Towards a resilient e-health system for monitoring and early detection of severity in hospitalized patients during a pandemic

Rodrigo S. Cañibano[1 [0000-0001-6992-5421]], Santino Castagno[1], Mariano Conchillo[1 [0000-0002-4241-3511]], Guillermo Chiarotto[1], Claudia Rozas[1 [0000-0002-9876-4724]], Claudio Zanellato[1], Cristina Orlandi[2], and Javier Balladini[1 ⋆ [0000-0002-9769-7830]]

[1] Universidad Nacional del Comahue, Neuquén, Argentina,
{rcanibano, stefano.castagno, mariano.conchillo, claudia.rozas, claudio.zanellato, javier.balladini}@fi.uncoma.edu.ar,
guillermo.chiarotto@est.fi.uncoma.edu.ar
[2] Hospital Área Programa General Roca, Río Negro, Argentina,
cristina.orlandi@gmail.com

Abstract. During a pandemic, the traditional methodology of doctors and nurses must be optimized to allow the care of a greater number of patients without reducing the quality of care. The triage method select and classify patients in different levels of severity to achieve a correct allocation of care. Although triage can be performed through manual calculations, the process would be error-prone and considerably increase the workload on health personnel, limiting the potential benefit of the method. Our research group designed and developed *COVINDEX*, a computer system aimed at monitoring and early detection of severity in hospitalized *COVID-19* patients. It implements an existing early warning system for patients hospitalized in general wards. Our system has a complex distributed architecture that is easy to deploy and maintain, and it provides high resilience to network and hardware failures. This work discusses some details of the system architecture.

Keywords: Early severity detection · e-health monitoring · fault tolerant distributed architecture.

1 Introduction

The *SARS-COV-2* pandemic has shown the health system's lack of preparation to face this type of situation. Health centers were overwhelmed by the sudden increase in demand. Specialized human resources became scarce, aggravated by contagions from the staff themselves, which led to them temporarily ceasing to provide services.

In this context, the traditional methodology of doctors and nurses must be optimized to allow the care of a greater number of patients without reducing the

⋆ Corresponding author.

quality of care. To organize and optimize resources in this critical situation it is possible to use the triage method. This method consists of the selection and classification of patients in different levels of severity to achieve a correct allocation of care. Triage was applied early in China [3] during the health emergency of *COVID-19* patients through the adaptation of one of the best known and validated *Early Warning Systems (SAT)* [2]. This allows staff performance to be improved, reducing controls in stable patients, increasing them in seriously ill patients, and reducing unexpected hospital mortality in general hospital areas.

Although triage can be performed through manual calculations, the process would be error-prone and considerably increase the workload on health personnel, limiting the potential benefit of the method.

To improve the quality of care [4], our research group designed and developed *COVINDEX*, a computer system aimed at monitoring and early detection of severity in hospitalized *COVID-19* patients. It implements the early warning system published in [1] for patients hospitalized in general wards, based on the knowledge of expert health personnel. It also allows the collection of data for later improvement of the prediction using machine learning techniques.

COVINDEX is a clinical decision support system that can be interconnected with a digital medical record. Our system has a complex distributed architecture that is resilient (the system continues to provide services in the event of network and device failures), and simple to deploy and maintain (it does not require the use of standard servers in the hospitals). To the best of our knowledge, no other systems exist with this properties.

2 Current development

COVINDEX performs an automatic real-time analysis of the clinical data of patients: vital signs, comorbidities, laboratories, and X-ray reports of the patient. Patients are classified according to four severity levels (according to what is expected to develop after 24 or 48 hours): low, moderate, high and critical (needs intensive care). The system issues alerts to health personnel if there is a change in severity, and organizes and coordinates the work of doctors and nurses in a patient-centered care process.

2.1 Architecture overview

The architecture design must meet the following characteristics:

- Take into account that hospital wards in Argentina are highly diverse in terms of the availability and reliability of their Internet connection. Health personnel must be able to record data related to patients in the ward where they are assigned, and the system must provide early warnings, even in the event of an interruption in the connectivity of local devices with system components located outside the hospital. Data consistency must be maintained, even after a communication interruption between system components.

– It should not require a standard server in the hospitals to simplify the deployment and maintenance of the system.
– A global and historical repository of data is required for the generation of statistical information and the extraction of new knowledge.

The system architecture can be summarized in three computing levels:

Edge: mobile devices used by doctors and nurses. They allow data to be entered into the system and the results to be observed. It can be directly connected to the fog level through a local area network or cloud level through the Internet. Cloud level connection is preferred to reduce overhead and power consumption of fog level devices.

Fog: mobile devices located inside the hospital. They maintain the primary copy (master copy) of a part of the global database. If connection to the cloud level is lost, it allows edge devices to be served through direct connections established over the local area network. This device can also be simultaneously used as an edge level device.

Cloud: maintains a complete and historical copy of all fog level device databases, and is a primary copy of hospital administrative data, and data shared between all hospitals. Also, it can reduce the overhead of fog level devices by taking care of edge devices.

2.2 Detailed description of the architecture

Three different types of domains are defined (shown as circles in figure 1):

Users domain: which contains the data of all users that are relevant to all hospitals, such as username, email, and the password data. There is only one instance of this domain.

Hospital domains: each one contains the specific data of a particular hospital, such as the list of users who are employees of the hospital, wards and beds. There is one instance of this domain per hospital.

Island domains: each one contains data specific to a group of wards of a particular hospital, such as vital signs, laboratories, and X-ray reports of patients. There are multiple instances of this domain on each hospital. Which wards belong to an island are defined by the connectivity between them and the capabilities of the fog device that manages them.

Each domain has a single leader device that holds the primary copy of the data, and client devices that hold replicas. When a device wants to store a data record for a specific domain, it must send a request to the leader of that domain. When the leader stores the record on its database, it sends a copy to all client devices in that domain, including the device that made the request. The leader might process the inserted record in conjunction with the historical data of the patient to generate new records (for example, related to triage).

To keep the devices up to date even after the connection to some other system component is lost, each domain instance has a virtual clock (represented as an

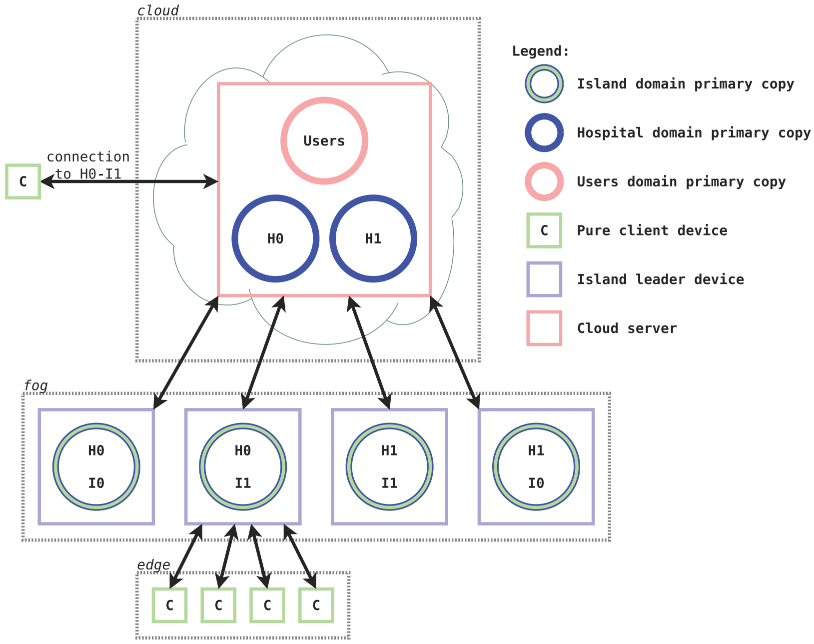

Fig. 1. Detailed architecture of the system

integer) that counts the number of transactions (insertions of new records). We call this virtual clock *sync_id*. Each record has the *sync_id* of the domain as an additional key. To keep a complete history of the inserted data, the records are never modified or deleted, and the most current version of the record is distinguished by the value of the logical clock. When a client device reconnects, it can ask the domain leader device for records with *sync_id* greater than the last one the device has stored.

Leader devices can be clients (maintain replicas) of each other, which allows them to serve as caches for other domains. Furthermore, they can be used by clients as bridges to leaders from other domains. This allows a leader device to reduce the workload on leaders with more limited capabilities and, if the network becomes disconnected, it can provide clients (still connected to the leader) with the latest data records stored before failure (properly indicating that this data may be out of date).

The current system defines two types of leaders and a pure client type (shown in figure 1 as squares):

Cloud server: is the leader of the users domain and of all the hospital domains, it is also a client of all the island domains. It maintains a complete and historical replica of all domains.

Island leader devices: these are fog level devices. Each one is a leader of an island domain, a client of a hospital domain, and a client of the users domain.

They are physically located within the hospital. Upon interruption of connectivity to the cloud server, clients can continue to operate by interfacing with the island leader device via a local area network.
In addition, each time a new record of a patient is inserted, a triage is performed (also using historical data of the patient), which determines a level of severity. If the severity of the patient has changed, an alert is issued.

Pure client devices: these are the edge level devices. They are clients of a single island domain, a single hospital domain, and the users domain. They keep a replica of these domains. In the event of an island leader device failure, one of the pure client devices can be promoted to island leader device.

3 Next steps

Since the database is distributed within the hospital and the devices can be easily replaced, the current system is highly resilient to connectivity and device failures (island leaders and pure clients). Meanwhile, the cloud server centralizes its services. To increase fault tolerance of this component, a spare node should be installed and a mechanism that allows the service to be automatically restored in the event of a failure should be implemented.

The *Ministry of Health of the Province of Río Negro, Argentina* is interested in deploying the system in its public hospitals. The system will be expanded for the following purposes: (1) Timely referral of critical patients from hospitals without intensive care units; (2) Unify the care process for hospitalized patients, and increase the quality of care; (3) Obtain statistical information from all hospitals in the province for decision making; (4) Extend the use to Non-COVID-19 patients.

Data is currently recorded into the system manually, and therefore data flow is limited. However, it may be desirable in the future to connect sensors at the edge level. For this the architecture will need to be tailored to support this new feature.

We still need to evaluate the system performance under a realistic load and improve the security of the data stored on edge and fog devices.

References

1. Huespe, I., Carboni Bisso, I., Gemelli, N.A., Terrasa, S.A., Di Stefano, S., Burgos, V., Sinner, J., Oubina, M., Bezzati, M., Delgado, P., et al.: Design and development of an early warning score for covid-19 hospitalized patients. MEDICINA (Buenos Aires) **81**(4) (2021)
2. Lee, J.R., Kim, E.M., Kim, S.A., Oh, E.G.: A systematic review of early warning systems' effects on nurses' clinical performance and adverse events among deteriorating ward patients. Journal of patient safety **16**(3), e104–e113 (2020)
3. Liao, X., Wang, B., Kang, Y.: Novel coronavirus infection during the 2019–2020 epidemic: preparing intensive care units—the experience in sichuan province, china. Intensive care medicine **46**(2), 357–360 (2020)
4. Rocco, C., Garrido, A.: Seguridad del paciente y cultura de seguridad. Revista Médica Clínica Las Condes **28**(5), 785–795 (2017)