Interoperating Infrastructures in Emergencies

Antoine Desmet and Erol Gelenbe *FIEEE*

Imperial College London
Department of Electrical and Electronic Engineering
Intelligent Systems and Networks Group
{a.desmet10,e.gelenbe}@imperial.ac.uk

Abstract. The great challenge in handling the security and resilience in emergency situations is that threats will typically affect more than one infrastructure. A fire is not only a direct hazard for people but it will also short the electrical system, cutting off the lights and possibly the communications and sensor infrastructure and even create more fires. The Tsunami in Japan in 2011 flooded the nuclear reactors but also cut of the pumps that were designed to respond to any flooding situations. This paper is part of a project that addresses these cascaded failures and studies them via simulation. To provide some quantitative estimates of the effect of such cascaded threats, we use the Distributed Building Evacuation Simulator (DBES) to represent the effect of a hazard (in this case a fire) which destroys the sensor system which is used to compute the best advice given to people that are evacuated during the fire. Our simulations compare the situation when the sensor system is intact, and also when it is compromised. As expected, some results highlight the poor overall system performance when the underlying infrastructures are damaged. However, in some scenarios, the degraded system appears to perform as well as the intact one. An analysis into the fault-tolerance of the system leads to some design guidelines which can be applied to design systems fault-tolerant systems.

Keywords: Cyber-physical systems, emergency navigation, interacting critical systems, wireless sensor networks

1 Introduction

Whenever an undesirable incident occurs such as a fire, a terrorist act, or an accident that requires the intervention of emergency services, critical infrastructures such as water and electricity, even if just within a single large building, must interact with sensor networks and distributed computation for decision making, and with each other in order to offer overall security for both human beings and goods and services. This paper addresses this precise issue of fusing inter-operable "networks" within the hypothetical situation of a need for emergency evacuation for illustration purposes. Thus emergency management serves as a vignette or test-case for an overall cyber-technical environment that exploits wireless technologies, micro-sensing and distributed decision making when incidents occur that implicate or damage several of the underlying infrastructures.

For instance, if the water sprinkling system in a building is actuated by the electrical system, while the sensor system also depends on electrical energy, a fire that is initially detected by the sensors will eventually cause short circuits and further fires in the electrical systems, which in turn will bring down both the sprinkling and the sensor network and the computerized distributed decision making system.

Networking and wireless sensing enable applications in environmental sensing [1], health monitoring [17], surveillance [21], intelligent transportation [26], smart tourism [5], and emergency response [22, 10]. During a fire, temperature and gas sensors are responsible for monitoring the spreading of hazards, cameras can track the spread of the fire and the movement of civilians, ultrasound detects obstacles in the environment, and monitors dynamic changes in built structures through destruction and debris. Intelligent evacuation scheduling can be carried out by cooperation between people with mobile devices, decision nodes, sensors, and civilians with mobile devices [8, 9] and in case of major breakdowns or catastrophes opportunistic communication can also be used [16]. Civilians with mobile devices will follow the best known paths, directions and distributed decisions will help select those paths, and signposting and mobile devices can be used for sharing the best advice with the evacuees. In this area there is an abundant literature; early work [20] assumes that there is only one exit, while *Artificial potential fields* that have long been used in mission planning [13, 18] can be used to compute evacuation paths in a distributed manner where the exit point creates an attractive force that pulls the evacuees (and their mobile devices) to the exit, while each obstacle and hazard generates a repulsive force pushing the away from obstacles and hazards. On the other hand, shortest paths to exits may lead the evacuees along paths that are close to hazards [20]. The work in [25] uses ideas from *multipath routing* in mobile ad hoc networks to guide people as far away from hazardous regions as possible, and in [23] this approach is extended to a 3D guide for people (downstairs to exits, or to rooftops when there are no obvious safe evacuation paths). Human congestion is considered in [6, **?**], where a distributed protocol is proposed to balance the number of evacuees among multiple navigation paths to different exits; each sensor is location-aware and capable of detecting the number of evacuees within its sensing coverage. Geometric approaches exploits the unique properties of geometric graphs to plan evacuation paths as far as possible from the hazards as in [4] where *Delaunay triangulations* [24] are used to partition a WSN into several triangular areas for planning area-to-area navigation paths in a distributed manner. Since location information regarding sensors and users may not always be available, in [19] a road map is used in each user device to compute navigation paths. Based on distances of sensors to the hazardous areas, the backbone of the road map [3], is created and a shortest path tree rooted at the exit is constructed so that each evacuee can avoid the hazardous areas. In [2], a system is presented which estimates how long a hazard takes to reach a given sensor, to compute evacuation paths which offer the longest safest time for evacuees to make their way out of the building. Concepts such as "hazard time" and the evacuation delay are also considered in [7].

2 Modeling the Concurrent Critical Infrastructures

The resiliency of coupled critical infrastructures is evaluated through simulation. The Distributed Building Evacuation Simulator (DBES) platform is chosen to simulate assistance and safety systems in the context of emergency building evacuation. DBES is an agent-based platform dedicated to building evacuation simulation. As such, DBES comes with generic and expandable models for buildings, occupants, hazards, evacuation policies and more. A distinctive feature of DBES is that it is completely run by autonomous agents. This organic structure allows DBES to operate as a fully decentralized simulation platform – with the notable exception of timekeeping and synchronization. An agent-based architecture also naturally lends itself to parallel processing: each agent is assigned to a particular simulation object or *entity*, and all agents run concurrently with some level of autonomy. Since the agent messaging system operates over most computer networks, the agent's physical location is only constrained by network performance. This allows DBES to extend beyond parallel processing and operate in a distributed fashion, over a pool of networked computers.

The following describes some of the fundamental aspects of the model implemented in DBES, along with some components which are of special relevance to the scenarios we present.

Building Model The physical building model in DBES consists of a coarse-grid graph, loaded from a user-configured XML file. DBES incorporates a model for path congestion which limits the number of building users allowed on a given edge at any given time. As such, some characteristic patterns of real-life emergency evacuation events occur in the simulation (bottlenecks, effects of the flow of large crowds and more).

Graphical User Interface The building graph can be represented on the simulator's Graphical User Interface (GUI), superimposed over a building floor plan. The GUI also shows the position and health of each user, the spread of the hazard and the condition of the sensor system.

Communications DBES provides a framework to model communication networks. As part of the scenario requirements, a communication device consisting of three networks layers is developed to model a simple broadcasting wireless communication network. The purpose of the wireless network is to support a mobile node localization system.

Wireless Localization Fixed-location RF *anchors* are disseminated across the building to serve as reference points when localizing the communication devices carried by every building user. The localization system is based on multilateration, and node-anchor distances are inferred from signal strength.

Utility Networks A generic *Utility Network* component provides the fundamentals to model any infrastructure network composed of nodes and vertices. The concept of Utility Network can be applied to water, power or gas distribution infrastructures, sensor networks, data communications, etc. The generic model can be extended to incorporate specific aspects and components of the network, such as electrical circuit breakers, valves, repeaters, sensors, routers, as required.

Fire Monitoring System Using the concept of *Utility Networks* introduced earlier, a fire monitoring system is implemented and deployed across the entire building. Networked fire sensors detect and measure the fire intensity in their vicinity. The model accounts for malfunctions induced by fire damage and the resulting disruptions in information supply.

Emergency Evacuation Assistance System This system amalgamates information from the fire monitoring system and the wireless localization component, to provide individualized optimal escape path recommendations to every building user. The escape path presented by the system begins at the user's sensed location, and provides a comprehensive step-by-step safest path to a building exit. It is assumed that each building user's communication node comes with some form of user interface, on which evacuation advice can be displayed.

3 Simulations

In the following simulations we model an emergency evacuation assistance system which relies on two interacting infrastructures : *a)* a user localization subsystem, which sets the origin of the recommended evacuation path, and *b)* the fire monitoring sub-system, which determines the areas that ought to be avoided and defines the safest path for each evacuee. This paper focuses on the high-level effects of fire-induced faults in the fire sensing system.

The assumptions, parameters and configurations used to produce the simulation results are as follows:

Building Parameters The building simulated is a 3-floor rectangular building of 60x24 meters. The upper floors mostly feature office space, while the ground floor is primarily a lobby area, as can be seen in Figure 1(b). The building comprises of two evacuation exits, located in close proximity to one another, at the Southern edge of the ground floor. Two stairways provide emergency access to the ground floor from upper levels, and are located at the N and N-E ends of the building. The two grey arrows on Figure 1(b) are near these two stairways and mark the main flow of users evacuating from the upper floors.

Building Users The number of building users is set to {10, 20, 30, 40} users *per floor*. For greater consistency, each simulation is iterated 10 times with a different set of initial user locations and designated working areas. The building users' normal behavior is to spend the majority of their time in a dedicated area – such as their office – and to make occasional short trips to other areas of the building. As soon as the fire alarm is triggered, the evacuee joins the path suggested by the assistance system, and strictly follows it until an exit is reached.

Tracking System Infrastructure A set of 8 tracking anchors is disseminated on every floor to provide comprehensive and redundant coverage of the building. The range of an anchor device is approximately 20 meters, and the resulting accuracy of the system is within 5 meters. These devices have an independent secondary power supply, and since they do not rely on a wired infrastructure network, their failure is not modeled. This is based on the assumption that their failure would only occur well after the building had been evacuated.

Fire Monitoring System The fire monitoring system consists of a dense deployment of fire sensors, from 40 to 60 sensors per floor. Each floor is divided into 5 to 8 *zones*, which correspond to areas of the building. Figures 1(b) and 1(a) show the topology of the first and second floor's fire monitoring sensor network. Different *line styles* are used to differentiate each zone on the Figures. Every zone is connected to the floor's fire panel, in a star topology. *Within a zone*, sensors are connected in a bus topology. Owing to this bus wiring, damage or short-circuiting of any wire or sensor leads to fault conditions which affect the entire zone. The subsequent effect is a loss of sensor information from the entire zone, even though the fire may have only damaged a minor part of it.

Evacuation Assistance System The evacuation assistance system is hosted on a central server, located in a safe area of the building. Its role is to *1)* gather the position of the building users, *2)* evaluate every possible egress path and *3)* select and transmit the safest one to the evacuee. The decision system favors path safety over shorter paths. Typically, if a shorter egress path requires the user to walk in close proximity to the fire, the system will instead recommend a safer but most likely longer path. Throughout the simulation, the fire sensing network will be affected by fire damage. Damaged sensor zones will no longer supply current information to the decision system, and the latter will eventually begin to base its decisions on outdated information, i.e. the last value supplied by the sensor. The coupling between the top-level application and its underlying infrastructure will be observed through the success rate of the assistance system. Success rate is derived from the number of building users which are able to evacuate the building.

3.1 Experimental Results

The results of three distinct simulation configurations are summarized in Figure 2(a) and 2(b). Each graph corresponds to a different fire outbreak location – each marked by a flame icon on Figures 1(b) and 1(a). These graphs present the percentage of building users evacuating the building for various densities of users per floor. The results shown are the average of ten simulations, and the error

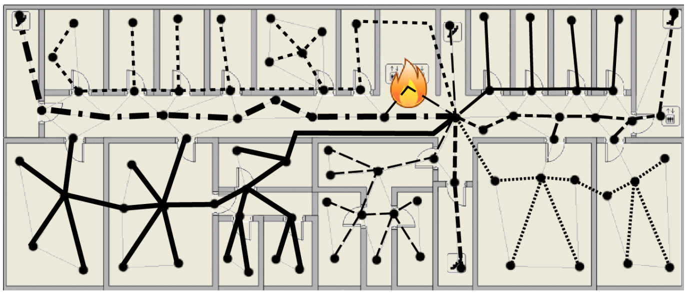

(a) Floor 2

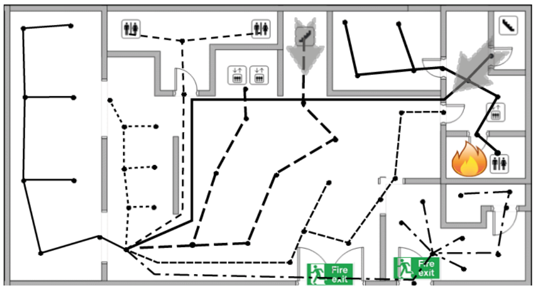

(b) Floor 1

Fig. 1. Floor plans of the building simulated, showing the sensor nodes (black dots) and the network layout, and the initial fire locations in different simulations (flames).

bars indicate the lowest and highest values obtained across the simulation iterations. Each bar graph has three series which correspond to different simulation models:

- The dark-gray bars serve as a control test, where occupants evacuate the building upon hearing the alarm – without any further form of assistance. This is referred to as the “autonomous” run.
- Bars in gray relate to simulations where fire-induced sensor failure is modeled. Simulation parameters are set so that sensors are able to record moderate fire intensities before eventually defaulting as the fire intensity grows. In turn, faulty fire sensors cause the fire sensor control panel to shutdown the entire zone on which the faulty sensor is located. This model is referred to as the “realistic” model.
- Bars in the light-gray shade represent the outcome of a simulation run where destruction of the fire sensing network does not occur. It is referred to as the “ideal” model, and shows the “best case” results that an ideal and fully fire-proof system could achieve.

Figure 2(a) shows somewhat expectable results. As the density of users in the building increases, the egress paths are subjected to greater congestion, and proportionally less users are able to evacuate the building in time, regardless of any

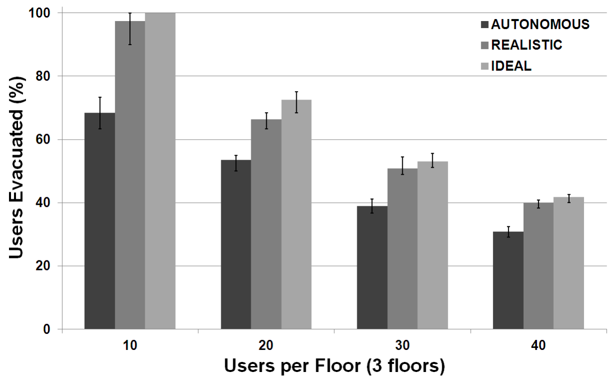

(a) Simulation results for fire starting on floor 1

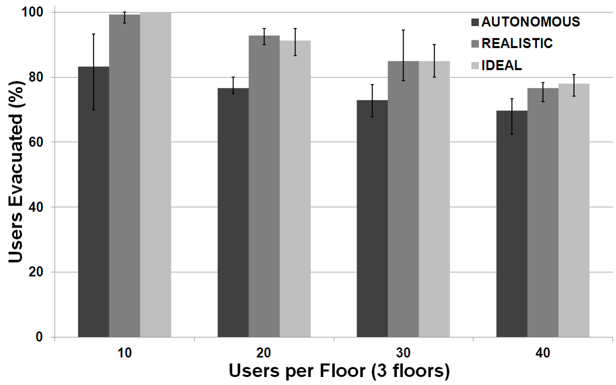

(b) Simulation results for fire starting on floor 2

Fig. 2. Histogram of building user's health, expressed as a percentage

assistance received. "Autonomous" users – which ignore the fire outbreak location and simply evacuate the building by following the shortest path they know of – often meet the fire on their way out and have to follow lengthy bypassing routes. This delays the evacuation and eventually causes users to become trapped inside the building, which explains the larger casualty count for this population. On the other hand, in the "ideal" model users are constantly kept informed of the safest path, and receive updates as the fire evolves. It is worth noting that casualties do occur in the "ideal" scenario, as the system favors the safest path and does not manage *overall* evacuation time. Often, the safest path is also the longest one, while shorter paths generally involve higher risks. As users are directed to this safe but lengthy path, some eventually become trapped in the building and perish. In the "realistic" scenario, the information provided is often inaccurate, as the system must resort to using out-of-date measurements as soon as sensors become inoperative. Overall, damage to the underlying fire sensor network increases casualty ratio by 5 to 10 points. This is aggravated, in this particular scenario, by the specific location of the fire outbreak. While the fire initially breaks out in a non-critical dead-end area of the building, it quickly destroys the area's fire monitoring systems and spreads, un-monitored, towards the critical evacuation path near the N-E staircase. This condition goes unnoticed until the fire eventually spreads far enough to reach an adjacent, healthy sensor zone. Only then does the evacuation assistance system become aware that the N-E staircase is no longer a safe route and begins re-directing users. By then, not only have a few users already perished by trying to cross this area, but the re-routing process also took its toll, by increasing the overall evacuation time. This experiment highlights how the high-level system can be affected by malfunctions in underlying infrastructures. In this particular scenario, the system's advice has shown to provide evacuees with misleading and downright life-threatening advice.

Overall, the evacuation ratios observed in Figure 2(a) are higher than those in Figure 2(b). This is only due to the fire initial location: an outbreak on the second floor leaves better alternatives for users to evacuate the building, and it also guarantees and unobstructed evacuation to all users on the floor below the fire outbreak. The order of magnitude, in terms of evacuation ratios between assisted and autonomous users, is comparable between Figures 2(a) and 2(b). While similarities exist between the two experiments, the most striking difference is that in Figure 2(b), the "ideal" and "realistic" models display near-identical results. This suggests that the information which is not captured by the defective sensor network is of little value, since the degraded system performs as well without it than the "indestructible" system does with it.

In order to explain this phenomenon, let us focus on the initial fire locations, and their particular implications on evacuation paths. The fire location on floor 1 is in a dead-end: a low-traffic area which would never form part of any evacuation path. On the other hand, the fire location on floor 2 is at the crossroads of several critical evacuation paths. From the evacuation assistance system's viewpoint, the outbreak of fire at the dead-end location on the first floor does not pose an immediate threat to users. As such, the system merely recommends each user to evacuate following the shortest path to the exits. As the fire grows, the local sensor network soon burns out, and stops reporting the progress of the fire as it reaches the critical N-E staircase evacuation path. Deprived from this crucial piece of information, the emergency evacuation assistance system fails its duty to advise users to avoid this area. Overall, it merely serves to trigger the evacuation process, and provides little more valuable assistance. On the other hand, when the fire occurs on floor 2, the value of the information acquired before the area's sensor network failed is much greater. Beyond detecting the initial fire outbreak, the sensor network is also able to report the presence of fire in a critical area of the building for evacuation purposes. Although the area's fire sensing system soon develops a fault, the evacuation assistance system has already received valuable information and is able to trigger the re-routing of the many paths going through this area, and effectively assists users with life-saving advice. In summary, the variations in success rate observed between the two scenarios is due to the differences in quantity and value of information acquired before the sensor network eventually defaulted.

4 Comments and Conclusions

This experiment highlights the fact that the *value* of hazard information is highly location-dependant. Clearly, as some areas of the building are critical for evacuation purposes, information relating to hazards in those areas is also of greater value compared to, for instance, information pertaining to a dead-end. Based on this observation, we can define two categories of information for emergency evacuations, and their value:

1. Initial outbreak: this information is of the highest value: it triggers the alarm, signalling every user to evacuate the building.
2. Hazard location: the value of this information depends directly on the location considered. The more evacuation paths tend to run through this location, the higher the information value.

To create a dependable and fault-tolerant system, the underlying sensor network must be designed to guarantee that the highest-value information will be acquired before defaulting. Information relating to the initial outbreak is not an issue: the fire is detected well before it become severe enough to start damaging the infrastructure. What most impacts the outcome of the evacuation process is the total *value* of the "hazard location" information acquired before the sensor network zones inevitably fail.

4.1 Design of Robust Infrastructures

The first step in designing a robust underlying infrastructure is to identify critical edges: those which are most likely to form part of evacuation paths, as this is directly related to hazard location information value. This can be done by simulating a series of evacuation drills, i.e. without any real hazard fire, and marking each edge which forms part of a shortest path to an exit. After several iterations with random sets of initial user locations, the "score" of each edge will represent their degree of criticality. The "zoning" and general layout of the network can then be optimized based on this information. A basic approach to designing such a wired sensor network is to divide the building using a rough grid pattern and assign one zone to each area of the building – as was done in this experiment. This design has proven to yield unpredictable and overall poor performance in degraded conditions. At the other end of the spectrum, one could envision wiring each sensor as an individual "zone", however this method is clearly not scalable. Instead, a seemingly optimal solution consists in assigning independent zones to each section of the critical exit paths identified. By grouping the "critical" sensors in one zone, high-value information would inherently be acquired before the sensors failed, i.e. if this zone has developed a fault, it can only mean that the fire has reached the corresponding major egress path and that through traffic should be re-routed. All remaining nodes which do not cover critical paths can form part of one single zone whose purpose is to detect the initial fire outbreak.

4.2 Future Direction

Having highlighted the flaws of a simplistic sensor network layout, and proposed a design method to increase the fault-tolerance of such networks, the next step in this research project will be to verify its effectiveness against other designs. In particular, the fault-tolerance of the overall system will be tested by running simulations with various fire outbreak locations, while expecting the performance of the "realistic" system to be as close as possible to the "ideal" model, regardless of the fire parameters. While in this paper we have focused on simulations to the study of inter-network dependencies and fault-tolerance in emergency management, in future work we may also call upon probability modeling used in computer systems, networks and distributed systems [11, 14, 15] and for models of uncertainty of information [12] to offer computationally fast predictions of overall performance prior to simulations.

References

1. Terrestrial ecology observing systems, center for embedded networked sensing, http://research.cens.ucla.edu/
2. Barnes, M., Leather, H., Arvind, D.K.: Emergency evacuation using wireless sensor networks. In: IEEE Conf. Local Computer Networks. pp. 851–857 (2007)
3. Bruck, J., Gao, J., Jiang, A.A.: MAP: medial axis based geometric routing in sensor networks. In: ACM Int'l Conf. Mobile Computing and Networking. pp. 88–102 (2005)
4. Chen, P.Y., Chen, W.T., Shen, Y.T.: A distributed area-based guiding navigation protocol for wireless sensor networks. In: IEEE Int'l Conf. Parallel and Distributed Systems. pp. 647–654 (2008)
5. Chen, P.Y., Chen, W.T., Tseng, Y.C., Huang, C.F.: Providing group tour guide by rfids and wireless sensor networks. IEEE Trans. Wireless Communications 8(2), 1536–1276 (2009)
6. Chen, W.T., Chen, P.Y., Wu, C.H., Huang, C.F.: A load-balanced guiding navigation protocol in wireless sensor networks. In: IEEE Global Telecomm. Conf. pp. 1–6 (2008)
7. Cherniak, A., Zadorozhny, V.: Towards adaptive sensor data management for distributed fire evacuation infrastructure. In: Int'l Conf. Mobile Data Management. pp. 151–156 (2010)
8. Filippoupolitis, A., Gelenbe, E.: A decision support system for disaster management in buildings. In: Summer Computer Simulation Conference. p. 141V147 (2009)
9. Filippoupolitis, A., Gelenbe, E.: A distributed decision support system for building evacuation. In: IEEE Int'l Conf. Human System Interactions. pp. 320–327 (2009)
10. Fischer, C., Gellersen, H.: Location and navigation support for emergency responders: A survey. IEEE Pervasive Computing 9(1), 38–47 (2009)

11. Gelenbe, E.: A unified approach to the evaluation of a class of replacement algorithms. IEEE Transactions on Computers C-22(6), 611–618 (June 1973)
12. Gelenbe, E., Hébrail, G.: A probability model of uncertainty in data bases. In: ICDE. pp. 328–333. IEEE Computer Society (1986)
13. Gelenbe, E., Hussain, K., Kaptan, V.: Simulating autonomous agents in augmented reality. Journal of Systems and Software 74(3), 255268 (February 2005)
14. Gelenbe, E., Sevcik, K.C.: Analysis of update synchronisation algorithms for multiple copy data bases. IEEE Transactions on Computers C-28(10), 737–747 (October 1979)
15. Gelenbe, E., Stafylopatis, A.: Global behavior of homogeneous random neural systems. Applied Mathematical Modelling 15
16. Gorbil, G., Gelenbe, E.: Opportunistic communications for emergency support systems. In: Int'l Conf. Ambient Systems, Networks and Technologies. pp. 1–9 (2011)
17. Hu, F., Xiao, Y., Hao, Q.: Congestion-aware, loss-resilient bio-monitoring sensor networking for mobile health applications. IEEE J. Selected Areas in Comm. 27(4), 450–465 (2009)
18. Kaptan, V., Gelenbe, E.: Fusing terrain and goals: agent control in urban environments. In: Multisource Information Fusion: Architectures, Algorithms, and Applications 2006. vol. 6242, p. 7179. SPIE (April 2006)
19. Li, M., Liu, Y., Wang, J., Yang, Z.: Sensor network navigation without locations. In: IEEE INFOCOM. pp. 2419–2427 (2009)
20. Li, Q., Rosa, M.D., Rus, D.: Distributed algorithms for guiding navigation across a sensor network. In: ACM Int'l Conf. Mobile Computing and Networking. pp. 313–325 (2003)
21. Liu, H., Wan, P., Jia, X.: Maximal lifetime scheduling for sensor surveillance systems with k sensors to one target. IEEE Trans. Parallel and Distributed Systems 17(12), 1526–1536 (2006)
22. Malan, D.J., Fulford-Jones, T.R., Nawoj, A., Clavel, A., Shnayder, V., Mainland, G., Welsh, M., Moulton, S.: Sensor networks for emergency response: Challenges and opportunities. IEEE Pervasive Computing 3(4), 16–23 (2004)
23. Pan, M.S., Tsai, C.H., Tseng, Y.C.: Emergency guiding and monitoring applications in indoor 3D environments by wireless sensor networks. Int'l J. Sensor Networks 1(1/2), 2–10 (2006)
24. Preparata, F.P., Shamos, M.I.: Computational Geometry: An Introduction. Springer-Verlag (1985)
25. Tseng, Y.C., Pan, M.S., Tsai, Y.Y.: Wireless sensor networks for emergency navigation. IEEE Computer 39(7), 55–62 (2006)
26. Tubaishat, M., Zhuang, P., Qi, Q., Shang, Y.: Wireless sensor networks in intelligent transportation systems. Wireless Comm. and Mobile Computing 9(3), 287–302 (2009)