Privacy-Preserving WebID Analytics on the Decentralized Policy-Aware Social Web

Yuh-Jong Hu
hu@cs.nccu.edu.tw

Emerging Network Technology (ENT) Lab.
Department of Computer Science
National Chengchi University, Taipei, Taiwan

October-28-2014

Seminar at TIGP-SNHCC, Academia Sinica, Taipei, Taiwan

Table of contents

1. Introduction
2. Background
3. Semantics-enabled Policies
4. Automatic WebID Analytics Pipeline
5. Privacy-Preserving WebID Analytics
6. Conclusion and Future Work

Problem with Today's Online Social Networks

Walled Gardens

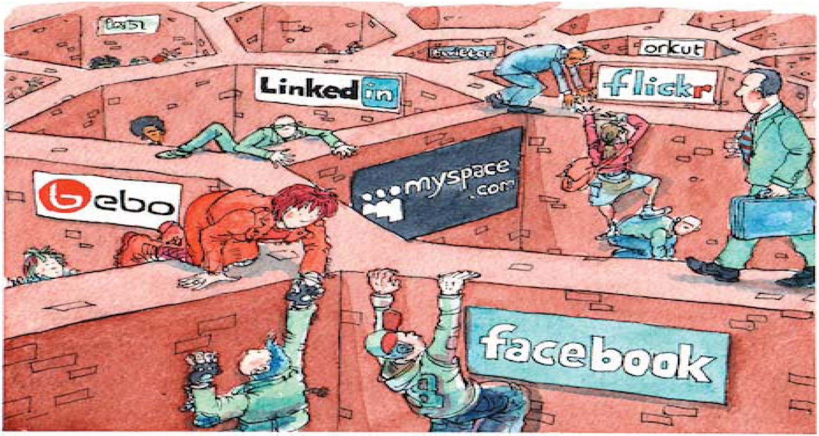

–Babbage: August-15th , 2012, The Economist Online

Problem with Today's Online Social Networks

Social Network Site Silos

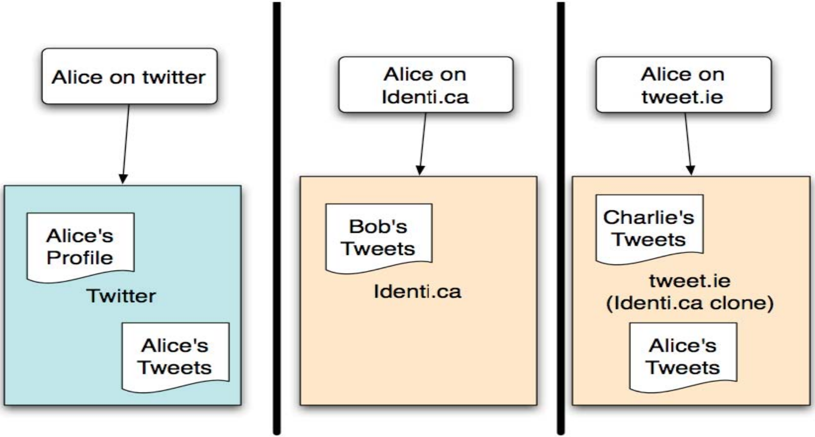

–Socially Aware Cloud Storage, Tim Berners-Lee

Problems with Centralized Online Social Networks

- Information silos on one site is not usable in the others.
- A user is “stuck”: migrating to another application is hard.
- Users cannot choose what Web applications to do with their data.
- New application must acquire a critical mass of data from scratch.
- Do not allow a user control over how his/her personal information is collected and disseminated, which results in potential privacy problems.

–Decentralization: The Future of Online Social Networking

Motivations

1. Centralized closed social networking sites are *walled gardens* that limit people activity only on a single site.
2. Big data analytics has been proposed for (centralized) online social networks, but the related privacy protection issue does not arise much attention.
3. Statistical Disclosure Control (SDC) methods have been well-developed for microdata protection, they are also possibly used for data disclosure control in online social networks.
4. Semantic Web technology, such as RDF(S), has been used for establishing a privacy-aware policy Web architecture to provide flexible and effective privacy-preserving data analytics services.

Motivations

1. Centralized closed social networking sites are *walled gardens* that limit people activity only on a single site.
2. Big data analytics has been proposed for (centralized) online social networks, but the related privacy protection issue does not arise much attention.
3. Statistical Disclosure Control (SDC) methods have been well-developed for microdata protection, they are also possibly used for data disclosure control in online social networks.
4. Semantic Web technology, such as RDF(S), has been used for establishing a privacy-aware policy Web architecture to provide flexible and effective privacy-preserving data analytics services.

Motivations

1. Centralized closed social networking sites are *walled gardens* that limit people activity only on a single site.
2. Big data analytics has been proposed for (centralized) online social networks, but the related privacy protection issue does not arise much attention.
3. Statistical Disclosure Control (SDC) methods have been well-developed for microdata protection, they are also possibly used for data disclosure control in online social networks.
4. Semantic Web technology, such as RDF(S), has been used for establishing a privacy-aware policy Web architecture to provide flexible and effective privacy-preserving data analytics services.

Motivations

1. Centralized closed social networking sites are *walled gardens* that limit people activity only on a single site.
2. Big data analytics has been proposed for (centralized) online social networks, but the related privacy protection issue does not arise much attention.
3. Statistical Disclosure Control (SDC) methods have been well-developed for microdata protection, they are also possibly used for data disclosure control in online social networks.
4. Semantic Web technology, such as RDF(S), has been used for establishing a privacy-aware policy Web architecture to provide flexible and effective privacy-preserving data analytics services.

Research Goals

1. We argue why we choose decentralized but not the centralized online social networking architecture for WebID analytics.
2. How can we provide privacy-preserving *batch* and *interactive* WebID analytics with effective and flexible data analytics services?
3. How can we call for privacy-preserving WebID analytics services through types of semantics-enabled policy enforcement?
4. How to provide an effective and flexible service platform for data analysts through integrating R+SPARQL for graph-parallel analytics and MapReduce paradigm for data-parallel analytics?

Research Goals

1. We argue why we choose decentralized but not the centralized online social networking architecture for WebID analytics.
2. How can we provide privacy-preserving *batch* and *interactive* WebID analytics with effective and flexible data analytics services?
3. How can we call for privacy-preserving WebID analytics services through types of semantics-enabled policy enforcement?
4. How to provide an effective and flexible service platform for data analysts through integrating R+SPARQL for graph-parallel analytics and MapReduce paradigm for data-parallel analytics?

Research Goals

1. We argue why we choose decentralized but not the centralized online social networking architecture for WebID analytics.
2. How can we provide privacy-preserving *batch* and *interactive* WebID analytics with effective and flexible data analytics services?
3. How can we call for privacy-preserving WebID analytics services through types of semantics-enabled policy enforcement?
4. How to provide an effective and flexible service platform for data analysts through integrating R+SPARQL for graph-parallel analytics and MapReduce paradigm for data-parallel analytics?

Research Goals

1. We argue why we choose decentralized but not the centralized online social networking architecture for WebID analytics.
2. How can we provide privacy-preserving *batch* and *interactive* WebID analytics with effective and flexible data analytics services?
3. How can we call for privacy-preserving WebID analytics services through types of semantics-enabled policy enforcement?
4. How to provide an effective and flexible service platform for data analysts through integrating R+SPARQL for graph-parallel analytics and MapReduce paradigm for data-parallel analytics?

Preliminary Contributions

1. Propose the concept of the semantic WebID analytics pipeline for automated privacy-preserving data analytics services.
2. Three types of semantics-enabled policy for access control, data handling, and data releasing, are designed and enforced to enable the effective and flexible privacy-preserving WebID analytics.
3. Data analysts can flexibly choose SDC techniques to enable data analytics processes on the large-scale privacy-aware Social Semantic Web.
4. We show how to effectively proceed anonymized WebIDs' collection and analysis but still ensure the data utility.

Preliminary Contributions

1. Propose the concept of the semantic WebID analytics pipeline for automated privacy-preserving data analytics services.
2. Three types of semantics-enabled policy for access control, data handling, and data releasing, are designed and enforced to enable the effective and flexible privacy-preserving WebID analytics.
3. Data analysts can flexibly choose SDC techniques to enable data analytics processes on the large-scale privacy-aware Social Semantic Web.
4. We show how to effectively proceed anonymized WebIDs' collection and analysis but still ensure the data utility.

Preliminary Contributions

1. Propose the concept of the semantic WebID analytics pipeline for automated privacy-preserving data analytics services.
2. Three types of semantics-enabled policy for access control, data handling, and data releasing, are designed and enforced to enable the effective and flexible privacy-preserving WebID analytics.
3. Data analysts can flexibly choose SDC techniques to enable data analytics processes on the large-scale privacy-aware Social Semantic Web.
4. We show how to effectively proceed anonymized WebIDs' collection and analysis but still ensure the data utility.

Preliminary Contributions

1. Propose the concept of the semantic WebID analytics pipeline for automated privacy-preserving data analytics services.
2. Three types of semantics-enabled policy for access control, data handling, and data releasing, are designed and enforced to enable the effective and flexible privacy-preserving WebID analytics.
3. Data analysts can flexibly choose SDC techniques to enable data analytics processes on the large-scale privacy-aware Social Semantic Web.
4. We show how to effectively proceed anonymized WebIDs' collection and analysis but still ensure the data utility.

The Spectrum of Online Social Networks

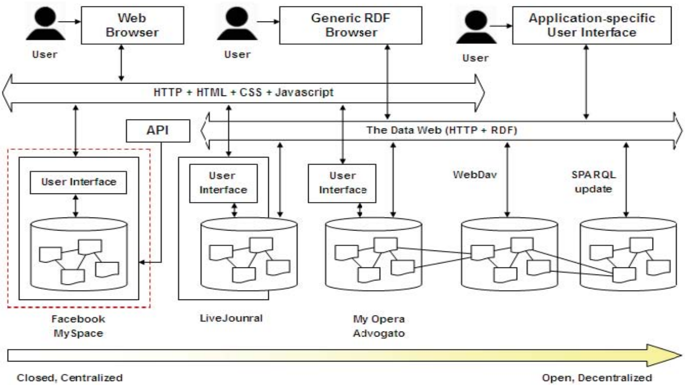

Figure 2: The spectrum of online social networking

–Decentralization: The Future of Online Social Networking

A Framework of Decentralized OSN

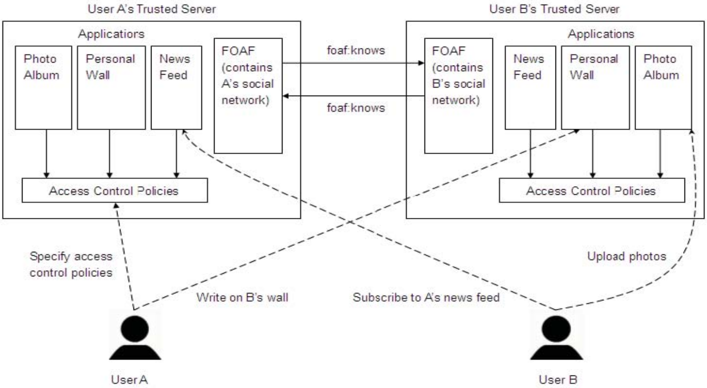

Figure 4: A framework of decentralized online social networking

–Decentralization: The Future of Online Social Networking

Decentralized Online Social Network

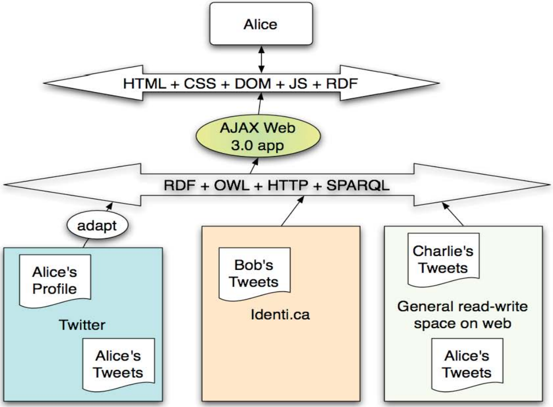

–Socially Aware Cloud Storage, Tim Berners-Lee

What are Decentralized Online Social Networks?

- Desire properties:
    - Decouple application from data
    - Give end-users control over their own data
    - Infrastructure (or platform) providers with IaaS or PaaS
    - Social network service developers with SaaS for users (or data owners)
    - Minimize the trust footprint for users to enforce their own data usage and control policies across applications.

–A World Wide Web Without Walls, Krohn, M., et al.

WebID As FOAF Ontology Representation

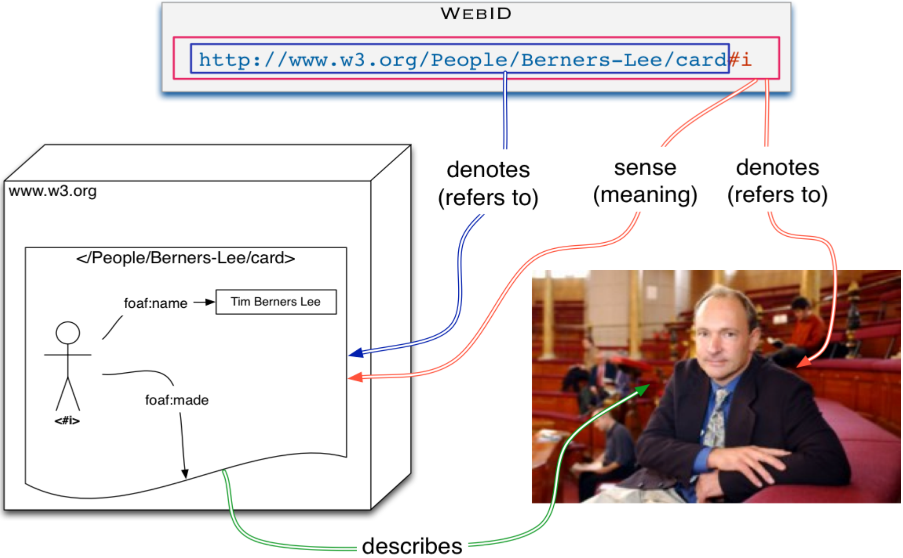

WebID

http://www.w3.org/People/Berners-Lee/card#i

denotes (refers to)

sense (meaning)

denotes (refers to)

www.w3.org

</People/Berners-Lee/card>

foaf:name

Tim Berners Lee

foaf:made

<#i>

describes

–Web Identity and Discovery, W3C Editor's Draft 19 October 2013

Decentralized Online Social Network with WebID Access Control

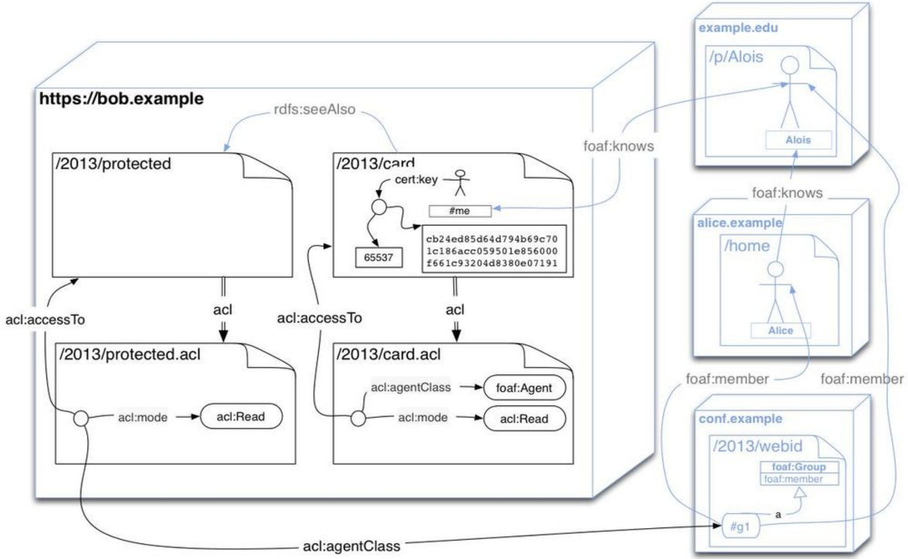

–WebAccessControl

Decentralized Online Social Network with WebID Access Control (conti.)

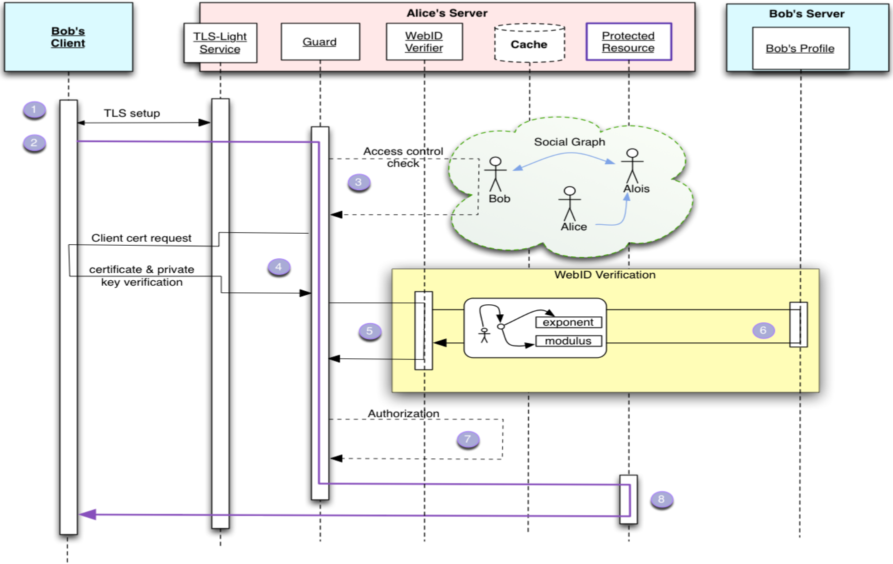

–WebID-TLS WebID Authentication over TLS, W3C Editor's Draft 21 October 2013

Decentralized Online Social Network with WebID Access Control (conti.)

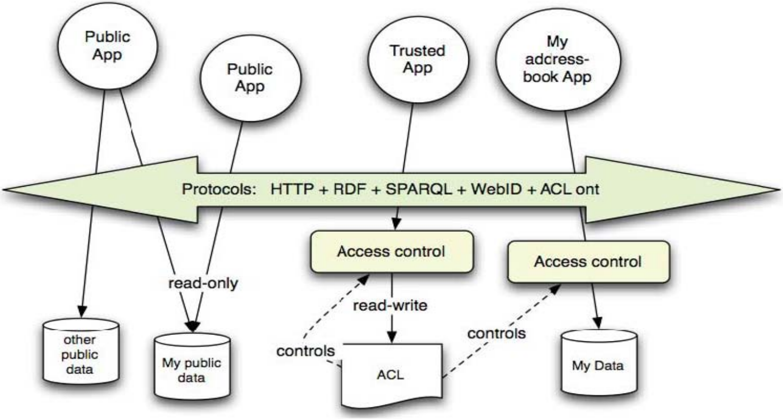

–Socially Aware Cloud Storage, Tim Berners-Lee

Privacy-Aware Social Semantic Web

- We need a new policy-oriented Social Semantic Web architecture.
- A profile management service that could be run in the browser or via a third-party website.
- Allow users to edit the attributes across multiple platforms and sites.

–A Standards-based, Open and Privacy-aware Social Web, W3C Incubator Group Report 6th Dec. 2010

A Super-Peer Domain (SPD) Data Cloud

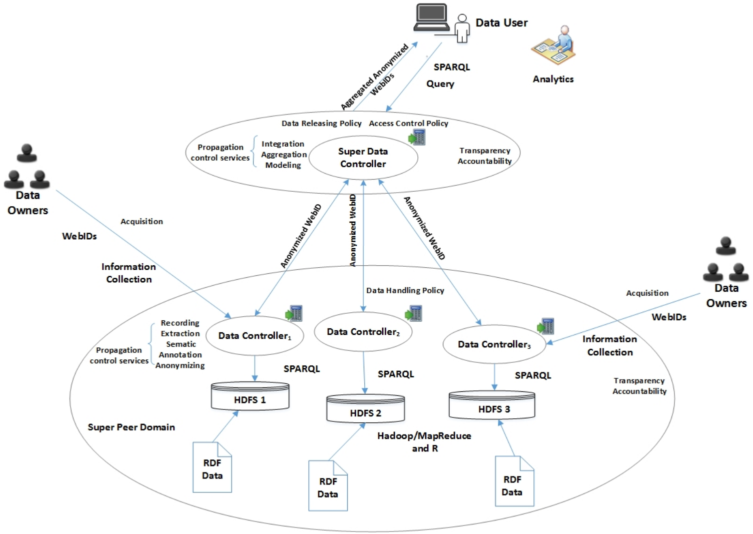

Policy Ontology for a Super-Peer Domain Cloud

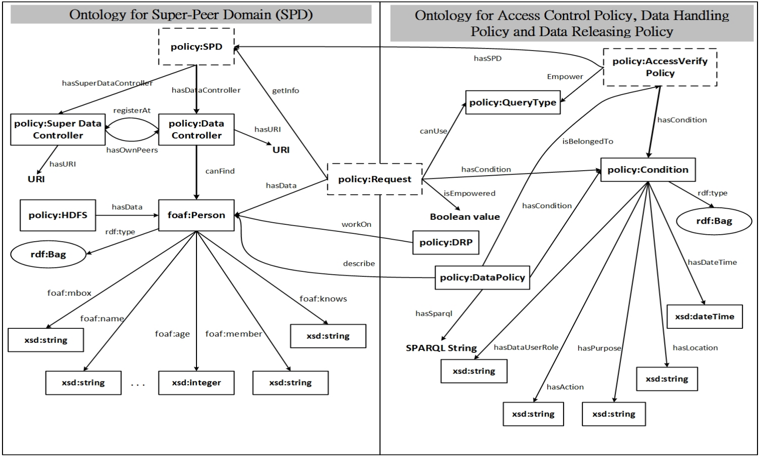

Semantics-enabled Policies

1. Semantics-enabled policies are composed of ontologies and queries, where ontologies describe the *concepts* of privacy-preserving WebID analytics services, and queries *enforce* the above privacy principles.
2. Semantics-enabled policies are correspond to query restriction, data manipulation/anonymization, and output perturbation:
    - Access Control Policy (ACP)
    - Data Handling Policy (DHP)
    - Data Releasing Policy (DRP)

Semantics-enabled Policies

1. Semantics-enabled policies are composed of ontologies and queries, where ontologies describe the *concepts* of privacy-preserving WebID analytics services, and queries *enforce* the above privacy principles.
2. Semantics-enabled policies are correspond to query restriction, data manipulation/anonymization, and output perturbation:
    - Access Control Policy (ACP)
    - Data Handling Policy (DHP)
    - Data Releasing Policy (DRP)

An Ontology for Access Control Policy (ACP)

Definition of ACP Ontology

The concept of a data user's request verifications is represented as an ACP ontology and enforced as a SPARQL query.

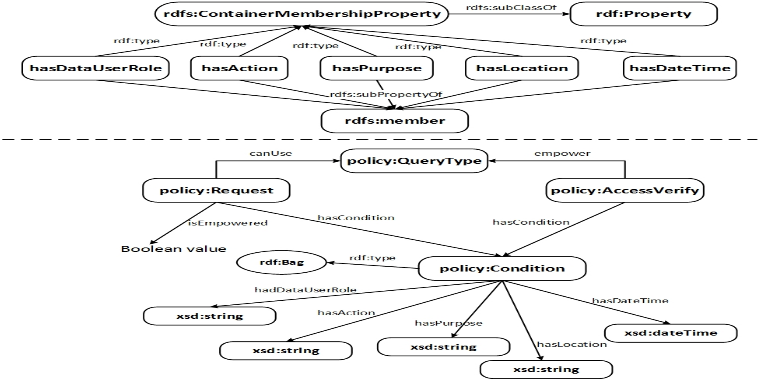

A Query for Access Control Policy (ACP)

An ACP SPARQL Query

```
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix policy: <http://nccu.edu.tw/policy> .
policy:QueryType rdf:type rdf:Class
policy:PBQ rdf:type policy:QueryType
policy:Condition [
hasDataUserRole “DataAnalyst”;
hasPurpose “Analytics”;
hasAction “Read”;
hasLocation “Taipei”;
hasDateTime “2013:12:25:15:00” ].
```

A Query for Access Control Policy (ACP) (conti.)

An ACP SPARQL Query (conti.)

```
Ask ?permit
From  < PeterRequest.rdf >
Where {?r policy : isEmpowered ?permit.
?r [ ?qt rdf : type policy : QueryType;
policy : hasCondition ?c [
hasDataUserRole ?role;
hasPurpose ?purpose;
hasAction ?action;
hasLocation ?location;
hasDateTime ?time ] ].}
```

An Ontology for Data Handling Policy (DHP)

Definition of DHP Ontology

A DHP describes which SDC techniques are used to anonymize WebIDs' profile attributes, but not yet social network structure, in an FOAF property graph.

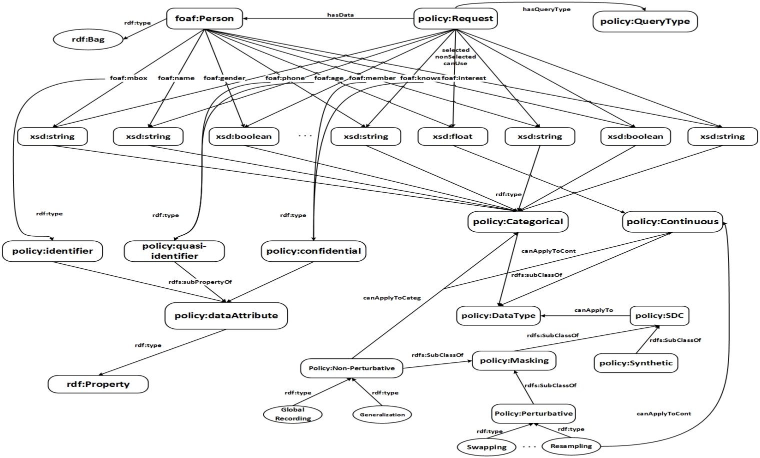

A Data Handling Policy (DHP)

A DHP calls for WebIds' profiles anonymizing

```
@prefix foaf : < http : //xmlns.com/foaf/0.1/ > .
@prefix policy : < http : //nccu.edu.tw/policy > .
< http : //nccu.edu.tw/j/foaf.rdf >  a
foaf : PersonalProfileDocument.
< http : //nccu.edu.tw/j/foaf.rdf > foaf : maker :me.
< http : //nccu.edu.tw/j/foaf.rdf >  foaf : primaryTopic :me.
:me a foaf : Person.
/ ∗ De − identification ∗ /
:me [ foaf : name "Yuh − Jong Hu";
foaf : homepage  < http : //nccu.edu.tw/j >;
foaf : mbox  < mailto : j@cs.nccu.edu.tw >;
/ ∗ Generalization ∗ /
foaf : phone  < tel : +886 − 2 − 29387620 >;
```

A Data Handling Policy (DHP)(conti.)

A DHP calls for WebIDs' profiles anonymizing

```
............
foaf : knows [ a foaf : Person;
/ * De − identification * /
foaf : name "Kua − Ping Cheng";
rdfs : seeAlso
/ * enhanced microdata protection techniques * /
< http : //nccu.edu.tw/k/foaf.rdf > ].
foaf : knows [ a foaf : Person;
/ * De − identification * /
foaf : name "Ya − Ling Huang";
rdfs : seeAlso
/ * enhanced microdata protection techniques * /
< http : //nccu.edu.tw/y/foaf.rdf > ].
............ ]
```

An Ontology for Data Releasing Policy (DRP)

Definition of DRP Ontology

Governs acceptable conditions to disclose anonymized WebID's attributes to ensure the compliance of privacy protection principle.

A DRP queries for anonymized WebID

```
Select  ?graph  ?gender  ?age  ?member  ?interest
From < http : //nccu.edu.tw/j/foaf.rdf >
From  named  graph???
Where{< http : //nccu.edu.tw/j/foaf.rdf#me >
foaf : knows  ?X.
{  ?X  rdfs : seeAlso  ?graph.
graph  ?graph  {[  a  foaf : Person.
/ * De − identification * /
foaf : mbox  ?mbox;
foaf : name  ?name;
foaf : gender  ?gender;
..................;
/ * Generalization * /
foaf : phone  ?phone;
/ * GlobalRecording * /
foaf : age  ?age;
foaf : member  ?member;
foaf : interest  ?interest;

foaf : knows  [  ?graph  ].  ]}}}
```

The Semantic WebID Analytics Pipeline

- In a six-stage lifecycle:
  1. Acquisition and recording
  2. Extraction, cleaning, and semantic annotation
  3. Representation, integration, and aggregation
  4. Modeling and analysis
  5. Query processing and disclosure for analytics
  6. Interpretation

The Semantic WebID Analytics Pipeline

- In a six-stage lifecycle:
  1. Acquisition and recording
  2. Extraction, cleaning, and semantic annotation
  3. Representation, integration, and aggregation
  4. Modeling and analysis
  5. Query processing and disclosure for analytics
  6. Interpretation

The Semantic WebID Analytics Pipeline

- In a six-stage lifecycle:
  1. Acquisition and recording
  2. Extraction, cleaning, and semantic annotation
  3. Representation, integration, and aggregation
  4. Modeling and analysis
  5. Query processing and disclosure for analytics
  6. Interpretation

The Semantic WebID Analytics Pipeline

- In a six-stage lifecycle:
  1. Acquisition and recording
  2. Extraction, cleaning, and semantic annotation
  3. Representation, integration, and aggregation
  4. Modeling and analysis
  5. Query processing and disclosure for analytics
  6. Interpretation

The Semantic WebID Analytics Pipeline

- In a six-stage lifecycle:
  1. Acquisition and recording
  2. Extraction, cleaning, and semantic annotation
  3. Representation, integration, and aggregation
  4. Modeling and analysis
  5. Query processing and disclosure for analytics
  6. Interpretation

The Semantic WebID Analytics Pipeline

- In a six-stage lifecycle:
  1. Acquisition and recording
  2. Extraction, cleaning, and semantic annotation
  3. Representation, integration, and aggregation
  4. Modeling and analysis
  5. Query processing and disclosure for analytics
  6. Interpretation

The Semantic WebID Analytics Pipeline (conti.)

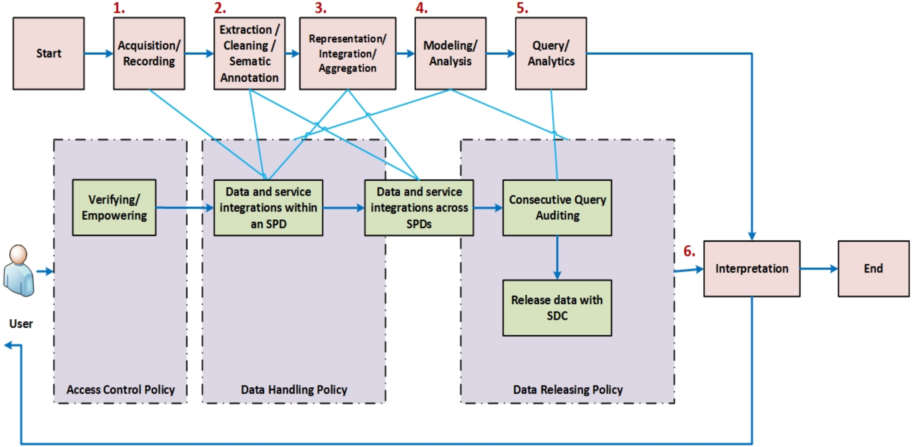

Figure: WebID Analytics Pipeline

Big Data Analytics Platforms

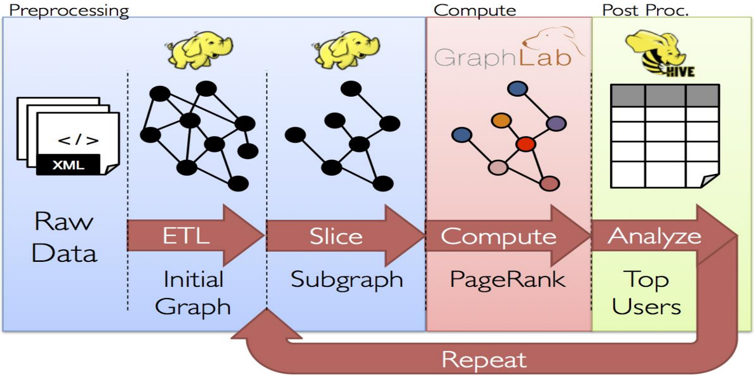

Figure: Berkeley AMP Lab. GraphX

Big Data Analytics Platforms

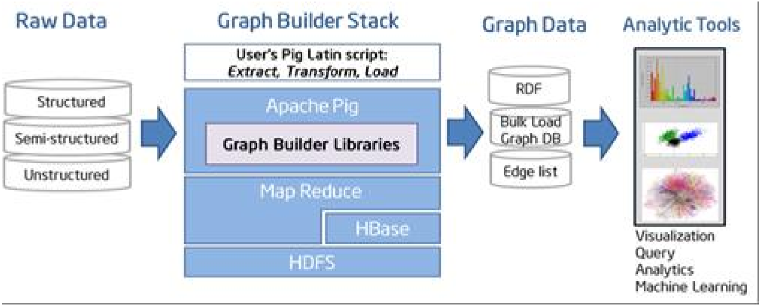

FIGURE:

Intel Lab GraphBuilder

Big Data Analytics Platforms

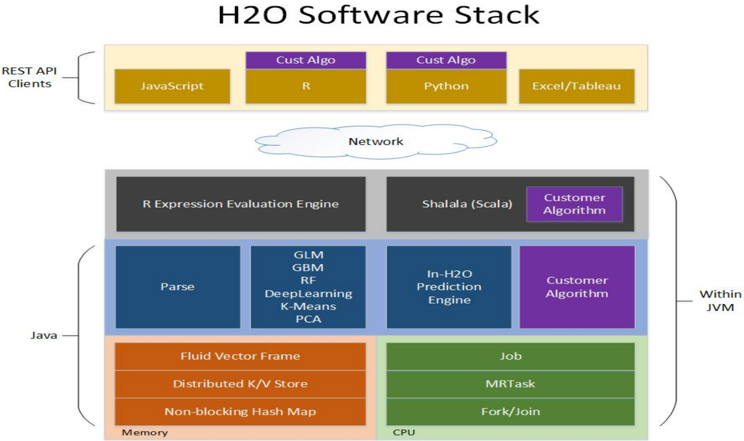

FIGURE:

H_2O Software Stack

Apply R+SPARQL and Hadoop/MapReduce for WebID Analytics

- MapReduce paradigm works for lightweight data-parallel analytics services.
- Distributed R for heavyweight graph-parallel WebID analytics services.
- Graph-parallel analytics services for discovering various social network's degree centralities, and data-parallel analytics services for WebID anonymizing and output perturbation.
- Integrating R+SPARQL and MapReduce brings heavyweight graph-parallel WebID analytics of R and lightweight data-parallel WebID analytics of MapReduce.

Apply R+SPARQL and MapReduce for WebID Analytics (conti.)

- We leverage the Semantic Web's open technologies for Social Semantic Web data representation and access.
- In future, possibly merging centralized social network data in JSON with JSON-Linked Data (JSON-LD) for integrated data analytics.
- JSON-LD are stored in the HDFS of cluster computers for *batch* and *interactive* data analytics in R.
- SPARQL queries and filters FOAF datasets in JSON-LD for R to enable heavyweight data analytics services.

Semantic Data Protection Protection and Analytics

- Improve the situation, where SDC enforcement is obliged to original data providers so a data analytics user lacks the flexibility to choose suitable SDC methods.
- Seek a balance between a data owner's right for privacy protection and a data user's need for data analytics through SDC methods selection.
- Semantics-enabled DHP and DRP call for feasible SDC methods and ensures maximum data utility with a tolerable data disclosure risk.

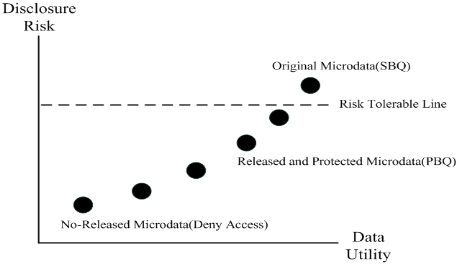

Semantic Data Protection Protection and Analytics

- Improve the situation, where SDC enforcement is obliged to original data providers so a data analytics user lacks the flexibility to choose suitable SDC methods.
- Seek a balance between a data owner's right for privacy protection and a data user's need for data analytics through SDC methods selection.
- Semantics-enabled DHP and DRP call for feasible SDC methods and ensures maximum data utility with a tolerable data disclosure risk.

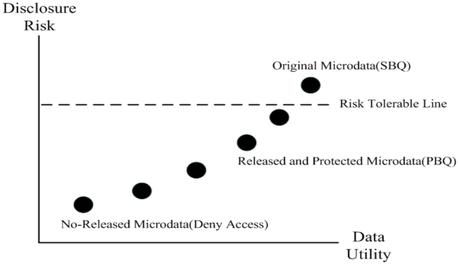

Semantic Data Protection Protection and Analytics

- Improve the situation, where SDC enforcement is obliged to original data providers so a data analytics user lacks the flexibility to choose suitable SDC methods.
- Seek a balance between a data owner's right for privacy protection and a data user's need for data analytics through SDC methods selection.
- Semantics-enabled DHP and DRP call for feasible SDC methods and ensures maximum data utility with a tolerable data disclosure risk.

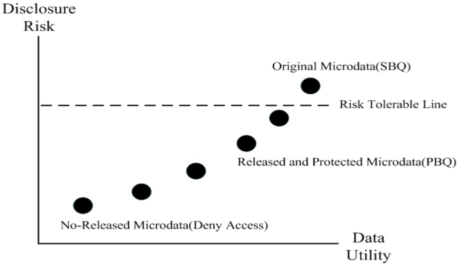

Conclusion and Future Works

- Preliminary Results:
  1. Semantics-enabled policies are proposed and verified to provide query restriction, data manipulation/anonymization, and output perturbation.
  2. The R+SPARQL (for graph-parallelism) and MapReduce (for data-parallelism) platform is establishing to enable a flexible and effective privacy-preserving WebID analytics.
  3. Billion Triples Challenge 2012 datasets have been used for FOAF/WebID analytics.
  4. A simple balance between data protection and utility through the semantics-enabled policies to call for suitable SDC methods.

Conclusion and Future Works

- Preliminary Results:
  1. Semantics-enabled policies are proposed and verified to provide query restriction, data manipulation/anonymization, and output perturbation.
  2. The R+SPARQL (for graph-parallelism) and MapReduce (for data-parallelism) platform is establishing to enable a flexible and effective privacy-preserving WebID analytics.
  3. Billion Triples Challenge 2012 datasets have been used for FOAF/WebID analytics.
  4. A simple balance between data protection and utility through the semantics-enabled policies to call for suitable SDC methods.

Conclusion and Future Works

- Preliminary Results:
  1. Semantics-enabled policies are proposed and verified to provide query restriction, data manipulation/anonymization, and output perturbation.
  2. The R+SPARQL (for graph-parallelism) and MapReduce (for data-parallelism) platform is establishing to enable a flexible and effective privacy-preserving WebID analytics.
  3. Billion Triples Challenge 2012 datasets have been used for FOAF/WebID analytics.
  4. A simple balance between data protection and utility through the semantics-enabled policies to call for suitable SDC methods.

Conclusion and Future Works

- Preliminary Results:
    1. Semantics-enabled policies are proposed and verified to provide query restriction, data manipulation/anonymization, and output perturbation.
    2. The R+SPARQL (for graph-parallelism) and MapReduce (for data-parallelism) platform is establishing to enable a flexible and effective privacy-preserving WebID analytics.
    3. Billion Triples Challenge 2012 datasets have been used for FOAF/WebID analytics.
    4. A simple balance between data protection and utility through the semantics-enabled policies to call for suitable SDC methods.

Conclusion and Future Works (conti.)

- Future Works:
  1. Fully enforcing the semantics-enabled policies on the R+SPARQL and MapReduce platform.
  2. Using differential privacy technique for output disclosure control.
  3. Applying bottom-up inductive reasoning, such as machine learning and top-down deductive reasoning, such as ontology and rule knowledge representation, for policy modelling and WebID analytics services.
  4. Crafting an optimized balance between WebID protection and utility.

Conclusion and Future Works (conti.)

- Future Works:
  1. Fully enforcing the semantics-enabled policies on the R+SPARQL and MapReduce platform.
  2. Using differential privacy technique for output disclosure control.
  3. Applying bottom-up inductive reasoning, such as machine learning and top-down deductive reasoning, such as ontology and rule knowledge representation, for policy modelling and WebID analytics services.
  4. Crafting an optimized balance between WebID protection and utility.

Conclusion and Future Works (conti.)

- Future Works:
  1. Fully enforcing the semantics-enabled policies on the R+SPARQL and MapReduce platform.
  2. Using differential privacy technique for output disclosure control.
  3. Applying bottom-up inductive reasoning, such as machine learning and top-down deductive reasoning, such as ontology and rule knowledge representation, for policy modelling and WebID analytics services.
  4. Crafting an optimized balance between WebID protection and utility.

Conclusion and Future Works (conti.)

- Future Works:
  1. Fully enforcing the semantics-enabled policies on the R+SPARQL and MapReduce platform.
  2. Using differential privacy technique for output disclosure control.
  3. Applying bottom-up inductive reasoning, such as machine learning and top-down deductive reasoning, such as ontology and rule knowledge representation, for policy modelling and WebID analytics services.
  4. Crafting an optimized balance between WebID protection and utility.

M. Krohn *et al.*, "A world wide web without walls," in *6th ACM Workshop on Hot Topics in Networking (Hotnets)*. ACM, 2007.

C. Yeung, A. *et al.*, "Decentralization: The future of online social networking," in *W3C Workshop on the Future of Social Networking*. W3C, 2009.

T. Berners-Lee, "Socially aware cloud storage," September 2011.

T. Inkster, H. Story, and B. Harbulot, "WebID-TLS: WebID authentication over TLS," W3C, Tech. Rep., October 2013.

D. Appelquist *et al.*, "A standard-based, open and privacy-aware social web," W3C Incubator Group Report, Tech. Rep., December 2010.

J. D. Weitzner *et al.*, "Creating a policy-aware web: Discretionary, rule-based access for the world wide web," in *Web and Information Security*, E. Ferrari and B. Thuraisingham, Eds. IGI, 2006, pp. 1–31.

V. Ciriani *et al.*, "Microdata protection," in *Secure Data Management in Decentralized Systems*, T. Yu and S. Jajodia, Eds. Springer, 2007, pp. 291–321.

A. Hundepool *et al.*, *Statistical Disclosure Control*. Wiley Series in Survey Methodology, 2012.

A. Labrinidis *et al.*, "Challenges and opportunities with big data," Computing Research Consortium (CSR), Tech. Rep., 2012.

- K. Liu *et al.*, "Privacy-preserving data analysis on graphs and social networks," in *Next Generation Data Mining*, H. Kargupta *et al.*, Eds. CRC Press, 2008, pp. 1–17.

- A. Narayanan *et al.*, "A critical look at decentralized personal data architectures," Cornell University Library, Tech. Rep., 2012.

- B. Carminati and E. Ferrari, "Privacy-aware access control in social networks: Issues and solutions," in *Privacy and Anonymity in Information Management Systems*, J. Nin and J. Herranz, Eds. Springer, 2010, pp. 181–195.

- E. Zheleva, E. Terizi, and L. Getoor, *Privacy in Social Networks*. Morgan&Claypool, 2012.

- S. D. C. d. Vimercati *et al.*, "Access control policies and languages in open environments," in *Secure Data Management in Decentralized Systems*, T. Yu and S. Jajodia, Eds. Springer, 2007, pp. 21–58.

- Y. J. Hu *et al.*, "Crafting a balance between big data utility and protection in the semantic data cloud," in *International Conference on Web Intelligence, Mining and Semantics (WIMS'13)*. ACM Press, June 2013.

- A. P. Bonatti, "Datalog for security, privacy and trust," in *Datalog 2010*, ser. LNCS 6702. Springer, 2011, pp. 21–36.

- D. Beckett *et al.*, "Turtle: Terse RDF triple language," W3C Candidate Recommendation, Tech. Rep., February 2013.

- J. Weaver and P. Tarjan, "Facebook linked data via the graph API," *Semantic Web - Interoperability, Usability, Applicability*, 2012.

- M. Spomy *et al.*, "JSON-LD 1.0," W3C Proposed Recommendation, Tech. Rep., November 2013.

- R. N. Adam and C. J. Worthmann, "Security-control methods for statistical databases: A comparative study," *ACM Computing Survey*, vol. 21, no. 4, pp. 515–556, 1989.