Computer Networks 55 (2011) 241–263

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

SATI: A scalable and traffic-efficient data delivery infrastructure for real-time sensing applications

Kyungmin Cho [a], Younghyun Ju [a], Sungjae Jo [a], Yunseok Rhee [b], Junehwa Song [a,*]

[a] Korea Advanced Institute of Science and Technology (KAIST), 373-1 Guseong-dong, Yuseong-gu, Daejeon 305-701, Republic of Korea
[b] Hankuk University of Foreign Studies, Mohyun, Yongin, Kyunggi-Do 449-791, Republic of Korea

ARTICLE INFO

Article history:
Received 23 June 2009
Received in revised form 4 August 2010
Accepted 17 August 2010
Available online 27 August 2010
Responsible Editor: A. Pitsillides

Keywords:
Network architecture
Overlay network
Delivery path management
Content-based publish/subscribe system

ABSTRACT

Upcoming ubiquitous technologies are expediting the advent of many real-time applications. Examples of such applications include physical world browsing, RFID-based supply chain management, city-wide road traffic monitoring, weather forecasting, and air pollution monitoring. These applications show different scales and characteristics in terms of sensing data delivery demands. They commonly demand a deep understanding on real-time data delivery from widely distributed data sources. Also, they have highly individualized and fine-grained delivery demands in terms of data and delay specifications, e.g., data value ranges of interest, spatial and temporal resolutions, and tolerable delay, etc. Due to the remarkable scale and complexity, however, existing data delivery systems cannot support such applications effectively. We present SATI (scalable and traffic-efficient data delivery infrastructure), a novel Internet-based sensing data delivery infrastructure that provides a common platform for data providers and consumers. Basically, it is comprised of a collection of proxy nodes forming an overlay network, where each proxy node conducts an in-network processing and efficient data delivery. It allows applications to specify their delivery requirements with intuitive and comprehensive delivery semantics. For scalable and efficient data delivery, SATI develops a novel delivery path management scheme based on an incremental relaxation method. The scheme enables SATI to construct and maintain efficient delivery paths satisfying a large number of delivery requests of high diversity. It fully exploits the diversity of delivery demands on both data and delay requirements, thus achieving a high level of service satisfaction and efficiency at the same time. The result from a large-scale simulation shows that SATI achieves a high level of scalability and bandwidth efficiency.

1. Introduction

The rapid advance of sensor technologies and mobile devices will lead to a new computing environment which is fully-instrumented with innumerable sensors and tiny devices. In such an environment, diverse sensor networks [11,19,1,29] will emerge and proliferate in a dynamic and ad hoc fashion beyond simply static ones. Leveraging these sensor networks, numerous real-time sensing applications, e.g., physical world browsing [31], RFID-based supply chain management, city-wide road traffic monitoring, weather forecasting, and air pollution monitoring [19,5], will be developed and used in various application domains. These applications will mostly require real-time understanding of sensing data, and have diverse, individualized needs on data according to the application users' personal or business purposes. Delivering such continuous and voluminous data to a large number of real-time sensing applications effectively and in a timely manner demands

* Corresponding author. Tel.: +82 42 350 3546.
E-mail addresses: kmcho@nclab.kaist.ac.kr (K. Cho), yhju@nclab.kaist.ac.kr (Y. Ju), sjjo@nclab.kaist.ac.kr (S. Jo), rheeys@hufs.ac.kr (Y. Rhee), junesong@nclab.kaist.ac.kr (J. Song).

1389-1286/$ - see front matter
doi:10.1016/j.comnet.2010.08.004

a scalable and efficient sensing data delivery infrastructure. Due to the remarkable scale and complexity, it is very costly and difficult for individual data providers owning sensor networks to conduct sensor data delivery from the bare ground without such infrastructural support.

In this paper, we present SATI (scalable and traffic-efficient data delivery infrastructure), a novel Internet-based sensing data delivery infrastructure which provides a common platform for real-time sensing applications. SATI enables the data providers to easily publish their data without considering the details of delivery operations. It also provides continuous and timely data delivery service to numerous consumers. It allows data consumers to search for and obtain desired sensing data in a convenient and unified way by supporting intuitive and comprehensive data delivery semantics such as data of interest and tolerable delay. Using the delivery semantics, consumers can specify their own delivery requirements in fine-granularity.

The basic architecture of SATI is based on a content-based publish/subscribe system [8,15,26]. It consists of a number of proxy nodes and mediates a large number of data providers and consumers through collaborative delivery and processing among the nodes. In SATI, one or more delivery trees are constructed, one for each provider, which connects proxy nodes so as to efficiently support all consumers interested in the data provider. In particular, the delivery trees are carefully organized to satisfy delivery demands of individual consumers while reducing redundant data delivery to efficiently utilize bandwidth resources. Based on the delivery trees, SATI makes full use of in-network processing at each node on a delivery path, where each node efficiently filters out unnecessary data that are not requested by the subsequent nodes while meeting delay requirements.

A number of efforts have explored data or content delivery such as application-level multicast (ALM) [12,4], content-based pub/sub (CBPS) [8,26], sensing data sharing infrastructure [32,17,31], and content delivery network (CDN) [30]. However, existing delivery systems are limited in considering the following two requirements, which are essential to effectively support real-time sensing applications. First, these applications should support individualized fine-grained needs on data in terms of *data value ranges of interest, and spatial and temporal resolutions*. For example, in a physical world browsing application such as Microsoft Sensor Map, users may select an area of interest, setting a value range to detect an occurrence of certain events such as a forest fire. To closely scrutinize the area, they also zoom in and out of the selected area at different spatial resolutions and adjust its sensing frequency for different temporal resolutions. Second, these real-time sensing applications require stable data delivery within a tolerable delay to perform timely actions. Needless to say, stale data will significantly deteriorate quality of and confidence in the applications.

However, it is not a simple problem to address the above requirements, especially when the number of data producers, data consumers, and proxy nodes greatly increases. To build a bandwidth-efficient delivery tree, a wide exploration of all the proxy nodes as well as providers and consumers is required. In addition, each delivery path should be maximally shared among as many nodes as possible. Meanwhile, the data specifications of consumers are fine-grained and different from each other, reflecting individualized interests. Thus, it is not easy to develop such an efficient data delivery tree while meeting all the data specifications of the consumers. Similarly, the data delivery tree should satisfy individual delay requirements while considering bandwidth constraints. For this purpose, careful investigation on possible path rearrangements according to network latencies among nodes is also required. Certainly, however, it demands a large amount of computation to find an optimal delivery tree satisfying the delay requirements of numerous consumers while maximizing bandwidth efficiency.

To address the above challenges, we propose a novel data delivery path management scheme where a data delivery tree is organized to support efficient bandwidth sharing by similarity-based data delivery, and also to satisfy delay requirements of real-time applications using an urgency-based path arrangement. For efficient and scalable tree construction, we develop an incremental relaxation method that efficiently supports a partial rearrangement of nodes in building a data delivery tree while considering path sharing and delay requirements. Our method attempts to add a new node to the delivery tree while keeping the current delivery tree structure as intact as possible. Only after it fails at the attempt, does it conduct a minimal and partial rearrangement of the tree structure to accommodate the new node. In fact, the rearrangement is conducted by iterations of detaching and rejoining nodes until it finds a delivery tree that can accommodate the new node as well as all the existing nodes. To minimize the cost of the rearrangements, our method incrementally relaxes the bandwidth and delay conditions in finding a proper parent for a new node. With this incremental relaxation method, SATI can efficiently construct and maintain a delivery tree that can satisfy the fine-grained and individualized data delivery requirements of a large number of consumers with efficient bandwidth utilization. In addition, each node of the tree is equipped with processing operators which perform selective data delivery down to its descendants. As a result, SATI efficiently supports real-time sensing applications through in-network processing considering the individual data requirements of consumers.

The contributions of our work can be summarized as follows. First, we characterize the semantics of data delivery specification for real-time sensing applications and enable the applications to specify their fine-grained needs on data, i.e., *data value ranges of interest, spatial* and *temporal resolutions*. SATI also allows the applications to specify *desired delay requirements* on the required data. Second, we propose a novel delivery path management scheme which fully considers the above fine-grained needs on data as well as delay requirements in constructing bandwidth-efficient delivery paths. Furthermore, SATI scales well with the number of consumers while satisfying their requirements. Through extensive experiments based on large-scale simulations, we demonstrate that SATI outperforms