International Journal of Control

ISSN: 0020-7179 (Print) 1366-5820 (Online) Journal homepage: http://www.tandfonline.com/loi/tcon20

Sensor selection for fault diagnosis in uncertain systems

Daniel Jung, Yi Dong, Erik Frisk, Mattias Krysander & Gautam Biswas

To cite this article: Daniel Jung, Yi Dong, Erik Frisk, Mattias Krysander & Gautam Biswas (2018): Sensor selection for fault diagnosis in uncertain systems, International Journal of Control, DOI: 10.1080/00207179.2018.1484171

To link to this article: https://doi.org/10.1080/00207179.2018.1484171

Accepted author version posted online: 09 Jun 2018.
Published online: 24 Jun 2018.

INTERNATIONAL JOURNAL OF CONTROL
https://doi.org/10.1080/00207179.2018.1484171

Taylor & Francis
Taylor & Francis Group

Sensor selection for fault diagnosis in uncertain systems

Daniel Jung [a], Yi Dong[b], Erik Frisk[a], Mattias Krysander[a] and Gautam Biswas[b]

[a]Department of Electrical Engineering, Linköping University, Linköping, Sweden; [b]Institute for Software Integrated Systems, Vanderbilt University, Nashville, TN, USA

ABSTRACT

Finding the cheapest, or smallest, set of sensors such that a specified level of diagnosis performance is maintained is important to decrease cost while controlling performance. Algorithms have been developed to find sets of sensors that make faults detectable and isolable under ideal circumstances. However, due to model uncertainties and measurement noise, different sets of sensors result in different achievable diagnosability performance in practice. In this paper, the sensor selection problem is formulated to ensure that the set of sensors fulfils required performance specifications when model uncertainties and measurement noise are taken into consideration. However, the algorithms for finding the guaranteed global optimal solution are intractable without exhaustive search. To overcome this problem, a *greedy stochastic search* algorithm is proposed to solve the sensor selection problem. A case study demonstrates the effectiveness of the greedy stochastic search in finding sets close to the global optimum in short computational time.

ARTICLE HISTORY
Received 6 November 2017
Accepted 28 May 2018

KEYWORDS
Fault diagnosis; fault detection and isolation; sensor selection

1. Introduction

In model-based diagnosis, mathematical models describing the monitored system are used to compare observed signals with the corresponding modelled signals to detect anomalies (Nyberg, 2002). Finding the optimal set of sensors to fulfil the fault detection and isolation requirements is important but can, in general, be computationally intractable due to exponential complexity properties. More sensors will give better fault diagnosis, i.e. fault detection and isolation, performance but also increase the sensor cost and require more space to fit all sensors (Bhushan, Narasimhan, & Rengaswamy, 2008). Therefore, the primary goal of this paper is to find, given a set of candidate sensors and a specified cost of using each sensor, a cheapest subset of sensors that fulfils the fault detection and isolation requirements. If the sensor cost is equal for all sensors the cheapest set is the minimum cardinality set.

A common approach to formulate fault diagnosability requirements when defining the sensor selection problem is to use fault detectability and isolability (see, e.g. Krysander & Frisk, 2008; Yassine, Ploix, & Flaus, 2008). Fault detectability and isolability are deterministic performance measures and describe whether faults can be detected and isolated or not in the ideal case, i.e. the measures can answer questions such as: *'Is it possible to detect a fault f_i?'* or *'Is a fault f_i isolable from another fault f_j?'* 'Is it possible to detect a fault f_i?' or 'Is a fault f_i isolable from another fault f_j?' (Eriksson, Frisk, & Krysander, 2013). One problem of using qualitative fault detectability and isolability to formulate performance requirements is that there is no way of specifying how easy it should be to detect or isolate different faults of different magnitudes.

An important factor when considering the sensor selection problem is the negative impact of model uncertainties and measurement noise on the performance of a diagnosis system. Large uncertainties complicate detection and isolation of small faults. Even if a set of sensors fulfils the deterministic fault detectability and isolability requirements, it is not certain that a diagnosis system based on these sensors will meet performance requirements when model uncertainties and measurement noise are taken into consideration. In Bhushan et al. (2008), it is emphasised that the importance of reducing cost or increasing robustness have a significant impact on the required number of sensors.

Another motivation is the increasing availability of cheap sensors. It can be more cost effective to use a large number of cheap sensors in the system instead of a few expensive ones with higher accuracy to achieve the same fault detection and isolation performance. With model uncertainties and measurement noise, the sensor selection problem needs to include a more realistic evaluation of diagnosability performance than, say, just finding a minimum cardinality sensor sets that implies fault detectability and isolability under ideal conditions. Typically, there are requirements on the diagnosis system, such as probability of false alarms and probability of missed detections given different faults. If these requirements can be translated to the sensor selection problem, a feasible solution would then assure that the corresponding diagnosis system can be developed with satisfactory performance.

A quantitative measure of diagnosability performance based on the Kullback–Leibler divergence, called *distinguishability*, is proposed in Eriksson et al. (2013). The same measure is also proposed in Harrou, Fillatre, and Nikiforov (2014).

CONTACT Daniel Jung daniel.jung@liu.se Department of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden

Distinguishability is used to quantify fault detection and isolation performance for a given model by taking measurement noise and model uncertainties into consideration. An important property of the distinguishability measure is that it is a model property and gives an upper bound of the maximum fault-to-noise ratio for any linear residual generator. Thus, by evaluating distinguishability of the model, the achievable performance of a diagnosis system can be predicted which is useful in the early diagnosis system design process. This is illustrated in Eriksson, Krysander, and Frisk (2012) where distinguishability is proposed when defining the sensor selection problem. It is shown that the distinguishability requirements have a significant impact on the number of sensors that are required to achieve the desired performance. In general, stricter distinguishability requirements result in higher solution cost since more sensors are needed. With respect to these previous works, a greedy stochastic search algorithm is proposed to solve the sensor selection problem. It is also shown how to include requirements on false alarm rates and missed detection rates when formulating the sensor selection problem using the distinguishability measure.

In Huber, Kopecek, and Hofbaur (2014), distinguishability is used to select a set of sensors which are good for estimating faults in an internal combustion (IC) engine using an extended Kalman filter. The nonlinear model of the IC engine was linearised at different operating points to capture distinguishability variations at different operating points. The number of available sensors was relatively small and, therefore, an exhaustive search was used. Thus, the work in Huber et al. (2014) did not explore the search problem which is a main topic here.

This work extends the results of the previous work in Eriksson et al. (2012), by analysing properties of the sensor selection problem and proposing an efficient algorithm for solving the problem. The main contribution is a formulation of the sensor selection problem where model uncertainties and measurement noise are taken into consideration when formulating the performance requirements using the distinguishability measure. It is shown how to formulate the quantitative fault detection and isolation performance requirements based on required false alarm rate and missed detection rate. This is important to assure that a diagnosis system can be designed, based on the selected sensor set, that fulfils the specified fault detection and isolation performance requirements. A second contribution is an analysis of the properties of the sensor selection problem. Based on the results from the analysis, a heuristic greedy stochastic search algorithm is proposed. A case study is used to show the effectiveness of the proposed algorithm with respect to other heuristic search algorithms.

2. Problem statement

Before formulating the sensor selection problem, a short discussion about modelling faults and the distinguishability measure is presented.

2.1 Modelling faults

The goal here is to solve the sensor selection problem, given a predetermined set of possible faults $\mathcal{F} = \{f_1, f_2, \ldots, f_{l_f}\}$ and a diagnosis requirement specification.

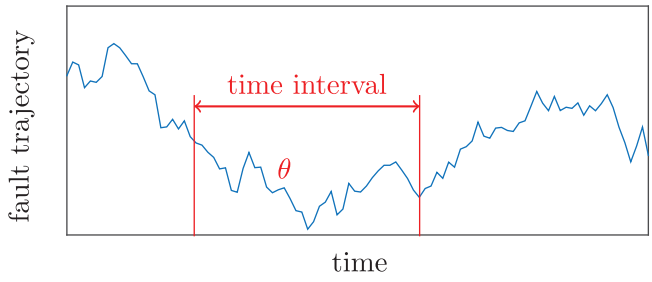

Figure 1. The fault time profile θ describes the fault trajectory during a given time interval.

A fault f_i is modelled as an unknown signal affecting the system (Eriksson et al., 2013) where $f_i = 0$ represents the fault-free case. It is assumed that each fault f_i can have any fault realisation, for example, a constant and a ramp, but also more complex realisations, as illustrated in Figure 1. The fault time profile θ is a vector describing the fault trajectory during a specific time interval. When formulating fault detection and isolation performance requirements, different fault time profiles θ can be used to represent different fault realisations that should be detectable and isolable.

2.2 Formulation of performance requirements for sensor sets

The sensor selection problem is formulated using a distinguishability measure that takes model uncertainties and measurement noise into consideration. The notion of the distinguishability measure (Eriksson et al., 2013) is described in detail in Section 4, but for the problem formulation it is sufficient to know that distinguishability, denoted $\mathcal{D}_{i,j}(\theta)$, quantifies the difficulty of isolating a fault f_i, with a given fault time profile θ, from another fault f_j. Each fault can have any fault realisation but the distinguishability measure is computed for a given fault scenario. Distinguishability can be computed for different fault time profiles to evaluate how the fault trajectory and magnitude affect detection and isolation performance. The distinguishability measure is non-negative where a larger value of $\mathcal{D}_{i,j}(\theta)$ corresponds to an easier fault isolation problem. The case $\mathcal{D}_{i,j}(\theta) = 0$ corresponds to the situation where f_i cannot be isolated from f_j.

Let $\mathcal{S} = \{y_1, y_2, \ldots, y_k\}$ be a set of k candidate sensors, where each sensor $y_i \in \mathcal{S}$ is assumed fault-free and the cost to use that sensor is denoted as c_i. The objective is to find a cheapest set of sensors $S \subseteq \mathcal{S}$, which fulfils a set of performance requirements defined by minimum required distinguishability $\mathcal{D}^{\text{req}}_{i,j}(\theta)$ for a predetermined set of fault realisations $\theta \in \Theta_i$ for each fault $f_i \in \mathcal{F}$. Each set Θ_i can be interpreted as a selected subset of fault time profiles a diagnosis system should be able to detect and isolate, i.e. the selected sensor set should fulfil the requirements $\mathcal{D}^{\text{req}}_{i,j}(\theta)$ for each $\theta \in \Theta_i$. The maximum achievable distinguishability is denoted as $\mathcal{D}^{\max}_{i,j}(\theta)$ and is the computed value of $\mathcal{D}_{i,j}(\theta)$ when all available candidate sensors are used. Then, each requirement $\mathcal{D}^{\text{req}}_{i,j}(\theta)$ is selected in the interval $0 \leq \mathcal{D}^{\text{req}}_{i,j}(\theta) \leq \mathcal{D}^{\max}_{i,j}(\theta)$ and defines a lower bound of $\mathcal{D}_{i,j}(\theta)$. Note that $\mathcal{D}^{\text{req}}_{i,j}(\theta)$ can be computed based on other performance requirements which are discussed in Section 5.

The reliability of each candidate sensor in $\mathcal{S}$ is not taken into consideration in the sensor placement problem. However, the increased lifetime cost by using a non-reliable sensor could be included when defining the sensor cost in the optimisation problem.

2.3 Sensor selection problem formulation

The sensor selection problem is formulated as follows.

Problem formulation 2.1: *Let $\mathcal{S} = \{y_1, y_2, \ldots, y_k\}$ be a set of k candidate sensors, $\mathcal{F} = \{f_1, f_2, \ldots, f_{l_f}\}$ a set of faults, and $\mathcal{D}^{\text{req}}_{i,j}(\theta)$ solution constraints. Then, the sensor selection problem is formulated as*

$$\begin{aligned} \min_{S\subseteq\mathcal{S}} \quad & \sum_{y_l\in S} c_l \\ \text{s.t.} \quad & \mathcal{D}^{S}_{i,j}(\theta) \geq \mathcal{D}^{\text{req}}_{i,j}(\theta), \quad \forall\theta \in \Theta_i \\ & \forall f_i, f_j \in \mathcal{F}, \end{aligned} \tag{1}$$

where $\sum_{y_l\in S} c_l$ is the solution cost and $\mathcal{D}^{S}_{i,j}(\theta)$ denotes the computed distinguishability for a given set of sensors S. The sensor candidates $\mathcal{S}$ are here assumed to be fault-free. However, the system can include already mounted sensors, not part of the sensor candidates $\mathcal{S}$, which can be faulty.

Note that distinguishability requirements regarding isolability from multiple faults can also be taken into consideration in Equation (1) by computing distinguishability from each fault to a set of other faults (see Eriksson, Frisk, & Krysander, 2012).

3. Related research

As mentioned in Section 1, many proposed sensor selection problems are formulated using fault detectability and isolability. In Frisk, Krysander, and Åslund (2009), an analytical approach for linear differential algebraic equation models is used to find all minimal sets of sensors that fulfil the required performance. In Dong and Biswas (2013) and Narasimhan, Mosterman, and Biswas (1998), a minimal solution to the sensor placement problem is found using the A* algorithm. In Casillas, Puig, Garza-Castañón, and Rosich (2013), a genetic algorithm is applied to select sensors for leakage detection in a water distribution network. In Krysander and Frisk (2008), an exhaustive search strategy is proposed which finds all minimal solutions. A greedy search strategy for finding a set of sensors fulfilling the diagnosability requirements is proposed in Raghuraj, Bhushan, and Rengaswamy (1999). Another greedy approach is proposed in Perelman, Abbas, Koutsoukos, and Amin (2016) which utilises the submodularity property of the sensor selection problem to significantly reduce computational time. In Rosich, Sarrate, and Nejjari (2009), the sensor placement problem is formulated as a binary integer linear programming problem. In Daigle, Roychoudhury, and Bregon (2014) and Travé-Massuyès, Escobet, and Olive (2006), the sensor selection problem is formulated as a test selection problem where the set of selected tests should use as few sensors as possible. In these previous works, fault diagnosability performance requirements are formulated in the sensor selection problem using fault detectability and isolability. For nonlinear models, a structural description of the system is used to evaluate fault detectability and isolability performance (see, e.g. Chi, Wang, & Zhu, 2015; Commault & Dion, 2007; Yassine et al., 2008).

There are also previous sensor selection research where the effects of different types of uncertainties are taken into consideration. In Bhushan et al. (2008), the probability of faults and sensor failures are taken into consideration when formulating the sensor placement problem. In Wu, Hsieh, and Li (2013), the cause–effect relations between faults and sensors are represented by a fuzzy graph and different quantitative factors are taken into consideration, such as sensor quality and sensitivity to different faults. In Namburu, Azam, Luo, Choi, and Pattipati (2007), diagnostic accuracy is taken into consideration in the selection problem where a data-driven diagnosis algorithm is designed. All mentioned works are closely related to this work, and a main contribution is that the effects of model uncertainties, measurement noise, fault realisations, and allowed time to detect are taken into consideration in addition to the search properties of the optimal selection problem.

4. Theoretical background

The quantitative diagnosability analysis used in the paper is here briefly reviewed. For detailed descriptions, the interested reader is referred to Eriksson et al. (2013).

4.1 Model

The class of models considered are linear descriptor models with additive faults represented as

$$\begin{aligned} Ex[t+1] &= Ax[t] + B_u u[t] + B_f f[t] + B_v v[t], \\ y[t] &= Cx[t] + D_u u[t] + D_f f[t] + D_\varepsilon \varepsilon[t], \end{aligned} \tag{2}$$

where t denotes time index, $x \in \mathbb{R}^{l_x}$ are unknown variables, $y \in \mathbb{R}^{l_y}$ are measured signals, $u \in \mathbb{R}^{l_u}$ are input signals, $f \in \mathbb{R}^{l_f}$ are modelled faults, and $v \sim N(0, \Lambda_v)$ and $\varepsilon \sim N(0, \Lambda_\varepsilon)$ are i.i.d. Gaussian random vectors, representing model uncertainties and noise, respectively, with zero mean and known symmetric positive definite covariance matrices $\Lambda_v \in \mathbb{R}^{l_v \times l_v}$ and $\Lambda_\varepsilon \in \mathbb{R}^{l_\varepsilon \times l_\varepsilon}$. The notation l_α denotes the number of elements in the vector α. All matrices are known and if l_q denotes the number of system equations, then $E \in \mathbb{R}^{l_q \times l_x}$, $A \in \mathbb{R}^{l_q \times l_x}$, $B_u \in \mathbb{R}^{l_q \times l_u}$, $B_f \in \mathbb{R}^{l_q \times l_f}$, $B_v \in \mathbb{R}^{l_q \times l_v}$, $C \in \mathbb{R}^{l_y \times l_x}$, $D_u \in \mathbb{R}^{l_y \times l_u}$, $D_f \in \mathbb{R}^{l_y \times l_f}$, and $D_\varepsilon \in \mathbb{R}^{l_y \times l_\varepsilon}$. Note that E can be singular.

Each sensor candidate $y_l \in \mathcal{S}$ measures one unknown variable x_i with additive i.i.d. Gaussian noise and is described by the model $y_l[t] = x_i[t] + \varepsilon_l[t]$, where $\varepsilon_l \sim \mathcal{N}(0, \sigma_l^2)$. Thus, when adding a set of candidate sensors S, represented by the vector $y^S[t]$, to model (2), the measurement equations are modified as follows:

$$\begin{aligned} \begin{pmatrix} y[t] \\ y^S[t] \end{pmatrix} &= \begin{pmatrix} C \\ C^S \end{pmatrix} x[t] + \begin{pmatrix} D_u \\ 0 \end{pmatrix} u[t] \\ &\quad + \begin{pmatrix} D_f \\ 0 \end{pmatrix} f[t] + \begin{pmatrix} D_\varepsilon & 0 \\ 0 & I \end{pmatrix} \begin{pmatrix} \varepsilon[t] \\ \varepsilon^S[t] \end{pmatrix}, \end{aligned} \tag{3}$$

where the uppercase S refers to the candidate sensors in S and C^S is a zero matrix with ones at positions corresponding to the measured variables, I is the identity matrix, and ε^S is a noise vector with the corresponding noise covariances.

Besides the magnitude and time profile of the fault, the difficulty to detect or isolate a fault also depends on allowed time to fault detection or time to fault isolation (see, e.g. Basseville & Nikiforov, 1993). The allowed time can also be seen as a design parameter in the sensor selection problem but is here considered to be fixed. Extended model (2)+(3) is rewritten as a time window model, or batch model, of length n to model the system behaviour during a given time window. Define the vectors

$$\begin{aligned} z &= (y[t-n+1]^{\mathrm{T}}, \ldots, y[t]^{\mathrm{T}}, u[t-n+1]^{\mathrm{T}}, \ldots, u[t]^{\mathrm{T}})^{\mathrm{T}}, \\ x &= (x[t-n+1]^{\mathrm{T}}, \ldots, x[t]^{\mathrm{T}}, x[t+1]^{\mathrm{T}})^{\mathrm{T}}, \\ f &= (f[t-n+1]^{\mathrm{T}}, \ldots, f[t]^{\mathrm{T}})^{\mathrm{T}}, \\ e &= (v[t-n+1]^{\mathrm{T}}, \ldots, v[t]^{\mathrm{T}}, \varepsilon[t-n+1]^{\mathrm{T}}, \ldots, \varepsilon[t]^{\mathrm{T}})^{\mathrm{T}}, \end{aligned} \tag{4}$$

where $z \in \mathbb{R}^{n(l_y+l_u)}$, $x \in \mathbb{R}^{(n+1)l_x}$, $f \in \mathbb{R}^{nl_f}$, and e is a random vector with a known distribution with zero mean and covariance matrix Λ_e.

A sliding window model of length n can then be written as

$$Lz = Hx + Ff + Ne, \tag{5}$$

where

$$L = \begin{pmatrix} 0 & 0 & \cdots & 0 & -B_u & 0 & \cdots & 0 \\ 0 & 0 & & 0 & 0 & -B_u & & 0 \\ \vdots & & \ddots & \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 & -B_u \\ I & 0 & \cdots & 0 & -D_u & 0 & \cdots & 0 \\ 0 & I & & 0 & 0 & -D_u & & 0 \\ \vdots & & \ddots & \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & I & 0 & 0 & \cdots & -D_u \end{pmatrix},$$

$$H = \begin{pmatrix} A & -E & 0 & \cdots & 0 \\ 0 & A & -E & & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & A & -E \\ C & 0 & 0 & \cdots & 0 \\ 0 & C & 0 & & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & C & 0 \end{pmatrix},$$

$$F = \begin{pmatrix} B_f & 0 & \cdots & 0 \\ 0 & B_f & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & B_f \\ D_f & 0 & \cdots & 0 \\ 0 & D_f & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & D_f \end{pmatrix},$$

$$N = \begin{pmatrix} B_v & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ 0 & B_v & & 0 & 0 & 0 & & 0 \\ \vdots & & \ddots & \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & B_v & 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & D_\varepsilon & 0 & \cdots & 0 \\ 0 & 0 & & 0 & 0 & D_\varepsilon & & 0 \\ \vdots & & \ddots & \vdots & \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & D_\varepsilon \end{pmatrix}.$$

The fault vector f in Equation (5) describes how a fault affects the system during the considered time interval, for example, a constant, a ramp, or an intermittent fault.

For the quantitative fault diagnosability analysis, it is assumed that no noise-free residuals can be generated. This corresponds to model (5) that fulfils the condition

$$\begin{pmatrix} H & N \end{pmatrix} \text{ has full row-rank.} \tag{6}$$

One sufficient criteria for Equation (2) to satisfy Equation (6) is that all sensors have measurement noise and the model has a unique solution for a given initial state (Kunkel & Mehrmann, 2006).

It proves useful to write Equation (5) in an input–output form where the unknowns, x, are eliminated, without losing information about the system behaviour (Eriksson et al., 2013). The input–output model can, in the general case, then be written as

$$\mathcal{N}_H Lz = \mathcal{N}_H Ff + \mathcal{N}_H Ne, \tag{7}$$

where the rows of $\mathcal{N}_H$ is an orthonormal basis for the left null space of H, i.e. $\mathcal{N}_H H = 0$ and $\mathrm{cov}(\mathcal{N}_H Ne) = \Sigma$.

4.2 Model-based diagnosis

Model-based diagnosis systems are based on a set of residual generators to monitor the system. Each residual generator $r(z)$ is a function of known variables and designed to monitor a specific part of the system. A residual is, ideally, zero in the fault-free case and is said to be *sensitive* to a fault f_i if a fault time profile $\theta \neq 0$ implies that $\mathbb{E}[r(z)] \neq 0$. Based on which faults each residual generator is sensitive to, it can be used to detect and isolate faults (Svärd, Nyberg, & Frisk, 2013).

Definition 4.1 (Fault detectability of a residual generator): A fault f is detectable if a residual generator is sensitive to that fault.

Definition 4.2 (Fault isolability of a residual generator): A fault f_i is isolable from another fault f_j if a residual generator is sensitive to f_i but not f_j.

Based on these definitions, fault detection and isolation performance of a diagnosis system can be evaluated by analysing the detection performance of the residual generators with corresponding fault sensitivities (Svärd et al., 2013).

4.3 Distinguishability

Distinguishability is a model property, based on the Kullback–Leibler divergence, and is a quantitative measure of fault detection and isolation performance (Eriksson et al., 2013). In order to compute distinguishability, model (5) is first transformed to simplify the computations. It is assumed without loss of generality that

$$\Sigma = \mathcal{N}_H N \Lambda_e N^{\mathrm{T}} \mathcal{N}_H^{\mathrm{T}} = I. \tag{8}$$

Note that any model in form (5), satisfying Equation (6), can be transformed into fulfilling $\Sigma = I$ by multiplying Equation (5) with an invertible transformation matrix T from the left. The choice of matrix T is non-unique and one possibility is

$$T = \begin{pmatrix} \Gamma^{-1}\mathcal{N}_H \\ T_2 \end{pmatrix}, \tag{9}$$

where Γ is non-singular and

$$\mathcal{N}_H N \Lambda_e N^{\mathrm{T}} \mathcal{N}_H^{\mathrm{T}} = \Gamma\Gamma^{\mathrm{T}} \tag{10}$$

is satisfied, and T_2 is any matrix ensuring invertibility of T. Matrix Γ can, for example, be computed by a Cholesky factorisation of the left-hand side of Equation (10).

Then, distinguishability can be computed according to the following lemma (Theorem 1 in Eriksson et al., 2013).

Lemma 4.3: *Distinguishability of a fault f_i with fault profile $\theta \in \mathbb{R}^n$ from another fault f_j, for sliding window model (5) with Gaussian distributed random vector e, with covariance matrix I, is given by*

$$\mathcal{D}_{i,j}(\theta) = \frac{1}{2}\|\mathcal{N}_{(H F_j)} F_i \theta\|^2, \tag{11}$$

where the rows of $\mathcal{N}_{(H F_j)}$ is an orthonormal basis for the left null space of the matrix $(H F_j)$ and the matrix $F_j \in \mathbb{R}^{n(l_q+l_y)\times n}$ contains the columns of F corresponding to the elements of f_j.

Distinguishability $\mathcal{D}_{i,j}(\theta) : \mathbb{R}^n \to \mathbb{R}$ is a quantitative measure to evaluate how difficult it is to isolate a fault f_i given model (5), if the fault realisation is given by the specific fault profile θ, from another fault f_j with an unknown fault profile. Note that $\mathcal{D}_{i,j}(\theta)$ increases with increasing fault magnitude θ illustrating that the larger fault is easier to detect or isolate. A fault f_i with fault time profile θ is isolable from a fault mode f_j if and only if

$$\mathcal{D}_{i,j}(\theta) > 0. \tag{12}$$

Distinguishability computed for a given model gives an upper limit of the fault-to-noise ratio that can be achieved by a linear residual generator

$$r(z) = \gamma \mathcal{N}_{(H F_j)} L z, \tag{13}$$

where γ is a row vector. The fault-to-noise ratio measures the change in residual mean, caused by a fault f_i with fault profile θ, normalised by the residual noise standard deviation. The upper limit is given by the following lemma (Theorem 2 in Eriksson et al., 2013).

Lemma 4.4: *For model (2), let ϕ be the optimal fault-to-noise ratio with respect to fault f_i with fault profile θ when fault f_j is decoupled. Then*

$$\mathcal{D}_{i,j}(\theta) = \frac{1}{2}\phi^2,$$

where $\phi = \mathcal{N}_{(H F_j)} F_i \theta$.

This means that any linear residual generator (13) designed to isolate f_i (with fault realisation θ) from f_j will have a fault-to-noise ratio that is equal to or lower than ϕ. Note that maximum fault-to-noise ratio is achieved if the residual generator is designed such that the vector γ is selected parallel to $(\mathcal{N}_{(H F_j)} F_i \theta)^{\mathrm{T}}$. Lemma 4.4 shows that it is possible to predict, by analysing the model, what fault detection and isolation performance can be achieved by a set of residual generators in a diagnosis system.

5. Reformulating diagnosis system requirements for sensor selection

Two common measures to quantify performance requirements of diagnosis tests, such as residuals, are probability of false alarms $p_{\mathrm{fa}}^{\mathrm{req}}$ and probability of missed detections $p_{\mathrm{md}}^{\mathrm{req}}$. The objective is then to find a set of sensors that makes it possible to design residual generators (13) with sufficient fault-to-noise ratio ϕ to fulfil these requirements. To illustrate this, a one-sided threshold J is considered (see Figure 2), where $p_{\mathrm{nom}}(r)$ represents the nominal residual pdf and $p_{\mathrm{f}}(r)$ the faulty case. However, it is simple to generalise for the two-sided case, i.e. when $|r| > J$. Note that since a fault in Equation (2) only affects the mean of residual generator (13), a normalised residual value is considered here such that the variance is one. Then, the fault-to-noise ratio is equal to the mean deviation of the residual output when the fault f_i occurs with realisation θ.

A lower bound on the required fault-to-noise ratio can be derived such that it is possible to select a threshold J that fulfils both $p_{\mathrm{fa}}^{\mathrm{req}}$ and $p_{\mathrm{md}}^{\mathrm{req}}$ (see Figure 2). Then, Lemma 4.4 can be used to formulate the required fault-to-noise ratio as required distinguishability $\mathcal{D}_{i,j}^{\mathrm{req}}(\theta)$. Note that requirements on fault isolation performance are treated in the same way as detection performance since ϕ in Lemma 4.4 refers to the fault-to-noise ratio of a residual sensitive to f_i but not f_j.

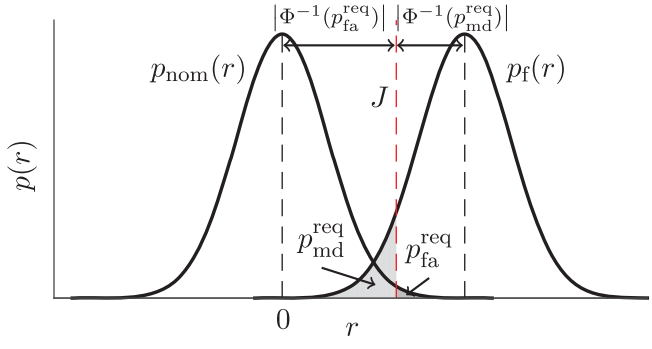

Figure 2. Requirements of maximum probability of false alarm $p_{\mathrm{fa}}^{\mathrm{req}}$ and missed detection $p_{\mathrm{md}}^{\mathrm{req}}$ can be translated to minimum required distance to the threshold J from the mean of the nominal pdf p_{nom} and the faulty case p_{f}, respectively.

To compute the lower bound, let the cumulative density function of the normal distribution with zero mean and variance one be denoted

$$\Phi(\psi) = \int_{-\infty}^{\psi} \frac{1}{\sqrt{2\pi}} e^{-\psi^2/2}\, d\psi = p_\psi \tag{14}$$

and the inverse $\psi = \Phi^{-1}(p_\psi)$. The value $|\psi|$ represents the distance from the distribution mean to a threshold J where the probability that ψ lies outside of the threshold is p_ψ. Thus, to fulfil $p_{\text{fa}}^{\text{req}}$ the distance from the mean of the nominal residual distribution $p_{\text{nom}}(r)$ to the threshold J must be greater than $|\Phi^{-1}(p_{\text{fa}}^{\text{req}})|$. Similarly, the distance from the mean of $p_{\text{f}}(r)$ to the threshold J must exceed $|\Phi^{-1}(p_{\text{md}}^{\text{req}})|$ to fulfil $p_{\text{md}}^{\text{req}}$ (see Figure 2). Thus, the fault-to-noise ratio ϕ must fulfil the inequality

$$\phi \geq |\Phi^{-1}(p_{\text{md}}^{\text{req}})| + |\Phi^{-1}(p_{\text{fa}}^{\text{req}})| \tag{15}$$

to satisfy both constraints.

Since computed distinguishability measure (11) gives the upper limit of achievable fault-to-noise ratio, it must exceed $\frac{1}{2}\phi^2$ as stated in Lemma 4.4, to assure that it is possible to design a set of residual generators that fulfils the requirements. Thus, the constraints $\mathcal{D}_{i,j}^{\text{req}}(\theta)$ can be computed as

$$\mathcal{D}_{i,j}^{\text{req}}(\theta) = \frac{1}{2}(|\Phi^{-1}(p_{\text{md}}^{\text{req}})| + |\Phi^{-1}(p_{\text{fa}}^{\text{req}})|)^2. \tag{16}$$

Note that Equation (16) reformulates the quantitative diagnosis system performance requirements to sensor selection problem (1) without designing a diagnosis system.

6. Analysis

In order to find a computationally efficient sensor placement algorithm, different properties of the problem are investigated. The search space grows exponentially with the number of available sensors which complicates the search for an optimal solution without an exhaustive search strategy. In order to find a global optimum, it is desirable to find some properties of the sensor placement problem that could be used to reduce the number of sensor combinations that needs to be evaluated.

6.1 A lattice representation of the search space

The set of all possible sensor combinations can be represented using a lattice. A lattice is a partially ordered set in which every two elements have a supremum and an infimum. Each element in the lattice set represents a sensor set and all sensor combinations of equal cardinality are positioned on the same level in the lattice. As a result, the lattice will be wider in the middle. Edges connect each set with its smallest supersets and its largest subsets, thus, representing the partial order. As an example, consider all combinations of four sensors $\mathcal{S} = \{y_1, y_2, y_3, y_4\}$. A lattice of all combinations of the four sensors is shown in Figure 3.

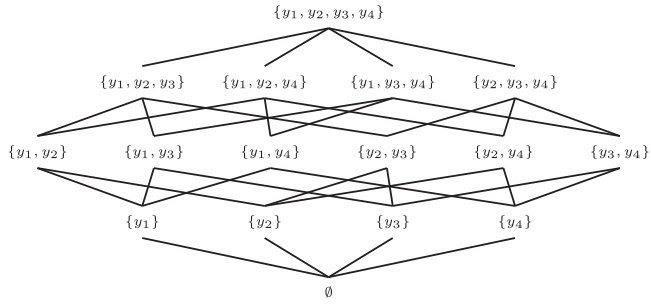

Figure 3. A lattice representing all combinations of four sensors $\{y_1, y_2, y_3, y_4\}$.

6.2 Distinguishability properties of partially ordered sensor sets

The distinguishability measure for each fault pair (f_i, f_j), i.e. isolating fault f_i from f_j, will never decrease when including more sensors to the already selected set of sensors. This means that if S_1 is consistent with the constraints in Equation (1) and $S_1 \subseteq S_2$ then S_2 is also consistent, i.e. all sensor combinations which are supersets of the feasible sensor set S_1 in the lattice are also feasible.

Theorem 6.1: *Consider sliding window model (5) and a fault pair (f_i, f_j). Let $\mathcal{D}_{i,j}^{S}(\theta)$ denote distinguishability when the sensor equations corresponding to the sensor set S are included in model (5). Let $S_1, S_2 \subseteq \mathcal{S}$ be two sets of sensors. If $S_1 \subseteq S_2$ then*

$$\mathcal{D}_{i,j}^{S_2}(\theta) \geq \mathcal{D}_{i,j}^{S_1}(\theta) \tag{17}$$

for all $f_i, f_j \in \mathcal{F}$.

Proof: Given model (5) where $\mathcal{D}_{i,j}(\theta) > 0$, there exists an optimal linear residual generator, to isolate the fault f_i with fault time profile θ from a fault f_j, that can be found using Theorem 3 in Eriksson et al. (2013). Since $S_1 \subseteq S_2$, the optimal residual generator given S_1 can also be generated with S_2. Thus, the optimal linear residual based on sensors S_2 is at least as good as for sensors S_1. Lemma 4.4 states that distinguishability gives the maximum fault-to-noise ratio of any linear residual generator based on the model, thus proving the inequality. ■

6.3 Distinguishability bounds and submodularity

An example of a candidate search algorithm to guarantee finding a globally optimal set of sensors is the A* algorithm (Russell, Norvig, Canny, Malik, & Edwards, 1995). To avoid evaluating all possible sensor combinations, the A* algorithm uses an admissible heuristic function to estimate the distance from a set of sensors $S \subseteq \mathcal{S}$ to the optimal solution. For the sensor selection problem, the heuristic function would define an upper bound of the distinguishability gain when adding a specific sensor to the solution set. However, it is important that the upper bound is not conservative in order for the search algorithm to be efficient.

Let

$$\Delta\mathcal{D}_{i,j}^{S\cup\{y_l\}}(\theta) = \mathcal{D}_{i,j}^{S\cup\{y_l\}}(\theta) - \mathcal{D}_{i,j}^{S}(\theta) \tag{18}$$

denote the increased distinguishability for a fault f_i with fault time profile θ from a fault f_j when adding a sensor y_l to a set S. For convenience, $\Delta\mathcal{D}_{i,j}^{S\cup\{y_l\}}(\theta)$ is also simply referred to the distinguishability gain when adding sensor y_l to S, since the fault pair (f_i, f_j) and fault time profile are given from the notation.

If the increase in distinguishability is lower when added to a larger solution set than a smaller, $\Delta\mathcal{D}_{i,j}^{S\cup\{y_l\}}(\theta)$ is a sub-modular

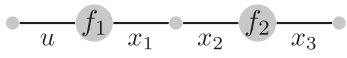

Figure 4. A schematic overview of the model described in Equation (21) with one known input, three measurable variables, and two faults.

set function. This property is utilised in, for example, Perelman et al. (2016) and Shamaiah, Banerjee, and Vikalo (2010), and if satisfied, a heuristic function would also be easy to compute. However, the amount of distinguishability that is increased for each fault pair $\Delta\mathcal{D}_{i,j}^{S\cup\{y_l\}}(\theta)$, when adding a sensor y_l, depends on the previous selected set of sensors S. If $S_1, S_2 \subseteq \mathcal{S} \setminus \{y_l\}$ are two sets of sensors such that $S_1 \subseteq S_2$, then

$$\Delta\mathcal{D}_{i,j}^{S_1\cup\{y_l\}}(\theta) \lesseqgtr \Delta\mathcal{D}_{i,j}^{S_2\cup\{y_l\}}(\theta), \tag{19}$$

i.e. it is not certain that more distinguishability is always gained by adding the sensor to a subset of sensors S_1 or a superset of sensors S_2. Thus, to compute the maximum distinguishability gain for a sensor $y_l \in \mathcal{S}$, means to solve the combinatorial optimisation problem

$$\max_{S\subseteq\mathcal{S}\setminus y_l} \Delta\mathcal{D}_{i,j}^{S\cup\{y_l\}}(\theta). \tag{20}$$

A small example is used to illustrate the difficulty of finding the global optimum to Equation (20) without using an exhaustive search strategy because of property (19).

Example 6.2: Consider a simple dynamic pipeline model shown in Figure 4. Let the inlet flow u be known, then a small flow model is given by

$$\begin{aligned} x_1[t+1] &= u[t] - f_1[t] + v_1[t], \\ x_2[t+1] &= x_1[t] + v_2[t], \\ x_3[t+1] &= x_2[t] - f_2[t] + v_3[t] \end{aligned} \tag{21}$$

with two faults f_1 and f_2 representing leakages in the pipeline. The process noise is i.i.d. Gaussian distributed as $v_1, v_2, v_3 \sim N(0,1)$. Each flow x_i can be measured by a sensor $y_i[t] = x_i[t] + \varepsilon_i[t]$, where $\varepsilon_i \sim N(0,1)$. For the analysis the time window is selected as $n=4$ and the analysed fault time profile is assumed constant with amplitude one.

Two cases of Equation (19) are shown where the distinguishability gain decreases with a larger set S in the first case, and increases in the second case. First, the distinguishability gain of adding y_1 to detect the fault f_1 is evaluated, i.e. $\Delta\mathcal{D}_1^{S\cup\{y_1\}}(\bar{1})$ where $\bar{1}$ denotes a vector with ones. All sensors in $\mathcal{S}$ can be used to detect f_1 since a residual generator sensitive to f_1 can be generated as long as we measure at least one of the flows, x_1, x_2, or x_3. If another sensor is already included in S when adding y_1, the total distinguishability value will increase since the knowledge about the system increases with more sensors. However, the distinguishability gain by adding y_1 will be smaller compared to if y_1 is selected first. The largest gain is achieved if y_1 is added before any other sensor, $\Delta\mathcal{D}_1^{\emptyset\cup\{y_1\}} = 1$, and the smallest gain if added last, $\Delta\mathcal{D}_1^{\{y_2,y_3\}\cup\{y_1\}} = 0.5$.

If instead considering detection of fault f_2 and maximising the distinguishability gain by adding y_3, the result is the opposite. To be able to generate a residual generator to detect f_2, it is necessary to use y_3 since it is the only sensor measuring a flow after the leakage f_2. If all three sensors are used, the distinguishability measure will be maximised. However, the distinguishability value will still be zero if y_3 is not added since y_3 is required to detect f_2. Thus, the highest distinguishability gain by adding y_3 is when y_1 and y_2 are already selected since the distinguishability measure will go from zero to maximum, $\Delta\mathcal{D}_2^{\{y_1,y_2\}\cup\{y_3\}} = 0.55$, while if y_3 is added before any other sensors, $\Delta\mathcal{D}_2^{\emptyset\cup\{y_3\}} = 0.13$.

7. Sensor placement algorithms

The analysis of the sensor placement problem shows that it is difficult to find the globally optimal set of sensors without the use of an exhaustive search strategy which is not feasible if the problem is too large. Therefore, a heuristic greedy stochastic search algorithm is proposed to solve the problem. The algorithm is designed to be computationally fast and still be able to find solutions, equal or close to global optimum, even for large search spaces. It uses multiple starting points, where each starting point is initialised with a random set of sensors that fulfils the requirements. Then, the algorithm iteratively removes sensors as long as the subset still fulfils the requirements, to guarantee that a feasible solution is always returned. Thus, the algorithm is sound.

7.1 Greedy stochastic search

A low-complexity heuristic search algorithm that can be used for solving the sensor selection problem is a greedy search algorithm (Eriksson et al., 2012). It solves a local optimisation problem in each iteration of the search and stops when the solution cannot be further improved. Since the greedy search is a deterministic algorithm, restarting the algorithm will always result in the same solution, that in the worst case could be far from the global optimum, which means that a better solution will never be found. To increase the chance of finding a better solution without significantly increasing the computational complexity, a greedy stochastic search algorithm is proposed. Greedy stochastic search algorithms have been successfully applied in similar problems when computing minimal diagnosis candidates (Feldman, Provan, & van Gemund, 2010). Because of the similarities between the two problems, greedy stochastic search is considered a suitable candidate search algorithm for the sensor selection problem. A general description of the algorithm in Feldman et al. (2010) is given here.

Instead of iteratively removing sensors that minimises a utility function, as for the case with the greedy search, a sensor is selected randomly from S and removed. Using this strategy, the algorithm finds different local optima each time the search is restarted. If the reduced sensor set, i.e. when the sensor is removed, is not feasible, another random sensor is selected instead. This is repeated maximally M times if no sensor to remove is found. The parameter $M \leq k$ is used to reduce the computational cost if k is large. If the reduced set of sensors is still feasible, another sensor is removed. If none of the M selected sensors can be removed, the remaining set of sensors is considered a local optimum and the search is stopped.

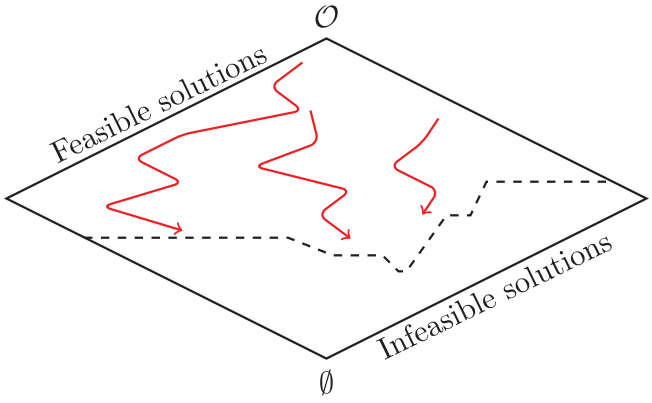

Figure 5. A lattice describing the greedy stochastic search algorithm. The algorithm is restarted from a randomly selected valid sensor set. A random sensor is iteratively removed as long as the subset is a valid solution.

To increase the probability of finding the global optimum, the algorithm uses N different starting points and the best solution of each search is stored. If S_i is a set of sensors found from run i, which fulfils the requirements, then the algorithm returns the cheapest set $S_{\min}$ as

$$S_{\min} = \arg \min_{1 \le i \le N} \sum_{y_l \in S_i} c_l. \tag{22}$$

Let X_i and $X_{\min}$ be random variables representing the distributions of $\sum_{y_l \in S_i} c_l$ and $\sum_{y_l \in S_{\min}} c_l$, respectively. If c^* is the cost of optimal solution, it can be shown that the discrete distribution $P(X_{\min} = c^*) \to 1$ when $N \to \infty$ (see Theorem 2 in Feldman et al., 2010). The probability that $\sum_{y_l \in S_{\min}} c_l = c^*$ can be written as

$$\begin{aligned} P(X_{\min} = c^*) &= 1 - P(X_1 > c^*, X_2 > c^*, \ldots, X_N > c^*) \\ &= 1 - \prod_{i=1}^{N} P(X_i > c^*), \end{aligned} \tag{23}$$

where the last equality follows from the fact that the set of sensors in each restart is found independently from the others. Since $P(X_i > c^*) < 1$, the probability of finding an optimal solution goes to one when $N \to \infty$.

If the number of sensor candidates is large, lots of sensors must be removed before reaching a local optimum. Therefore to speed up the algorithm, the starting point of each run is randomly generated among the feasible solutions, i.e. sensor sets fulfilling the requirements, to start closer to a local optimum (see Figure 5). Thus, the algorithm does not always start from $S = \mathcal{S}$ in each run. Here, a feasible set of sensors is generated by starting with an empty set $S = \emptyset$ and then iteratively adding a random subset of sensors $S' \subseteq \mathcal{S} \setminus S$ to S until the starting point is feasible as described in Algorithm 1. The function RandValidSensSet adds each sensor $s \in R$ to the set S' with probability p_{add}.

The algorithm StochSearch is described in Algorithm 2 and takes a model $\mathcal{M}$ in form (5), available sensors $\mathcal{S}$, requirements $\mathcal{D}^{\text{req}}_{i,j}$, and the parameters N and M as inputs. Each starting point is returned by a function RandValidSensSet which generates a feasible set of sensors as described above.

Algorithm 1: Generate random feasible set of sensors.

function $S = \text{RandValidSensSet}(\mathcal{M}, \mathcal{S}, \mathcal{D}^{\text{req}}_{i,j}(\theta))$;
  $R := \mathcal{S}$;
  $S := \emptyset$;
  while $\mathcal{D}^{S}_{i,j}(\theta) < \mathcal{D}^{\text{req}}_{i,j}(\theta), \forall \theta \in \Theta_i, \forall f_i, f_j \in \mathcal{F}$ **do**
    $S' := \text{SelectRandomSubset}(R)$;
    $S := S \cup S'$;
    $R = R \setminus S'$;
  end
return S

The computational complexity of the greedy stochastic search algorithm is $\mathcal{O}(kMN)$. Since the algorithm is stochastic, there is no guarantee that it will find a solution within a certain quality range from the global optimum, i.e. how close the cost of the found solution will be from the global optimum.

If considering the multiple-fault case when formulating sensor selection problem (1), distinguishability can be computed using the results in Eriksson et al. (2012).

Algorithm 2: Greedy stochastic search

function $S_{\text{opt}} = \text{StochSearch}(\mathcal{M}, \mathcal{S}, \mathcal{D}^{\text{req}}_{i,j}(\theta), N, M)$;
  $S_{\text{opt}} := \mathcal{S}$;
  for $n = 1, \ldots, N$ **do**
    $S := \text{RandValidSensSet}(\mathcal{M}, \mathcal{S}, \mathcal{D}^{\text{req}}_{i,j}(\theta))$;
    $m := 0$;
    while $m < M$ **do**
      $S' := \text{RemoveRandomSensor}(S)$;
      if $\mathcal{D}^{S'}_{i,j}(\theta_i) \ge \mathcal{D}^{\text{req}}_{i,j}(\theta), \forall \theta \in \Theta_i, \forall f_i, f_j \in \mathcal{F}$ **then**
        $S := S'$;
        $m := 0$;
      end
      else
        $m := m + 1$;
      end
    end
    if $\sum_{y_l \in S} c_l < \sum_{y_l \in S_{\text{opt}}} c_l$ **then**
      $S_{\text{opt}} := S$;
    end
  end
return S_{opt}

8. Case study

Here, the greedy stochastic search algorithm is evaluated on a system with 24 measurable variables with different sensor costs. Two heuristic search methods are also evaluated to compare the results, a greedy search (Eriksson et al., 2012) and a genetic algorithm (Deep, Singh, Kansal, & Mohan, 2009). A depth-first search algorithm (Russell et al., 1995) is also used as a reference to evaluate the performance of the search algorithms.

8.1 Model description

A schematic of the system is shown in Figure 6. Since there is no principal difference of computing distinguishability between static and dynamic systems, because a dynamic system is written in static form (5) for a given time window, a static system is considered here. Here, the window length is chosen as $n = 1$ and the fault time profiles for each of the faults are $\theta = 1$. The

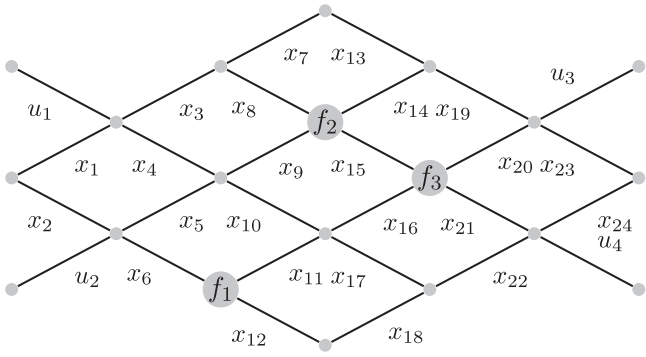

Figure 6. A schematic overview of a model describing a static flow through a number of nodes. The inputs u are known and the flows through the branches x_i, $i = 1, 2, \ldots, 24$, are unknown, and f_1, f_2, f_3 are three additive faults.

sensors have different costs so the goal is to minimise total cost (1).

The system is under-determined and described by the following set of equations:

$$
\begin{aligned}
x_1 + x_2 &= v_1, \quad x_4 + x_3 = x_1 + u_1 + v_2,\\
x_6 + x_5 &= x_2 + u_2 + v_3, \quad x_8 + x_7 = x_3 + v_4,\\
x_{10} + x_9 &= x_4 + x_5 + v_5, \quad x_{12} + x_{11} = x_6 + f_1 + v_6,\\
x_{13} &= x_7 + v_7, \quad x_{15} + x_{14} = x_9 + x_8 + f_2 + v_8,\\
x_{17} + x_{16} &= x_{11} + x_{10} + v_9, \quad x_{18} = x_{12} + v_{10},\\
x_{19} &= x_{14} + x_{13} + v_{11},\\
x_{21} + x_{20} &= x_{16} + x_{15} + f_3 + v_{12},\\
x_{22} &= x_{18} + x_{17} + v_{14}, \quad x_{23} + u_3 = x_{20} + x_{19} + v_{15},\\
x_{24} + u_4 &= x_{22} + x_{21} + v_{16}, \quad 0 = x_{24} + x_{23} + v_{17},
\end{aligned}
$$

where $v_k \sim \mathcal{N}(0, 0.01)$ for $k = 1, 2, \ldots, 17$. Each unknown variable x_l can be measured by a sensor $y_l = x_l + \varepsilon_l$ where $\varepsilon_l \sim \mathcal{N}(0, 1)$ for $l = 1, 2, \ldots, 24$. The cost of using each sensor is given by

$$
c_l = \begin{cases} 1 & \text{for } l \in \{6, 8, 9, 11, 12, 14, 15, 16, 20, 21\}, \\ 0.7 & \text{for } l \in \{3, 4, 5, 10, 19, 22, 23\}, \\ 0.4 & \text{for } l \in \{1, 2, 7, 13, 17, 18, 24\}, \end{cases} \tag{24}
$$

and the number of sensor combinations is $2^{24} = 16\,777\,216$.

8.2 Evaluation

Maximum distinguishability $\mathcal{D}^{\max}_{i,j}(\theta)$ for each fault pair when using all sensors is shown in Table 1. The problem is solved for $\mathcal{D}^{\text{req}}_{i,j}(\theta) = \alpha \mathcal{D}^{\max}_{i,j}(\theta)$ for all $f_i, f_j \in \mathcal{F}$, where $\alpha = \{0, 0.01, 0.02, \ldots, 1\}$ is a scaling factor. First, the greedy stochastic search is parameterised as $N = 10$ and $M = 4$ which are relatively small values compared to the number of candidate sensors. The algorithm is evaluated 100 times for each set of requirements to estimate the mean and standard deviation of the computed cost.

The search algorithms are evaluated on a standard desktop computer with an Intel I5 processor. The greedy search and the greedy stochastic search algorithms take a couple of seconds

Table 1. Maximum achievable distinguishability for the case study model in Figure 6.

$\mathcal{D}^{\max}_{i,j}(1)$	NF	f_1	f_2	f_3
f_1	3.26	0	0.48	0.44
f_2	3.28	0.47	0	0.27
f_3	3.28	0.43	0.27	0

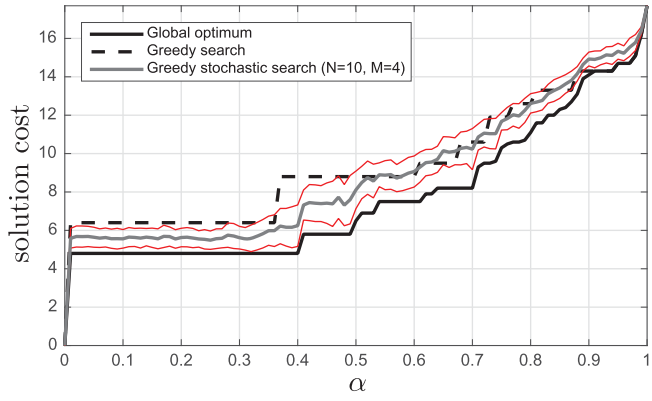

Figure 7. Evaluation of the greedy search algorithm, greedy stochastic search ($N = 10, M = 4$), and depth-first search for different requirements.

for each run. The depth-first search algorithm takes a couple of minutes for $\alpha = 0.9$ and the time increases to around 10 h when $\alpha = 0.5$. Therefore, the depth-first search algorithm was not evaluated for lower requirements. A global search, evaluating all sensor combinations, was estimated to take around 83 h and was instead performed on a server using 9 cores which reduced the computational time to around 12 h.

The result of the evaluation for different requirements is shown in Figure 7. The cost of the solutions of the different search algorithms are compared. The result of the greedy stochastic search is presented by the mean value (grey thick solid line) and standard deviation (thinner solid lines) of the 100 evaluations. The greedy stochastic search shows a noticeable better result closer to global optimum compared to the pure greedy search, even though the parameters N and M are selected relatively restrictive compared to the size of the problem, except when the requirements are chosen in the interval $\alpha \in [0.9, 1]$, i.e. close to $\mathcal{D}^{\max}_{i,j}$. This is expected since there are not many sensors that can be eliminated while still fulfilling the high requirements.

To see how sensitive the greedy stochastic search algorithm is to different parameters the system is evaluated for $(N = 10, M = 24)$ and $(N = 50, M = 10)$, see Figures 8 and 9, respectively. Each run of the greedy stochastic search took around 2 s for $(N = 10, M = 4)$, 5 s for $(N = 10, M = 24)$, and around 24 s for $(N = 50, M = 10)$. In this case study, different values of M were tested and for $M > 10$, the performance did not change significantly except computational time which increased slightly. This could be explained by the fact that in many cases, the cardinality of the solution set is seldom significantly larger than 10, meaning that the number of sensors evaluated in each step of the search is often close to the number of sensors in S. For lower values of M the performance is slightly worse, i.e. the mean and variance of the solution cost are higher, which is visible when comparing Figures 7 and 8.

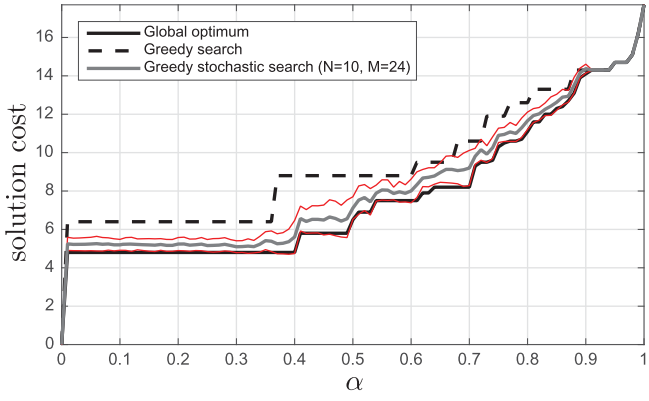

Figure 8. Evaluation of the greedy search algorithm, greedy stochastic search ($N = 10, M = 24$), and depth-first search for different requirements.

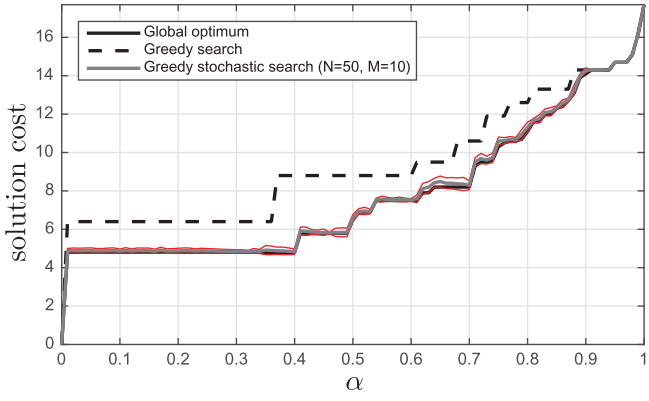

Figure 9. Evaluation of the greedy search algorithm, greedy stochastic search ($N = 50, M = 10$), and depth-first search for different requirements.

Figure 9 shows that the solution of the greedy stochastic search is significantly improved with increasing N. In this case, the greedy stochastic search algorithm finds solutions which in average are very close to the global optimum (within 1–3% of global optimum). For $(N = 200, M = 10)$, the search took a little more than a minute and the found solution is in average around 0.2–0.4% from the global optimum. This shows that the global optimum is found in almost all cases.

To compare the quality of the computed solutions of Figures 7– 9, the mean costs relative to the global optimum are compared (see Figure 10). The greedy search finds solutions that are more expensive compared to the mean value of the greedy stochastic searches. When $\alpha \leq 0.5$, the increased cost varies between 30% and 90% higher than the optimal solution.

It is also visible for the greedy stochastic search that higher M and N increase the probability of finding better solutions. However, since there is no point in selecting $M > k$, the main tuning parameter is N to improve the performance. That is, the more starting points the greedy stochastic search algorithm can explore, the higher is the probability of finding the global optimum.

To evaluate the greedy stochastic search, an off-the-shelf genetic algorithm is used to see if the performance is comparable given similar running time. The mixed integer genetic algorithm (Deep et al., 2009) is implemented using the function `ga` in MATLAB's Global Optimization Toolbox with standard

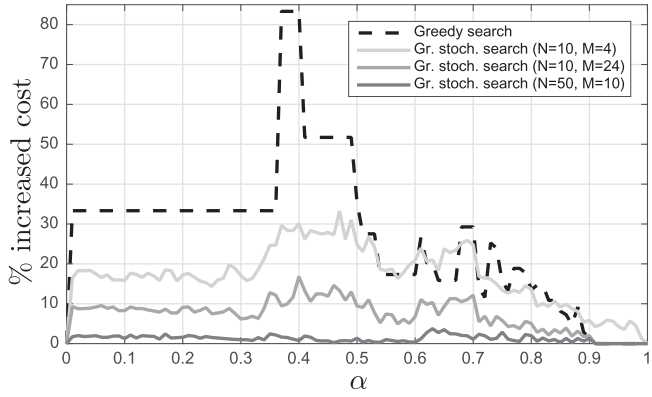

Figure 10. Relative mean cost of optimal solutions found by the different algorithms.

Table 2. A summary of the standard configurations used in MATLAB for the mixed integer genetic algorithm.

Parameter	Function	Value
Population size		200
Elite selection		10
Selection to mate	Tournament	4
Crossover	Laplace	160
Mutation	Power	30

Note: The functions are described in Deep et al. (2009).

configurations (see Table 2). Since the greedy stochastic search can be considered a plug-and-play solver without any necessary tuning, no specific tuning was made of the genetic algorithm (see Deep et al., 2009; MATLAB & Toolbox, 2013 for further details). In this comparison, the number of generations was set to a large value and a time limit of the genetic algorithm was set to 24 s which corresponds to the runtime of the greedy stochastic search with parameters $(N = 50, M = 10)$. The genetic algorithm was run 100 times for each requirement. The genetic algorithm gave a cost of the found solution around 10% higher than global optimum compared to the greedy stochastic search which gave around 2% higher cost. Thus, the greedy stochastic search shows a good result in comparable computational time and at the same time requires no tuning. The result of the genetic algorithm was slightly worse than the greedy stochastic search with parameters $(N = 10, M = 24)$ as shown in Figure 10. Note that the genetic algorithm did not always have time to converge within the time limit of 24 s. The computational complexity of the genetic algorithm depends on the number of generations G and the population size P in each generation as $\mathcal{O}(GP)$.

Based on this case study, the greedy stochastic search appears as a good candidate for finding solutions close to global optimum while using limited computational time and no complicated tuning. In the case study, this means from a couple of seconds to a minute per run compared to the depth-first search which can take from a couple of hours to days. However, each algorithm can be parallelised to further speed up the computations, especially for the greedy stochastic search algorithm.

9. Conclusions

The sensor selection problem is formulated by taking measurement noise and model uncertainties into consideration to

get a more realistic evaluation of fault diagnosability performance. Fault detection and isolation performance is evaluated using a measure called distinguishability which gives an upper bound of achievable fault-to-noise ratio for any linear residual generator based on the model of the system. By reformulating diagnosis system requirements for the sensor selection problem means that it is possible to take more realistic fault diagnosis performance requirements into consideration early in the system design process. It also gives an intuitive interpretation of the distinguishability measure in terms of possible false alarm rate and missed detection rate. The properties of the sensor selection problem are analysed and examples show that the amount of distinguishability gained by using a specific sensor depends on previously selected sensors. This complicates the use of informed search strategies, such as the A* algorithm, since it is difficult to estimate the distance to the optimal solution. The case study shows that the proposed greedy stochastic search algorithm is able to find solutions close to the global optimum in relatively short time compared to other heuristic search algorithms.

Disclosure statement

No potential conflict of interest was reported by the authors.

ORCID

Daniel Jung http://orcid.org/0000-0003-0808-052X

References

Basseville, M., & Nikiforov, I. V. (1993). *Detection of abrupt changes: Theory and application* (Vol. *104*). Englewood Cliffs, NJ: Prentice Hall.

Bhushan, M., Narasimhan, S., & Rengaswamy, R. (2008). Robust sensor network design for fault diagnosis. *Computers & Chemical Engineering*, *32*(4), 1067–1084.

Casillas, M. V., Puig, V., Garza-Castañón, L. E., & Rosich, A. (2013). Optimal sensor placement for leak location in water distribution networks using genetic algorithms. *Sensors*, *13*(11), 14984–15005.

Chi, G., Wang, D., & Zhu, S. (2015). An integrated approach for sensor placement in linear dynamic systems. *Journal of the Franklin Institute*, *352*(3), 1056–1079.

Commault, C., & Dion, J. M. (2007). Sensor location for diagnosis in linear systems: A structural analysis. *IEEE Transactions on Automatic Control*, *52*(2), 155–169.

Daigle, M., Roychoudhury, I., & Bregon, A. (2014). *Diagnosability-based sensor placement through structural model decomposition*. Proceedings of European conference of the prognostics and health management society, Nantes, France.

Deep, K., Singh, K. P., Kansal, M., & Mohan, C. (2009). A real coded genetic algorithm for solving integer and mixed integer optimization problems. *Applied Mathematics and Computation*, *212*(2), 505–518.

Dong, Y., & Biswas, G. (2013). *Comparison for sensor placement algorithms*. Proceedings of 24th international workshop on principles of diagnosis (DX-13), Jerusalem, Israel.

Eriksson, D., Frisk, E., & Krysander, M. (2012). *A sequential test selection algorithm for fault isolation*. Proceedings of 10th European workshop on advanced control and diagnosis, Copenhagen, Denmark.

Eriksson, D., Frisk, E., & Krysander, M. (2013). A method for quantitative fault diagnosability analysis of stochastic linear descriptor models. *Automatica*, *49*(6), 1591–1600.

Eriksson, D., Krysander, M., & Frisk, E. (2012). *Using quantitative diagnosability analysis for optimal sensor placement*. Proceedings of IFAC safeprocess, Mexico City, Mexico.

Feldman, A., Provan, G., & van Gemund, A. (2010). Approximate model-based diagnosis using greedy stochastic search. *Journal of Artificial Intelligence Research*, *38*, 371–00.

Frisk, E., Krysander, M., & Åslund, J. (2009). Sensor placement for fault isolation in linear differential-algebraic systems. *Automatica*, *45*(2), 364–371.

Harrou, F., Fillatre, L., & Nikiforov, I. (2014). Anomaly detection/detectability for a linear model with a bounded nuisance parameter. *Annual Reviews in Control*, *38*(1), 32–44.

Huber, J., Kopecek, H., & Hofbaur, M. (2014). *Sensor selection for fault parameter identification applied to an internal combustion engine*. Proceedings of IEEE conference on control applications (CCA), Juan Les Antibes, France.

Krysander, M., & Frisk, E. (2008). Sensor placement for fault diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, *38*(6), 1398–1410.

Kunkel, P., & Mehrmann, V. (2006). *Differential-algebraic equations: Analysis and numerical solution*. Zürich: European Mathematical Society.

Namburu, S. M., Azam, M. S., Luo, J., Choi, K., & Pattipati, K. R. (2007). Data-driven modelling, fault diagnosis and optimal sensor selection for hvac chillers. *IEEE Transactions on Automation Science and Engineering*, *4*(3), 469–473.

Narasimhan, S., Mosterman, P. J., & Biswas, G. (1998). A systematic analysis of measurement selection algorithms for fault isolation in dynamic systems. In Proceedings of 9th international workshop on principles of diagnosis (DX-98) (pp. 94–101), Cape Cod, MA, USA.

Nyberg, M. (2002). Criterions for detectability and strong detectability of faults in linear systems. *International Journal of Control*, *75*(7), 490–501.

Perelman, L. S., Abbas, W., Koutsoukos, X., & Amin, S. (2016). Sensor placement for fault location identification in water networks: A minimum test cover approach. *Automatica*, *72*, 166–176.

Raghuraj, R., Bhushan, M., & Rengaswamy, R. (1999). Locating sensors in complex chemical plants based on fault diagnostic observability criteria. *AIChE Journal*, *45*(2), 310–322.

Rosich, A., Sarrate, R., & Nejjari, F. (2009). *Optimal sensor placement for fdi using binary integer linear programming*. Proceedings of 20th international workshop on principles of diagnosis, Stockholm, Sweden.

Russell, S. J., Norvig, P., Canny, J. F., Malik, J. M., & Edwards, D. D. (1995). *Artificial intelligence: A modern approach* (Vol. *2*). Englewood Cliffs, NJ: Prentice Hall.

Shamaiah, M., Banerjee, S., & Vikalo, H. (2010). Greedy sensor selection: Leveraging submodularity. In 49th IEEE conference on decision and control (CDC), 2010 (pp. 2572–2577), Atlanta, Giorgia, USA.

Svärd, C., Nyberg, M., & Frisk, E. (2013). Realizability constrained selection of residual generators for fault diagnosis with an automotive engine application. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *43*(6), 1354–1369.

MATLAB, & Toolbox, G. O. (2013). *Version 8.1.0.604 (r2013a)*. Natick, MA: The MathWorks Inc.

Travé-Massuyès, L., Escobet, T., & Olive, X. (2006). Diagnosability analysis based on component-supported analytical redundancy relations. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, *36*(6), 1146–1160.

Wu, Z., Hsieh, S., & Li, J. (2013). Sensor deployment based on fuzzy graph considering heterogeneity and multiple-objectives to diagnose manufacturing system. *Robotics and Computer-Integrated Manufacturing*, *29*(1), 192–208.

Yassine, A., Ploix, S., & Flaus, J. M. (2008). A method for sensor placement taking into account diagnosability criteria. *International Journal of Applied Mathematics and Computer Science*, *18*(4), 497–512.