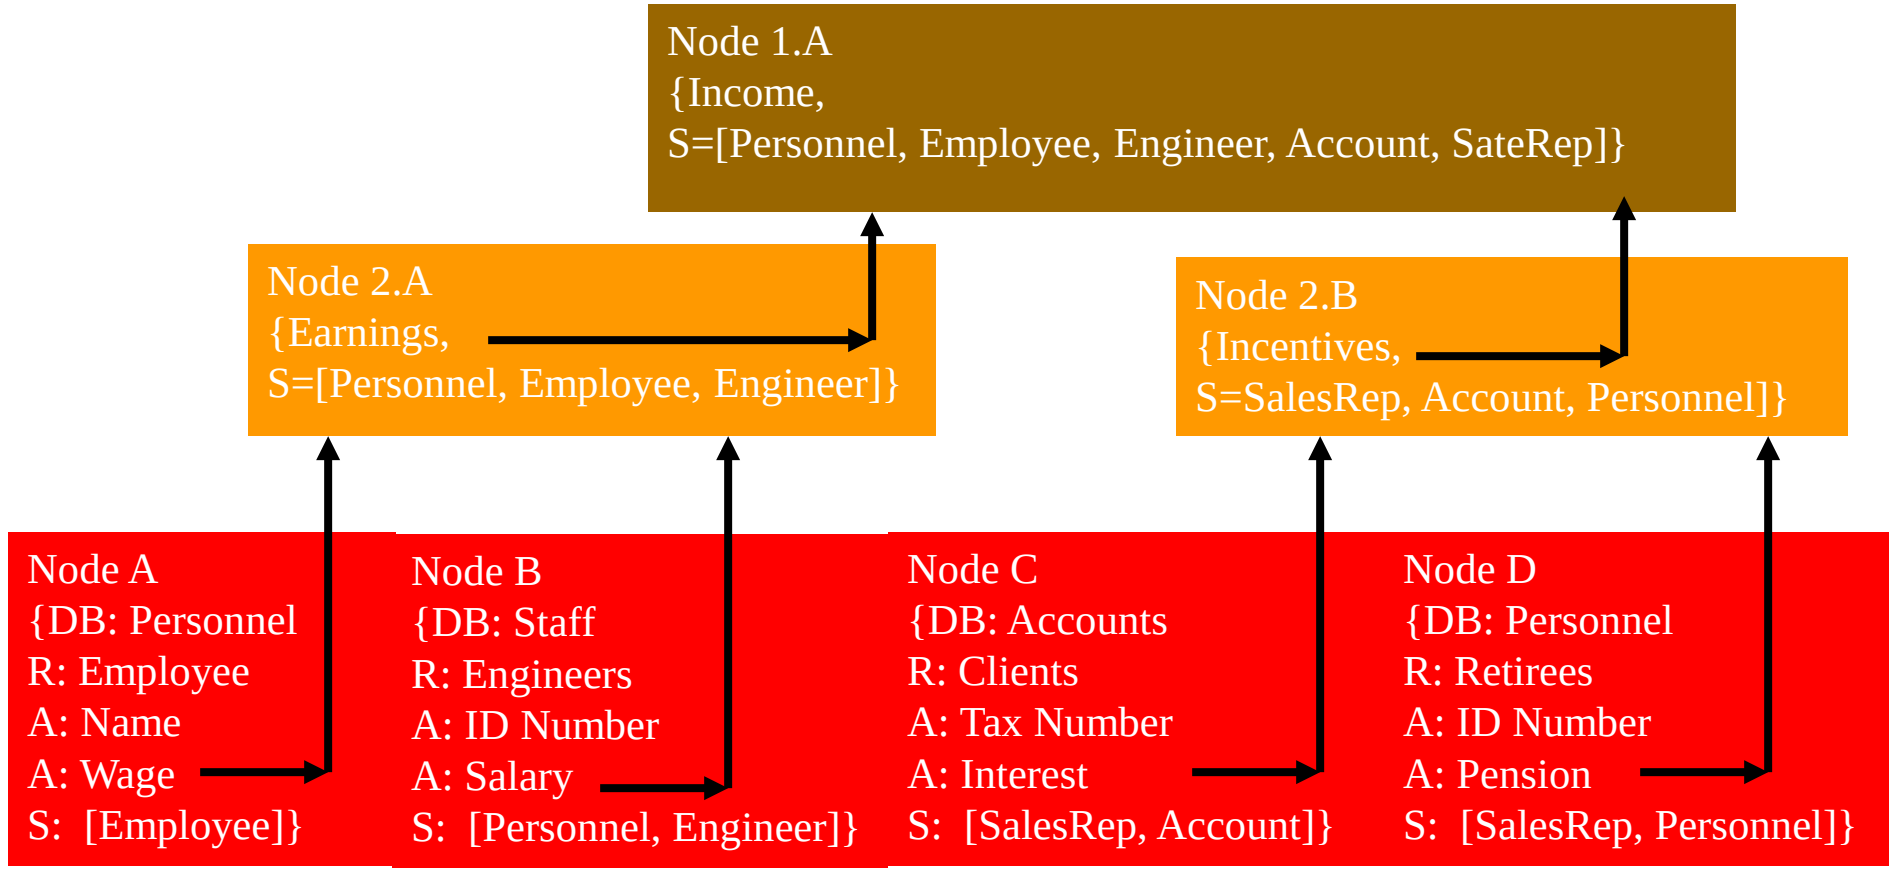

Heterogeneous Distributed Databases

- MultiDatabase Systems – Security in SSM
  - The following assumptions are made:

Assumptions:

(i) An imprecise query is submitted at any node in the SSM hierarchy ,

(ii) If one access term is rejected due to insufficient authority, the whole query is rejected.

(iii) At the query origin node, a query is parsed to identified access terms, a submitted query is also tagged with a valid global role of the user.

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authorized SSM

```
1.  FOR each imprecise term in the query Do;
2.        Compute SDM of each term,
3.        IF a match is found and accessible by the user role,
4.        THEN  IF this is a local node,
5.                  THEN replace imprecise term with the corresponding
                              precise term,
6.                  ELSE send it to lower node and continue at line 2,
7.        ELSE   IF a match is found, but inaccessible by the user role,
8.                 THEN reject the whole query,
9.                  ELSE IF this is the root of the SSM hierarchy,
10.                       THEN reject the query,
11.                       ELSE send it to higher node and continue at line 1,
```

Heterogeneous Distributed Databases

- Mu
  -

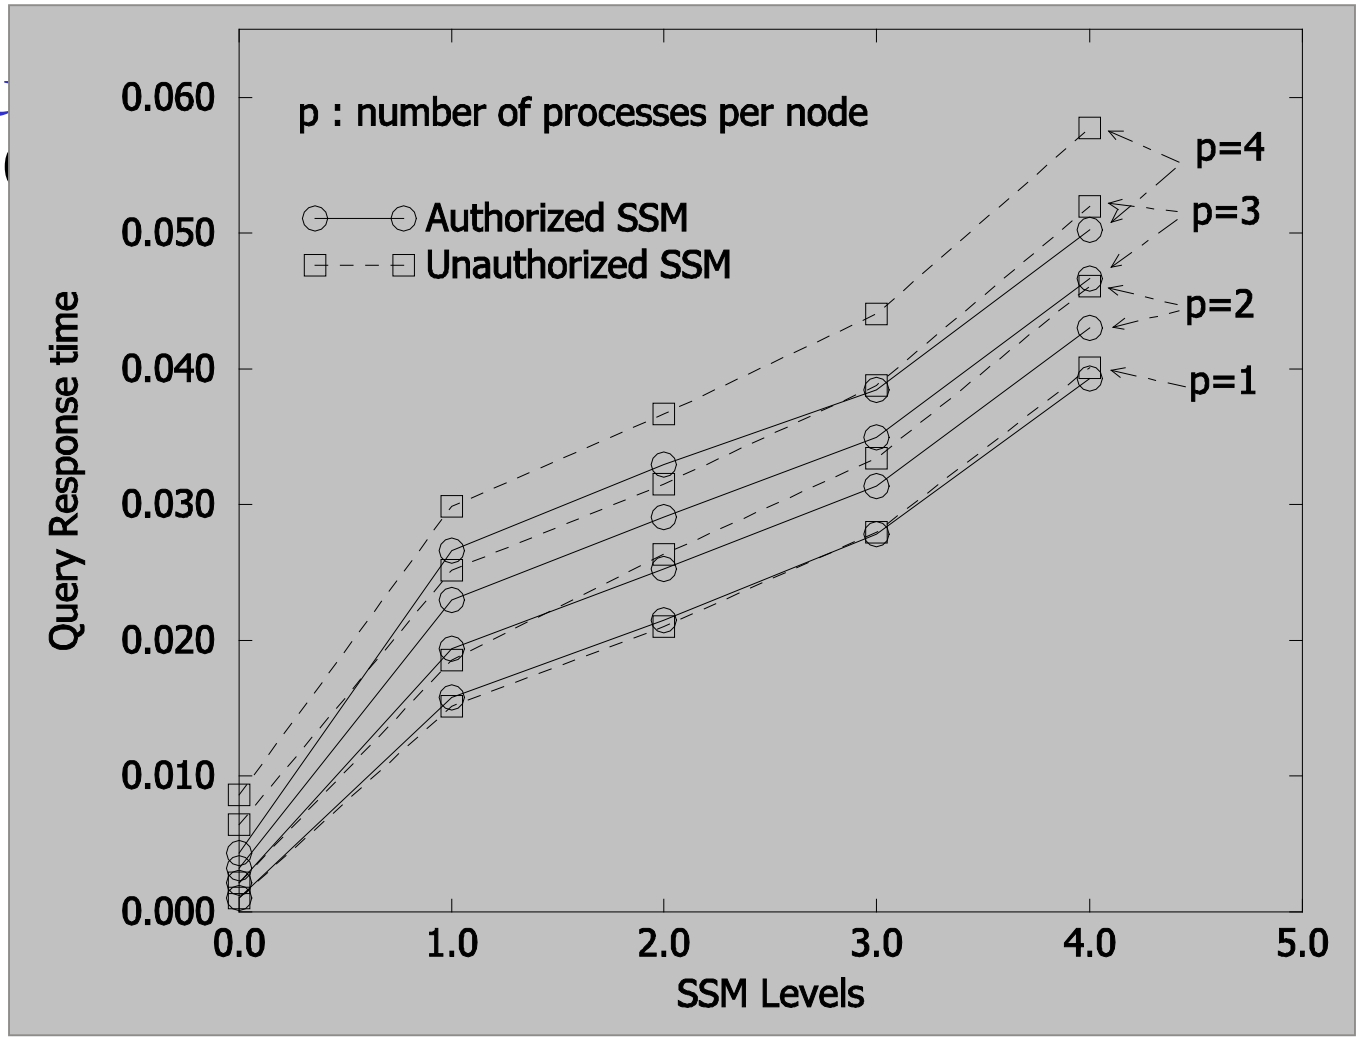

Heterogeneous Distributed Databases

- Multi
  - Re

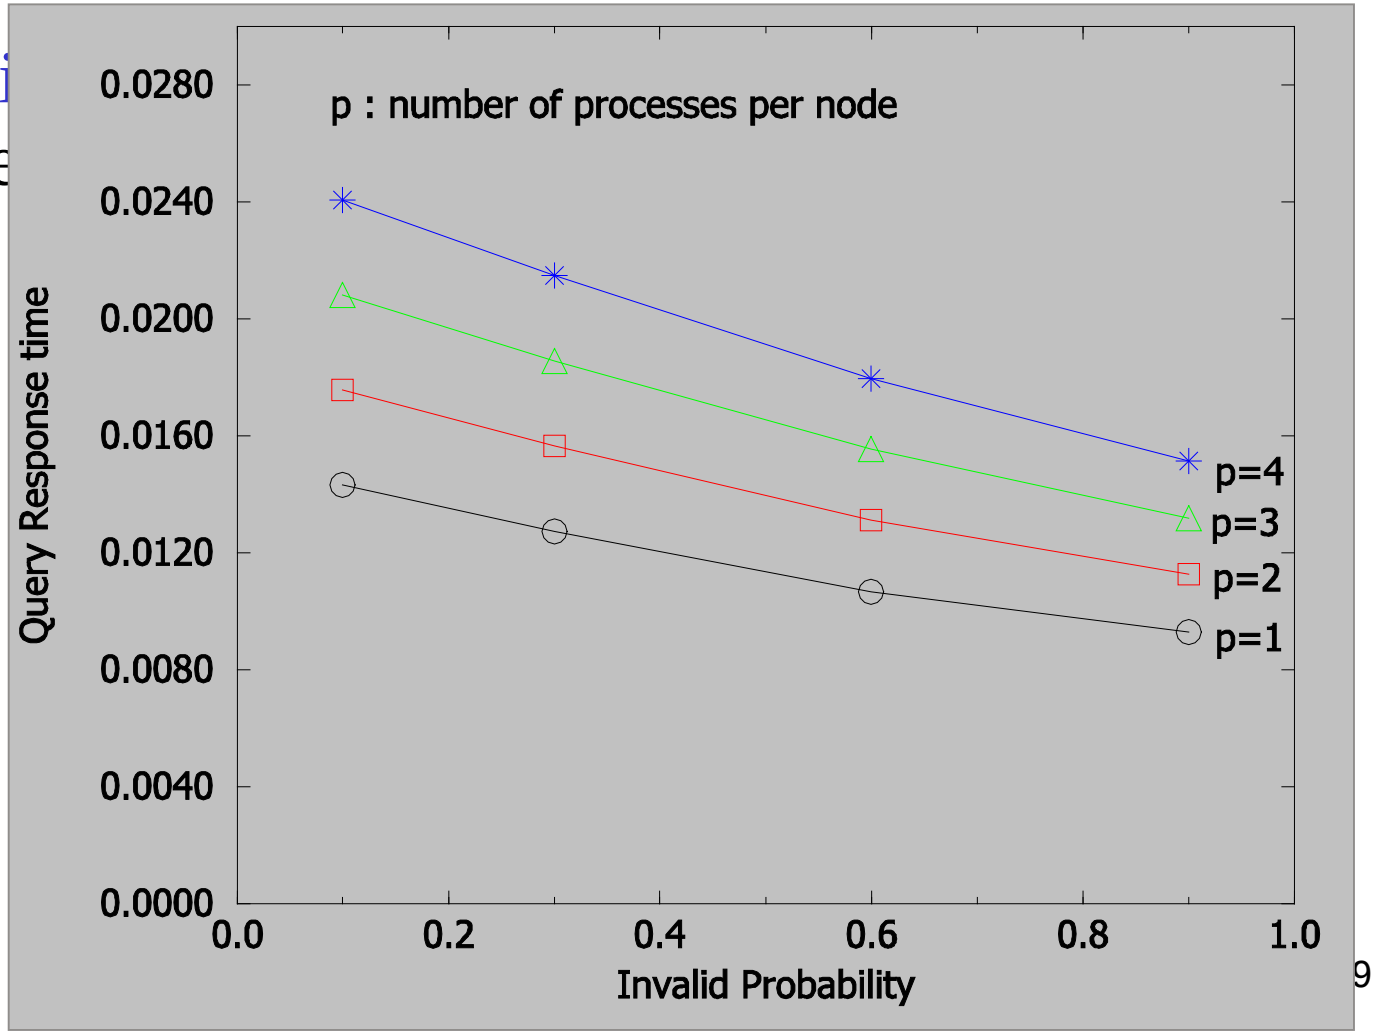

Heterogeneous Distributed Databases

- MultiDatabase Systems — Authorized SSM

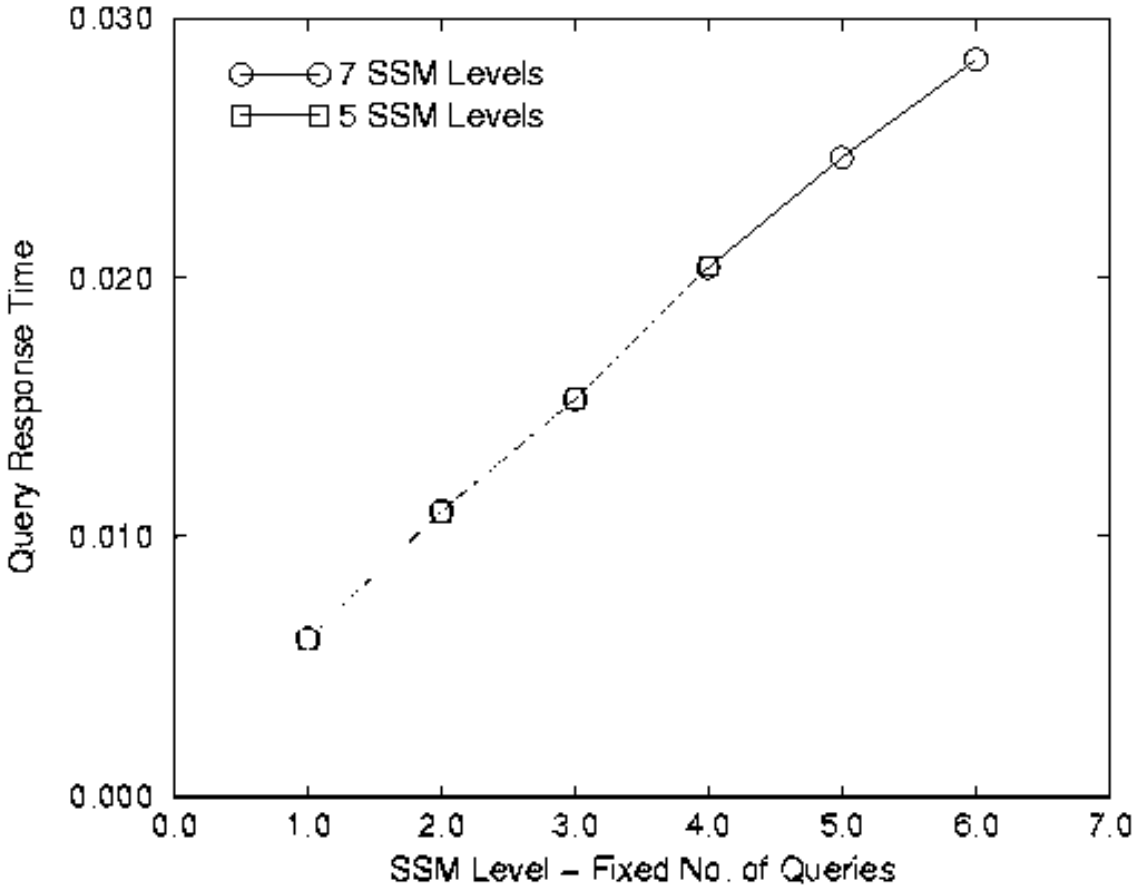

Heterogeneous Distributed Databases

- MultiDatabase Systems – Authorized SSM