Censorship-resistant Publishing

Emiran Curtmola, Alin Deutsch
UC San Diego
{ecurtmola,deutsch}@cs.ucsd.edu

K.K. Ramakrishnan, Divesh Srivastava
AT&T Labs Research
{kkrama,divesh}@research.att.com

ABSTRACT

As the web evolves, it is becoming easier to form communities based on shared interests, and to create and publish data on a wide variety of topics. With this "democratization of information creation" comes the natural desire to make one's data accessible for querying within the community and also be able to query the global collection that is the union of all local data collections of others within the community.

In order to fully deliver on the promise of free data exchange, any community-supporting infrastructure needs to enforce the key requirement of being resistant to censorship by third parties, be they of governmental, corporate, of other special interest nature. Censorship resistance precludes some obvious approaches that reuse and build on existing centralized technologies, e.g., search engines, hosted online communities, etc.

We propose a 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 prevents third parties from pinpointing which publisher advertises what data (without extensively colluding with or attacking community members).

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 subtree.

We experimentally evaluate design tradeoffs, and demonstrate that UQDT can maximize throughput by preventing any overlay network node from becoming a bottleneck.

Technical Report. University of California, San Diego. March 2008.
.

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