Computer Networks 56 (2012) 1215–1235

Contents lists available at SciVerse ScienceDirect

Computer Networks

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

On data dissemination for large-scale complex critical infrastructures

Marcello Cinque [a], Catello Di Martino [a], Christian Esposito [a,b,*]

[a] Dipartimento di Informatica e Sistemistica (DIS), Università di Napoli Federico II, via Claudio 25, Napoli 80125, Italy
[b] Laboratory iTem "Carlo Savy" of the Consorzio Inter-universitario Nazionale per l'Informatica (CINI), Campus Monte S. Angelo, Napoli 80125, Italy

ARTICLE INFO

Article history:
Received 20 May 2011
Received in revised form 2 October 2011
Accepted 18 November 2011
Available online 8 December 2011

Keywords:
Publish/subscribe middleware
Peer-to-peer systems
Reliability assessment
Stochastic Activity Networks

ABSTRACT

Middleware plays a key role for the achievement of the mission of future large scale complex critical infrastructures, envisioned as federations of several heterogeneous systems over Internet. However, available approaches for data dissemination result still inadequate, since they are unable to scale and to jointly assure given QoS properties. In addition, the best-effort delivery strategy of Internet and the occurrence of node failures further exacerbate the correct and timely delivery of data, if the middleware is not equipped with means for tolerating such failures.

This paper presents a peer-to-peer approach for resilient and scalable data dissemination over large-scale complex critical infrastructures. The approach is based on the adoption of epidemic dissemination algorithms between peer groups, combined with the semi-active replication of group leaders to tolerate failures and assure the resilient delivery of data, despite the increasing scale and heterogeneity of the federated system. The effectiveness of the approach is shown by means of extensive simulation experiments, based on Stochastic Activity Networks.

1. Introduction

Large-scale complex critical infrastructures (LCCIs) [1] are emerging as a new paradigm to build future world-wide monitor and control systems (MCSs) [2]. An LCCI consists of an Internet-scale interconnection of heterogeneous, sometimes already existing, systems, glued together into a federated and open system by a data distribution middleware. Concrete examples are the novel framework for Air Traffic Management under development in Europe by EUROCONTROL called "*Single European Sky ATM Research*" (SESAR) [3], and the collaborative effort of the US Department of Energy (DOE) and the North American Electric Reliability Corporation (NERC) called "*North American Synchro-Phasor Initiative*" (NASPI) [4]. In these cases, it is not practical to deploy a dedicated and proprietary network between Air Traffic Management (ATM) entities in SESAR, or Phasor Measurements Units (PMU) in NASPI. Therefore, communication is realized by means of available network infrastructures and by using IP-based protocols. As a concrete example, interaction among ATM entities will be realized on top of Pan-European Network Service (PENS) [5], which aligns with SESAR Implementing Rules and industry standard services by providing a common IP-based network service across the European region by means of several virtual private networks (VPNs) with Gold Class of Service.

The novel requirements imposed by such world-wide interconnected systems of systems cannot be fulfilled adopting the traditional architectural model of MCSs. Currently, large MCSs are conceived as monolithic and "closed world" architectures, where the overall system is fragmented in several islands of control, each focused on an assigned portion of the infrastructure and with limited collaboration with the other ones. Since each fragment takes control decisions without considering the state of

* Corresponding author at: Dipartimento di Informatica e Sistemistica (DIS), Università di Napoli Federico II, via Claudio 25, Napoli 80125, Italy. Tel.: +39 081 7683874/676770.

E-mail addresses: macinque@unina.it (M. Cinque), catello.dimartino@unina.it (C. Di Martino), christian.esposito@unina.it (C. Esposito).

1389-1286/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2011.11.016

the others, the final control decisions are not optimal. Therefore, the novel LCCI perspective pursues the possibility to let fragments orchestrating their control decisions by an intensive information sharing conveyed by the Internet.

However, the federated architectural model envisioned for LCCIs pushes the frontiers of current technologies by posing new challenging requirements:

- High scalability: the ultra large scale of the infrastructure in terms of generated traffic load or interconnected entities must not compromise the dissemination quality;
- Interoperability: heterogeneous entities must be able to communicate among each other;
- Resiliency: messages must be delivered to all the interested destinations, even in the presence of node and/or link failures [6].

It is still unclear what underlying data distribution model should be adopted in LCCIs, since none of the current solutions are able to satisfy all the mentioned requirements. Today, a promising solution to obtain scalability is represented by middleware infrastructures adopting the publish/subscribe interaction model [7], characterized by natural decoupling properties among interacting parties. At the same time, interoperability issues are being faced by adopting standardized solutions, such as the OMG *Data Distribution Service* (DDS) for publish/subscribe services [8], as demonstrated by its track record of industrial deployments in mission- and business-critical systems. The *Real Time Publish/Subscribe* (RTPS), provides a standardized message exchanging protocol in the DDS standard. For instance, EUROCONTROL has selected DDS as reference technology for the SESAR project. Despite its advantages, DDS is not a viable solution for LCCIs since it is not able to guarantee both scalability and resiliency in the context of large-scale federated systems. In fact, RTPS presents the following issues: (i) it uses an *Automatic Repeat reQuest* (ARQ) scheme for guaranteeing data delivery resiliency by means of message buffering and retransmission, which presents scalability issues, such as implosion and exposure, and does not provide full resiliency due to buffer overflow problems [9]; and (ii) it adopts a decentralized unbrokered architecture, based on IP Multicast, which is known to exhibit severe deployment limitations over the Internet [10,11].

The contribution of this paper is twofold. First, we present TODAI (Two-tier Organization for DAta dissemination Infrastructure), a novel data dissemination scheme for DDS-compliant middleware, which addresses the mentioned issues, and contributes to the on-going discussion about the improvement of RTPS. Second, we develop a set of detailed performability models [12] for Publish/Subscribe LCCIs, used to compare TODAI and RTPS in terms of their performance and resiliency to node and link failures. The paper extends our previous research [13] providing further details on TODAI, proposing the performability models and the failure assumptions defined to conduct the analysis, and providing additional experimental results.

The proposed approach is based on a peer-to-peer (P2P) model, which is able to provide high data delivery resiliency in Internet-scale infrastructures, such as LCCIs, while inheriting the attractive scalability properties of publish/subscribe services and interoperability characteristics of the DDS. Specifically, the proposed paradigm (i) adopts a *super-peer* architecture, to handle the LCCI federated structure; (ii) it implements a *semi-active replication strategy*, to reduce the probability that a group of peers is unreachable due to node crashes; and (iii) it uses an *epidemic algorithm* to implement a proactive and reliable Internet-scale multicasting service. In fact, it has been illustrated how epidemic forwarding can be used to build scalable and reliable communication systems [14]; while, it has been empirically proved that epidemic multicast causes balanced overhead distribution among receiving peers and it is scalable as group size, publishing rate and network failure rate increase [15].

TODAI and RTPS are modeled using the Stochastic Activity Networks (SANs) [16] formalism. The proposed models allow to evaluate several figures of interests such as: (i) super-peer availability, (ii) resiliency to node crashes and link failures during the data delivery process, (iii) dissemination latency, and (iv) dissemination overhead, while varying several parameters, such as the failure rate, the gossiping fan-out, the data publishing rate, and the number of nodes. For instance, the proposed models allow us to point out that TODAI is able to deliver a resiliency level up to 99.999% over a year, while keeping a delivery latency of 87 ms, against a resiliency of 99.9% and a latency of 62 ms in the case of RTPS, in the same simulated scenario. Finally, the performed simulations allow us to investigate the limits of both solutions, and they confirm the limitations of RTPS when applied to LCCIs, in terms of buffer overflow issues due to the adopted ARQ scheme.

The paper is structured as follows: Section 2 introduces assumptions and middleware requirements for LCCIs and it includes a description of the background, with a brief presentation of the RTPS protocol; Section 3 describes the problem statement, by highlighting the limitations of RTPS when applied to LCCIs; Section 4 describes in details the proposed approach; Section 5 illustrates key concepts of SANs and presents the models we have developed to perform the simulation based performance analysis, while Section 6 describes the obtained results; Section 7 presents the related work; last Section 8 concludes the paper with final remarks and future directions for the presented work.

2. Background and requirements

2.1. LCCI

An LCCI, depicted in Fig. 1(a), is composed of several systems interconnected by means of wide-area networks, such as the Internet; therefore, it represents an example of a *Ultra Large Scale* (ULS) system [17]. LCCIs represent a solution to overcome the unsuitability of traditional architectures, and to satisfy the urgent need for a more integrated control architecture. In fact, the federation that