Mobile Triage Management in Disaster Area Networks Using Decentralized Replication

Dominik Meißner, Benjamin Erb*, Rens van der Heijden*, Kristin Lange‡, Frank Kargl*
*Institute of Distributed Systems
‡Institute of Psychology and Education, Dept. Human Factors
Ulm University, Germany
{firstname}.{lastname}@uni-ulm.de

ABSTRACT

In large-scale disaster scenarios, efficient triage management is a major challenge for emergency services. Rescue forces traditionally respond to such incidents with a paper-based triage system, but technical solutions can potentially achieve improved usability and data availability. We develop a triage management system based on commodity hardware and software components to verify this claim. We use a single-hop, ad-hoc network architecture with multi-master replication, a tablet-based device setup, and a mobile application for emergency services. We study our system in cooperation with regional emergency services and report on experiences from a field exercise. We show that state-of-the-art commodity technology provides the means necessary to implement a triage management system compatible with existing emergency service procedures, while introducing additional benefits. This work highlights that powerful real-world ad-hoc networking applications do not require unreasonable development effort, as existing tools from distributed systems, such as replicating NoSQL databases, can be used successfully.

CCS Concepts

•**Networks** → **Mobile ad hoc networks;**
•**Human-centered computing** → *Mobile computing;*

Keywords

triage management, disaster area network, field study, multi-master replication

1. INTRODUCTION

Emergency services are presented with a plethora of challenges when large-scale disasters take place. Large-scale disasters are characterized by a confusing and chaotic situation, a high level of uncertainty, and missing knowledge about casualties. Emergency responders arriving at the scene should proceed in a calculated and systematic way. Triage management has become one of several established procedures for emergency forces to counter the chaos of a disaster. Triage is the process of assigning casualties with priorities based on condition and severity, in order to allocate limited resources.

For emergency services, triage represents the first step casualty management, which also includes subsequent treatments and off-site evacuations. Triage tags are special labels for managing this entire process, which contain treatment and status information about the patient it is issued to. While this approach is straightforward and easy-to-use, the sole usage of a physical tag comes with a major drawback—the locality of information. Information is physically bound to the patient, which works well for the individual treatment, but makes it difficult to give an overview of the entire operation. In a pen-and-paper approach, this results in the administration of additional lists with information gathered from the triage tags, and frequent manual copying of this information. Outdated and stale information or incomplete reports can skew the insights for operational control and yield misconceptions and even wrong decisions. Radio dispatches are also not well-suited for the centralized collection of a complete situation report due to capacity limitations. Another limitation of triage tags is that they can easily be lost or modified by non-cooperative patients.

Previous research in mobile ad-hoc networks has not only addressed infrastructural considerations for so-called disaster area networks [4, 6, 7]. These projects also demonstrated potential applications for disaster relief, such as the use of digital triage tags, and involved the design of highly customized mobile systems and network equipment.

In this paper, we want to share our experience with a different approach. We aim to determine whether research has progressed to a state where an augmented triage system can be designed and implemented using only off-the-shelf hardware components and open software components. Therefore, we cooperated extensively with local units from the German Red Cross to design and test such a system. We collected their requirements for an augmented triage management for large-scale disaster scenarios, as summarized in Section 2. Subsequently, we designed and implemented a distributed disaster relief application, based on commodity hardware and open source software (Section 3). To corroborate our hypothesis, we took part in a joint field exercise and evaluated our system against the traditional pen-and-paper approach in a mixed-methods evaluation, described in Section 4. We compare this to earlier work in Section 5 and summarize the experiences and lessons learned in Section 6.

CHANTS'16, October 03-07 2016, New York City, NY, USA

ISBN 978-1-4503-4256-8/16/10...$15.00

DOI: http://dx.doi.org/10.1145/2979683.2979689

2. REQUIREMENTS

When augmenting standardized operations with digital assistance, it is important to adhere to the governing processes already in place. A digital triage tool should primarily supplement the existing communication and information infrastructure for the rescue forces. This includes a faster dissemination of information and the provision of an additional, independent information channel, while still maintaining the general idea of triage management. Therefore, we briefly review existing processes before deriving our requirements.

2.1 Disaster Response Processes

The following gives an overview of standardized German disaster response procedures for a small to medium sized mass casualty incident. The basic process of a disaster response involves three phases: setup, triage and treatment, and tear down. In the first phase, first responders and emergency rescue services relocate patients to a safe distance; once disaster response teams arrive, they set up the treatment area, while patients in the waiting area receive preliminary triage and if resources allow, registered with name and where the patient was found.

The main focus of our work is the second phase, where triage is performed; in this phase, patients are moved to the treatment area, receive initial treatment, are registered, and distributed to nearby hospitals as soon as they can be transported and a hospital has resources available. Patients are prioritized for treatment and transport by their triage category, which indicates the severity of their injuries. After all patients are processed and transported to a hospital, the tear down phase begins. In this phase, the treatment sites are torn down; here it is important to preserve collected information about patients, in order to inform relatives.

Based on these procedures, the area is separated into four operation sections: incident site, casualties treatment area, transport zone, and hospital zone (see Figure 1). A hierarchical structure [9] is used to organize the operation, consisting of operational control, and several operational sections led by a section leader and one or more quick response units. Operational control consists of a chief emergency doctor and an organizational commander responsible for the entire operation. The section leaders coordinate leaders of the quick response units, which typically consist of 8 to 16 aid workers with medical training and various specializations.

During and after the incident, it is very important that current records are available on where patients are, what their current status is, what medication they have received, and where they are transported to. This information, as well as the result of the triage process that classifies the patient, is usually documented on a paper triage tag, which is carried on the patient's person, often with some hand-written notes. One of the primary use cases of our application is to replace this tag, allowing easy dissemination of patient information within the incident area.

2.2 System Requirements

Network Aspects. The system should be independent of public infrastructure, which can fail due to a natural or human-made disaster. Additionally, cellular networks are vulnerable to congestion in accidents scenarios, where both emergency workers and affected persons communicate at the same time. The system should be tolerant of delay and disruptions within the network, maximizing the availability of information. The system should be self-healing; able to recover from node and transmission failures. Finally, extending the system with additional nodes and integration with different response teams should be possible; the network architecture should be scalable.

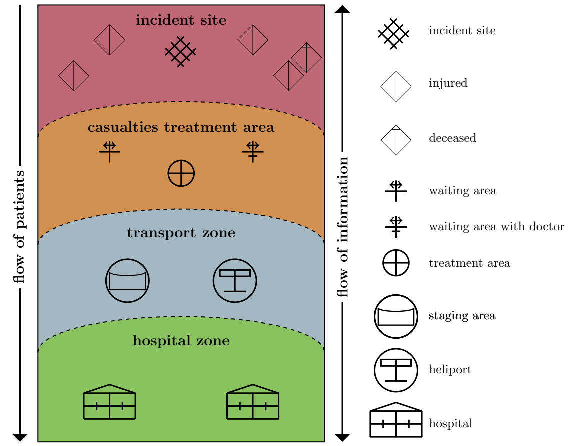

Figure 1: Segmentation of the incident area, based on [9]. The tactical symbols in this figure are used according to the recommendations of the German civil protection [3].

Deployment Aspects. The system should be rapidly deployable, because disastrous events require fast response and occur without warning. Because the system will be essential for effective triage management, configuration and deployment time should be minimal [8]. In addition, the system should be cost-effective enough to be deployed by regular local Red Cross units.

Features. **Operational overview**. For proper decision-making, the superiors and operational control need information about every tactical unit, including their location, status, and assigned task. Similar information should be recorded about vehicles and other assets. This and other relevant information should be presented on a topographic map for a situational overview of the entire operation. This allows operational control to efficiently coordinate between different relief units, redirecting and reorganizing them as more information about the incident becomes available.

Operation documentation. The operation documentation is the chronological attestation of all information and every decision made by operational control, required in case of subsequent inquiries to the operation.

Registration of affected people. During and after the triage process, it is important to keep track of the patients, so that their relatives can easily receive this information. In existing systems with triage tags, this process is known to be cumbersome and incur a large delay. If patients are transported to hospitals, it is often difficult to track them down afterwards: the system should avoid this weakness through central administration.

3. CONCEPT & IMPLEMENTATION

To meet the requirements from Section 2, we identify an ad-hoc networking approach as the most suitable, due to the flexible and self-healing and self-managing nature of these networks. For robustness and several of the required features, we replicate information to as many different nodes as possible, leading us to a single-hop communication infrastructure, i.e., without explicit routing. This is an integral difference to many related systems, which either rely on a central sink for all information or allow nodes to subscribe to specific objects in a publish-subscribe manner [6]. Our approach directly meets the needs of disaster response procedures, where any aid worker may require access to any piece of information, such as patient data, at any time.

Each aid worker will have access to a node consisting of a local data storage, a *disasterBeacon* service and a *disasterApp* service, as shown in Figure 2. Our prototype implementation consists of a Raspberry Pi equipped with an external battery pack (10 000 mA h for about 20 hours of power), two WLAN USB dongles to connect to the ad-hoc network and to create an access point. The access point is used to connect a mobile device that can display web pages, which allows the user to interact with *disasterApp*.

Data storage is responsible for storage and consistency of collected information, such as patient data. This information is bundled in self-contained documents that are given a globally unique, immutable identifier [5]. By only replicating changes to these documents, we can protect against loss of data. Further challenges arising from our peer-to-peer architecture can be solved through the use of multi-version concurrency control and implementing application-based conflict resolution. The advantage of using a replicated local data storage is that our system is tolerant of partitions and has guaranteed availability [10]. In our prototype, we use the eventually consistent, document-oriented database CouchDB, which has built-in support of multi-master replication and concurrent updates. The unidirectional master-master replication system of CouchDB [1], can be used to synchronize two replicas of the same database—locally or remotely. CouchDB performs replication when the state of replicas diffe by assigning a new version number and sending a stream of modification and deletion events over HTTP. Each of these events is atomic, which ensures the database stability if the connection fails.

disasterBeacon is an ad-hoc service application, which is responsible for neighbor discovery and replication target selection. It is implemented using the event-driven, non-blocking JavaScript runtime Node.js. Every 10 seconds, a beacon in form of a UDPv6 all-nodes-multicast is transmitted. Received beacons are used to create a neighbor table with the latest known remote data storage states. Periodically, disasterBeacon weights the neighbors and selects the optimal replication partner from this table, and triggers a replication at the local storage. The weight of a neighbor is computed to be $\mathrm{lsr}(x) \cdot \alpha_{\Delta t}(x) \cdot 2^{-\mathrm{utd}(x)} \cdot 2^{\epsilon(x)-1}$, where $\mathrm{lsr}(x)$ represents the time since the last successful replication to x, $\alpha(x)$ is ratio of beacons received from x in the last Δt minutes, $\mathrm{utd}(x)$ represents the up-to-dateness of x and $\epsilon(x)$ is the ratio of failed replications with x. Replication occurs in direct communication range only, i.e., no routing is performed.

disasterApp is the user interface component of our system, which is a web application running on our network nodes. It is responsible for submitting changes the user makes to the local data storage and providing direct access to the information required by aid workers. The web application, which we select due to its platform independence and suitability for user interface design, is served by CouchDB and developed using a mobile UI framework based on AngularJS and Bootstrap 3. Whenever the database receives an update, it pushes this information to the web application that uses AJAX to subscribe itself to a long polling based changes feed of CouchDB. disasterApp therefore provides near real-time information updates to the user.

Most information that was traditionally gathered using different paper-based lists, forms, and tags can be gathered with the patients module. The input options are derived from the traditional system and was improved at some points. For example, a user can provide a significantly more exact short diagnose for a patient (see Figure 3). Input fields, which could not be filled out intuitively in the traditional solution, were improved with drop down selection fields that show the user all possible input options.

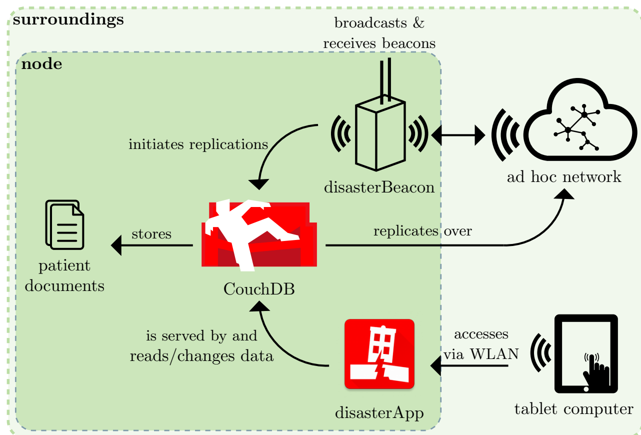

Figure 2: Node architecture in relation to its surroundings.

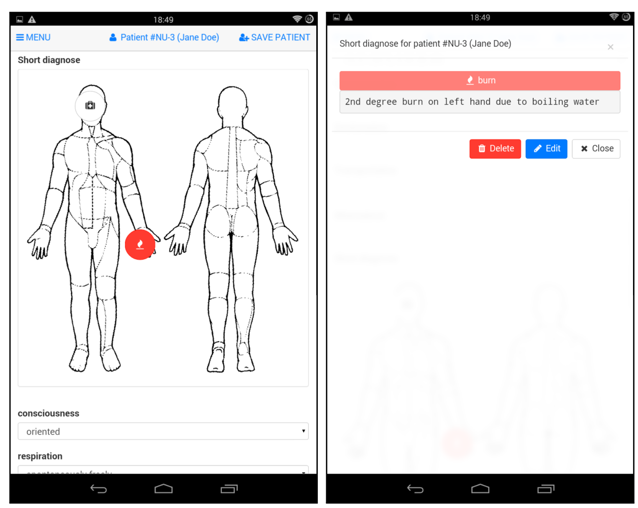

Figure 3: Screenshots of a patient's short diagnose.

4. FIELD STUDY

In order to evaluate the design and implementation of our application in a realistic setting, we cooperated with two local district chapters of the German Red Cross during a joint

Figure 4: disasterApp during the field study.

exercise, where two groups engaged in a simulated major incident. The setting was replicated, hence one team used the traditional pen-and-paper solution while the other team made use of our disasterApp solution. We used a mixed-methods approach for our analysis, relying on verbose logging on all nodes, a quantitative empirical user study, semi-structured interviews, and video to document our study.

4.1 Exercise Setting

The exercise has been conducted simultaneously for two groups, one using *disasterApp* and a control group using the traditional pen-and-paper solution. A typical scenario was tested: an incident caused by a sudden thunderstorm during a street festival. Patients have to be transported to the treatment area, since the incident site is not easily accessible. Seven persons (later referred to as patients) were moderately to seriously injured during the incident, who are to be triaged and treated. Each patient with the same injuries was played by two people, who wore make-up reflecting their injuries. They were given a numbered wristband to ensure that both groups have a comparable set of patients.

Exercise Script. Each group has a separate organizational hierarchy. After the arrival on-site, a triage team is sent to the incident site, coordinated by a group leader, who is also responsible for communication with operational command. The digital group is equipped with tablets and special triage tags that only include the patients name, a preliminary triage category, and the patient number; the other group used traditional triage tags. Patients from both groups were delivered to the same triage tent, which was divided into a digital and an analog section, and have been processed independently. After triage, the patients were moved to the digital or analog treatment tent, depending on which group they are part of, where they received initial treatment and their data was collected. As soon as a patient was in a transportable condition and transport was available, they were escorted to the transport area. One casualty was instructed to move off-site and disappear to test the communication between the response team and operational control.

Exercise Report. Our system was sufficiently robust; some nodes were rebooted and seamlessly re-integrated into the network. In the digital group, the missing patient was swiftly recognized, due to the improved overview offered by *disasterApp*, while in the analog group, registration only reported 6 patients. This discrepancy was not responded to by their operational command. However, they did finish treatment and transport slightly faster, which we attribute to the missing patient and time required to input information on tablets. One interesting result was that radio communication was noticeably reduced in the digital group.

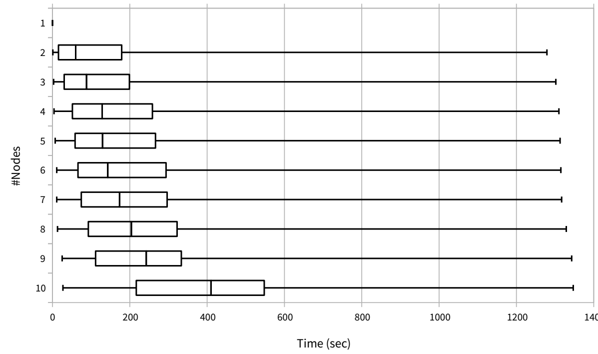

Figure 5: Per-hop distribution of the change propagation times when disseminated through the network.

4.2 Technical Evaluation

To provide an analysis of the overall replication behavior of our system, we performed measurements during the field study, logging the following: beacon broadcast events, beacon reception events, replication attempts, received replication attempts, the outcome of each replication attempt, locally introduced changes, and successful replication events. Each of these log entries includes a time stamp based on the local system time, which we cross-reference to determine directional link probabilities (Table 1), based on the beacon broadcast and reception events. In addition, we compute the time necessary to disseminate a change through the network by finding the shortest replication sequence between each node and the node that introduced the change locally.

Table 1: Fraction of beacons were received per node (senders on top, receivers on the left).

	N1	N2	N3	N4	N5	N6	N7	N8	N9	N10	N11
N1		.47	.45	.37	.44	.44	.45	.21	.23	.16	.19
N2	.44		.89	.72	.81	.84	.85	.44	.45	.46	.36
N3	.45	.92		.75	.85	.87	.87	.44	.43	.45	.39
N4	.39	.80	.83		.85	.86	.83	.52	.57	.55	.40
N5	.46	.92	.88	.85		.93	.92	.51	.52	.50	.39
N6	.46	.93	.93	.84	.91		.94	.49	.50	.51	.38
N7	.46	.93	.92	.86	.90	.95		.48	.50	.51	.35
N8	.42	.91	.87	.90	.92	.90	.91		.89	.81	.44
N9	.42	.82	.77	.88	.84	.84	.84	.77		.83	.42
N10	.24	.70	.67	.75	.68	.71	.70	.64	.75		.31
N11	.46	.84	.88	.85	.80	.85	.80	.64	.57	.50	

Results. Table 1 shows several effects related to the way different nodes were deployed. For example, node 11 was used by operational control, and it was within a vehicle for most of the exercise. Similar effects were caused by high mobility: node 1 was highly mobile and wandered throughout the network, leading to poor connectivity. Nodes 8, 9 and 10 were deployed with first responders, who located casualties while other team members set up tents, which explains why many beacons they sent were not received by the other nodes. All the other nodes were well-connected, because they were in the casualty treatment area throughout the exercise.

The results for propagation of the changes through the network are as expected; in the best cases, the propagation took only 27 seconds for the entire network. The high maxima, up to 23 minutes, can be explained due to individual nodes that were partially in the process of transportation. The median propagation is high, but is significantly faster than standard propagation through hand-written paper.

4.3 Empirical User Study

We conducted a user study to assess the user experience of the rescue forces with our application. The study included a general application assessment, as well as a pre- and post-test for the group using the application during the exercise.

Participants & Procedure. The participants stem from two local German Red Cross groups, which took part in the joint exercise. In total, 24 subjects (7 of which were female) participated in the empirical study, where $n_p = 10$ used the paper version, and $n_t = 14$ used tablets with disasterApp. For team efficacy reasons, we opted against a reshuffling of teams, and randomly allocated the tablet condition to one of the existing quick response units. The age of the participants was comparable in both groups (paper: $M_p = 32.5$, $SD_p = 16.5$; tablet: $M_t = 31.8$, $SD_t = 10.8$). On average, participants had about 10 years experience in emergency medical services, with a minimum of at least 1 year. The tablet group showed mixed experience levels using smart devices. While 100% possessed a smartphone, only 57% of them owned a tablet personally. Average device usage (in hours per week) was also varied strongly (smartphone: $M = 22.5$, $SD = 28.1$; tablet: $M = 7$; $SD = 10.7$).

Before the briefing of the exercise, we gave an introduction into our tablet-based application and provided a short, hands-on tutorial for all participants of both teams. After the tutorial, we asked all participants to fill in a survey about their preferences towards a future usage of such a tablet-based system. We also asked them for an estimation of usefulness in comparison to the traditional paper-based version. Furthermore, we asked for an estimation whether the system will improve patients' care. These four attitudes were covered by a set of items using 5-point Likert scales ranging from absolute disagreement to complete agreement. Some items were reversed and required recoding prior to the analysis. In order to control for confounding variables, we also asked the participants to fill out an additional TA-EG [2] questionnaire, an established German inventory measuring general affinity with technical devices.

After the exercise, but before tear-down of the treatment area, the group in tablet condition was asked to fill in the survey for their preferences towards a future usage once again. Hence, we had pre- and post-tests for this group before and after using the application in the exercise.

Results & Discussion. After the short introduction and briefing, the participants ($N = 18$; as 6 surveys had to be removed due to incompleteness) generally agreed on the positive impact of the tablet-based application. For the application, the usefulness, innovativeness and preference for future operations were rated significantly positive ($t_{use}(19) = 4.36$, $t_{inno}(19) = 6.66$, $t_{pref}(18) = 9.80$, $p < .01$).

A similar trend was shown—at least descriptively—for the statement that the application allows for a better patient-centered care. However no significant difference from zero (no agreement in any direction) was found. This rating was also unchanged in the post condition after the exercise.

In general, participants rated their technical affinity positively ($M = 0.21$; $SD = 0.82$). This result was not significant, due to the large variance. Participants were significantly showing a positive valuation of their own competence ($M = 0.19$; $SD = 0.38$; $t(23) = 2.43, p < .05$), a positive attitude towards technical devices ($M = 0.46$; $SD = 0.30$; $t(23) = 7.53$, $p < .01$) and a rejection of negative outcomes caused by new technologies ($M = -0.22$; $SD = 0.45$; $t(23) = -2.35, p < .05$).

An analysis of the relation between affinity towards technical devices and statements on the future usage revealed a negative medium correlation between the outcome of new technologies and the opinion about a better patient-centered care due to the technical application ($r_s = -.47$, $p = .05$). Users afraid of the impact of new technologies do not believe that our system can improve patient-centered care.

We are confident that these results confirm our initial application design: The overall feedback was consistently positive before, and even after using the application during a realistic disaster relief exercise. We assume that our application aligned with the users' requirements and expectations. The study also showed that it is important to specifically address users that are generally critical of technical solutions.

4.4 Interview with Operation Leaders

After the regular debriefing and user study test, the operation leaders of both rescue units were surveyed using a semi-structured interview, which we summarize here.

The exercise script contained a *lost patient*, who disappeared during the mission. In the analog group, 7 patients where registered during the initial triage. Later, the group leader reported 6 patients in treatment. Their operational control failed to detect the mismatch and overlooked the lost patient. In the digital group, the group leader quickly realized the mismatch and found the missing patient.

Communication between operational control and rescue forces was an aspect heavily affected by the application usage. In the analog group, the operational control had intensive radio communication with the forces on-site, yielding the feeling of a good situational awareness. Once the first patients arrived at the treatment area, the group leader was steadily involved in radio communication. For the digital group, the experience was very different. Neither the operational control nor the group leader were frequently radioing, preferring the application's messages feature. Also, state updates arrived directly in the application. This was an unfamiliar experience, but the benefit of re-reading messages was highlighted by the team.

Concerning *up-to-dateness*, the digital solution was favored, because operational control did not have to wait for updated lists to arrive by foot. Regarding information *trustworthiness*, operational control of the digital team sometimes had difficulties to assess the actual up-to-dateness of information provided by the application. *Batteries usage* was not considered a problem for longer deployments: they could be recharged at mobile diesel generators. As a *backup solution* for increased reliability, the usage of a special node with a printer attached was suggested, which can periodically print physical backup lists. For the *implementation* of such a system, the operation leaders highlighted the importance of practical exercises. They also emphasized the need for a free and open solution that is usable by anyone. Finally, a limited, but focused feature set was considered to be important for those with low technical affinity.

5. RELATED WORK

Many existing systems utilize multiple wireless communication technologies, such as IEEE 802.11 WLAN or cellular networks [6]. A reliance on infrastructure-oriented communication entails significant risk, due to potential congestion and the central point of failure. To avoid this risk, research adopted ideas from mobile ad-hoc networks to disaster scenarios in the form of peer-to-peer architectures. The initial designs utilized a traditional client-server architecture with the constraint that nodes had to be in communication range of a certain root node or were bound to the knowledge of multi-hop end-to-end routes. The dynamic nature of emergency scenarios heavily exacerbate the management of such end-to-end routes. Due to the movement of emergency workers, vehicles, and their equipment, the topology of the wireless is subject to constant change and therefore network partitions may occur frequently. We describe several examples in the following; unlike these examples, our work shows a small-scale project with commodity soft- and hardware components can also develop an effective solution.

Gao et al. [4] focus on the use case of patient monitoring during the treatment process: they introduce electronic triage tags for advanced health monitoring and disaster aid networks. This wearable tag includes several biomedical sensors to monitor the patient's vitals directly; the sensor can be managed and tracked with laptops or PDAs. This management and tracking occurs over a ZigBee network, which is also used to collect the sensor data to the devices for processing. ZigBee is suitable for this type of application, where a sink node (such as a laptop) is available. In our work, we do not attach sensors or information to the patients directly, but we rather focus on the dissemination and sharing of information between the different rescue forces. Our solution could be enhanced by including biomedical sensors, but this would require custom hardware; one of our goals was to rely on commercial-off-the-shelf hardware only. Their field study was comparable to ours, as where their results; they also observed a reduction in radio communication.

Another research team has developed Wireless Information System for Medical Response in Disasters (WIISARD) [6], which is a system based on 802.11 mesh networks. Nodes connect to this mesh network, consisting of battery-powered APs, and apply a publish/subscribe service approach to disseminate data within the network. In this architecture, clients are responsible for dealing with connection losses by caching local data where relevant; replication is performed as soon as connectivity is restored. As with our system, WIISARD was evaluated during a field study, where it was also shown that the digital solution reduced the amount of radio communication. In contrast to our design, the WIISARD system is strongly centralized, and individual nodes do not have access to the complete data set. In case of a server failure are the local changes not replicated to other nodes, even if they have a working connection. Despite having identified these issues, the authors have further evaluated the system in a larger field study, with over 100 participants [7].

Interestingly, both of these major field studies have neglected the analysis of system usability and user acceptance, although Gao et al. [4] did include a brief post-test analysis similar to ours. Our work has cost-effectively replicated the results discussed in these works, and additionally includes a thorough study of the systems' usability.

6. CONCLUSION

In this paper we have presented our experience in deploying mobile communication technologies in a real-world environment using commercial-off-the-shelf hardware and open source software. We analyzed the behavior and usability of our system in a realistic field study, in cooperation with local Red Cross units. The main result is that it is feasible for a small team (1-2 full time developers) to develop this application type with commercial-off-the-shelf hardware within a few months, applying research results for challenged networks in a real-world scenario. Our usability analysis has shown that the resulting application has a significant acceptance from users, despite this comparatively brief development time. The interviews conducted with operation leaders indicated that the system is considered useful for future deployments. In summary, we conclude that mobile ad-hoc network research has reached a point where commercial-off-the-shelf hardware and open source software is suitable to deploy applications that apply these research results.

7. REFERENCES

[1] J. C. Anderson, J. Lehnardt, and N. Slater. *CouchDB: the definitive guide*. O'Reilly Media, Inc., 2010.

[2] C. Bruder, C. Clemens, C. Glaser, and K. Karrer-Gauß. TA-EG – Fragebogen zur Erfassung von Technikaffinität. Technical report, FG Mensch-Maschine Systeme TU Berlin, 2009.

[3] Bundesamt für Bevölkerungsschutz und Katastrophenhilfe. Empfehlungen für Taktische Zeichen im Bevölkerungsschutz, 2012.

[4] T. Gao, T. Massey, L. Selavo, D. Crawford, B.-r. Chen, K. Lorincz, V. Shnayder, L. Hauenstein, F. Dabiri, J. Jeng, et al. The advanced health and disaster aid network: A light-weight wireless medical system for triage. *Biomedical Circuits and Systems, IEEE Transactions on*, 1(3):203–216, 2007.

[5] P. Leach, M. Mealling, and R. Salz. A Universally Unique IDentifier (UUID) URN Namespace. RFC 4122 (Proposed Standard), July 2005.

[6] L. Lenert, T. C. Chan, W. Griswold, J. Killeen, D. Palmer, D. Kirsh, R. Mishra, and R. Rao. Wireless internet information system for medical response in disasters (wiisard). In *AMIA annual symposium proceedings*, volume 2006, page 1192. American Medical Informatics Association, 2006.

[7] L. A. Lenert, D. Kirsh, W. G. Griswold, C. Buono, J. Lyon, R. Rao, and T. C. Chan. Design and evaluation of a wireless electronic health records system for field care in mass casualty settings. *Journal of the American Medical Informatics Association*, 18(6):842–852, 2011.

[8] N. Marchetti. *Telecommunications in Disaster Areas*. River Publishers, 2010.

[9] T. Mitschke and H. Peter. *Handbuch für Schnell-Einsatz-Gruppen*. Verlagsgesellschaft Stumpf & Kossendey, 3. überarbeitete Auflage edition, 2001.

[10] W. Vogels. Eventually consistent. *Communications ACM*, 52(1):40–44, Jan. 2009.