

EXPLOITING ATTRIBUTE POPULARITY DISTRIBUTION SKEW TO ENHANCE THE PERFORMANCE OF PEER TO PEER PUBLISH/SUBSCRIBE SYSTEMS

CHYOUHWA CHEN, CHIA-LIANG TSAI AND SHI-JINN HORNG*

Department of Computer Science and Information Engineering
National Taiwan University of Science and Technology
No. 43, Sec. 4, Keelung Road, Taipei 106, Taiwan
{similar; M9615041}@mail.ntust.edu.tw

*Corresponding author: horngsj@yahoo.com.tw

Received February 2010; revised July 2010

ABSTRACT. *An inherent phenomenon in real-world publish/subscribe services is the skew in attribute popularity, where some attributes are much more popular than others. None of the many publish/subscribe systems in the literature proposed effective means to address this phenomenon, which results in workload imbalance among peers, excessive total computation cost, inefficient event delivery or all of the above. To combat these problems, we propose two fundamental architectural designs for these systems. First, we propose to exploit the attribute popularity distribution skew in the subscription installation process. A number of novel subscription installation strategies are presented in this paper that exploits the popularity of attributes to spread the load across a multitude of peers so that the peer workload both in subscription storage and in event matching computation is much more balanced. However, wide spreading of subscription storage and event computation workload usually result in inefficient event delivery. As our second contribution, we propose an architecture design where the subscription storage and event delivery components are decoupled. However, the two components are also seamlessly integrated so that event delivery remains extremely efficient. Using an extensive set of experiments, our proposed approach is shown to require only a modest level of subscription replication, require significantly less total computation cost, achieve better load balance among peers in subscription storage and event matching, and deliver events in an extremely efficient manner. In addition, our approach is shown to be scalable in the number of events, subscriptions, schema attributes and network peers.*

Keywords: Publish/subscribe services, Subscription, Event, Structured P2P networks, Distributed hash table, Rendezvous-based mapping function

1. Introduction. Content-based publish/subscribe (pub/sub) systems are becoming increasingly popular for information dissemination from event publishers to event subscribers [2,3,8-10,15,22,24,28-30]. Subscribers register their interests in events through subscriptions using a set of predicates over attributes in a subscription language. An event instance published by a publisher specifies constant values of a set of attributes using predicates involving equality operator only. Upon receiving an event, the pub/sub system is responsible for finding the subscriptions that match the event and deliver the event to the matched subscribers. Therefore, subscriptions act as filters for subscribers to select desired events out of all possible events published.

The subscription language is made up from a pub/sub schema, denoted as $S = \{A_1, A_2, \dots, A_n\}$, where each A_i is an attribute. Each attribute $A_i \in S$ has a name, a type and a domain, and is denoted by a tuple [name, type, A_{\min} , A_{\max}]. The type could be integer, float, string, etc. The A_{\min} and A_{\max} values define the range of values that can be taken by