Automatica 50 (2014) 2470–2481

Contents lists available at ScienceDirect

Automatica

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

Error probability bounds for nuclear detection: Improving accuracy through controlled mobility ☆

Chetan D. Pahlajani [a,1], Jianxin Sun [b], Ioannis Poulakakis [b], Herbert G. Tanner [b]

[a] Department of Mathematical Sciences, University of Delaware, Newark, DE 19716, United States

[b] Department of Mechanical Engineering, University of Delaware, Newark, DE 19716, United States

ARTICLE INFO

Article history:
Received 3 July 2013
Received in revised form
3 January 2014
Accepted 19 May 2014
Available online 10 September 2014

Keywords:
Nuclear detection
Mobile sensor networks
Inhomogeneous Poisson processes
Chernoff bounds

ABSTRACT

A collection of static and mobile radiation sensors is tasked with deciding, within a fixed time interval, whether a moving target carries radioactive material. Formally, this is a problem of detecting weak time-inhomogeneous Poisson signals (target radiation) concealed in another Poisson signal (naturally occurring background radiation). Each sensor locally processes its observations to form a likelihood ratio, which is transmitted once—at the end of the decision interval—to a fusion center. The latter combines the transmitted information to optimally (in the Neyman–Pearson sense) decide whether the measurements contain a radiation signal, or just noise. We provide a set of analytically derived upper bounds for the probabilities of false alarm and missed detection, which are used to design threshold tests without the need for computationally intensive Monte Carlo simulations. These analytical bounds couple the physical quantities of interest to facilitate planning the motion of the mobile sensors for minimizing the probability of missed detection. The network reconfigures itself in response to the target motion, to allow more accurate collective decisions within the given time interval. The approach is illustrated in numerical simulations, and its effectiveness demonstrated in experiments that emulate the statistics of nuclear emissions using a pulsed laser.

1. Introduction

This paper proposes a theoretical framework for network-based decision making, tailored to the problem of detecting nuclear material in transit within a given time interval, using a network of small and inexpensive static and mobile radiation sensors. This is an instance of a general problem of detecting a signal buried in noise, which is found for either single sensor or sensor network settings in a surprisingly rich application domain, from nuclear detection (Nemzek, Dreicer, Torney, & Warnock, 2004; Pahlajani, Poulakakis, & Tanner, 2013) and optical communications (Teich & Rosenberg, 1973), to radar (Yang, Blum, & Sadler, 2009) and acoustic (Kreucher & Shapo, 2011) surveillance, to medical sensing (Estes et al., 2003) and neuroscience (Gold & Shadlen, 2007), to natural disaster early warning systems (Faulkner et al., 2011), and to high-energy experimental physics (Cranmer & Plehn, 2007).

A network approach to deploying and managing data from radiation sensors can be one out of several layers in a comprehensive, integrative system for nuclear detection (Byrd et al., 2005; Srikrishna, Chari, & Tisch, 2005). The sensor of choice is a Geiger counter; larger and more sophisticated sensors (providing spectroscopy information) are prohibitively expensive to be deployed on a large scale (Sundaresan, Varshney, & Rao, 2007) and too big to be mounted on mobile platforms. In addition, any active (e.g. X-ray) interrogation technology cannot be used to check vehicles that carry passengers or livestock (Srikrishna et al., 2005).

A first challenge in detecting the presence of radioactive material with these types of sensors is that such a detector not only picks up the signal coming from the material, but also another one from ubiquitous cosmic and naturally occurring background radiation. From the sensor's perspective, the two signals are of identical nature and once superimposed, it is impossible to tell them apart. A second challenge relates to attenuation: although

☆ This work was supported in part by ARL MAST CTA W911NF-08-2-0004. The material in this paper was partially presented at the 52nd IEEE Conference on Decision and Control (CDC), December 10–13, 2013, Florence, Italy. This paper was recommended for publication in revised form by Associate Editor Michele Basseville under the direction of Editor Torsten Söderström.

E-mail addresses: chetan@math.udel.edu, cdpahlajani@iitgn.ac.in (C.D. Pahlajani), jxsun@udel.edu (J. Sun), poulakas@udel.edu (I. Poulakakis), btanner@udel.edu (H.G. Tanner).

[1] Now at: Indian Institute of Technology Gandhinagar, VGEC Campus, Chandkheda, 382 424 Ahmedabad, Gujarat, India.

http://dx.doi.org/10.1016/j.automatica.2014.08.025
0005-1098/

a kilogram of Highly Enriched Uranium (HEU) can emit as many as 4×10^7 gamma rays per second (Byrd et al., 2005), shielding and attenuation (Nemzek et al., 2004) limit the effective detection range to a few feet and require detection times that can range from minutes to hours (Srikrishna et al., 2005). In fact, a study performed on the radiation emitted by actual nuclear missiles (Fetter et al., 1990) concluded that for the type of warheads containing HEU, the gamma-ray emission just 25 cm away from the warhead casing is comparable to background. In a similar study (Fetter, Cochran, Grodzins, Lynch, & Zucker, 1990), it was concluded that a nuclear cruise-missile is practically undetectable by portable gamma-ray detectors at a distance of more than 5 m. And although remote count-based detection with vehicle-mounted sensors is possible (Fetter et al., 1990), the required sensor sensitivity and resolution are beyond that of common portable detectors which would most likely form the basis of a mobile sensor network. The problem is exacerbated when the source of the signal that needs to be detected is in motion. Not only does the stochastic process describing the signal become time-inhomogeneous from the detector's perspective (Nemzek et al., 2004), but the detector(s) only have a limited amount of time to make a decision before the potential target disappears from sight. They are called to detect within a small time interval a Poisson signal, buried inside another comparable Poisson signal.

Networks of spatially distributed sensors have been recognized as an important component of a multi-layered approach to nonproliferation, security and defense (Byrd et al., 2005). The potential of static sensor networks for the detection of stationary (Chandy, Pilotto, & McLean, 2008; Chin, Yau, & Rao, 2011; Rao, Chin, Yau, Ma, & Madan, 2010) or moving (Nemzek et al., 2004) radiation sources has been examined. In this context, the value of *sensor* mobility in nuclear measurement has recently been recognized (Cortez et al., 2008; Klimenko, Priedhorsky, Tanner, Borozdin, & Hengartner, 2006; Ma, Yau, Yip, Rao, & Chen, 2009), but the problem setup has been different from the one in this paper. Either sensor motion was random and the objective was network coverage (Ma et al., 2009), or detector motion was controlled and the objective was radiation mapping (Cortez et al., 2008; Cortez, Tanner, Lumia, & Abdallah, 2011), or sensor paths were predetermined (Kumar, Tanner, Klimenko, Borozdin, & Priedhorsky, 2006), or the source was static (Klimenko et al., 2006; Ristic & Gunatilaka, 2008; Ristic, Morelande, & Gunatilaka, 2010).

In a general setting, detection is a decision problem between two alternative hypotheses (source plus background versus background only) and a fair amount of literature in signal processing exists (Boel, Varaiya, & Wong, 1975; Brémaud, 1981; Davis & Andreadakis, 1977; Geraniotis & Poor, 1985; Hibey, Snyder, & van Schuppen, 1978; Verdú, 1986a). A likelihood ratio test (LRT) is a common approach, according to which a certain ratio computed based on collected sensor data is compared to a constant threshold; if this ratio is above the threshold, we decide that a source is present (else, we decide a source is absent). Nuclear emission is modeled as a Poisson process, and the approaches for detecting Poisson signals using networks of detectors generally follow either a Bayesian (Brennan, Mielke, & Torney, 2005; Chamberland & Veeravalli, 2003; Morelande, Ristic, & Gunatilaka, 2007; Nemzek et al., 2004) or a sequential formulation (DeLucia & Poor, 1997; Kazakos & Papantoni-Kazakos, 1980).[2] Typically, data collected at individual sensors is assumed to be independent identically distributed (i.i.d.) (a notable exception is Sundaresan et al., 2007). Sequential Probability Ratio Test (SPRT) approaches (Chin et al., 2010; Jarman, Smith, & Carlson, 2004) are not comparable to the one presented here for the following reason. In the SPRT setting, data is typically collected until such time as a decision can be made with sufficient accuracy. In our setting, however, the data is only available over a fixed time interval (while the target is within sensing range), at the end of which a decision necessarily has to be made. Networked Bayesian formulations, on the other hand, tend to be computationally intensive to the point that current computing power would limit the scale of networks that can implement them in real-time to single-digit network sizes (Brennan et al., 2005). Neyman–Pearson formulations (Kailath & Poor, 1998; Viswanathan & Varshney, 1997) can be an alternative to Bayesian approaches, but have not yet been adapted to the case of time-inhomogeneous Poisson processes like the ones resulting from relative motion between sensor and source.

Sensor mobility changes the dynamics of nuclear measurement. We now understand how and why closing the distance between sensor and source affects the information content of the sensor measurement (Nemzek et al., 2004): the signal-to-noise ratio scales with the inverse square of the distance. In this sense, bringing a small sensor closer to the source has an effect equivalent to that of using a much bigger sensor at a greater distance. Sensor mobility can be exploited (Cortez et al., 2008; Klimenko et al., 2006; Kumar et al., 2006; Ma et al., 2009) in the context of nuclear measurement, but it is not entirely clear what exact purpose it should serve. For example, in Klimenko et al. (2006) and Kumar et al. (2006) the variance of the assumed mean count rate at each spatial bin was taken as a performance measure, while Cortez et al. (2011), Ristic and Gunatilaka (2008) and Ristic et al. (2010) used various information-theoretic measures. Although these choices are intuitive, they may be considered equally arbitrary from the perspective of the decision maker. What is more, it is not always clear how the performance metric depends explicitly on sensor mobility, and how the latter can *optimally* be utilized.

This paper formulates the fixed-interval detection problem of a mobile source by a reconfigurable sensor network as an LRT developed within the Neyman–Pearson framework (Pahlajani, Poulakakis, & Tanner, 2014). In such a test, two types of errors can occur: the first is to decide that a source is present when there is not, and this constitutes a false alarm; the second is to decide that there is nothing when in fact there is a source, which is a case of missed detection. The Neyman–Pearson test is designed to minimize the probability of missed detection for a given acceptable probability of false alarm. The contribution of the paper is in showing explicitly how the relative distance between sensor and source affects the error probabilities in a fixed-interval LRT, formulated for detecting weak, time-inhomogeneous Poisson processes buried in Poisson background noise. Since the error probabilities cannot be analytically computed, however, this hinders their use in a sensor motion optimization scheme that would aim directly at improving the accuracy of the LRT. The paper addresses this problem by deriving appropriate Chernoff bounds as proxies for these error probabilities and utilizes an optimal (motion) control approach to steer the sensors, as well as to derive the optimal threshold values for the decision test, as a function of sensor and source trajectories. Chernoff bounds (Evans, 1974; Hibey et al., 1978; Newman & Stuck, 1979; Snyder, 1975) are used in robust detection for scalar Poisson processes with uncertain intensities in Geraniotis and Poor (1985). The bounds are also employed in performance evaluation of communication systems (Nelson & Poor, 1995; O'Reilly & da Rocha, 1987; Prabhu, 1982; Verdú, 1986b). Connections between Chernoff bounds and large deviations are explored in Bucklew and Sadowsky (1993), Kazakos (1991) and Sadowsky (1987).

Section 2 recalls from Pahlajani et al. (2014) the (Neyman–Pearson) optimal decision rule for our problem of interest. Section 3 derives analytical probability bounds for the fixed-interval detection test, tightens them, and validates them by

[2] Arguably, both have common theoretical underpinnings (Kailath & Poor, 1998; Viswanathan & Varshney, 1997).