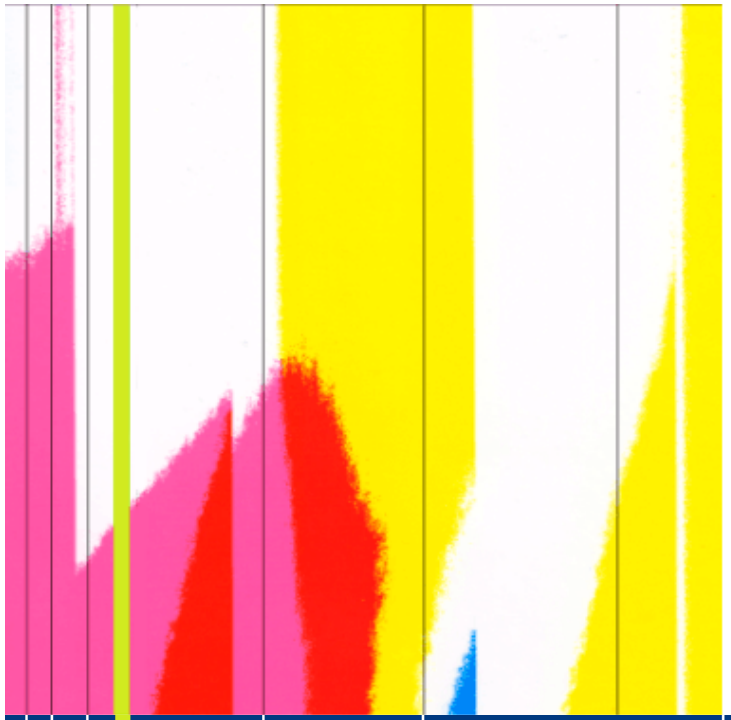

Watermarking ontologies, or how to create plausible facts

Serge Abiteboul, David Gross-Amblard, Fabian Suchanek
WebDam - INRIA Saclay

March 4 2011

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

INRIA

centre de recherche SACLAY - ÎLE-DE-FRANCE

Motivation

- Building (big/useful) ontologies with individual entities: hard work (Yago, dbpedia, …)
- High cost, high price
- Ontology producers want:
  - (proprietary model): to limit illegal resales of the ontology
  - (open model): to be properly referenced when used
- Need for ownership proof methods on suspect ontologies
- Malevolent resaler is not necessarily stupid

Attacks to evade detection

- Renaming things (Paris_City becomes cityOfParis)
- Reselling only subsets, or mix with other ontologies
- Random alteration of facts
- …
- (your proposition here)

An approach : Watermarking

Watermarking techniques

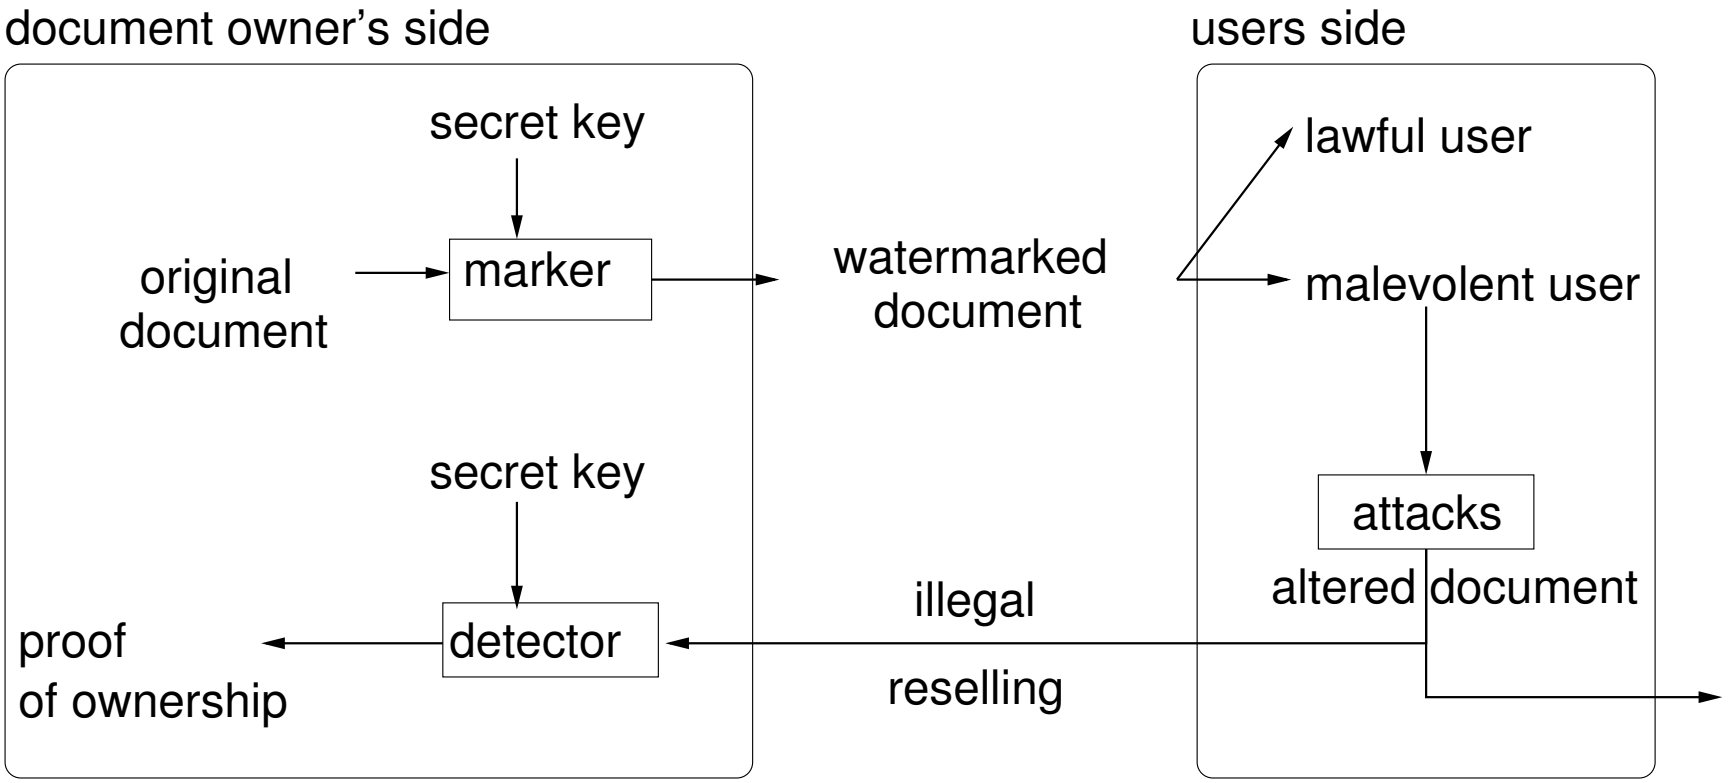

Existing methods

- For databases or XML
  - Adding errors in numerical/categorical/spatial/geometrical attributes
  - Ok if errors are allowed
  - A priori knowledge of attacker is not taken into account now
- For ontologies
  - Syntactical rewriting (<a> exchanged with <a/>, fake spaces, empty attributes)
  - Easily removed
- Adding fake data
  - So natural…but not considered so far (one short student paper)
  - How to watermark an ontology using fake facts ?

Fake facts requirements

- Populate the ontology: subset attack
- Being commonplace: the attacker should
  - Compare with existing datasets (easy) and ground-truth (Wikipedia/hard)
  - Delete strange facts
    - A city with a « OwnershipProof » relation ?
  - Delete outliers
    - A city with 2 inhabitants ?

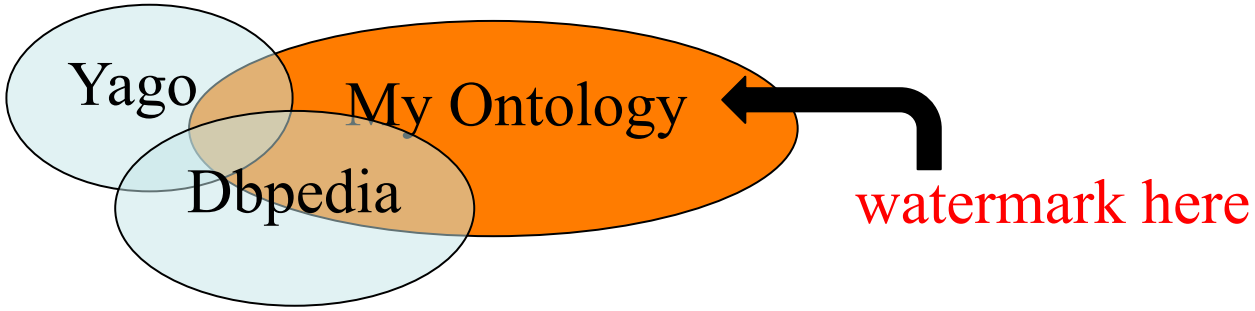

Open Questions

- How can we deal with other attacks?
- How can we integrate DL constraints?
- How can we better define statistical invisibility?

Thanks.