

Private and Trusted Interactions*

Bharat Bhargava, Leszek Lilien, and Dongyan Xu
{bb, llilien, [dxu](mailto:dxu@cs.purdue.edu)}@cs.purdue.edu)

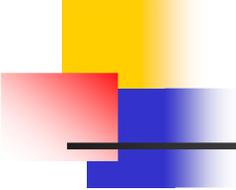
Department of Computer Sciences, CERIAS[†] and CWSA[‡]
Purdue University

in collaboration with Ph.D. students and postdocs in the Raid Lab
Computer Sciences Building, Room CS 145, phone: 765-494-6702
www.cs.purdue.edu/homes/bb

* Supported in part by NSF grants IIS-0209059, IIS-0242840, ANI-0219110, and Cisco URP grant. More grants are welcomed!

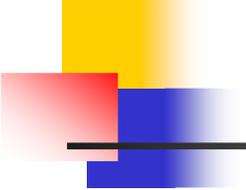
[†] Center for Education and Research in Information Assurance and Security (Executive Director: Eugene Spafford)

[‡] Center for Wireless Systems and Applications (Director: Catherine P. Rosenberg)



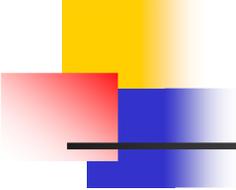
Motivation

- Sensitivity of personal data [Ackerman *et al.* '99]
 - 82% willing to reveal their favorite TV show
 - Only 1% willing to reveal their SSN
- Business losses due to privacy violations
 - Online consumers worry about revealing personal data
 - This fear held back \$15 billion in online revenue in 2001
- Federal Privacy Acts to protect privacy
 - E.g., Privacy Act of 1974 for federal agencies
 - Still many examples of privacy violations even by federal agencies
 - JetBlue Airways revealed travellers' data to federal gov't
 - E.g., Health Insurance Portability and Accountability Act of 1996 (HIPAA)



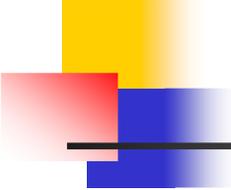
Privacy and Trust

- Privacy Problem
 - Consider computer-based interactions
 - From a simple transaction to a complex collaboration
 - Interactions involve *dissemination of private data*
 - It is voluntary, “pseudo-voluntary,” or required by law
 - Threats of privacy violations result in lower trust
 - Lower trust leads to isolation and lack of collaboration
- Trust must be established
 - Data – provide quality and integrity
 - End-to-end communication – sender authentication, message integrity
 - Network routing algorithms – deal with malicious peers, intruders, security attacks



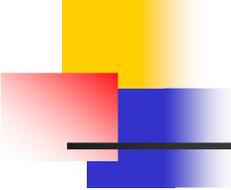
Fundamental Contributions

- Provide measures of privacy and trust
- Empower users (peers, nodes) to control privacy in ad hoc environments
 - Privacy of user identification
 - Privacy of user movement
- Provide privacy in data dissemination
 - Collaboration
 - Data warehousing
 - Location-based services
- Tradeoff between privacy and trust
 - *Minimal* privacy disclosures
 - Disclose private data absolutely necessary to gain a level of trust required by the partner system



Proposals and Publications

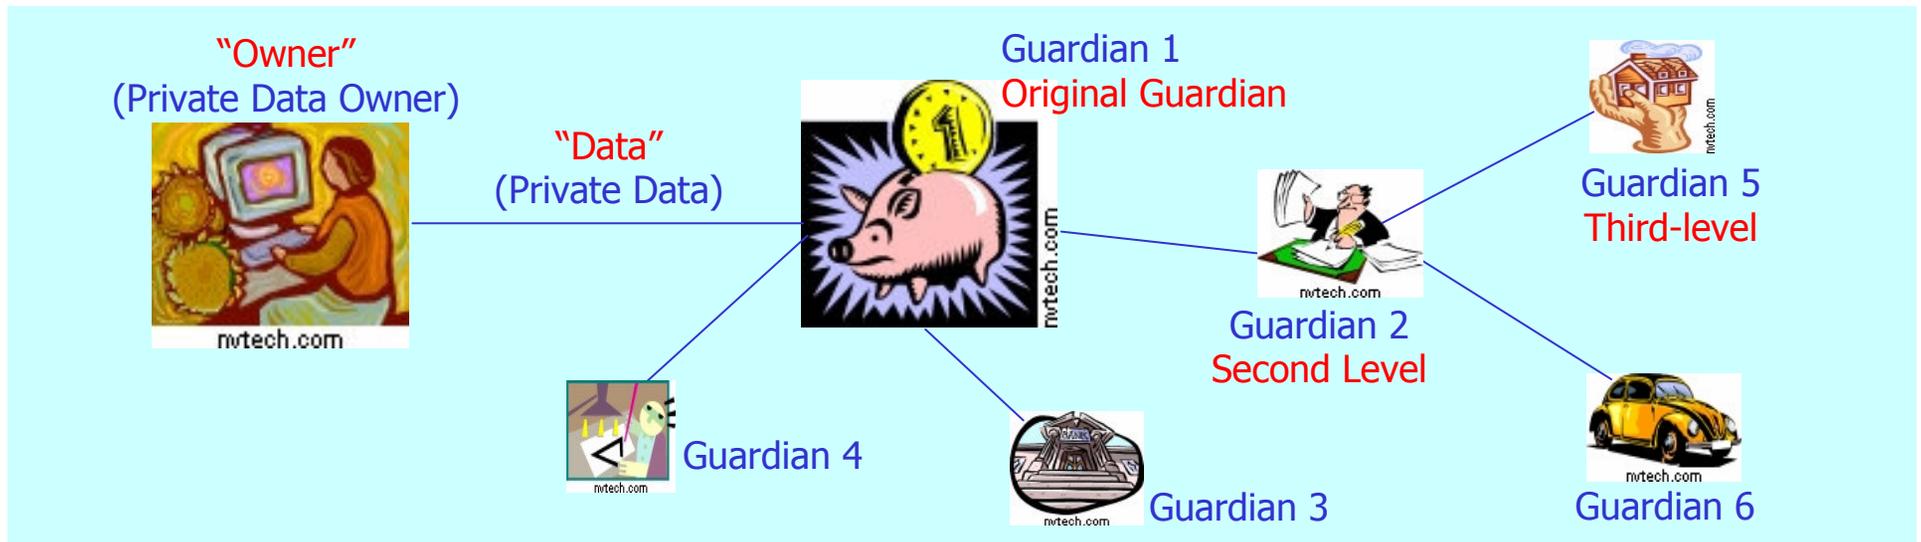
- Submitted NSF proposals
 - “Private and Trusted Interactions,” by B. Bhargava (PI) and L. Lilien (co-PI), March 2004.
 - “Quality Healthcare Through Pervasive Data Access,” by D. Xu (PI), B. Bhargava, C.-K.K. Chang, N. Li, C. Nita-Rotaru (co-PIs), March 2004.
- Selected publications
 - “On Security Study of Two Distance Vector Routing Protocols for Mobile Ad Hoc Networks,” by W. Wang, Y. Lu and B. Bhargava, Proc. of IEEE Intl. Conf. on Pervasive Computing and Communications (PerCom 2003), Dallas-Fort Worth, TX, March 2003.
<http://www.cs.purdue.edu/homes/wangwc/PerCom03wangwc.pdf>
 - “Fraud Formalization and Detection,” by B. Bhargava, Y. Zhong and Y. Lu, Proc. of 5th Intl. Conf. on Data Warehousing and Knowledge Discovery (DaWaK 2003), Prague, Czech Republic, September 2003. <http://www.cs.purdue.edu/homes/zhong/papers/fraud.pdf>
 - “Trust, Privacy, and Security. Summary of a Workshop Breakout Session at the National Science Foundation Information and Data Management (IDM) Workshop held in Seattle, Washington, September 14 - 16, 2003” by B. Bhargava, C. Farkas, L. Lilien and F. Makedon, CERIAS Tech Report 2003-34, CERIAS, Purdue University, November 2003.
<http://www2.cs.washington.edu/nsf2003> or
https://www.cerias.purdue.edu/tools_and_resources/bibtex_archive/archive/2003-34.pdf
 - “e-Notebook Middleware for Accountability and Reputation Based Trust in Distributed Data Sharing Communities,” by P. Ruth, D. Xu, B. Bhargava and F. Regnier, Proc. of the Second International Conference on Trust Management (iTrust 2004), Oxford, UK, March 2004.
<http://www.cs.purdue.edu/homes/dxu/pubs/iTrust04.pdf>
 - “Position-Based Receiver-Contention Private Communication in Wireless Ad Hoc Networks,” by X. Wu and B. Bhargava, submitted to the Tenth Annual Intl. Conf. on Mobile Computing and Networking (MobiCom’04), Philadelphia, PA, September - October 2004.
http://www.cs.purdue.edu/homes/wu/HTML/research.html/paper_purdue/mobi04.pdf



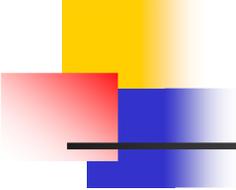
Outline

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies

1. Privacy in Data Dissemination

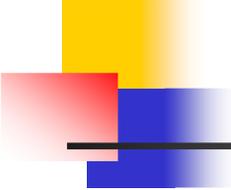


- **"Guardian:"**
Entity entrusted by private data owners with collection, storage, or transfer of their data
 - owner can be a guardian for its own private data
 - owner can be an institution or a system
- Guardians allowed or required by law to share private data
 - With owner's explicit consent
 - Without the consent as required by law
 - research, court order, etc.



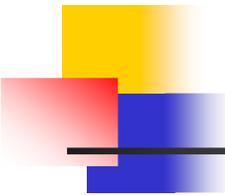
Problem of Privacy Preservation

- Guardian passes private data to another guardian in a data dissemination chain
 - Chain within a graph (possibly cyclic)
- Owner privacy preferences *not* transmitted due to neglect or failure
 - Risk grows with chain length and milieu fallibility and hostility
- If preferences lost, receiving guardian unable to honor them



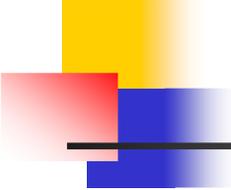
Challenges

- Ensuring that owner's metadata are never decoupled from his data
 - Metadata include owner's privacy preferences
- Efficient protection in a hostile milieu
 - Threats - examples
 - Uncontrolled data dissemination
 - Intentional or accidental data corruption, substitution, or disclosure
 - Detection of data or metadata loss
 - Efficient data and metadata recovery
 - Recovery by retransmission from the original guardian is most trustworthy



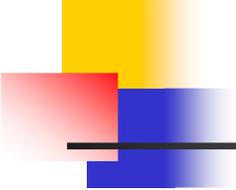
Related Work

- Self-descriptiveness
 - Many papers use the idea of self-descriptiveness in diverse contexts (meta data model, KIF, context-aware mobile infrastructure, flexible data types)
- Use of self-descriptiveness for data privacy
 - The idea briefly mentioned in [Rezgui, Bouguettaya, and Eltoweissy, 2003]
- Securing mobile self-descriptive objects
 - Esp. securing them via apoptosis, that is clean self-destruction [Tschudin, 1999]
- Specification of privacy preferences and policies
 - Platform for Privacy Preferences [Cranor, 2003]
 - AT&T Privacy Bird [AT&T, 2004]



Proposed Approach

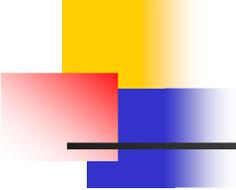
- A. Design self-descriptive private objects
- B. Construct a mechanism for apoptosis of private objects
apoptosis = clean self-destruction
- C. Develop proximity-based evaporation of private objects



A. Self-descriptive Private Objects

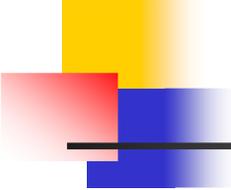
- Comprehensive metadata include:

- owner's privacy preferences How to read and write private data
- guardian privacy policies For the original and/or subsequent data guardians
- metadata access conditions How to verify and modify metadata
- enforcement specifications How to enforce preferences and policies
- data provenance Who created, read, modified, or destroyed any portion of data
- context-dependent and other components Application-dependent elements
Customer trust levels for different contexts
Other metadata elements



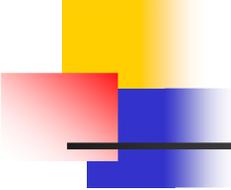
Notification in Self-descriptive Objects

- Self-descriptive objects simplify notifying owners or requesting their permissions
 - Contact information available in the *data provenance* component
- Notifications and requests sent to owners immediately, periodically, or on demand
 - Via pagers, SMSs, email, mail, etc.



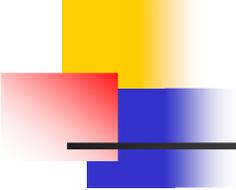
Optimization of Object Transmission

- Transmitting *complete* objects between guardians is inefficient
 - They describe all foreseeable aspects of data privacy
 - For any application and environment
- Solution: prune transmitted metadata
 - Use application and environment semantics along the data dissemination chain



B. Apoptosis of Private Objects

- Assuring privacy in data dissemination
 - In benevolent settings:
use *atomic* self-descriptive object with retransmission recovery
 - In malevolent settings:
when attacked object threatened with disclosure, use *apoptosis* (clean self-destruction)
- Implementation
 - Detectors, triggers, code
 - False positive
 - Dealt with by retransmission recovery
 - Limit repetitions to prevent denial-of-service attacks
 - False negatives

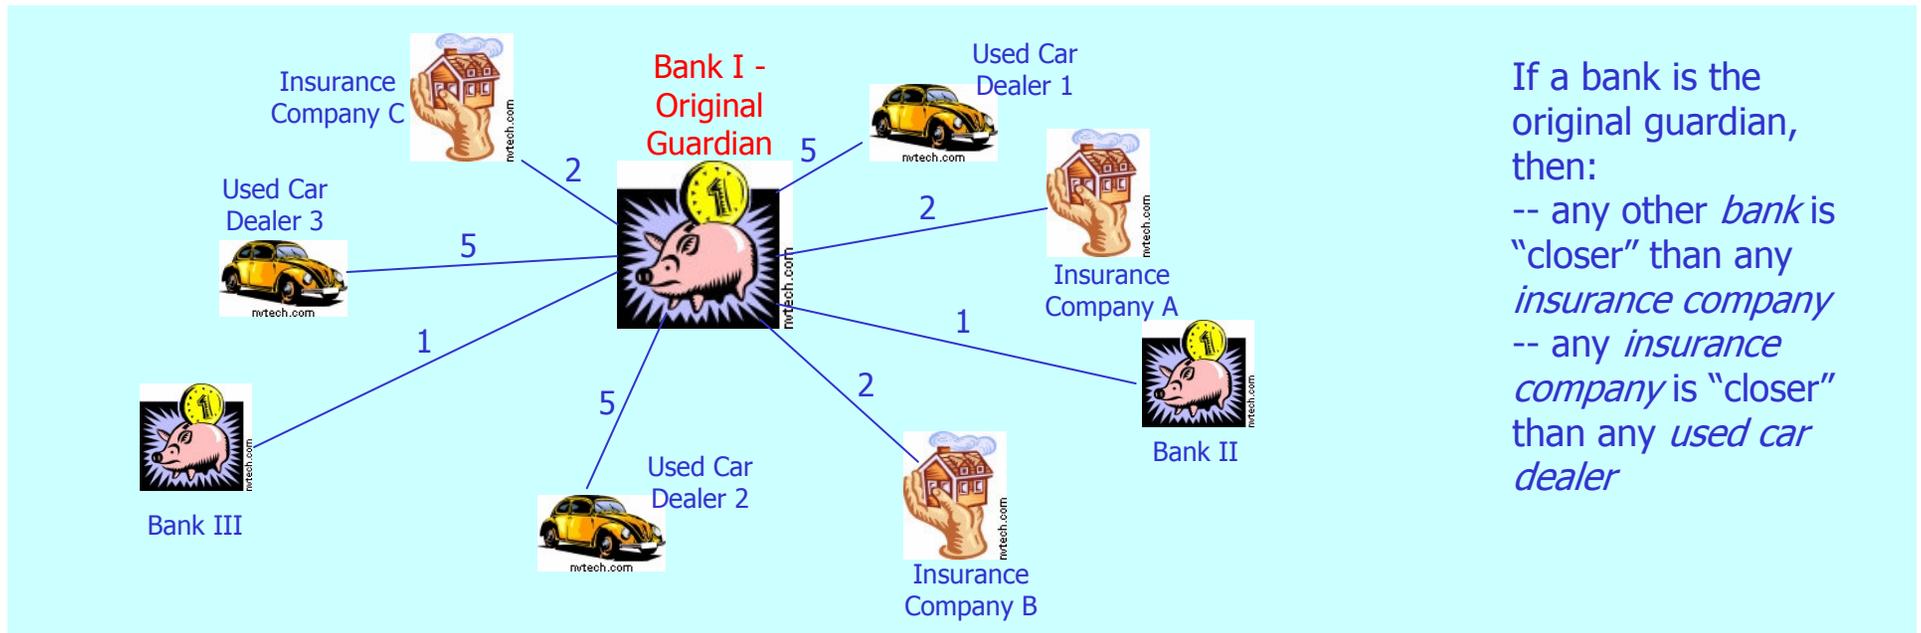


C. Proximity-based Evaporation of Private Data

- Perfect data dissemination not always desirable
 - Example: Confidential business data shared within an office but *not outside*
- Idea: Private data *evaporate* in proportion to their “distance” from their owner
 - “Closer” guardians trusted more than “distant” ones
 - Illegitimate disclosures more probable at less trusted “distant” guardians
 - Different distance metrics
 - Context-dependent

Examples of Metrics

- Examples of one-dimensional distance metrics
 - Distance \sim business type



- Distance \sim distrust level: more trusted entities are "closer"
- Multi-dimensional distance metrics
 - Security/reliability as one of dimensions

Evaporation Implemented as Controlled Data Distortion

- Distorted data reveal less, protecting privacy
- Examples:

accurate

more and more distorted

250 N. Salisbury
Street
West Lafayette, IN

Salisbury Street
West Lafayette, IN

somewhere in
West Lafayette, IN

250 N. Salisbury
Street
West Lafayette, IN
[home address]

250 N. University
Street
West Lafayette, IN
[office address]

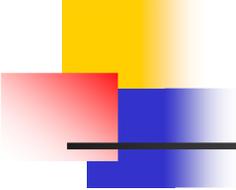
P.O. Box 1234
West Lafayette, IN
[P.O. box]

765-123-4567
[home phone]

765-987-6543
[office phone]

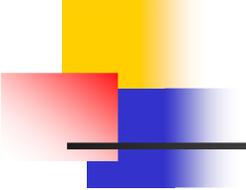
765-987-4321
[office fax]





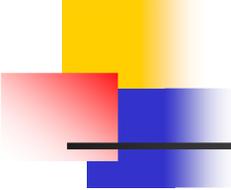
Evaporation as Apoptosis Generalization

- Context-dependent apoptosis for implementing evaporation
 - Apoptosis detectors, triggers, and code enable context exploitation
- Conventional apoptosis as a simple case of data evaporation
 - Evaporation follows a step function
 - Data self-destructs when proximity metric exceeds predefined threshold value



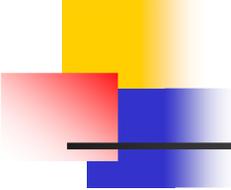
Application of Evaporation for DRM

- Evaporation used for digital rights management
 - Objects self-destruct when copied onto “foreign” media or storage device



Outline

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies



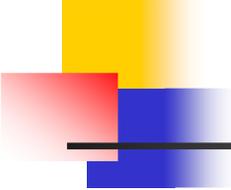
2. Privacy-trust Tradeoff

- Problem

- To build trust in open environments, users provide digital credentials that contain private information
- How to gain a certain *level of trust* with the least *loss of privacy*?

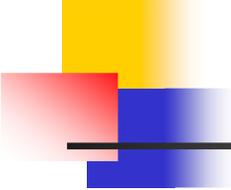
- Challenges

- Privacy and trust are fuzzy and multi-faceted concepts
- The amount of privacy lost by disclosing a piece of information is affected by:
 - Who will get this information
 - Possible uses of this information
 - Information disclosed in the past



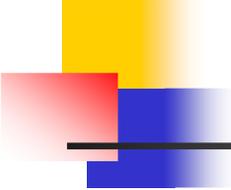
Related Work

- Automated trust negotiation (ATN) [Yu, Winslett, and Seamons, 2003]
 - Tradeoff between the length of the negotiation, the amount of information disclosed, and the computation effort
- Trust-based decision making [Wegella et al. 2003]
 - Trust lifecycle management, with considerations of both trust and risk assessments
- Trading privacy for trust [Seigneur and Jensen, 2004]
 - Privacy as the linkability of pieces of evidence to a pseudonym; measured by using *nymity* [Goldberg, thesis, 2000]



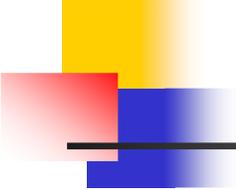
Proposed Approach

- A. Formulate the privacy-trust tradeoff problem
- B. Estimate privacy loss due to disclosing a set of credentials
- C. Estimate trust gain due to disclosing a set of credentials
- D. Develop algorithms that minimize privacy loss for required trust gain



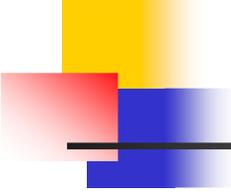
A. Formulate Tradeoff Problem

- Set of private attributes that user wants to conceal
- Set of credentials
 - Subset of *revealed* credentials R
 - Subset of *unrevealed* credentials U
- Choose a subset of credentials NC from U such that:
 - NC satisfies the requirements for trust building
 - $\text{PrivacyLoss}(NC + R) - \text{PrivacyLoss}(R)$ is minimized



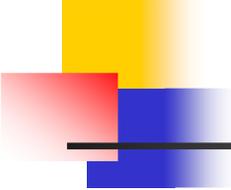
Formulate Tradeoff Problem - cont.1

- If multiple private attributes are considered:
 - Weight vector $\{w_1, w_2, \dots, w_m\}$ for private attributes
 - Privacy loss can be evaluated using:
 - The weighted sum of privacy loss for all attributes
 - The privacy loss for the attribute with the highest weight



B. Estimate Privacy Loss

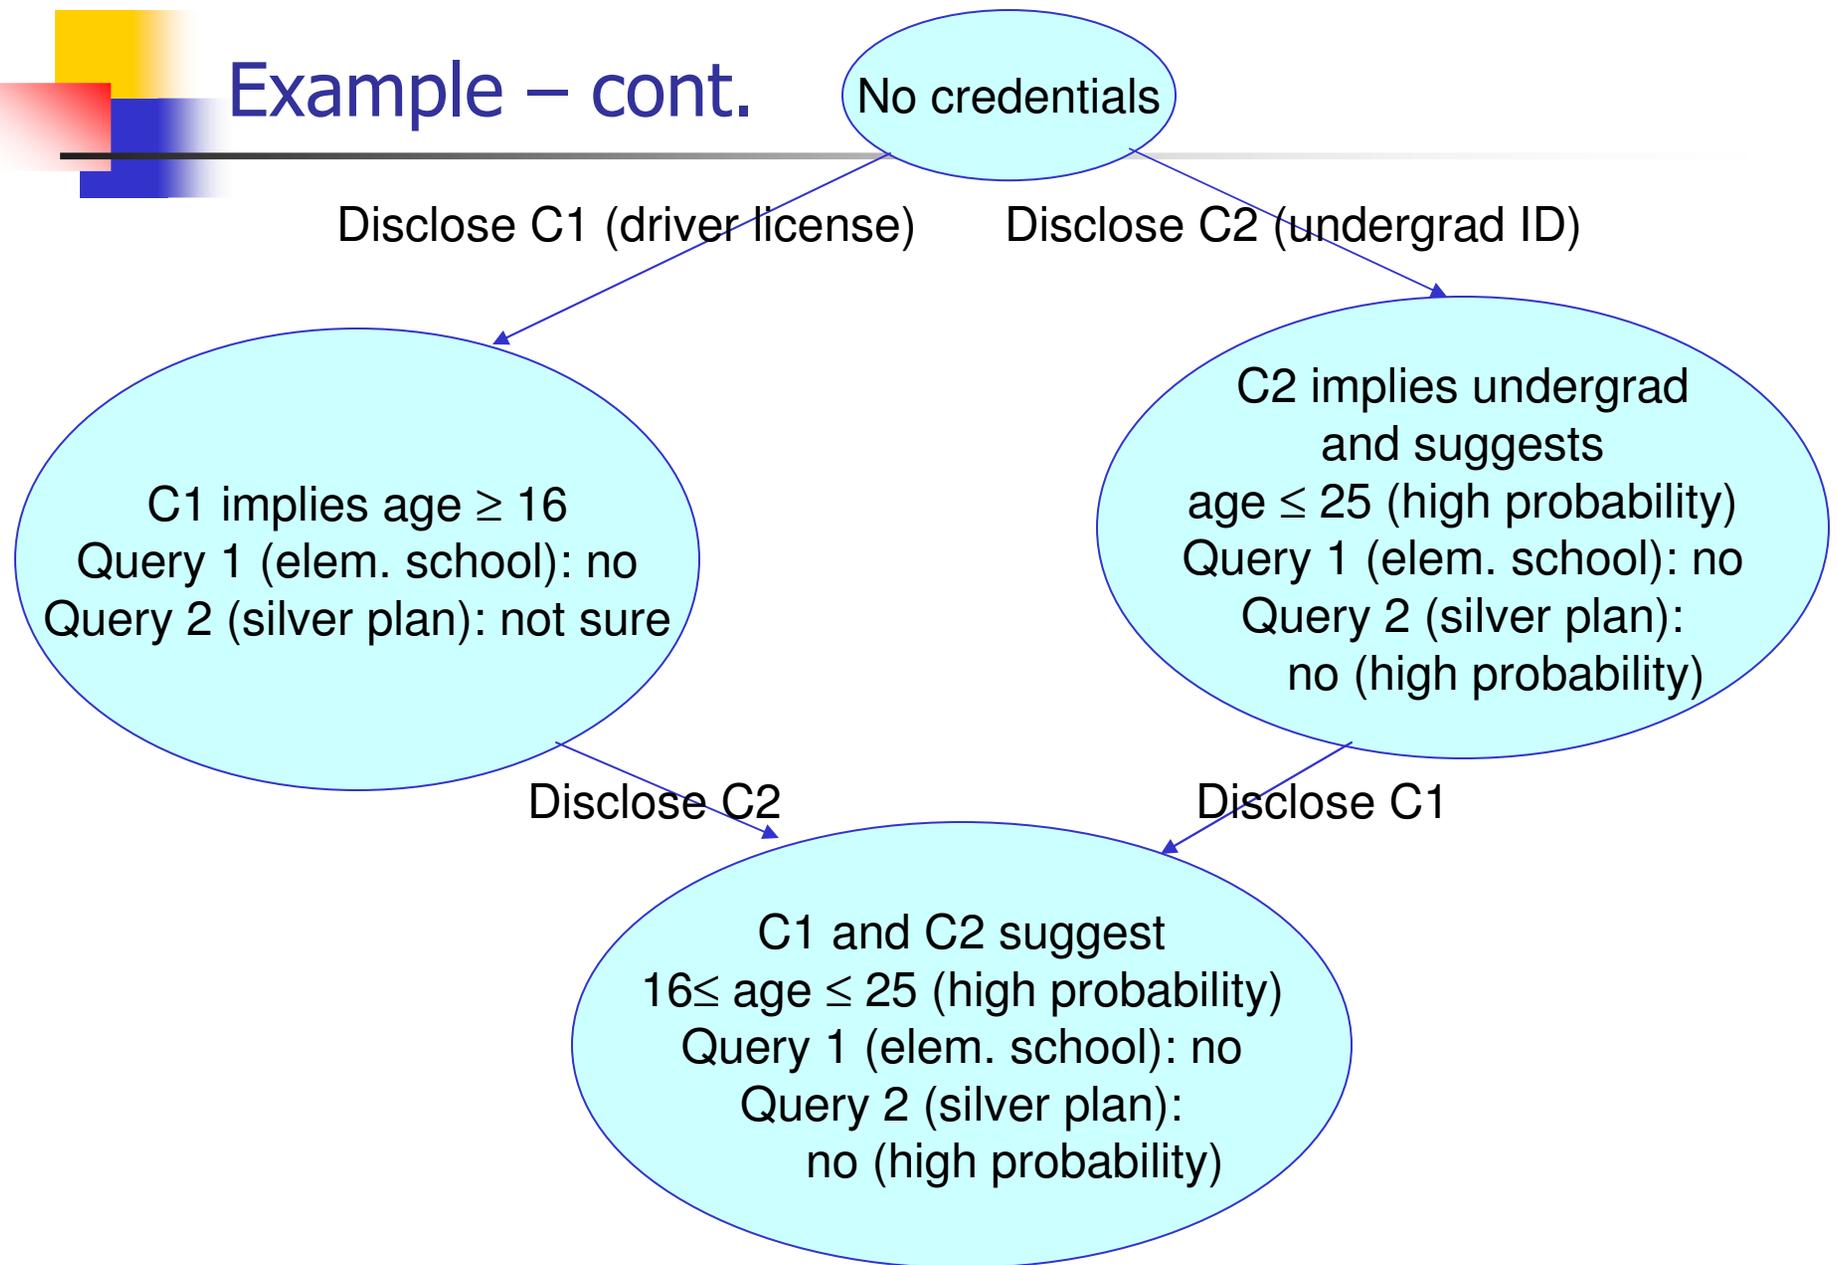
- Query-independent privacy loss
 - Provided credentials reveal the value of a private attribute
 - User determines her private attributes
- Query-dependent privacy loss
 - Provided credentials help in answering a specific query
 - User determines a set of potential queries that she is reluctant to answer

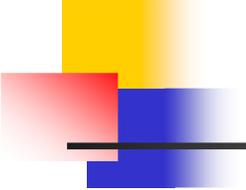


Privacy Loss Example

- Private attribute
 - age
- Potential queries:
 - (Q1) Is Alice an elementary school student?
 - (Q2) Is Alice older than 50 to join a silver insurance plan?
- Credentials
 - (C1) Driver license
 - (C2) Purdue undergraduate student ID

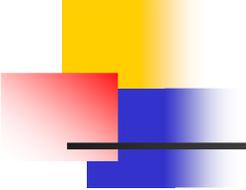
Example – cont.





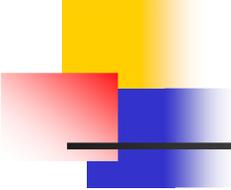
Example - Observations

- Disclose license (C1) and then undergrad ID (C2)
 - Privacy loss by disclosing *license*
 - low query-independent loss (wide range for age)
 - 100% loss for Query 1 (elem. school student)
 - low loss for Query 2 (silver plan)
 - Privacy loss by disclosing *ID* after license
 - high query-independent loss (narrow range for age)
 - zero loss for Query 1 (because privacy was lost by disclosing license)
 - high loss for Query 2 (“not sure” → “no - high probability”)
- Disclose undergrad ID (C2) and then license (C1)
 - Privacy loss by disclosing *ID*
 - low query-independent loss (wide range for age)
 - 100% loss for Query 1 (elem. school student)
 - high loss for Query 2 (silver plan)
 - Privacy loss by disclosing *license* after ID
 - high query-independent loss (narrow range of age)
 - zero loss for Query 1 (because privacy was lost by disclosing ID)
 - zero loss for Query 2



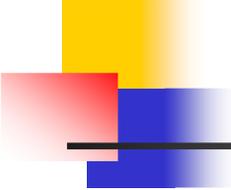
Example - Summary

- High query-independent loss does not necessarily imply high query-dependent loss
 - e.g., disclosing *ID* after *license* causes
 - high query-independent loss
 - zero loss for Query 1
- Privacy loss is affected by the order of disclosure
 - e.g., disclosing *ID* after *license* causes different privacy loss than disclosing *license* after *ID*



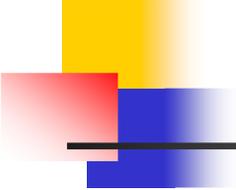
Privacy Loss Estimation Methods

- Probability method
 - Query-independent privacy loss
 - Privacy loss is measured as the difference between entropy values
 - Query-dependent privacy loss
 - Privacy loss for a query is measured as difference between entropy values
 - Total privacy loss is determined by the weighted average
 - Conditional probability is needed for entropy evaluation
 - Bayes networks and kernel density estimation will be adopted
- Lattice method
 - Estimate query-independent loss
 - Each credential is associated with a tag indicating its privacy level with respect to an attribute a_j
 - Tag set is organized as a lattice
 - Privacy loss measured as the *least upper bound* of the privacy levels for candidate credentials



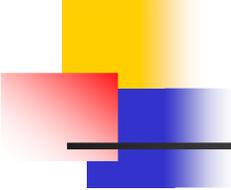
C. Estimate Trust Gain

- Increasing trust level
 - Adopt research on trust establishment and management
- Benefit function $B(\text{trust_level})$
 - Provided by service provider or derived from user's utility function
- Trust gain
 - $B(\text{trust_level}_{new}) - B(\text{trust_level}_{prev})$



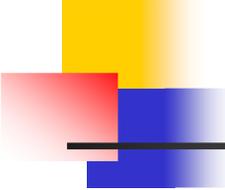
D. Minimize Privacy Loss for Required Trust Gain

- Can measure privacy loss (B) and can estimate trust gain (C)
- Develop algorithms that minimize privacy loss for required trust gain
 - User releases more private information
 - System's trust in user increases
 - How much to disclose to achieve a target trust level?



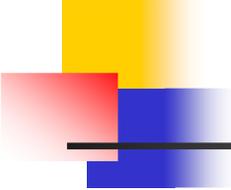
Outline

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies



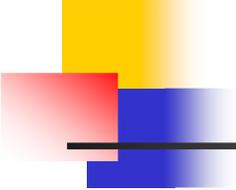
3. Privacy Metrics

- Problem
 - How to determine that certain degree of data privacy is provided?
- Challenges
 - Different privacy-preserving techniques or systems claim different degrees of data privacy
 - Metrics are usually ad hoc and customized
 - Customized for a user model
 - Customized for a specific technique/system
 - Need to develop uniform privacy metrics
 - To confidently compare different techniques/systems



Requirements for Privacy Metrics

- Privacy metrics should account for:
 - Dynamics of legitimate users
 - How users interact with the system?
E.g., repeated patterns of accessing the same data can leak information to a violator
 - Dynamics of violators
 - How much information a violator gains by watching the system for a period of time?
 - Associated costs
 - Storage, injected traffic, consumed CPU cycles, delay



Related Work

- Anonymity set without accounting for probability distribution [Reiter and Rubin, 1999]
- An entropy metric to quantify privacy level, assuming static attacker model [Diaz *et al.*, 2002]
- Differential entropy to measure how well an attacker estimates an attribute value [Agrawal and Aggarwal 2001]



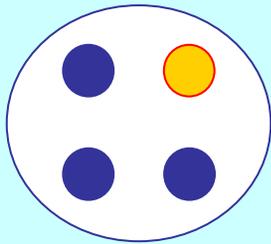
Proposed Approach

- A. Anonymity set size metrics
- B. Entropy-based metrics

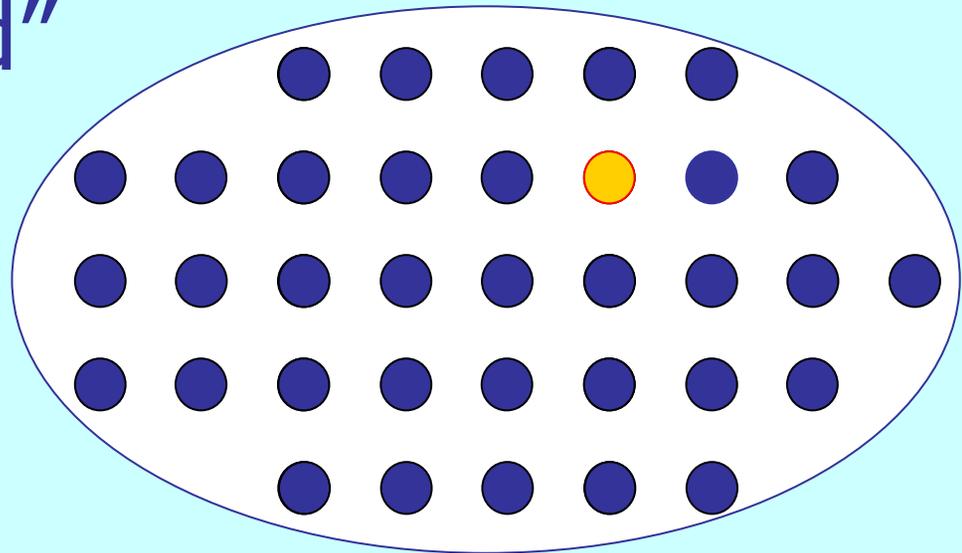
A. Anonymity Set Size Metrics

- The larger set of indistinguishable entities, the lower probability of identifying any one of them
 - Can use to "anonymize" a selected private attribute value within the domain of its all possible values

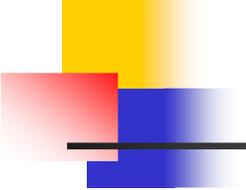
"Hiding in a crowd"



"Less" anonymous ($1/4$)



"More" anonymous ($1/n$)

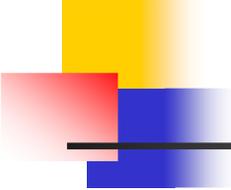


Anonymity Set

- Anonymity set A

$$A = \{(s_1, p_1), (s_2, p_2), \dots, (s_n, p_n)\}$$

- s_i : subject i who might access private data
or: i th possible value for a private data attribute
- p_i : probability that s_i accessed private data
or: probability that the attribute assumes the i th possible value

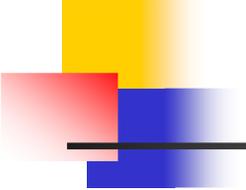


Effective Anonymity Set Size

- Effective anonymity set size is

$$L = |A| \sum_{i=1}^{|A|} \min(p_i, 1/|A|)$$

- Maximum value of L is |A| iff all p_i 's are equal to $1/|A|$
- L below maximum when distribution is skewed
 - skewed when p_i 's have different values
- Deficiency:
L does not consider violator's *learning* behavior

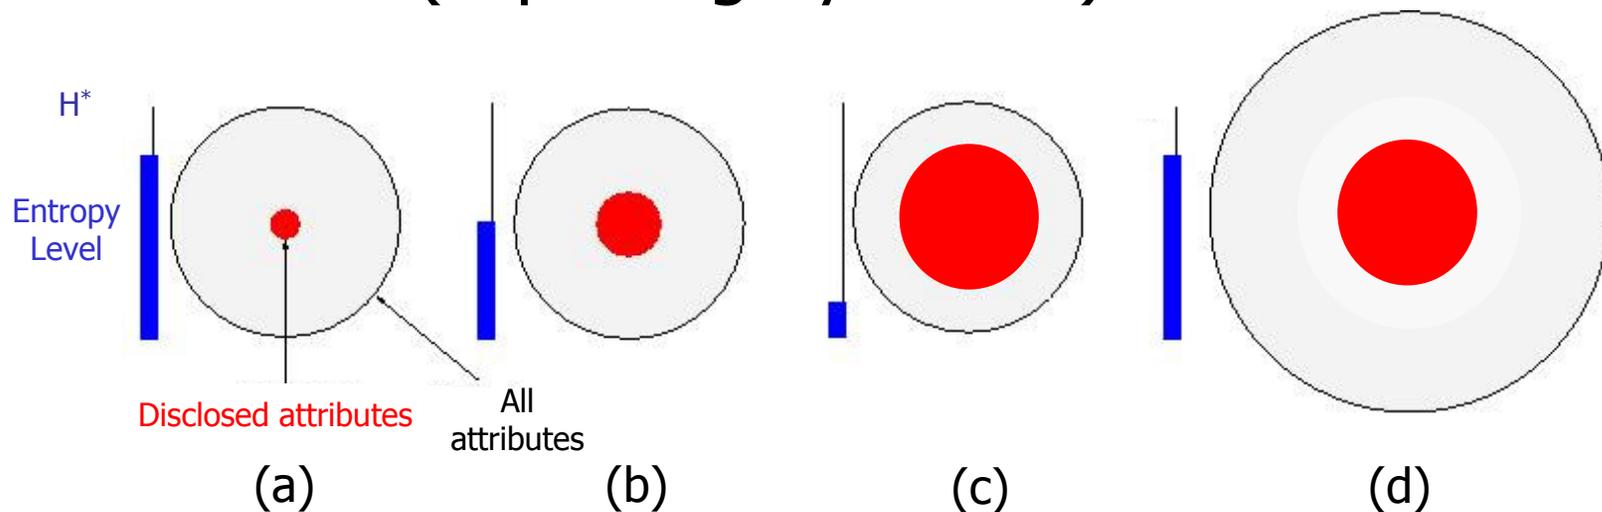


B. Entropy-based Metrics

- Entropy measures the randomness, or uncertainty, in private data
- When a violator gains more information, entropy decreases
- Metric: Compare the current entropy value with its maximum value
 - The difference shows how much information has been leaked

Dynamics of Entropy

- Decrease of system entropy with attribute disclosures (capturing dynamics)



- When entropy reaches a threshold (b), *data evaporation* can be invoked to increase entropy by controlled data distortions
- When entropy drops to a very low level (c), *apoptosis* can be triggered to destroy private data
- Entropy increases (d) if the set of attributes grows or the disclosed attributes become less valuable – e.g., obsolete or more data now available

Quantifying Privacy Loss

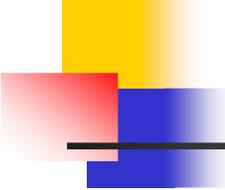
- Privacy loss $D(A,t)$ at time t , when a subset of attribute values A might have been disclosed:

$$D(A,t) = H^*(A) - H(A,t)$$

- $H^*(A)$ – the maximum entropy
 - Computed when probability distribution of p_i 's is uniform
- $H(A,t)$ is entropy at time t

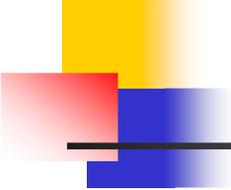
$$H(A,t) = \sum_{j=1}^{|A|} w_j \left(\sum_{\forall i} (-p_i \log_2(p_i)) \right)$$

- w_j – weights capturing relative privacy “value” of attributes



Using Entropy in Data Dissemination

- Specify two thresholds for D
 - For triggering evaporation
 - For triggering apoptosis
- When private data is exchanged
 - Entropy is recomputed and compared to the thresholds
 - Evaporation or apoptosis may be invoked to enforce privacy



Entropy: Example

- Consider a private phone number: $(a_1 a_2 a_3) a_4 a_5 a_6 - a_7 a_8 a_9 a_{10}$
- Each digit is stored as a value of a separate attribute
- Assume:
 - Range of values for each attribute is [0—9]
 - All attributes are equally important, i.e., $w_j = 1$
- The maximum entropy – when violator has no information about the value of each attribute:
 - Violator assigns a *uniform* probability distribution to values of each attribute
 - e.g., $a_j = i$ with probability of 0.10 for each i in [0—9]

$$H^*(A) = \sum_{j=0}^9 \left(w_j \sum_{i=1}^{10} (-0.1 \log_2(0.1)) \right) = 33.3$$

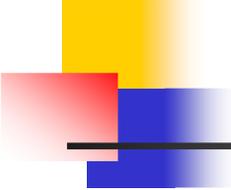
Entropy: Example – cont.

- Suppose that after time t , violator can figure out the state of the phone number, which may allow him to learn the three leftmost digits
- Entropy at time t is given by:

$$H(A, t) = 0 + \sum_{j=4}^{10} w_j \left(\sum_{i=0}^9 (-0.1 \log_2(0.1)) \right) = 23.3$$

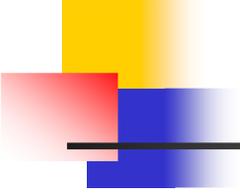
- Attributes a_1, a_2, a_3 contribute 0 to the entropy value because violator knows their correct values
- Information loss at time t is:

$$D(A, t) = H^*(A) - H(A, t) = 10.0$$



Outline

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies



4a. Application: Privacy in LBRS for Wireless Networks

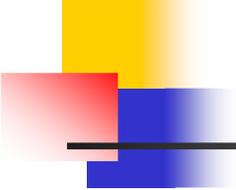
LBRS = location-based routing and services

■ Problem

- Users need and want LBRS
- LBRS users do not want their stationary or mobile *locations* widely known
- Users do not want their *movement patterns* widely known

■ Challenge

- Design mechanisms that preserve location and movement privacy while using LBRS

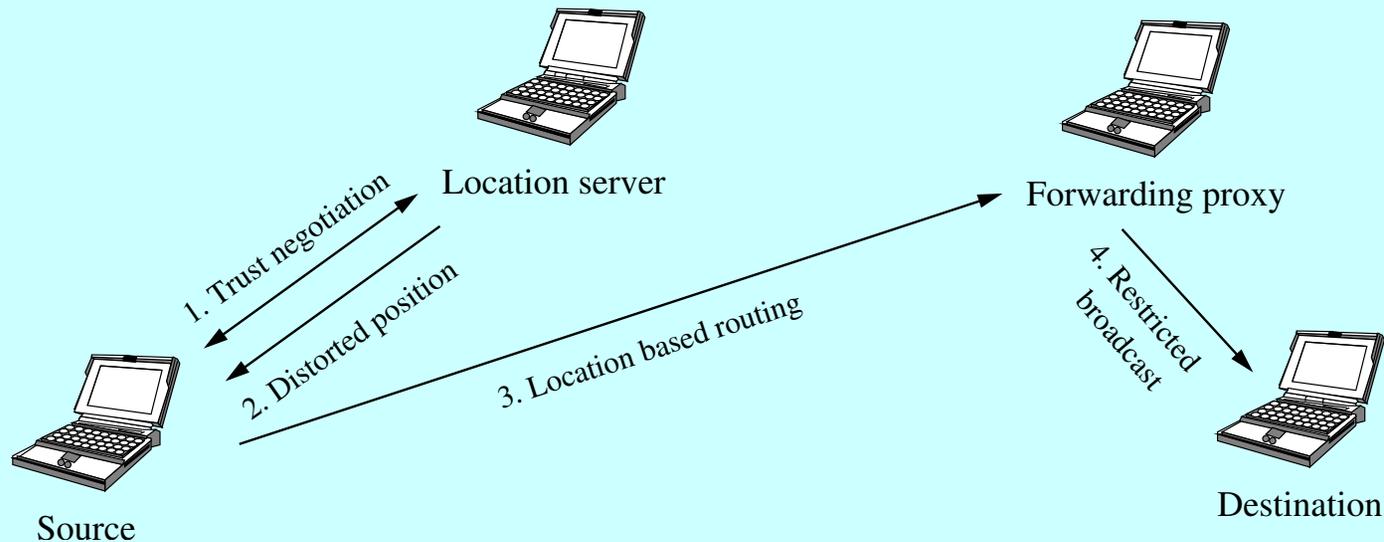


Related Work

- Range-free localization scheme using Point-in-Triangulation [He *et al.*, MobiCom'03]
- Geographic routing without exact location [Rao *et al.*, MobiCom'03]
- Localization from connectivity [Shang *et al.*, MobiHoc 03]
- Anonymity during routing in ad hoc networks [Kong *et al.*, MobiHoc'03]
- Location uncertainty in mobile networks [Wolfson *et al.*, Distributed and Parallel Databases'99]
- Querying imprecise data in mobile environments [Cheng *et al.*, TKDE'04]

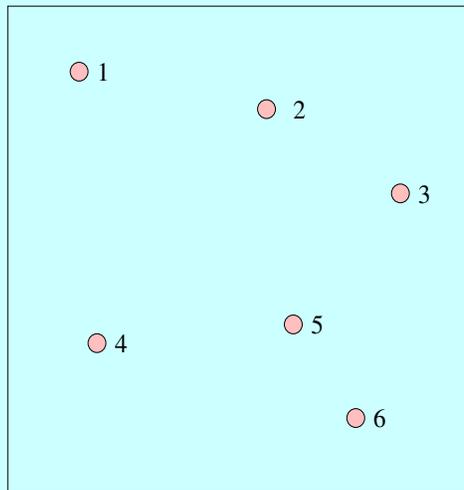
Proposed Approach: Basic Idea

- Location server distorts actual positions
 - Provide approximate position (stale or grid)
 - Accuracy of provided information is a function of the *trust level* that location server assigns to the requesting node
- Send to forwarding proxy (FP) at approximate position
Then apply restricted broadcast by FP to transmit the packet to its final destination

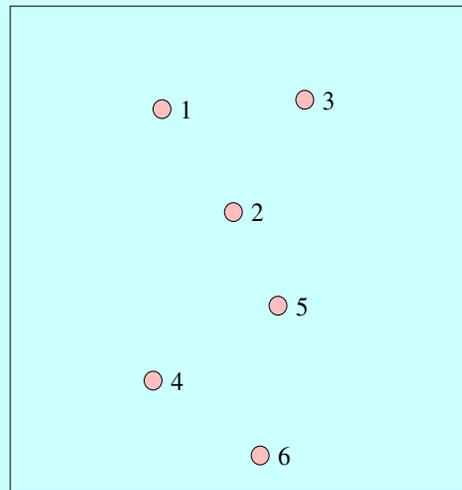


Trust and Data Distortion

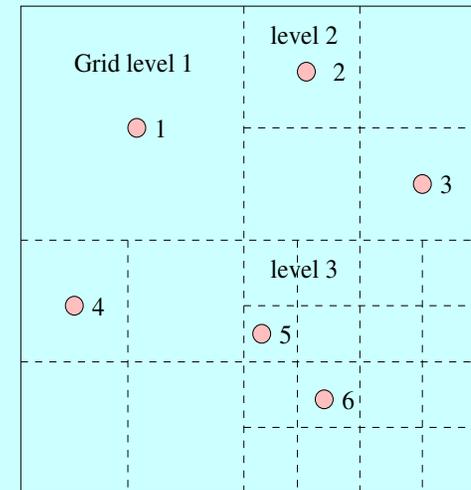
- Trust negotiation between source and location server
 - Automatic decision making to achieve tradeoff between privacy loss and network performance
- Dynamic mappings between trust level and distortion level
 - Hiding destination in an anonymity set to avoid being traced



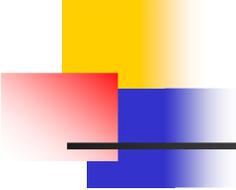
(a) Current topology



(b) Time-based distortion method: The figure shows the topology 5 minutes ago.

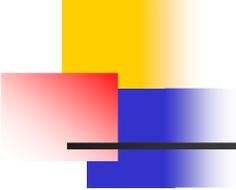


(c) Grid-based distortion method: The position is reported as the center of the grid. Three grid levels are shown in the figure.



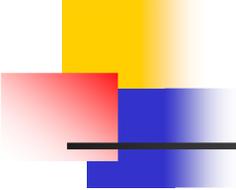
Trust Degradation and Recovery

- Identification and isolation of privacy violators
 - Dynamic trust updated according to interaction histories and peer recommendations
- Fast degradation of trust and its slow recovery
 - This defends against smart violators



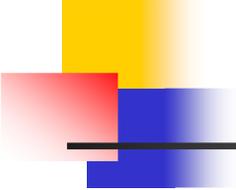
Contributions

- More secure and scalable routing protocol
- Advances in QoS control for wireless networks
- Improved mechanisms for privacy measurement and information distortion
- Advances in privacy violation detection and violator identification



Outline

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies



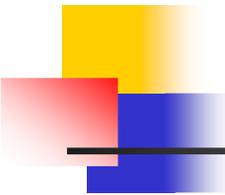
4b. Application: Privacy in e-Supply Chain Management Systems

- Problem

- Inadequacies in privacy protection for e-supply chain management system (e-SCMS) hamper their development

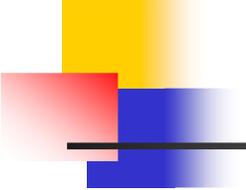
- Challenges

- Design privacy-related components for privacy-preserving e-SCMS
 - When and with whom to share private data?
 - How to control their disclosures?
 - How to accommodate and enforce privacy policies and preferences?
 - How to evaluate and compare alternative preferences and policies?



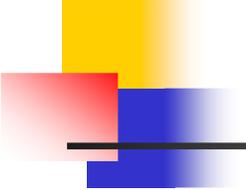
Related Work

- Coexistence and compatibility of e-privacy and e-commerce [Frosch-Wilke, 2001; Sandberg, 2002]
 - Context: electronic customer relationship management (e-CRM)
 - e-CRM includes e-SCMS
- Privacy as a major concern in online e-CRM systems for providing personalization and recommendation services [Ramakrishnan, 2001]
- Privacy-preserving personalization techniques [Ishitani *et al.*, 2003]
- Privacy preserving collaborative filtering systems [Mender project, <http://www.cs.berkeley.edu/~jfc/mender/>]
- Privacy-preserving data mining systems [Privacy, Obligations, and Rights in Technologies of Information Assessment
<http://theory.stanford.edu/~rajeev/privacy.html>]



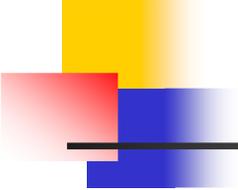
Proposed Approach

- Intelligent data sharing
 - Implementation of privacy preferences and policies at data warehouses
 - Evaluation of credentials and requester trustworthiness
 - Evaluation of cost benefits of privacy loss vs. trust gain
- Controlling misuse
 - Automatic enforcement via private objects
 - Distortion / summarization
 - Apoptosis
 - Evaporation



Proposed Approach – cont.

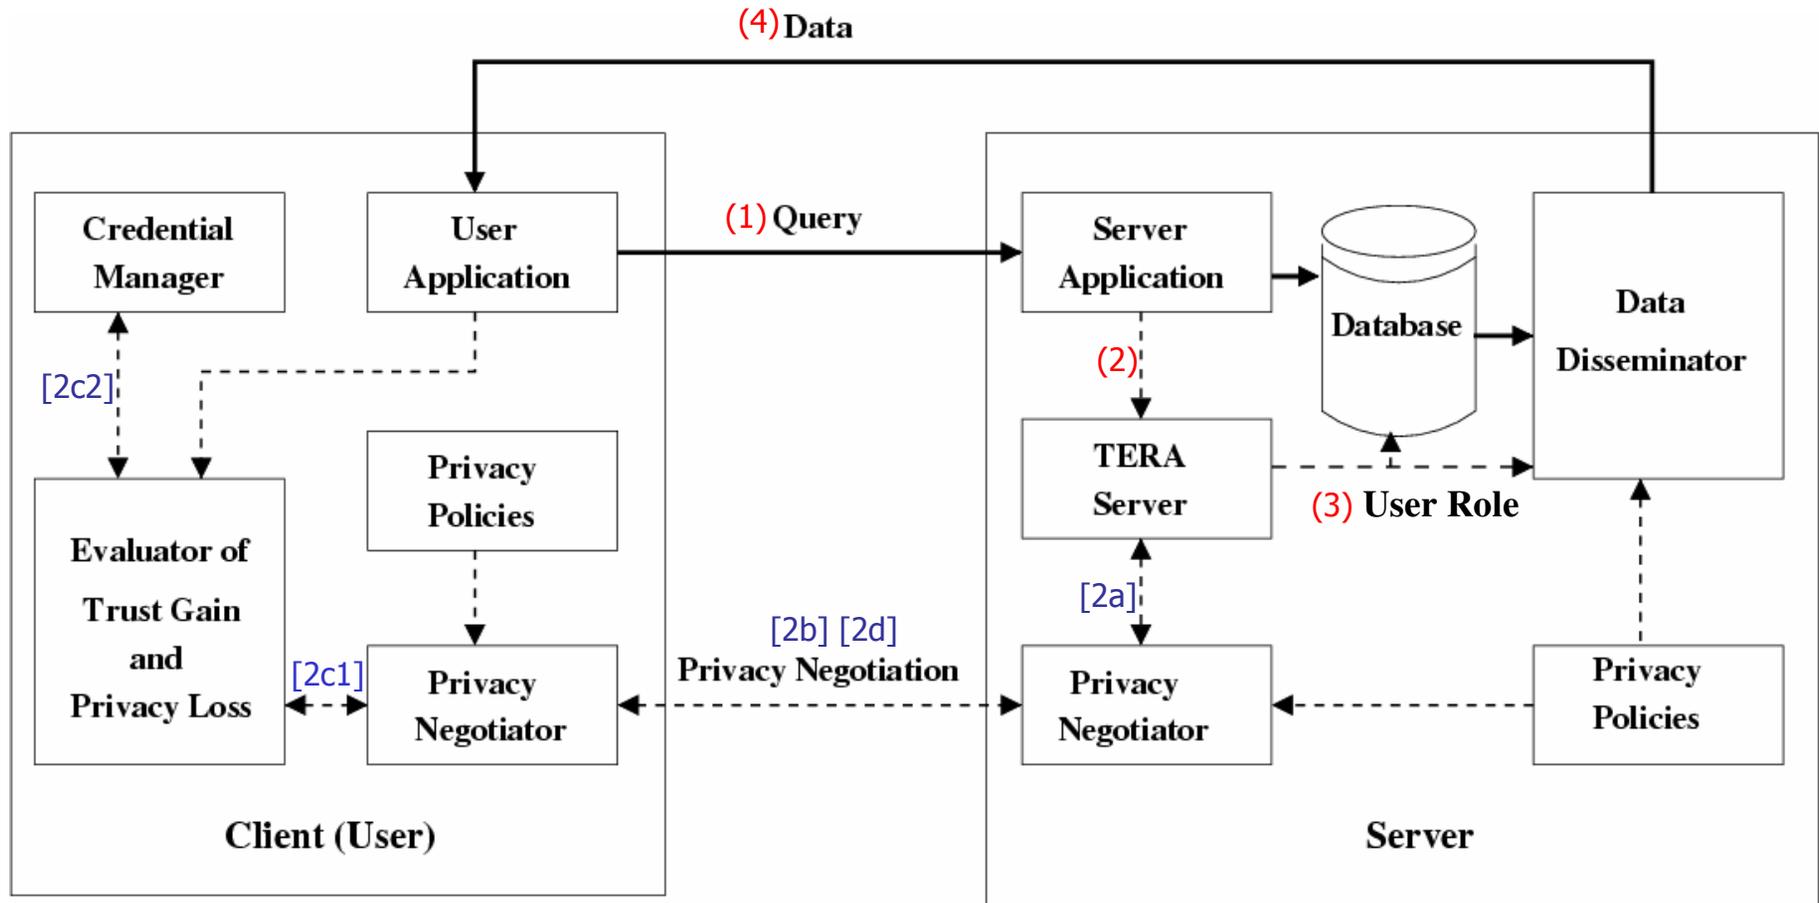
- Enforcing and integrating privacy components
 - Using privacy metrics for policy evaluation before its implementation
 - Integration of privacy-preservation components with e-SCMS software
 - Modeling and simulation of privacy-related components for e-SCMS
 - Prototyping privacy-related components for e-SCMS
 - Evaluating the effectiveness, efficiency and usability of the privacy mechanisms on PRETTY prototype
 - Devising a privacy framework for e-SCMS applications



Outline

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies

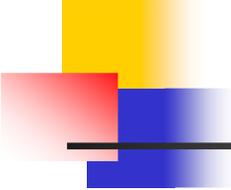
5. PRETTY Prototype for Experimental Studies



(<nr>) – unconditional path

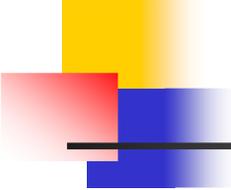
[<nr>]– conditional path

TERA = Trust-Enhanced Role Assignment



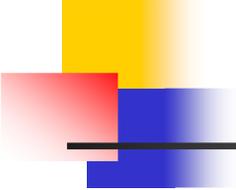
Information Flow for PRETTY

- 1) User application sends query to server application.
- 2) Server application sends user information to TERA server for trust evaluation and role assignment.
 - a) If a higher trust level is required for query, TERA server sends the request for more user's credentials to privacy negotiator.
 - b) Based on server's privacy policies and the credential requirements, privacy negotiator interacts with user's privacy negotiator to build a higher level of trust.
 - c) Trust gain and privacy loss evaluator selects credentials that will increase trust to the required level with the least privacy loss. Calculation considers credential requirements and credentials disclosed in previous interactions.
 - d) According to privacy policies and calculated privacy loss, user's privacy negotiator decides whether or not to supply credentials to the server.
- 3) Once trust level meets the minimum requirements, appropriate roles are assigned to user for execution of his query.
- 4) Based on query results, user's trust level and privacy policies, data disseminator determines: (i) whether to distort data and if so to what degree, and (ii) what privacy enforcement metadata should be associated with it.



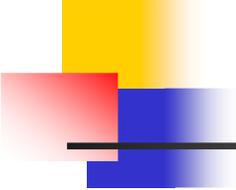
Example Experimental Studies

- Private object implementation
 - Validate and evaluate the cost, efficiency, and the impacts on the dissemination of objects
 - Study the apoptosis and evaporation mechanisms for private objects
- Tradeoff between privacy and trust
 - Study the effectiveness and efficiency of the probability-based and lattice-based privacy loss evaluation methods
 - Assess the usability of the evaluator of trust gain and privacy loss
- Location-based routing and services
 - Evaluate the dynamic mappings between trust levels and distortion levels



Private and Trusted Interactions - Summary

1. Assuring privacy in data dissemination
2. Privacy-trust tradeoff
3. Privacy metrics
4. Example applications to networks and e-commerce
 - a. Privacy in location-based routing and services in wireless networks
 - b. Privacy in e-supply chain management systems
5. Prototype for experimental studies



Bird's Eye View of Research

- Research integrates ideas from:
 - Cooperative information systems
 - Collaborations
 - Privacy, trust, and information theory
- General privacy solutions provided
- Example applications studied:
 - Location-based routing and services for wireless networks
 - Electronic supply chain management systems
- Applicability to:
 - Ad hoc networks, peer-to-peer systems
 - Diverse computer systems
 - The Semantic Web

