

UC San Diego

Technical Reports

Title

Censorship-resistant Publishing

Permalink

<https://escholarship.org/uc/item/86c1m73c>

Authors

Curtmola, Emiran
Deutsch, Alin
Ramakrishnan, K. K.
[et al.](#)

Publication Date

2010-03-20

Peer reviewed

Censorship-resistant Publishing

EMIRAN CURTMOLA

University of California, San Diego

and

ALIN DEUTSCH

University of California, San Diego

and

K.K. RAMAKRISHNAN

AT&T Labs – Research

and

DIVESH SRIVASTAVA

AT&T Labs – Research

As the web evolves, it is becoming easier to form communities based on shared interests, and to create, publish, and query data on a wide variety of topics. In order to fully deliver on the promise of free data exchange, any community-supporting infrastructure needs to enforce the key requirement to preserve privacy of the association of content providers with potential sensitive published information. This privacy preserving publishing requirement prevents censorship, harassment, or discrimination of users by third parties.

We propose a novel privacy-preserving distributed infrastructure in which data resides only with the publishers owning it. The infrastructure disseminates user queries to publishers, who answer them at their own discretion. The infrastructure enforces a *publisher k-anonymity* guarantee, which prevents leakage of information about which publishers are capable of answering a certain query. Given the virtual nature of the global data collection, we study the challenging problem of efficiently locating publishers in the community that contain data items matching a specified query. We propose a distributed index structure, UQDT, that is organized as a union of Query Dissemination Trees (QDTs), and realized on an overlay (i.e., logical) network infrastructure. Each QDT has data publishers as its leaf nodes, and overlay network nodes as its internal nodes; each internal node routes queries to publishers, based on a summary of the data advertised by publishers in its subtrees. We experimentally evaluate design tradeoffs, and demonstrate that UQDT can maximize throughput by preventing any overlay network node from becoming a bottleneck.

Categories and Subject Descriptors: H.2.4 [Database Management]: Systems—*Distributed databases*

General Terms: Design, Performance

Additional Key Words and Phrases: Distributed query dissemination, privacy-aware communities, k-anonymity publishing

Author's addresses: Emiran Curtmola, Department of Computer Science at University of California San Diego, La Jolla, CA 92093; email: ecurtmola@cs.ucsd.edu; Alin Deutsch, Department of Computer Science at University of California San Diego, La Jolla, CA 92093; email: deutsch@cs.ucsd.edu; K.K. Ramakrishnan, AT&T Labs – Research, Florham Park, NJ 07932; email: kkrama@research.att.com; Divesh Srivastava, AT&T Labs – Research, Florham Park, NJ 07932; email: divesh@research.att.com;

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Request

To obtain a copy of this UCSD technical report please send an email or a letter request to:

Emiran Curtmola ecurtmola@cs.ucsd.edu

University of California, San Diego - CSE
9500 Gilman Dr.
La Jolla, California 92093, USA