The Journal of Systems and Software 86 (2013) 2574–2593

Contents lists available at ScienceDirect

The Journal of Systems and Software

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

Supporting end-to-end quality of service properties in OMG data distribution service publish/subscribe middleware over wide area networks

Akram Hakiri[a,b,*], Pascal Berthou[a,b], Aniruddha Gokhale[c], Douglas C. Schmidt[c], Thierry Gayraud[a,b]

[a] CNRS, LAAS, 7 Avenue du Colonel Roche, F-31400 Toulouse, France
[b] Univ de Toulouse, UPS, LAAS, F-31400 Toulouse, France
[c] Institute for Software Integrated Systems, Department of EECS, Vanderbilt University, Nashville, TN 37212, USA

ARTICLE INFO

Article history:
Received 24 September 2012
Received in revised form 20 April 2013
Accepted 20 April 2013
Available online 15 May 2013

Keywords:
DDS services
Schedulability
QoS Framework
DiffServ

ABSTRACT

Assuring end-to-end quality-of-service (QoS) in distributed real-time and embedded (DRE) systems is hard due to the heterogeneity and scale of communication networks, transient behavior, and the lack of mechanisms that holistically schedule different resources end-to-end. This paper makes two contributions to research focusing on overcoming these problems in the context of wide area network (WAN)-based DRE applications that use the OMG Data Distribution Service (DDS) QoS-enabled publish/subscribe middleware. First, it provides an analytical approach to bound the delays incurred along the critical path in a typical DDS-based publish/subscribe stream, which helps ensure predictable end-to-end delays. Second, it presents the design and evaluation of a policy-driven framework called Velox. Velox combines multi-layer, standards-based technologies—including the OMG DDS and IP DiffServ—to support end-to-end QoS in heterogeneous networks and shield applications from the details of network QoS mechanisms by specifying per-flow QoS requirements. The results of empirical tests conducted using Velox show how combining DDS with DiffServ enhances the schedulability and predictability of DRE applications, improves data delivery over heterogeneous IP networks, and provides network-level differentiated performance.

1. Introduction

Current trends and challenges. Distributed real-time and embedded (DRE) systems, such as video surveillance, on-demand video transmission, homeland security, on-line stock trading, and weather monitoring, are becoming more dynamic, larger in topology scope and data volume, and more sensitive to end-to-end latencies (White et al., 2012). Key challenges faced when fielding these systems stem from how to distribute a high volume of messages per second while dealing with requirements for scalability and low/predictable latency, controlling trade-offs between latency and throughput, and maintaining stability during bandwidth fluctuations. Moreover, assuring end-to-end quality-of-service (QoS) is hard because end-system QoS mechanisms must work across different access points, inter-domain links, and within network domains.

Over the past decade, standards-based middleware has emerged that can address many of the DRE system challenges described above. In particular, the OMG's *Data Distribution Service* (DDS) (OMG-DDS, 2013) provides real-time, data-centric publish/subscribe (pub/sub) middleware capabilities that are used in many DRE systems. DDS's rich QoS management framework enables DRE applications to combine different policies to enforce desired end-to-end QoS properties.

For example, DDS defines a set of *network scheduling policies* (*e.g.*, end-to-end network latency budgets), *timeliness policies* (*e.g.*, time-based filters to control data delivery rate), *temporal policies* to determine the rate at which periodic data is refreshed (*e.g.*, deadline between data samples), *network priority policies* (*e.g.*, transport priority is a hint to the infrastructure used to set the priority of the underlying transport used to send data in the DSCP field for DiffServ), and other policies that affect how data is treated once in transit with respect to its reliability, urgency, importance, and durability.

Although DDS has been used to develop many scalable, efficient and predictable DRE applications, the DDS standard has several limitations, including:

- **Lack of policies for processor scheduling.** DDS does not define policies for processor-level packet scheduling *i.e.*, it provides no

* Corresponding author at: CNRS, LAAS, 7 Avenue du Colonel Roche, F-31400 Toulouse, France. Tel.: +33 634927039.
E-mail addresses: hakiri@laas.fr, akram.hakiri@uha.fr (A. Hakiri), berthou@laas.fr (P. Berthou), a.gokhale@vanderbilt.edu (A. Gokhale), d.schmidt@vanderbilt.edu (D.C. Schmidt), gayraud@laas.fr (T. Gayraud).

0164-1212/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jss.2013.04.074

standard means to designate policies to schedule IP packets. It therefore lacks support for analyzing end-to-end latencies in DRE systems. This limitation makes it hard to assure real-time and predictable performance of DRE systems developed using standard-compliant DDS implementations.

- **End-to-end QoS support.** Although DDS policies manage QoS between publisher and subscribers, its control mechanisms are available only at end-systems. Overall response time and pub-sub latencies, however, are also strongly influenced by network behavior, as well as end-system resources. As a result, DDS provides no standard QoS enforcement when a DRE system spans multiple different interconnected networks, *e.g.*, in wide-area networks (WANs).

Solution approach → End-system performance modeling and policy-based management framework to ensure end-to-end QoS. This paper describes how we enhanced DDS to address the limitations outlined above by defining mechanisms that (1) coordinate scheduling of the host and network resources to meet end-to-end DRE application performance requirements (Kartik and Kyoung-Don, 2007) and (2) provision end-to-end QoS over WANs composed of heterogeneous networks comprising networks with different transmission technologies over different links managed by different service providers that support different technologies (such as wired and wireless network links). In particular, we focus on the end-to-end timeliness and scalability dimensions of QoS for this paper, while referring to these properties simply and collectively as "QoS."

To coordinate scheduling of host and network resources, we developed a performance model that calculates each node's local latency and communicates it to the DDS data space. This latency is used to model each end-system as a schedulable entity. This paper first defines a pub/sub system model to verify the correctness and effectiveness of our performance model and then validates this model via empirical experiments. The parameters found in the performance model are injected in the framework to configure the latency budget DDS QoS policies.

To provision end-to-end QoS over WANs composed of heterogeneous networks, we developed a QoS policy framework called *Velox* to deliver end-to-end QoS for DDS-based DRE systems across the Internet by supporting QoS across multiple heterogeneous network domains. Velox propagates QoS-based agreements among heterogeneous networks involving the chain of inter-domain service delivery. This paper demonstrates how those different agreements can be used together to assure end-to-end QoS service levels:: the QoS characterization is done from the application, and notifies the upper layer about its requirements, which adapt the middleware's service to them using the DDS QoS settings. Then, the middleware negotiates the network QoS with Velox on behalf of the application. Fig. 1 shows the high-level architecture of our solution.

We implemented the two mechanisms described above into the Velox extension of DDS and then used Velox to evaluate the following issues empirically:

- How DDS scheduling overhead contributes to processing delays, which is described in Section 3.2.2.
- How DDS real-time mechanisms facilitate the development of predictable DRE systems, which is described in Section 3.2.4.
- How DDS QoS mechanisms impact bandwidth protection in WANs, which is described in Section 3.3.2.
- How customized implementations of DDS can achieve lower end-to-end delay, which is described in Section 3.3.3.

The work presented in this paper differs from our prior work on QoS-enabled middleware for DRE systems in several ways. Our most recent work (Hakiri et al., 2011a,b) only focused on

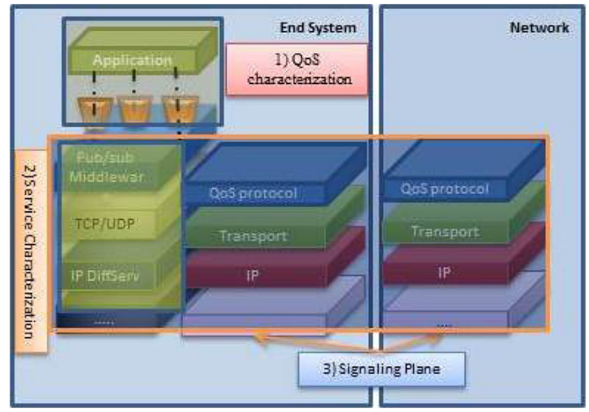

Fig. 1. End-to-end architecture for guaranteeing timeliness in OMG DDS.

bridging OMG DDS with the *Session Initiation Protocol* (SIP) to assure end-to-end timeliness properties for DDS-based application. In contrast, this paper uses the Velox framework to manipulate network elements to use mechanisms, such as DiffServ, to provide QoS properties. Other earlier work (Schantz et al., 2003) described how priority- and reservation-based OS and network QoS management mechanisms could be coupled with CORBA-based distributed object computing middleware to better support dynamic DRE applications with stringent end-to-end real-time requirements in controlled LAN environments. In contrast, this paper focuses on DDS-based applications running WANs.

We focused this paper on DDS and WANs due to our observation that many network service providers allow clients to use MPLS over DiffServ to support their traffic over the Internet, which also is also the preferred approach to support QoS over WANs. We expect our Velox technique is general enough to support end-to-end QoS for a range of communication infrastructure, including CORBA and other service-oriented and pub/sub middleware. We emphasize OMG DDS in this paper since prior studies have showcased DDS in LAN environments, so our goal was to extend this existing body of work to evaluate DDS QoS properties empirically in WAN environments.

Paper organization. The remainder of this paper is organized as follows: Section 2 conducts a scheduling analysis of the DDS specification and describes how the Velox QoS framework manages both QoS reservation and the end-to-end signaling path between remote participants; Section 3 analyzes the results of experiments that evaluate our scheduling analysis models and the QoS reservation capabilities of Velox; Section 4 compares our research on Velox with related work; and Section 5 presents concluding remarks and lessons learned.

2. The Velox modeling and end-to-end QoS management framework

This section describes the two primary contributions of this paper:

- **The performance model of DDS scheduling**. This contribution describes the end-system that hosts the middleware itself and analyzes its capabilities and drawbacks in terms of scheduling capabilities and timeliness used by DDS on the end-system and across the network.
- **The Velox policy-based QoS framework**. This contribution performs the QoS negotiation and the resource reservation to fulfill participants QoS requirements across WANs.