

Distributed Model-Invariant Detection of Unknown Inputs in Networked Systems *

James Weimer
Department of Computer and
Information Science
School of Engineering and
Applied Science
University of Pennsylvania
Philadelphia, PA, USA
weimerj@seas.upenn.edu

Damiano Varagnolo
ACCESS Linnaeus Centre
School of Electrical
Engineering
KTH Royal Institute of
Technology
Stockholm, Sweden
damiano@kth.se

Karl Henrik Johansson
ACCESS Linnaeus Centre
School of Electrical
Engineering
KTH Royal Institute of
Technology
Stockholm, Sweden
kallej@kth.se

ABSTRACT

This work considers hypothesis testing in networked systems under severe lack of prior knowledge. In previous work we derived a centralized Uniformly Most Powerful Invariant (UMPI) approach to testing unknown inputs in unknown Linear Time Invariant (LTI) networked dynamics subject to unknown Gaussian noise. The detector was also shown to have Constant False Alarm Rate (CFAR) properties. Nonetheless, in large-scale systems, centralized testing may be infeasible or undesirable. Thus, we develop a distributed testing version of our previous work that utilizes a statistic that is maximally invariant to the unknown parameters and the non-local/neighbor measurements. Similar to the centralized approach, the distributed test is shown to have CFAR properties and to have performance that asymptotically approaches that of the centralized test. Simulation results illustrate that the performance of the distributed approach suffers marginal performance degradation in comparison to the centralized approach. Insight to this phenomena is provided through a discussion.

*The research leading to these results has received funding from the European Union Seventh Framework Programme [FP7/2007-2013] under grant agreement n°257462 HYCON2 Network of Excellence. The research was also supported by the Swedish Research Council and the Knut and Alice Wallenberg Foundation. Additionally, this material is based on research sponsored by DARPA under agreement number FA8750-12-2-0247. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HiCoNS'13, April 9–11, 2013, Philadelphia, Pennsylvania, USA.
Copyright 2013 ACM 978-1-4503-1961-4/13/04 ...\$15.00.

Categories and Subject Descriptors

H.1.1 [Information Systems]: Probability and Statistics-Models and Principles[Systems and Information Theory]

Keywords

Invariant Testing, Networked Systems

1. INTRODUCTION

Driven by the possibility of augmenting the flexibility and the reconfiguration capabilities of very complex systems, in many applications the current trend is to exploit multitudes of sensors and actuators, as in environmental monitoring [1], building energy management [2, 3], wireless communications [4] and power grids [5, 6]. The trend, however, comes with drawbacks: the high number of devices induces an increased possibility of faults with potentially disruptive ripple effects, like extended blackouts in power systems. There is thus a factual need for distributed fault detection algorithms.

We then consider that in every system, including dynamically networked ones such as the smart grid and building thermal dynamics, fault detection algorithms undoubtedly benefit from the knowledge of accurate models [6, 1, 3]. However, obtaining accurate models is often difficult or unrealistic due to the complexity of the system itself or the effects of environmental disturbances. For instance, in the smart grid security domain, it is common to assume the admittance of a transmission line is known [6]; however, the power line admittance is known to change with the temperature, humidity, and power flow, which leads to inaccurate models. Similarly, in building thermal dynamic modeling, even the simplest first-order heat equation model requires the knowledge of inter air-mass interactions, which change with the state of windows and doors (open or closed), the prevailing winds, the temperature, and the humidity. Thus, it is necessary to design fault detection schemes robust to these complex interactions.

If one were to consider large-scale networked systems, centralized approaches which apply model identification techniques in cascade with hypothesis testing may not be feasible. Similarly, when there are limited measurements, these identification and testing approaches tend to yield unexpected results, primarily due to the lack of information suit-

able for accurate parameter identification, see, e.g., [7, Example 1, page 46]. In this situation, distributed testing approaches that are designed to be invariant to the actual model parameters can result in better performance. In this paper we thus analyze if it is possible to derive distributed decision rules that do not depend on the model parameters and that are, in some sense to be defined, optimal with respect to the available information.

Literature review. Centralized classical hypothesis testing approaches usually use Generalized Likelihood Ratio (GLR) strategies, relying on obtaining Maximum Likelihood (ML) estimates of the unknown parameters under the various hypotheses and then testing their likelihood ratios. Maximally Invariant (MI) tests [8, Sec. 4.8] instead perform some additional preliminary operations so that the test is not influenced by the nuisance parameters. If MI tests are Uniformly Most Powerful Invariant (UMPI), then when the Signal to Noise Ratio (SNR) tends to infinity (e.g., when the number of measurements approaches infinity, see [9]), Generalized Likelihood Ratio (GLR) and UMPI strategies are asymptotically equivalent. When small datasets are available, nonetheless, MI tests can outperform GLR approaches [10].

Invariant strategies have been used in several applications, like detection of structural changes in linear regression models [11] or in spectral properties of disturbances [12]. The literature focuses mainly on finding invariant methods in linear models with unknown or partially known covariance matrices [13, 14, 15, 16, 17], with efforts specially in finding tests that exploit maximally invariant statistics and that have Constant False Alarm Rate (CFAR) properties.

Recently, there has been substantial research in distributed GLR tests for networked systems, e.g., in environmental monitoring, smart grid fault detection, and building HVAC failure detection and diagnostics applications. While all these approaches yield asymptotically accurate results as the number of measurements increases, their performance under limited measurements is sporadic and unpredictable. This motivates the need for distributed testing techniques which have predictable performance regardless of the number of measurements.

In our previous work [18], we considered the centralized detection of unknown inputs in unknown dynamically networked Linear Time Invariant (LTI) Gaussian systems and developed a UMPI test with CFAR properties. This work not only showed the existence on a UMPI test, but also established an upper bound on the performance of any distributed detection scheme.

Statement of contributions. here we again focus on LTI-Gaussian models, but reduce the prior information to be the smallest possible. More precisely, we assume the knowledge of *just* the fact that the system dynamics is networked, LTI with Gaussian driving noises and, furthermore, a weak knowledge on the structure of the input fault. We thus develop a distributed CFAR test that is invariant to the unknown parameters and the non-local/neighborhood measurements describing the system. The distributed test is then numerically evaluated against the centralized test developed in [18] as well as the best case (assuming a known model) and the worst case (assuming no model) scenarios, where

it is shown empirically that the distributed test approaches the performance of the centralized UMPI test.

Structure of the paper. Section 2 reports the needed basic results and definitions on invariant hypothesis testing. Section 3 formulates precisely the problem considered. We propose our testing technique along with its statistical characterization in Section 4. Section 5 numerically compares the performance of the distributed detector against the performance of the centralized UMPI detector in [18] and strategies endowed with more prior information and no prior information for different operating points and systems. Finally, Section 6 reports some concluding remarks and proposes future extensions.

Notation. we use plain lower case italic fonts to indicate scalars or functions with scalar range, bold lower case italic fonts to indicate vectors or functions with vector range, and plain upper case italic fonts to indicate matrices. We also use \otimes to denote Kronecker products, and $\mathbf{e}_{i,j}$ to denote the elementary vector of dimension i consisting of all zeros with a single unit entry in the j -th position.

2. HYPOTHESIS TESTING PRELIMINARIES

Commiserate with [8], we recall definitions and methodology employed in designing UMPI tests. Let \mathbf{y} be a r.v. with probability density $f(\mathbf{y}; \mathbf{d}, \boldsymbol{\delta})$ parametrized in \mathbf{d} , $\boldsymbol{\delta}$. We define \mathbf{d} to be the set of parameters of interest, and thus $\boldsymbol{\delta}$ to be the set of nuisance parameters, which induce a *transformation group* G , i.e., a set of endomorphisms g on the space of the realizations \mathbf{y} [8, Sec. 4.8]. This group of transformations partitions the measurement space into equivalence classes (or orbits) where points are considered equal if there exist $g, g' \in G$ mapping the first into the second and vice versa.

Definition 1 (Maximally Invariant Statistic [8]) A statistic $T[\mathbf{y}]$ is said to be maximally invariant w.r.t. a transformation group G if it is:

$$\text{invariant: } T[g(\mathbf{y})] = T[\mathbf{y}], \quad \forall g \in G$$

$$\text{maximal: } T[\mathbf{y}'] = T[\mathbf{y}''] \Rightarrow \exists g \in G \text{ s.t. } \mathbf{y}'' = g(\mathbf{y}').$$

A statistical test, ϕ , based on an invariant statistic is said to be an invariant test:

Definition 2 (Invariant Test [8, Sec. 4.8]) Let G be a transformation group, $T[\mathbf{y}]$ a statistic and $\phi(\cdot)$ a hypothesis test. ϕ is said to be invariant w.r.t. G if

$$\phi(T[g(\mathbf{y})]) = \phi(T[\mathbf{y}]) \quad (1)$$

for every $g \in G$.

The statistical performance of an invariant test ϕ is measured in terms of its *size* and *power*, where an invariant test is desired to be Uniformly Most Powerful Invariant (UMPI):

Definition 3 (Uniformly Most Powerful Invariant (UMPI) Test [8, Sec. 4.8]) Let G be a transformation group, $T[\mathbf{y}]$ a statistic and $\phi(\cdot)$ a test for deciding between H_0 and H_1 that is invariant w.r.t. G . Then $\phi(T[\mathbf{y}])$ is said to be an *uniformly most powerful invariant* (UMPI) test of size α if for every competing invariant test $\phi'(T[\mathbf{y}])$ it holds that

$$\begin{aligned} \text{(size)} \quad & \sup_{\mathbf{d}, \boldsymbol{\delta} \text{ under } H_0} \Pr[\phi(T[\mathbf{y}]) = H_1 \mid \mathbf{d}, \boldsymbol{\delta}] = \alpha; \\ & \sup_{\mathbf{d}, \boldsymbol{\delta} \text{ under } H_0} \Pr[\phi'(T[\mathbf{y}]) = H_1 \mid \mathbf{d}, \boldsymbol{\delta}] \leq \alpha; \end{aligned} \quad (2)$$

$$\begin{aligned} \text{(power)} \quad & \Pr[\phi(T[\mathbf{y}]) = H_1 \mid \mathbf{d}, \boldsymbol{\delta} \text{ under } H_1] \geq \\ & \Pr[\phi'(T[\mathbf{y}]) = H_1 \mid \mathbf{d}, \boldsymbol{\delta} \text{ under } H_1]. \end{aligned} \quad (3)$$

As a remark, thanks to the Karlin-Rubin theorem [8, Sec. 4.7, page 124], a scalar maximally invariant statistic whose likelihood ratio is monotone can be used to construct an UMPI test.

3. PROBLEM FORMULATION AND NOTATION

This section introduces a distributed hypothesis testing problem for deciding whether a signal, driven by unknown LTI networked Gaussian dynamics, lies also in a given subspace. Specifically, we consider a system of M interconnected nodes for which there exists an underlying interconnection graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, between the M nodes, where $\mathcal{V} := \{1, \dots, M\}$ is the vertex set, with $i \in \mathcal{V}$ corresponding to node i , and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the edge set of the graph. The undirected edge $\{i, j\}$ is incident on vertices i and j if nodes i and j share an interconnection, such that the neighborhood of node i , \mathcal{N}_i , is defined as

$$\mathcal{N}_i := \{j \in \mathcal{V} \mid \{i, j\} \in \mathcal{E}\} \quad (4)$$

The inter-node dynamics are governed by discrete-time LTI-Gaussian dynamics

$$\begin{aligned} x_j(k+1) &= x_j(k) + m_j \sum_{i \in \mathcal{N}_j} a_{ji} (x_i(k) - x_j(k)) \\ &\quad + b_j d_j(k) + w_j(k) \\ y_j(k) &= x_j(k) + v_j(k) \end{aligned} \quad (5)$$

where:

- $k = 0, \dots, T$ is the time index (T even for notational simplicity¹);
- $j = 1, \dots, M$ is the agent index;
- the states $x_j(k)$'s, measurements $y_j(k)$'s and inputs $d_j(k)$'s are scalar;

¹For ease of notation and without loss of generality we assume that the available measurements are over a given period whose length is fixed *ex ante*.

- $m_j a_{ji} = m_j a_{ij} \in \mathbb{R}$ and $b_j \in \mathbb{R}$ denote respectively the gains between $x_i(k)$ and $x_j(k+1)$, and between $d_j(k)$ and $x_j(k+1)$;
- $w_j(k), v_j(k) \in \mathbb{R}$ are uncorrelated i.i.d. Gaussian process noise and measurement noise with moments

$$\mathbb{E}[w_j(k)] = \chi_{j,w} \quad \mathbb{E}[v_j(k)] = \chi_{j,v},$$

$$\mathbb{E}[(w_j(k) - \bar{w}_j)^2] = \sigma_{j,w}^2 \quad \mathbb{E}[(v_j(k) - \bar{v}_j)^2] = \sigma_{j,v}^2.$$

To compact the notation we let, for $j = 1, \dots, M$,

$$\begin{aligned} A &:= [\alpha_{ij}] \\ \alpha_{ij} &:= \begin{cases} 1 - m_j \sum_{n \in \mathcal{N}_j} a_{nj} & \text{if } i = j \\ m_j a_{ij} & \text{if } i \in \mathcal{N}_j, \quad i \neq j \\ 0 & \text{otherwise} \end{cases} \\ B &:= \text{diag}[b_1, \dots, b_M] \\ \mathbf{y}_j &:= [y_j(0), \dots, y_j(T)]^\top \\ \mathbf{d}_j &:= [d_j(0), \dots, d_j(T)]^\top. \end{aligned}$$

Additionally, we consider the following quantities: let $\mathcal{N}_j = \{i_1, \dots, i_J\}$ be the sorted list of neighbors of agent j . Then

$$\begin{aligned} \vec{\alpha}_j &:= [\alpha_{i_1 j}, \dots, \alpha_{i_J j}]^\top \\ \vec{y}_j(k) &:= [y_{i_1}(k), \dots, y_{i_J}(k)]^\top \\ \vec{\mathbf{y}}_j &:= [\mathbf{y}_{i_1}^\top, \dots, \mathbf{y}_{i_J}^\top]^\top, \end{aligned}$$

i.e., $\vec{y}_j(k)$ is the set of the measurements of agent j and its neighbors (sorted lexicographically) at time k , while $\vec{\mathbf{y}}_j$ is the set of *all* the measurements of agent j and its neighbors (again sorted lexicographically).

Consider then a *specific* agent $\ell \in \{1, \dots, M\}$. The structure of the input \mathbf{d}_ℓ is assumed to be as follows:

- $\mathbf{u}_\ell := [u_\ell(0), \dots, u_\ell(T)]^\top$ is a *desired* and *known* input signal;
- $\mathbf{s}_\ell^f := [s_\ell^f(0), \dots, s_\ell^f(T)]^\top$, $f = 1, \dots, N_\ell$ are some *known* signals defining the space of signals $\text{span}\langle \mathbf{s}_\ell^1, \dots, \mathbf{s}_\ell^{N_\ell} \rangle$ (with $S_\ell := [\mathbf{s}_\ell^1, \dots, \mathbf{s}_\ell^{N_\ell}]$ being a shorthand for the \mathbf{s}_ℓ^f 's);
- $\boldsymbol{\theta}_\ell \in \mathbb{R}^{N_\ell}$ is an unknown (but constant) signal selection parameter.

Then

$$\mathbf{d}_\ell = S_\ell \boldsymbol{\theta}_\ell + \mu_\ell \mathbf{u}_\ell \quad (6)$$

where the scalar μ_ℓ is an unknown parameter.

Summarizing, the information owned by agent ℓ is either *available* or *unavailable* as follows:

Assumption 4 Available information:

- the time-series measurements \vec{y}_ℓ
- the local desired input signal \mathbf{u}_ℓ ;

- the local nuisance subspace S_ℓ ;
- the local weight m_ℓ ;
- the fact that the state dynamics are LTI-Gaussian, constant in time, and with $b_\ell \neq 0$.

Assumption 5 Unavailable information:

- all the time-series measurements but \vec{y}_j ;
- all the local desired input signals but \mathbf{u}_ℓ ;
- all the local nuisance subspaces but S_ℓ ;
- all the local weights but m_ℓ ;
- the weights A and B ;
- the moments of the process and measurement noises $\chi_{j,w}, \chi_{j,v}, \sigma_{j,w}^2, \sigma_{j,v}^2, j = 1, \dots, M$;
- the parameters θ_j and μ_j ;
- the initial conditions $x_1(0), \dots, x_M(0)$;
- the input signals $\mathbf{d}_1, \dots, \mathbf{d}_M$.

We then assume the unknown μ_ℓ to be either 0 or 1 and pose the following binary hypothesis testing problem:

Assumption 6 Structure of the fault μ_ℓ satisfies either one of the two following hypotheses:

$$\begin{aligned} H_0 \text{ (null hypothesis):} & \quad \mu_\ell = 0 \\ H_1 \text{ (alternative hypothesis):} & \quad \mu_\ell = 1 \end{aligned}$$

In words, both hypotheses assume the actual \mathbf{d}_ℓ to be unknown, since θ_ℓ is unknown, but with a fixed and known functional structure. H_1 additionally assumes the presence of a known input \mathbf{u}_ℓ .

*Our aim is thus: develop a distributed test that considers a **specific** agent $\ell \in \{1, \dots, M\}$, and decides among the hypotheses H_0 vs. H_1 in Assumption 6 using only the information in Assumption 4 and, at the same time, being invariant to the unavailable information in Assumption 5.*

We note that the problem formulated in this section is fundamentally different from the problem formulated in [18]. Indeed, the novel test should be computable distributedly *and* should be invariant also to the non-local measurements (in addition to all the unavailable information in [18]).

We thus aim to find a test that detects whether node ℓ has a fault independently of whether a fault exists at any other node $j \neq \ell$ (fault isolation) *and* maximizes the probability of detection (power) for any probability of false alarm (size), i.e., we require the detector to be UMPI. Formally, thus, we aim to solve the following:

Problem 7

1. find a statistic $T[\vec{y}_\ell]$ that satisfies Definition 1 (maximal invariance) w.r.t. the transformation group induced by nuisance parameters in Assumption 5;
2. find a test $\phi(T[\vec{y}_\ell])$ that satisfies Definition 3 (UMPI test) w.r.t. to the class of tests based on the previously introduced maximal invariant statistic $T[\vec{y}_\ell]$.

4. DISTRIBUTED INVARIANT TESTING

In this section we solve the previously posed problem and develop a distributed UMPI test that uses only local and neighboring measurements. The algorithm is based on the following novel result, solving the first part of Problem 7:

Theorem 8 A maximally invariant statistic that solves Problem 7-1 is

$$T[\mathbf{z}_\ell] = \frac{\mathbf{z}_\ell^\top P_\ell \mathbf{z}_\ell}{\frac{1}{N_\ell - 1} \mathbf{z}_\ell^\top (I_{N_\ell} - P_\ell) \mathbf{z}_\ell} \quad (7)$$

with

$$\begin{aligned} \mathbf{z}_\ell &:= F_\ell Q \mathbf{y}_\ell \\ P_\ell &:= \frac{F_\ell Q \mathbf{u}_\ell \mathbf{u}_\ell^\top Q^\top F_\ell^\top}{\mathbf{u}_\ell^\top Q^\top F_\ell^\top F_\ell Q \mathbf{u}_\ell} \\ N_\ell &:= \frac{k}{2} - \|\mathcal{N}_\ell\|_0 \end{aligned} \quad (8)$$

and where the exploited quantities satisfy

$$\begin{aligned} F_\ell^\top F_\ell &= I_{\frac{k}{2}} - \vec{Y}_\ell (\vec{Y}_\ell^\top \vec{Y}_\ell)^{-1} \vec{Y}_\ell^\top \\ Q &= I_{\frac{k}{2}} \otimes [0 \quad 1] \\ \vec{Y}_\ell &= \begin{bmatrix} \vec{y}_\ell^\top(0) & (s_\ell^f(0))^\top & 1 \\ \vec{y}_\ell^\top(2) & (s_\ell^f(2))^\top & 1 \\ \vec{y}_\ell^\top(4) & (s_\ell^f(4))^\top & 1 \\ \vdots & \vdots & \vdots \\ \vec{y}_\ell^\top(T) & (s_\ell^f(T))^\top & 1 \end{bmatrix} \end{aligned} \quad (9)$$

PROOF. The proof for Theorem 8 is provided in the appendix.

We observe that the maximally invariant statistic in (7) can be equivalently written as a ratio of independent chi-square random variables. This particular ratio is known to follow an F -distribution, which has a monotone likelihood ratio [8]. Thus we solve the second part of Problem 7 by applying the Karlin-Rubin theorem, obtaining directly the following:

Corollary 9 A distributed UMPI test of size α for Problem 7-2 is

$$\phi_\ell(\mathbf{z}_\ell) = \begin{cases} H_0 & \text{if } T_\ell[\mathbf{z}_\ell] < \mathcal{F}_{1, N_\ell-1}^{-1}(\alpha) \\ H_1 & \text{otherwise.} \end{cases} \quad (10)$$

where $\mathcal{F}_{n,m}^{-1}(\alpha)$ is the inverse central cumulative F -distribution of dimensions n and m .

We remark that, w.r.t. the algorithm proposed in [18], test (10) can be performed in parallel and it is invariant to the non-local measurements. This comes with a price: the test exploits only about half of the available measurements (either local or from neighbors). The remaining local and neighbors' measurements are in fact lost in the attempt of obtaining invariance. Since the data set is smaller than the one exploited in [18], it is expected that the novel test will perform worse. In the following section we then numerically evaluate this loss.

5. NUMERICAL EXAMPLES

We perform three Monte-Carlo characterizations as follows:

1. we fix a desired probability of false alarms α (0.01, 0.1 and 0.25);
2. we randomly generate 500 stable networked systems of 10 agents like (5) as described in Table 1 (i.e., we discarded the unstable realizations);
3. for each of the 500 systems (5) we generated exactly one realization $y_j(1), \dots, y_j(500)$, $j = 1, \dots, 10$;
4. for each $T = 1, \dots, 500$ and each of the 500 systems (5) we executed the following four tests, all with the same desired probability of false alarms α :
 - (a) full information test: assume the perfect knowledge of the weights A and B ; the moments of the process and measurement noises $\chi_{j,w}$, $\chi_{j,v}$, $\sigma_{j,w}^2$, $\sigma_{j,v}^2$; the parameters θ_j ; the initial conditions $x_1(j)$ ($j = 1, \dots, 10$). Then design the Uniformly Most Powerful (UMP) test for testing H_0 vs. H_1 given all this information;
 - (b) centralized UMPI test: the UMPI test developed in [18], which is provided in the appendix using the notation introduced within this work;
 - (c) distributed UMPI test (DUMPI): our test (10);
 - (d) no information test: perform a weighted coin flip s.t. the desired probability of false alarms α is met.

The outcomes are then summarized in the following Figures 1, 2 and 3, that plot for each test and each T the average correct detection rate reached over the 500 considered realizations of system 5.

From the previous graphics we draw the following conclusions. Before the number of measurements (proportional to T) passes the threshold $\frac{T}{2} - N_\ell - M + 1$ (independent of the chosen α), both the centralized and distributed UMPI tests

$a_j, b_j \sim \mathcal{U}[-0.5, 0.5]$	$m_j \sim \mathcal{U}[1, 2]$
$\chi_{j,w}, \chi_{j,v} \sim \mathcal{N}(0, 1)$	$\sigma_{j,w}^2, \sigma_{j,v}^2 \sim \mathcal{U}[0.1, 1]$

Table 1: Random extraction mechanisms for the generation of the systems (5). \mathcal{N} indicates Gaussian distributions, \mathcal{U} uniform distributions. All the quantities are extracted independently.




Figure 1: Monte-Carlo characterization of the detection tests given $\alpha = 0.01$.




Figure 2: Monte-Carlo characterization of the detection tests given $\alpha = 0.1$. Legend as in Figure 1.




Figure 3: Monte-Carlo characterization of the detection tests given $\alpha = 0.25$. Legend as in Figure 1.

are equivalent to a coin flipping (since the amount of information is insufficient to take meaningful decisions). After that threshold, instead, the two test start increasing their correct detection rates (with different speeds, depending on the selected probability of false alarms), discerning better and better. Eventually they reach the same performance of the full information-based test, i.e., the best one might desire. We then notice that the difference in the correct detection rates between the centralized and distributed approaches starts small and vanishes quickly. This indicates that, from practical purposes, the distributed strategy performs well. The reason for such a similar performance between the centralized and distributed approaches lies in that the centralized approach from [18] (also provided in the appendix of this extended abstract), effectively disregards half of the measurements to achieve maximal invariance. In the distributed approach, the same measurements that are discarded by the centralized approach are employed to provide invariance to the local inter-node dynamics.

6. DISCUSSION AND FUTURE WORKS

We considered fault detection in networked Linear Time Invariant-Gaussian systems. More precisely, we defined a hypothesis testing problem over the structure of the inputs of the agents, and then derived a distributed Uniformly Most Powerful Invariant detector with Constant False Alarm Rate properties that is invariant to most of the parameters of the systems. We address the situation where there is little prior information available, and develop a distributed test starting from our previous centralized results described in [18]. Remarkably we obtain a distributed algorithm that has some capability of detecting faults even if knowledge of the overall system is really uncertain and the number of measurements is limited.

As in the centralized case, tests that exploit information of the system have better performance in terms of false positives / negatives rates. Nonetheless, the more measurements that are taken the more the distributed detector is shown to be perform better, achieving performance of its centralized counterpart quickly.

The value of the proposed strategy relies in its optimality properties, being in fact based on a maximally invariant statistic and being uniformly most powerful. This implies that in a certain sense it characterizes the performance that can be achieved when testing the posed hypotheses under the severe lack of knowledge assumed here.

The main future direction is thus to compare the developed strategy, both from practical and theoretical aspects, with the distributed fault detection algorithm that are based on dynamically identified systems. It is in fact necessary to understand if there are conditions s.t. the invariant test developed here is guaranteed to perform better than algorithms that start identifying the test and then perform tests on the identified model. Additionally, extensions of this work to feedback control applications where stabilizing control is desired in the presence of unknown parameters and disturbances is planned.

7. REFERENCES

- [1] J. Weimer, B. Sinopoli, and B. Krogh, "An approach to leak detection using wireless sensor networks at carbon sequestration sites," *International Journal of Greenhouse Gas Control*, vol. 9, pp. 243–253, 2012.
- [2] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, no. 0, pp. 15 – 27, 2012.
- [3] J. Weimer, S. Ahmadi, J. Araujo, F. Mele, D. Papale, I. Shames, H. Sandberg, and K. Johansson, "Active actuator fault detection and diagnostics in hvac systems," in *ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys)*, 2012.
- [4] T. Arampatzis, J. Lygeros, and S. Manesis, "A survey of applications of wireless sensors and wireless sensor networks," in *IEEE International Symposium on Intelligent Control, MCCA*, 2005, pp. 719–724.
- [5] S. Bolognani, A. Carron, A. Di Vittorio, D. Romeres, L. Schenato, and S. Zampieri, "Distributed multi-hop reactive power compensation in smart micro-grids subject to saturation constraints," in *51st IEEE Conference on Decision and Control*, 2012, pp. 785–790.
- [6] I. Shames, A. Teixeira, H. Sandberg, and K. Johansson, "Distributed fault detection for interconnected systems," *Automatica*, vol. 47, no. 12, pp. 2757–2764, 2011.
- [7] V. N. Vapnik, *Statistical learning theory*. New York: Wiley, 1998.
- [8] L. L. Scharf, *Statistical Signal Processing, Detection, Estimation, and Time Series Analysis*. Addison-Wesley Publishing Company Inc., Reading, Massachusetts, 1991.
- [9] J. Gabriel and S. Kay, "On the relationship between the glrt and umpi tests for the detection of signals with unknown parameters," *Signal Processing, IEEE Transactions on*, vol. 53, no. 11, pp. 4194–4203, 2005.
- [10] H. Kim and A. Hero III, "Comparison of glr and invariant detectors under structured clutter covariance," *Image Processing, IEEE Transactions on*, vol. 10, no. 10, pp. 1509–1520, 2001.
- [11] M. Hušková, "Some invariant test procedures for detection of structural changes," *Kybernetika*, vol. 36, no. 4, pp. 401–414, 2000.
- [12] N. Begum and M. L. King, "Most mean powerful invariant test for testing two-dimensional parameter spaces," *Journal of Statistical Planning and Inference*, vol. 134, no. 2, pp. 536 – 548, 2005.
- [13] S. Bose and A. Steinhardt, "A maximal invariant framework for adaptive detection with structured and unstructured covariance matrices," *Signal Processing, IEEE Transactions on*, vol. 43, no. 9, pp. 2164–2175, 1995.
- [14] K. Noda and H. Ono, "On ump invariant f-test procedures in a general linear model," *Communications in Statistics-Theory and Methods*, vol. 30, no. 10, pp. 2099–2115, 2001.
- [15] E. Conte, A. De Maio, and C. Galdi, "Cfar detection of multidimensional signals: an invariant approach," *Signal Processing, IEEE Transactions on*, vol. 51, no. 1, pp. 142–151, 2003.
- [16] A. De Maio, "Rao test for adaptive detection in gaussian interference with unknown covariance

matrix,” *Signal Processing, IEEE Transactions on*, vol. 55, no. 7, pp. 3577–3584, 2007.

- [17] A. De Maio and E. Conte, “Adaptive detection in gaussian interference with unknown covariance after reduction by invariance,” *Signal Processing, IEEE Transactions on*, vol. 58, no. 6, pp. 2925–2934, 2010.
- [18] J. Weimer, D. Varagnolo, M. Stankovic, and K. Johansson, “Model-invariant detection of unknown inputs in networked systems,” in *European Control Conference (under review)*, 2013.

Appendix

This appendix provides a proof for Theorem 8. To identify a maximally invariant statistic requires identifying the the group of transformations induced by the unknown parameters. Identifying this group is achieved writing the measurement dynamics for node ℓ in 5 as

$$y_\ell(k+1) = y_\ell(k) + m_\ell \sum_{i \in \mathcal{N}_\ell} a_{\ell i} (y_i(k) - y_\ell(k)) + b_\ell d_\ell(k) + n_\ell(k) \quad (11)$$

where

$$n_\ell(k) = w_\ell(k) + v_\ell(k+1) - \left(1 - m_\ell \sum_{i \in \mathcal{N}_\ell} a_{\ell i}\right) v_\ell(k) - m_\ell \sum_{i \in \mathcal{N}_\ell} a_{\ell i} v_i(k). \quad (12)$$

We write the time-series concatenation of the measurements as

$$\mathbf{y}_\ell = \tilde{H}_\ell \boldsymbol{\rho} + b_\ell \mu_\ell \mathbf{u}_\ell + \mathbf{n}_\ell \quad (13)$$

where

$$\tilde{H}_\ell = \begin{bmatrix} \tilde{\mathbf{y}}_\ell^\top(0) & (s_\ell^f(0))^\top & 1 \\ \tilde{\mathbf{y}}_\ell^\top(1) & (s_\ell^f(1))^\top & 1 \\ \vdots & \vdots & \vdots \\ \tilde{\mathbf{y}}_\ell^\top(T) & (s_\ell^f(T))^\top & 1 \end{bmatrix} \quad (14)$$

$$\mathbf{n}_\ell = [n_\ell(0), n_\ell(1), \dots, n_\ell(T)]^\top$$

$$\text{Cov}[\mathbf{n}_\ell] = \sigma_0^2 I + \sigma_1^2 \sum_{i=0}^{\frac{T}{2}} \left(\mathbf{e}_{2i} \mathbf{e}_{2i+1}^\top + \mathbf{e}_{2i+1} \mathbf{e}_{2i}^\top \right)$$

and $\boldsymbol{\rho}$ is a vector of unknown parameters such that each row of the time-series measurements is equivalent to the dynamics in (11). The unknown parameters induce a group of transformations on the measurements

$$G = \left\{ g \mid g(\mathbf{y}_\ell) = \sigma_0 \left(I - \sum_{i=0}^{\frac{T}{2}} c_i \mathbf{e}_{2i} \mathbf{e}_{2i+1}^\top \right) \mathbf{y}_\ell + \tilde{H}_\ell \boldsymbol{\rho} + \mu_\ell b_\ell \mathbf{u}_\ell \right\} \quad (15)$$

where $c_i \in \mathbb{R}$ is an unknown gain induced by the unknown correlation and noise realization (which varies with time). It then follows that Theorem 8 is maximally invariant to the transformation group induced by the unknown parameters.

PROOF. Invariance: Observing the following:

$$Q \left[\sigma_0 \left(I - \sum_{i=0}^{\frac{T}{2}} c_i \mathbf{e}_{2i} \mathbf{e}_{2i+1}^\top \right) \mathbf{y}_\ell + \tilde{H}_\ell \boldsymbol{\rho} + \mu_\ell b_\ell \mathbf{u}_\ell \right] \quad (16)$$

$$= \sigma_0 Q \mathbf{y}_\ell + \tilde{Y}_\ell \boldsymbol{\rho} + \mu_\ell b_\ell Q \mathbf{u}_\ell$$

and

$$F_\ell \left(\sigma_0 Q \mathbf{y}_\ell + \tilde{Y}_\ell \boldsymbol{\rho} + \mu_\ell b_\ell Q \mathbf{u}_\ell \right) = \sigma_0 \mathbf{z}_\ell + \mu_\ell b_\ell F_\ell Q \mathbf{u}_\ell \quad (17)$$

then

$$T[g(\mathbf{y}_\ell)] = \frac{\sigma_0^2 \mathbf{z}_\ell^\top P_\ell \mathbf{z}_\ell}{\sigma_0^2 \frac{1}{N_\ell - 1} \mathbf{z}_\ell^\top (I_{N_\ell} - P_\ell) \mathbf{z}_\ell} = \frac{\mathbf{z}_\ell^\top P_\ell \mathbf{z}_\ell}{\frac{1}{N_\ell - 1} \mathbf{z}_\ell^\top (I_{N_\ell} - P_\ell) \mathbf{z}_\ell} = T[\mathbf{y}_\ell] \quad (18)$$

Maximality: We observe that

$$T[\hat{\mathbf{z}}_\ell] = T[\mathbf{z}_\ell]$$

$$\rightarrow \frac{\mathbf{z}_\ell^\top P_\ell \mathbf{z}_\ell}{\mathbf{z}_\ell^\top (I - P_\ell) \mathbf{z}_\ell} = \frac{\hat{\mathbf{z}}_\ell^\top P_\ell \hat{\mathbf{z}}_\ell}{\hat{\mathbf{z}}_\ell^\top (I - P_\ell) \hat{\mathbf{z}}_\ell} \quad (19)$$

$$\rightarrow \hat{\mathbf{z}}_\ell^\top \left(P_\ell - I \frac{\mathbf{z}_\ell^\top P_\ell \mathbf{z}_\ell}{\mathbf{z}_\ell^\top \mathbf{z}_\ell} \right) \hat{\mathbf{z}}_\ell = 0$$

$$\rightarrow \mathbf{u}_\ell^\top Q^\top F_\ell^\top \hat{\mathbf{z}}_\ell = c \mathbf{u}_\ell^\top Q^\top F_\ell^\top \mathbf{z}_\ell, \quad \exists c \in \mathbb{R}.$$

and complete the proof for maximality as

$$\mathbf{u}_\ell^\top Q^\top F_\ell^\top \hat{\mathbf{z}}_\ell = c \mathbf{u}_\ell^\top Q^\top F_\ell^\top \mathbf{z}_\ell$$

$$\rightarrow \hat{\mathbf{y}}_\ell = c \mathbf{y}_\ell + (I - P_\ell) (c \mathbf{y}_\ell - \hat{\mathbf{y}}_\ell) \quad (20)$$

$$\rightarrow \hat{\mathbf{y}}_\ell = g(\mathbf{y}_\ell), \quad \exists g \in G$$