Automatica 52 (2015) 160–172

Contents lists available at ScienceDirect

Automatica

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

Randomized methods for design of uncertain systems: Sample complexity and sequential algorithms[☆]

Teodoro Alamo[a], Roberto Tempo[b], Amalia Luque[a], Daniel R. Ramirez[a]

[a] Departamento de Ingeniería de Sistemas y Automática, Universidad de Sevilla, Escuela Superior de Ingenieros, Camino de los Descubrimientos s/n, 41092 Sevilla, Spain

[b] CNR-IEIIT, Politecnico di Torino, Corso Duca degli Abruzzi 24, Torino 10129, Italy

ARTICLE INFO

Article history:
Received 18 March 2013
Received in revised form
18 July 2014
Accepted 27 October 2014
Available online 10 December 2014

Keywords:
Randomized and probabilistic algorithms
Uncertain systems
Sample complexity

ABSTRACT

In this paper, we study randomized methods for feedback design of uncertain systems. The first contribution is to derive the sample complexity of various constrained control problems. In particular, we show the key role played by the binomial distribution and related tail inequalities, and compute the sample complexity. This contribution significantly improves the existing results by reducing the number of required samples in the randomized algorithm. These results are then applied to the analysis of worst-case performance and design with robust optimization. The second contribution of the paper is to introduce a general class of sequential algorithms, denoted as Sequential Probabilistic Validation (SPV). In these sequential algorithms, at each iteration, a candidate solution is probabilistically validated, and corrected if necessary, to meet the required specifications. The results we derive provide the sample complexity which guarantees that the solutions obtained with SPV algorithms meet some pre-specified probabilistic accuracy and confidence. The performance of these algorithms is illustrated and compared with other existing methods using a numerical example dealing with robust system identification.

1. Introduction

The use of randomized algorithms for systems and control has matured thanks to the considerable research efforts made in recent years. Key areas where we have seen convincing developments include uncertain and hybrid systems (Tempo, Calafiore, & Dabbene, 2013; Vidyasagar, 1997). A salient feature of this approach is the use of the theory of rare events and large deviation inequalities, which suitably bound the tail of the probability distribution. These inequalities are crucial in the area of statistical learning theory (Bousquet, Boucheron, & Lugosi, 2004; Vapnik, 1998), which has been utilized for feedback design of uncertain systems (Vidyasagar, 2001).

Design in the presence of uncertainty is of major relevance in different areas, including mathematical optimization and robustness (Ben-Tal & Nemirovski, 1998; Petersen & Tempo, 2014). The goal is to find a feasible solution which is optimal in some sense for all possible uncertainty instances. Unfortunately, the related semi-infinite optimization problems are often NP-hard (examples of NP-hard problems in systems and control can be found in Blondel & Tsitsiklis, 1997, 2000), and this may seriously limit their applicability from the computational point of view. There are two approaches to resolve this NP-hard issue. The first approach is based on the computation of deterministic relaxations of the original problem, which are usually polynomial time solvable. However, this might lead to overly conservative solutions (Scherer, 2006). An alternative is to assume that a probabilistic description of the uncertainty is available. Then, a randomized algorithm may be developed to compute, in polynomial time, a solution with probabilistic guarantees (Tempo et al., 2013; Vidyasagar, 1997). Stochastic programming methods (Prékopa, 1995) are similar in spirit to the methods studied in this paper and take advantage that, for random uncertainty, the underlying probability distributions are known or can be estimated. The goal is to find a solution that is feasible for almost all possible uncertainty realizations and maximizes the expectation of some function of the decisions variables.

☆ This work was supported by the MCYT-Spain under projects DPI2010-21589-C05-01, DPI2013-48243-C2-2-R and by the European Commission (FP7-257462). The work of Roberto Tempo was supported by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257462 HYCON2 Network of Excellence. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Valery Ugrinovskii under the direction of Editor Ian R. Petersen.

E-mail addresses: alamo@cartuja.us.es (T. Alamo), roberto.tempo@polito.it (R. Tempo), amalia@cartuja.us.es (A. Luque), danirr@cartuja.us.es (D.R. Ramirez).

http://dx.doi.org/10.1016/j.automatica.2014.11.004
0005-1098/

The field of probabilistic methods (Calafiore, Dabbene, & Tempo, 2011; Tempo et al., 2013; Tempo & Ishii, 2007) has received growing attention in the systems and control community. Two complementary approaches, non-sequential and sequential, have been proposed. A classical approach for non-sequential methods is based upon statistical learning theory (Vapnik, 1998; Vidyasagar, 1997). Subsequent work along this direction includes Alamo, Tempo, and Camacho (2009), Chamanbaz, Dabbene, Tempo, Venkatakrishnan, and Wang (2014), Koltchinskii, Abdallah, Ariola, Dorato, and Panchenko (2000), Vidyasagar (2001) and Vidyasagar and Blondel (2001). Furthermore, in Alamo, Tempo, and Luque (2010a,b) and Luedtke and Ahmed (2008) the case in which the design parameter set has finite cardinality is analyzed. The advantage of these methods is that the problem under attention may be non-convex. For convex optimization problems, a non-sequential paradigm, denoted as the scenario approach, has been introduced in Calafiore and Campi (2005, 2006), see also Alamo et al. (2010b), Calafiore (2010) and Campi and Garatti (2008, 2011) for more advanced results, and Prandini, Garatti, and Vignali (2014) and Vayanos, Kuhn, and Rustem (2012) for recent developments in the areas of stochastic hybrid systems and multi-stage optimization, respectively. Finally, we refer to de Farias and Van Roy (2003) for a randomized approach to solve approximate dynamic programming.

In non-sequential methods, the original robustness problem is reformulated as a single optimization problem with sampled constraints, which are randomly generated. A relevant feature of these methods is that they do not require any validation step and the sample complexity is defined a priori. The main result of this line of research is to derive explicit lower bounds to the required sample size. However, the obtained explicit sample bounds can be overly conservative because they rely on a worst-case analysis and grow (at least linearly) with the number of decision variables.

For sequential methods, the resulting iterative algorithms are based on stochastic gradient (Calafiore & Polyak, 2001; Polyak & Tempo, 2001), ellipsoid iterations (Kanev, De Schutter, & Verhaegen, 2003; Oishi, 2007); or analytic center cutting plane methods (Calafiore & Dabbene, 2007; Wada & Fujisaki, 2009), see also Alamo, Tempo, Ramirez, and Camacho (2007) and Chamanbaz, Dabbene, Tempo, Venkataramanan, and Wang (2013) for other classes of sequential algorithms. Convergence properties in finite-time are one of the focal points of these papers. Various control problems have been solved using these sequential randomized algorithms, including robust LQ regulators (Polyak & Tempo, 2001), switched systems (Liberzon & Tempo, 2004) and uncertain linear matrix inequalities (LMIs) (Calafiore & Polyak, 2001). Sequential methods are often used for uncertain convex feasibility problems because the computational effort at each iteration is affordable. However, they have been studied also for non-convex problems, see Alamo et al. (2009) and Ishii, Basar, and Tempo (2005).

The common feature of most of these sequential algorithms is the use of the validation strategy presented in Dabbene, Shcherbakov, and Polyak (2010) and Oishi (2007). The candidate solutions provided at each iteration of these algorithms are tested using a validation set which is drawn according to the probability measure associated to the uncertainty. If the candidate solution satisfies the design specifications for every sampled element of this validation set, then it is classified as probabilistic solution and the algorithm terminates. The main point in this validation scheme is that the cardinality of the validation set increases very mildly at each iteration of the algorithm. The strategy guarantees that, if a probabilistic solution is obtained, then it meets some probabilistic specifications.

In this paper, we derive the sample complexity for various analysis and design problems related to uncertain systems. In particular we provide new results which guarantee that the tail of the binomial distribution is bounded by a pre-specified value. These results are then applied to the analysis of worst-case performance and constraint violation. With regard to design problems, we consider the special cases of finite families and robust convex optimization problems. This contribution improves the existing results by reducing the number of samples required to solve the design problem. We remark that the results we have obtained are fairly general and the assumptions on convexity and on finite families appear only in Section 4 which deals with probabilistic analysis and design.

The second main contribution of this paper is to propose a sequential validation scheme, denoted as Sequential Probabilistic Validation (SPV), which allows the candidate solution to violate the design specifications for one (or more) of the members of the validation set. The idea of allowing some violations of the constraints is not new and can be found, for example, in the context of system identification (Bai, Cho, Tempo, & Ye, 2002), chance-constrained optimization (Campi & Garatti, 2011; Nemirovski & Shapiro, 2006) and statistical learning theory (Alamo et al., 2009). This scheme makes sense in the presence of soft constraints or when a solution satisfying the specifications for all the admissible uncertainty realizations cannot be found. In this way, we improve the existing results with this relaxed validation scheme that reduces the chance of not detecting the solution even when it exists. Furthermore, we also show that a strict validation scheme may not be well-suited for some robust design problems.

This paper is based on the previous works of the authors Alamo, Luque, Ramirez, and Tempo (2012) and Alamo et al. (2010b). However, some results are completely new (Property 4) and others (Theorem 2, Property 1 and Property 3 and their proofs) are significant improvements of the preliminary results presented in the conference papers. Furthermore, the unifying approach studied here, which combines sample complexity results with SPV algorithms, was not present in previous papers. Finally, the numerical example in Section 8, which compares various approaches available in the literature, is also new. The rest of the paper is organized as follows. In the next section, we first introduce the problem formulation. In Section 3, we provide bounds for the binomial distribution which are used in Section 4 to analyze the probabilistic properties of different schemes involving randomization. In Section 5, we introduce the proposed family of probabilistically validated algorithms. The sample complexity of the validating sets is analyzed in Section 6. A detailed comparison with the validation scheme presented in Oishi (2007) is provided in Section 7. A numerical example where different schemes are used to address a robust identification problem is presented in Section 8. The paper ends with a section of conclusions and an Appendix which contains some auxiliary properties and proofs that are used in the previous sections.

2. Problem statement

We assume that a probability measure $\Pr_{\mathcal{W}}$ over the sample space $\mathcal{W}$ is given. Given $\mathcal{W}$, a collection of N independent identically distributed (i.i.d.) samples $\mathrm{w} = \{w^{(1)}, \ldots, w^{(N)}\}$ drawn from $\mathcal{W}$ belongs to the Cartesian product $\mathcal{W}^N = \mathcal{W} \times \cdots \times \mathcal{W}$ (N times). Moreover, if the collection w of N i.i.d. samples $\{w^{(1)}, \ldots, w^{(N)}\}$ is generated from $\mathcal{W}$ according to the probability measure $\Pr_{\mathcal{W}}$, then the *multisample* w is drawn according to the probability measure $\Pr_{\mathcal{W}^N}$. The scalars $\eta \in (0, 1)$ and $\delta \in (0, 1)$ denote probabilistic parameters called accuracy and confidence, respectively. Furthermore, $\ln(\cdot)$ is the natural logarithm and e is the Euler number. For $x \in \mathbb{R}$, $x \geq 0$, $\lfloor x \rfloor$ denotes the largest integer smaller than or equal to x; $\lceil x \rceil$ denotes the smallest integer greater than or equal to x. For