Facing reliability requirements for timely information sharing in future crisis management systems

Marcello Cinque, Domenico Cotroneo, Mario Fiorentino
Laboratorio ITEM Carlo Savy - Consorzio Interuniversitario Nazionale per l'Informatica
Via Cinthia 21, 80126, Naples, Italy. Contact author: mario.fiorentino@consorzio-cini.it

***Abstract*—During the management of a crisis, the reliable and timely sharing of information can help to save lives. This fast abstract identifies the main reliability requirements for federated crisis management systems and proposes DESTRIERO, a novel architectural solution to face them.**

I. Rationale

Today, people are threatened by crisis situations, both from natural origin, such as earthquakes and floods, or caused by malicious acts, such as terrorist attacks. In such situations emergency management, but also Post-Crisis Damage and Needs Assessment (PDNA) and Reconstruction and Recovery Planning (RRP), is usually coordinated by local authorities or dedicated civil protection organizations, with the support of a variety of different national and international relief organizations acting relatively autonomously.

For these reasons, there has been a growing interest in Crisis Management Systems (CMS) focussing on filling existing interoperability gaps, improving collaborative decision making among involved actors, approaching the problem of the integration of heterogeneous information sources and improving data visualization and aggregation through GIS technologies and sensors' data. Despite the interest of the research community, there is still a lack in proper technologies that could facilitate timely and reliable cross-boarder information sharing for joint situation awareness and cross-agency resource management.

All the reliability concerns are typically demanded to the underlying middleware used to enable the interoperability among systems. In several commercial solutions, currently used to mange real crisis scenarios, no evidences have been found regarding the reliability mechanisms and approaches put in practice in order to tolerate faults. These federated systems usually focus on requirements about crisis coordination and resource allocation, with limited attention on reliability concerns. For a practical example, EDEN, a CMS realized by the SAHANA foundation (http://sahanafoundation.org), is entirely focussed on the optimization of functionalities for the management of the crisis providing a minor attention on security aspects (mainly authorization and authentications controls) rather than approaching reliability concerns. Another key example is ResilienceDirect, an online private network which enables civil protection practitioners to work together across geographical and organizational boundaries, during the preparation, response and recovery phases of an emergency. In this case, the solution only deals with the resilience of the telecommunications adopted, via the High Integrity Telecommunications System (HITS) that provides a resilient communications backbone between Strategic Co-ordination Centers (SCCs) in police force areas across England and Wales, and central government crisis management facilities. It is evident that due to the importance of such complex federated systems, reliability became a primary key technological aspect to be taken into account when developing a CMS.

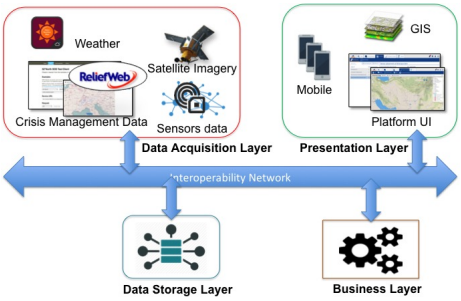

Fig. 1: Schematic Architecture of a federated CMS.

In this paper we identify a set of reliability requirements to be satisfied by a CMS, and, based on these, we propose a novel solution for reliable and timely information sharing in federated crisis information systems. The tool here proposed is based on hybrid (active/passive) replication, coupled with a lightweight reliable multicast communication support, and it is under development within the EU-funded project named DESTRIERO. Overall, the project aims to (i) offer a faster and better damage assessment for planning and monitoring of progress of recovery, (ii) facilitate user access to visualize the dynamic "common operational picture", during the planning and reconstruction period and (iii) improve the management of the information in relation to PDNA and RRP.

II. Reliability Requirements

Figure 1 shows the entities composing a typical node in a federated crisis management system. The schema identifies for each of the node (i) a presentation layer for data visualization, (ii) a business layer in charge of processing shared information, (iii) a data acquisition layer whose responsibility mainly concern the gathering of data from different information sources (*e.g.* sensors displaced on the field, meteorological information, satellite and aerial imagery etc.), (iv) a Data Storage Layer for persistency operations on acquired data (v) and, finally, an interoperability network for boosting data sharing among the nodes. In such a complex scenario, faults can be activated in processes running on each of the node of the federation, or they can occur on the network. In this context, we identify the following key high level requirements to assure reliable and timely information sharing in CMS federations:

R1: CMSes have to be designed with a high degree of redundancy so that, in case of crash of a node, another one is able to replace it.

R2: Proactive recovery means must be taken into account for facing network faults and for assuring that messages are received in the correct manner, without affecting delivery times.

R3: Fault-tolerance techniques must not affect the overall performances of the system, assuring near real-time information sharing among all the organizations that are taking part in the field operations.

R4: The system has to be monitored to assure that it is able to apply mitigation actions in case it is not behaving as expected.

III. The Proposed Solution

To satisfy the above mentioned requirements, the proposed solution, named DESTRIERO, has been designed (i) to assure a high degree of nodes replication without affecting system performance for all the information that is distributed among the organizations (to satisfy R1 and R3 in the list above), (ii) with a proper resilient multicast communication protocol to assure that information reaches its destination in a timely and reliable way (for R2). Replication and multicast are two key functionalities of the DESTRIERO node, shown in Figure 2. The node is conceived as a software artifact to be used by the CMS layers shown in Figure 1 (such as, presentation and/or 3rd party systems for data acquisition) to access the interoperability network through so called methodology services. These, in turn, invoke core services, in order to transparently handle replicas and use resilient multicast services. Finally, the CMS can be seen as a federation of distributed (and transparently replicated) DESTRIERO nodes.

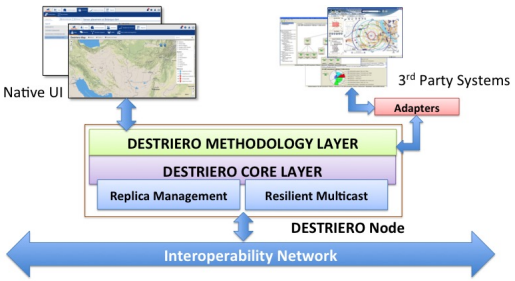

Fig. 2: Internals of a DESTRIERO node.

Two replication schemas have been investigated: passive and active. In a passive replication schema there is a single master replica, called *primary replica*, whose responsibility is to receive requests from clients and return responses. In this schema, all the backup replicas should interact with the primary one in order to maintain consistent their states. Although a low computational processing, this schema has one main drawback [1], in case of failure of the primary replica another replica, among the backup ones, should be elected. In this case, if the primary replica crashes before sending an answer to its clients, they will time-out. The usage of this schema is therefore discouraged in the case of critical systems. In a different manner, in an active replication schema all the replicas have an independent lifecycle but the schema ensures that all of them receive requests in the same order to maintain aligned their states. This has the advantage to ensure a low response time when clients submit their requests but with a high resources consumption. Moreover, requests should be processed in a deterministic way [1].

The adopted redundancy solution is based on a hybrid schema in which active and passive replication have been taken into account as a tradeoff to satisfy R3. In particular, when a write request is submitted all the nodes process and store the data so to maintain states aligned. When the write operation finishes an entry in a log file is inserted with all the operations that have been processed. In case of a node failure (*e.g.* due to a crash) state can be restored exchanging messages containing missing information among the nodes in the federation.
The proposed hybrid replication schema is then coupled with multicast communication, for what concerns R2. In literature there are two main approaches for tolerating faults in a multicast communication infrastructure: spatial and temporal redundancy. In the first case messages are sent over multiple paths through the network and error-correction codes are applied to check that no losses have affected the communication. The temporal redundancy instead, makes use of proper time-outs for checking losses on the communication channels and retransmissions of lost messages are done in such cases. Given the strict requirement of crisis information systems to timely distribute new data among all the organizations our attention has been focussed on spatial redundancy algorithms such as Path Redundancy (PR) and Forward Error Correction (FEC) [2]. Due to well known drawbacks related to FEC and PR [3], [4] (*e.g.* PR is effective only if path diversity is guaranteed) we focussed our attention on a hybrid approach combining the FEC and PR algorithm in a unique communication protocol defined in [4] where diverse path trees are built selecting paths from any new node to its parent that expose the lowest measure of diversity. Paths from parents to their children maintain a value of diversity that is closer to the one they had before a new child node is added to the tree. A preliminary study has demonstrated that this protocol is able to reach a level of diversity near to the intrinsic diversity of the topology at a network level assuring a robust delivery of multicast messages with nearly no overhead.

IV. Future Work

Future activities will encompass thorough experimentation of the proposed solution on the real DESTRIERO prototype, when available. Monitoring techniques will be included as well so to timely identify failures and their causes (to satisfy R4).

Acknowledgment

This work has been supported by the European Commission under the Collaborative Project "A DEcision Support Tool for Reconstruction and recovery and for the IntEroperability of international Relief units in case Of complex crises situations, including CBRN contamination risks" (DESTRIERO, http://www.destriero-fp7.eu/ - Grant agreement no: 312721).

References

[1] Xavier Dfago, Andr Schiper, Nicole Sergent, *Semi-Passive Replication*, Proceedings of the 17th IEEE Symposium on Reliable Distributed Systems, 1998.

[2] L. Rizzo and L. Vicisano, *RMDP: an FEC-based reliable multicast protocol for wireless environments*, ACM SIGMOBILE Mobile Computing and Communications Review, 1998.

[3] C. Perkins, O. Hodson, and V. Hardman, *A Survey of Packet Loss Recovery Techniques for Streaming Audio*, IEEE Network 12, 1998.

[4] C. Esposito, D. Cotroneo, and A. Gokhale, *Reliable publish/subscribe middleware for time-sensitive internet-scale applications*, Proceedings of the Third ACM International Conference on Distributed Event-Based Systems (DEBS), 2009.