FROM FAULT DIAGNOSIS AND PREDICTIVE MAINTENANCE TO CONTROL RECONFIGURATION

Digital twins are increasingly being developed for industrial high-tech systems. In this article, we outline the exciting opportunities that digital twins have to offer for fault diagnosis, predictive maintenance, and controller reconfiguration. The proposed solutions increase economic value by minimising downtime through nonintrusive diagnostics.

KOEN CLASSENS, JEROEN VAN DE WIJDEVEN, MAURICE HEEMELS AND TOM OOMEN

Introduction

The economic value of high-tech production equipment is largely determined by its productivity, which is, in turn, heavily related to its uptime. Without maintenance, it is not a question whether a machine will fail, but rather when it will fail. Unexpected failure often results in downtime, leading to a severe loss of productivity. These unexpected breakdowns can be attributed to various factors, such as defects, the ageing of system components, and wear and tear, among others.

Ideally, critical faults are detected in an early stage and handled such that downtime is minimised. Traditionally, faults have been addressed through preventive maintenance strategies or reactive responses. In sharp contrast to these traditional approaches, predictive maintenance offers a way to further minimise equipment downtime [1,2]; see Figure 1. This is achieved by detecting faults within the equipment and precisely pinpointing their origin, a process known as fault detection and isolation (FDI) [3,4].

Digital twins are increasingly being developed for industrial high-tech systems [5-8]. In sharp contrast to the popularity that artificial intelligence (AI) based solutions for fault diagnosis have gained [9-11], we argue that model-based approaches should form the core foundation of digital twins for FDI in high-tech systems [12]. Typically, models of the system have been created prior to commissioning a machine. These models range in complexity from simple first-principles modelling to data-enriched finite-element modelling (FEM), or they can be identified during the system integration phase. These models, developed during the machine's design and integration phases, are at the heart of digital twins for FDI and constitute a system of interconnected systems. Despite the trend towards more AI-based solutions, we argue that the building blocks for digital twins for fault diagnosis were readily available far before the digital twin concept gained prominence in the early 2000s [13].

Surprisingly, after system integration and controller design, the developed models are often left unused. In sharp contrast, we propose to repurpose these models, since these models have a predictive power and can be harnessed in the form of a digital counterpart that is continuously informed with real-time data through already existing sensors and actuators.

AUTHORS' NOTE

Koen Classens (Ph.D. candidate), Maurice Heemels (professor) and Tom Oomen (professor) are all associated with the Control Systems Technology group in the department of Mechanical Engineering at Eindhoven University of Technology, the Netherlands.
Jeroen van de Wijdeven is a senior researcher at ASML, Veldhoven (NL).

k.h.j.classens@tue.nl
www.tue.nl/cst
www.asml.com

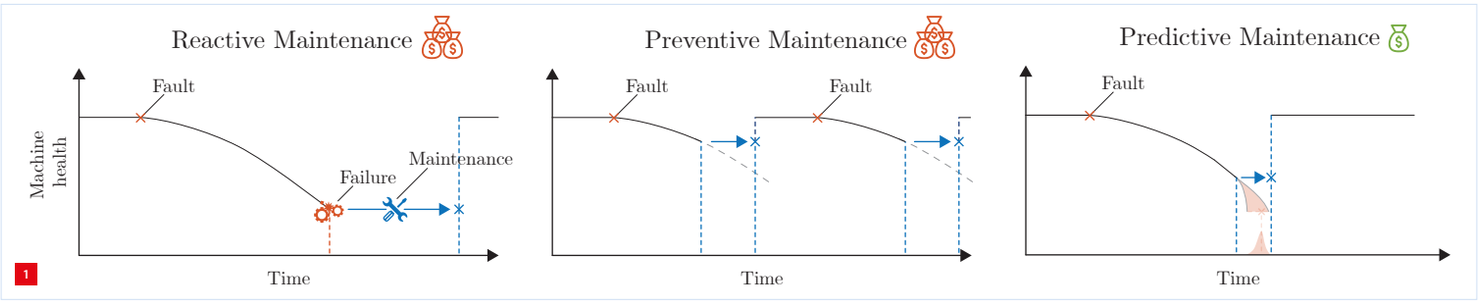

1

Illustration of three different methods to perform maintenance. Responding reactively leads to unexpected downtime and typically a long repair time. Preventive strategies lead to more interruptions than strictly necessary. Predictive maintenance is preferred as it exploits data and models to anticipate and prevent equipment failures, improving overall operational efficiency.

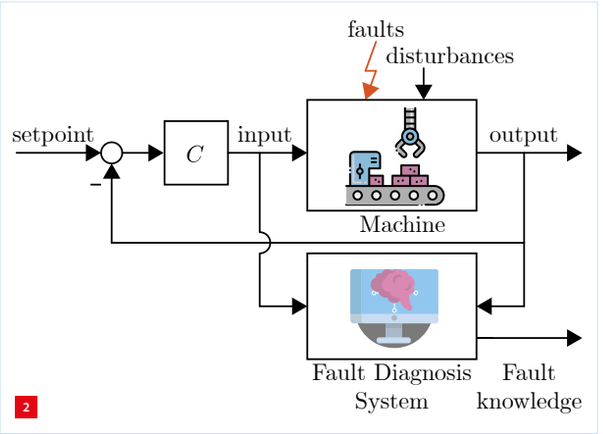

Schematic overview of a closed-loop controlled mechatronic system subjected to faults and disturbances. Data is extracted from the controlled system, which serves as input to the fault-diagnosis system that generates signals revealing knowledge about potential faults.

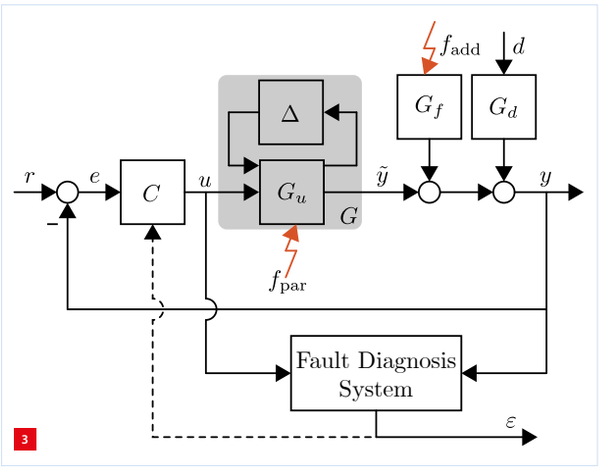

Block diagram of a closed-loop controlled system with a plant subjected to modelling uncertainty, disturbances and faults. The closed-loop controlled system is augmented with a fault-diagnosis system generating residual signals ε. The information from the fault-diagnosis system may be fed back into the controller, e.g., for reconfiguration and self-healing.

Properly designed digital counterparts have the capability to accurately predict the behaviour of the machine. Consequently, detection of an increasing mismatch between the model and the machine can indicate an upcoming failure. This yields the opportunity for scheduled instead of unscheduled service actions. As such, they provide major opportunities for enhancing uptime [12] [V1].

As the demand for higher performance increases, machines are evolving into more intricate systems. Their complexity becomes apparent through the increasing presence of actuators, sensors and components, which, in turn, generates a vast volume of data that can be harnessed for multiple purposes. Firstly, this data can be used to enhance the predictive capability of the digital counterpart. Secondly, the data can be used for monitoring and detection of anomalies. Moreover, after successful detection additional actuators and sensors might allow addressing the fault through effective self-healing. For instance, by reallocating the control effort among the remaining healthy actuators, so-called control reconfiguration.

As opposed to the opportunities arising from the abundance of actuators and sensors, this article also highlights several challenges that come with the growing complexity of systems in the context of fault diagnosis. Firstly, the increasing number of components gives rise to a greater number of potential fault scenarios, complicating the task to isolate the root cause. Secondly, for successful fault detection, it must be guaranteed that faults are distinguished from external disturbances, inherently present in any system, while simultaneously accounting for uncertainty in the system model. Thirdly, feedback controllers in mechatronic systems are designed to minimise the effects of disturbances and anomalies, making fault detection more difficult [14]. At the same time, faults in closed-loop controlled systems may be particularly hazardous as these affect machine performance and stability margins [15].

In this article, we first show the basic functionality of a model-based fault-diagnosis system and illustrate how to exploit the available models for basic FDI filter design. Subsequently, we illustrate how to give robustness guarantees and demonstrate the estimation of changing system dynamics. Finally, it is shown how to exploit fault information to enhance uptime through control reconfiguration.

Model-based fault-diagnosis system

Every fault-diagnosis system relies on data sourced from the system, such as actuator, controller and sensor data. This data can be processed offline, but also online during normal production. Fault-diagnosis systems processing this data can be designed in different ways [16-19], including data analyses, and machine-learning algorithms.

Model-based approaches acquire and process the data concurrently with the control algorithm and thus in real time, as schematically depicted in Figure 2. The output of the fault-diagnosis system is a signal or data stream that aids in identifying the presence of faults or anomalies within the system. Analysing these specifically designed signals helps diagnose and locate the specific fault in the system. So, summarising, faults that are difficult to detect directly in original data sources and/or suppressed by a feedback control loop will become visible in the output of the fault-diagnosis system.

In contrast to offline data processing, a significant advantage of online methods, such as model-based methods, is their ability to enable nearly instantaneous fault detection. One potential drawback of model-based approaches lies in their reliance on the accuracy of the underlying model. Fortunately, when it comes to mechatronic systems, precise models are readily available even before a machine is commissioned. The complexity of

the required model depends on the specifics of the faults. The most precise models are identified using experimental data, ideally accompanied by a certificate of accuracy. For other faults, even first-principle models provide sufficient accuracy for diagnostics.

The mathematical descriptions of these models are especially useful when they accurately encapsulate the system's relevant physical behaviour in an interpretable manner. A high accuracy is important to detect small deviations caused by faults at an early stage, even before they lead to significant deviations in system behaviour. Early fault detection can prevent minor problems from escalating into major problems.

The interpretability provides a profound understanding of how the system should behave under normal circumstances and proves to be invaluable for pinpointing the origin of anomalous behaviour, i.e., fault isolation, and quantifying its severity for closed-loop performance and stability. It is extremely valuable to isolate the specific component or subsystem responsible for a fault. Namely, this information directs service teams to the exact source of the problem, reducing diagnostic time and improving the accuracy of corrective actions.

Figure 3 shows a typical block diagram for the design and synthesis of model-based fault-diagnosis systems. The plant model of the plant is denoted by $G: u \to \tilde{y}$, which is a linear fractional transformation of a nominal model G_u and corresponding modelling uncertainty Δ. The feedback controller is denoted by C. Without loss of generality, the control algorithm can be extended with, e.g., a feedforward filter. Disturbances are inherently present in any control loop and are denoted by d. Disturbances may be coloured. For example, the sensor noise may have a specific frequency spectrum. This is modelled through the filter G_d.

Faults are often categorised into two parts [3], namely additive contributions f_{add}, modelled by the filter G_f, and parametric contributions f_{par}, also referred to as multiplicative faults. Sensor bias and drift, bearing wear or load variations, are examples of additive faults. Amplifier degradation and shifting resonance dynamics are examples of parametric faults. Depending on the nature of the fault, an appropriate fault-diagnosis strategy must be selected. The control input u, the output y, including effects from d, f_{add}, and f_{par}, and possibly the reference signal r, are fed into the fault-diagnosis system. How these models are used for a basic fault-diagnosis system design is described in detail below.

Fault detection and isolation

As mentioned, faults may be modelled as additive contributions; see Figure 3. In practise, many faults can be represented or approximated by additive contributions, making this a reasonable fault representation in many engineering applications. Modelling faults as additive contributions is attractive due to the mathematical simplicity and the alignment with linear system theory. This makes it easier to integrate fault detection, isolation and control into existing linear and numerically reliable frameworks for larger-scale systems. First, a basic design strategy is illustrated [18,19], followed by a case study on a prototype wafer stage.

Consider Figure 3 without multiplicative faults f_{par} and without modelling uncertainty Δ. A fault-diagnosis system to detect and isolate additive faults consists of a stable linear time-invariant residual generator, $[Q_y \quad Q_u]: (y, u) \to \varepsilon$, i.e.,

$$\varepsilon = [Q_y \quad Q_u]\begin{bmatrix} y \\ u \end{bmatrix} = [Q_y \quad Q_u]\begin{bmatrix} G_u u + G_d d + G_f f \\ u \end{bmatrix}.$$

This can be rewritten as:

$$\begin{bmatrix} \varepsilon_1 \\ \vdots \\ \varepsilon_q \end{bmatrix} = \underbrace{\begin{bmatrix} Q_y^{(1)} & Q_u^{(1)} \\ \vdots & \vdots \\ Q_y^{(q)} & Q_u^{(q)} \end{bmatrix}\begin{bmatrix} G_u & G_d & G_f \\ I_{m_u} & 0 & 0 \end{bmatrix}}_{[R_u \quad R_d \quad R_f]}\begin{bmatrix} u \\ d \\ f \end{bmatrix}.$$

The aim is to design the filter Q such that:

1. $R_u = 0$;
2. $R_d \approx 0$;
3. $R_f \neq 0$ and is structured.

The first objective, $R_u = 0$, aims to decouple the effect of the control input in the residual. The second objective, $R_d \approx 0$, aims to have a minimal effect of disturbances. With the imposed structure on R_f, effective fault isolation can be achieved. Structured residuals can be achieved by redistributing a distinct subset of faults as disturbances for every row in Q_y and Q_u.

A residual generator is typically synthesised in a few sequential steps [19]. Two steps are illustrated next. To accomplish $R_u = 0$, we define $Q_y = Q_2 Q_{y1}$ and $Q_u = Q_2 Q_{u1}$. Hence, to achieve

$$[Q_y \quad Q_u]\begin{bmatrix} G_u \\ I_{m_u} \end{bmatrix} = 0,$$

we can take the immediate solution $Q_{y1} = I$, $Q_{u1} = -G_u$, or alternatively take a so-called left coprime factorisation of G_u as a solution.

Subsequently Q_2 is used for fault-to-disturbance optimisation [20], i.e., maximise detection of f, given an allowable level of d. Hence, given $\gamma > 0$, determine $\beta > 0$ and find a stable proper filter Q_2 such that:

$$\beta = \max_{Q_2} \left\{ \left|\left|Q_2 Q_{y1} G_f\right|\right|_{\infty-} \Big| \left|\left|Q_2 Q_{y1} G_d\right|\right|_{\infty} \leq \gamma \right\}$$

with

$$\left|\left|Q_2 Q_{y1} G_f\right|\right|_{\infty-} := \min_{1 \leq i \leq q} \left|\left|Q_2 Q_{y1} G_{f_i}\right|\right|_{\infty}.$$

Simply stated, maximise the sensitivity to faults and simultaneously minimise the sensitivity to disturbances. A higher ratio between these sensitivities allows for earlier detection of faults. This optimisation step can be solved efficiently using a Ricatti equation [21,22].
The following case study illustrates a fault detection and isolation filter that allows to detect and isolate 17 distinct potential faults in a prototype wafer stage.

Case study: Fault detection and isolation for a next-generation prototype wafer stage

Prototype next-generation wafer stage with 13 actuators and four sensors in z-direction, which are all prone to faults. The set-up is currently installed at the Mechanical Engineering department of Eindhoven University of Technology.

The overactuated test-rig (OAT) is a prototype lightweight motion stage that can be positioned in all six degrees of freedom. The stage is levitated with four gravitation compensators in the corners and contains four sensors to measure the position in z-direction; see Figure 4. The position is measured using linear incremental encoders with a resolution of 1 nm. The stage is actuated by Lorentz actuators, of which 13 actuate in the z-direction; see Figure 5.

For this case study, it was assumed that all sensors and actuators in z-direction are prone to faults and as a result may malfunction. To detect and isolate these 17 distinct faults, a fault-diagnosis system has been designed. First, a frequency response measurement was conducted on the fault-free system and a 30th-order modal $G_u(s)$ model was estimated; see Figure 6 for a snapshot of the first four inputs.

Based on this model, a residual generator, $[Q_y\ Q_u]$, was synthesised following the method described in the section on fault detection and isolation, where a structure

$$R_f = \begin{bmatrix} 0 & (\cdot) & (\cdot) & (\cdot) & \cdots \\ (\cdot) & 0 & (\cdot) & (\cdot) & \cdots \\ (\cdot) & (\cdot) & 0 & (\cdot) & \cdots \\ (\cdot) & (\cdot) & (\cdot) & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

is enforced, and $(\cdot)$ implies nonzero transfer functions. An experiment was conducted where the 17 faults were artificially triggered by introducing signals f_i. In Figure 7, the first four residuals ε_i are shown, where the different faults can be clearly distinguished using the enforced structure above. For example, during the first seconds ε_1 is zero while the other signals are triggered, concluding that the first fault is occurring, which in this case is actuator 1.

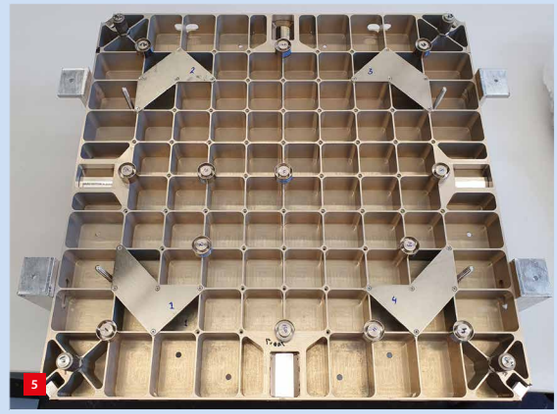

Bottom of the prototype wafer stage showing the 13 spatially distributed actuators and four gravitation compensators.

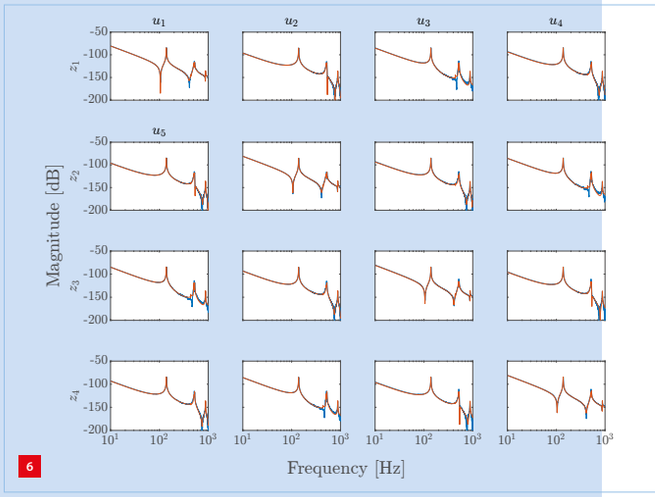

4 × 4 section of the 4 × 13 wafer-stage model. The nonparametric estimate (blue) and the parametric modal model (red), used for the fault diagnosis design, are depicted.

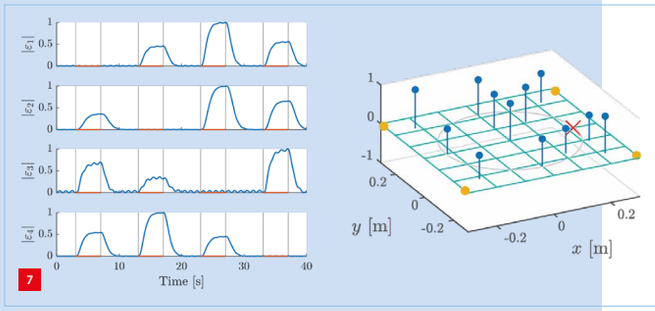

Four of the 17 residual signals ε_i (left). The structure imposed on R_f is clearly visible, allowing to isolate the root cause of the anomalous behaviour. The faults are applied in the highlighted red regions. A digital representation of the counterpart is shown (right), where the red cross marks the faulty component in real-time.

Robustness guarantees

One of the major challenges in fault-diagnosis system design is that it should be able to distinguish faults from disturbances and modelling uncertainty. Disturbances are inherently present in any control system, e.g., sensor noise and external vibrations. The modelling uncertainty can be attributed to limited estimation accuracy, simplifications and assumptions, or system variability. It is widely recognised that a compromise is inevitable between robustness to these disturbances and model uncertainties, and sensitivity to faults.

A framework to optimise the fault sensitivity despite disturbances has been presented in the previous section. This idea can be extended to account for the modelling uncertainty Δ. Recently, a H_-/H_∞ approach has been developed for models with an uncertain system representation, denoted by $G(\Delta)$, where the uncertainty Δ lies in a norm-bounded set $\mathbf{\Delta}$. This is exactly a model set that is often available for precision mechatronics, as it is considered in their control design.

Let us now consider the uncertain input-output relation:

$$y = G(\Delta)u + G_d d + G_f f.$$

Similar to the design of Q_{y1} and Q_{u1} in the previous section, the control input, generated by the controller, is largely cancelled through factorising the nominal plant, i.e., $G_u = M_u^{-1} N_u$. Hence, take $Q_{y1} = M_u$ and $Q_{u1} = -N_u$. The remaining effect through modelling uncertainty can be incorporated in an augmented disturbance matrix $\bar{G}_d$, while a similar Ricatti-based algorithm is solved to maximise the fault sensitivity, subjected to modelling uncertainty and disturbances.

Figure 8 shows an example of the transfer functions from the augmented disturbances and faults to the residual, i.e., $R_{\bar{d}}$ and R_f, for realisations within the set of uncertain systems. It can be concluded that the effect of the disturbance and modelling uncertainty lies below the specified dashed line for all different uncertainty realisations and the worst-case fault sensitivity is quantified. This approach guarantees that the optimal filter truly detects faults and is invariant to disturbances and implications of a poor model.

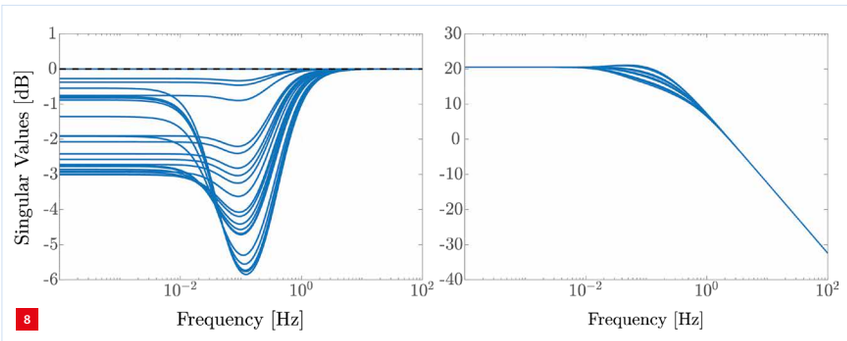

Illustration of the responses from generalised disturbance $\bar{d}$ (left) and fault f (right) to residual ε for all $\Delta \in \Delta$. The transfer from generalised disturbance to residual is clearly bounded by the specified dashed line.

Parametric faults in closed-loop controlled systems

Ageing-induced wear is a typical phenomenon that causes the system characteristics to change. Even slight changes in terms of stiffness or damping, that is, resonance characteristics, are key indicators of an increased risk of failure [15]. This is an example of a multiplicative fault f_{par}. First, the importance of detecting these faults is illustrated, followed by a case study.

Multiplicative faults are particularly dangerous for closed-loop controlled systems, as the controller has often been designed without considering the effect of potential parametric faults f_{par}. Consider the block diagram in Figure 2, without additive faults, i.e., $f_{add} = 0$, and without modelling uncertainty, i.e., $\Delta = 0$. In that case, the error signal is equal to:

$$e = S(f_{\text{par}})r - S(f_{\text{par}})G_d d.$$

Here, the fault-dependent sensitivity function is:

$$S(f_{\text{par}}) = \left(1 + G(f_{\text{par}})C\right)^{-1}.$$

As a result, performance is directly affected by faults f_{par}. The feedback controller often contains dedicated notch filters to attenuate the effect of specific resonances. A slightly changing G due to f_{par} then not only will deteriorate

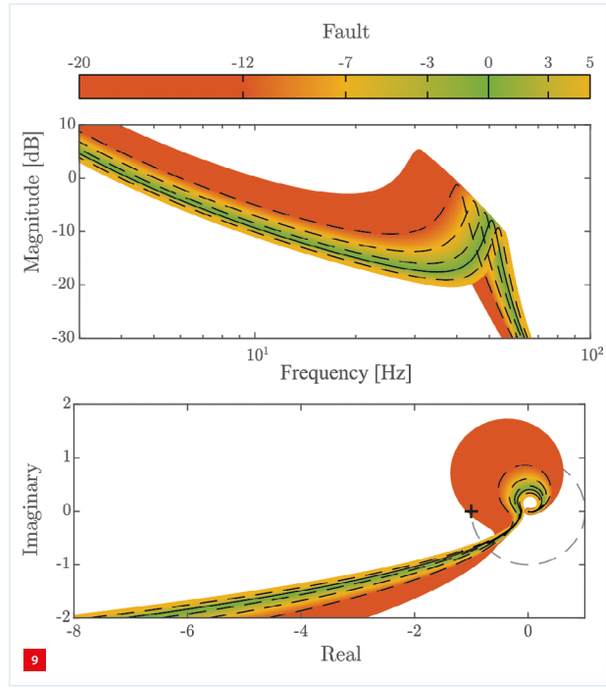

Bode diagram of $G(f_{par})$ and Nyquist diagram of $G(f_{par})C$ illustrating the effect of a parametric fault f_{par} affecting the stability margin. The gain and modulus margin have been drastically affected. The green area indicates the healthy state of the system. The red indicates a poor stability margin.

the closed-loop performance, but also drastically affect traditional stability margins such as the modulus margin $\max_\omega |S(j\omega)|$. An example of a spring characteristic changing due to a fault and its influence on stability margins is depicted in Figure 9. As illustrated in this example, monitoring changes in the system itself is highly desirable. To this end, inexpensive and online estimation algorithms are particularly useful [3]. However, off-the-shelf algorithms do not always provide the desired interpretability. For fault isolation in mechatronic applications, a newly developed and more interpretable algorithm is better suited, as illustrated by the following case study.

Case study: Online estimation of modal models for mechanical systems

Mechanical systems are typically described as [23]:

$$G(s) = \sum_{i=1}^{n_{\mathrm{rb}}} \frac{c_i b_i^T}{s^2} + \sum_{i=n_{\mathrm{rb}}+1}^{n_s} \frac{c_i b_i^T}{s^2 + 2\zeta_i \omega_i s + \omega_i^2}.$$

Here, n_{rb} is the number of rigid-body modes and n_s the total number of modes; the vectors c_i, b_i are associated with the mode shapes, and ζ_i, ω_i are the damping ratio and natural frequency, respectively.

Typically, for single-input single-output systems, a single unfactored transfer function $G(s) = B(s)/A(s)$ is estimated. However, we aim to estimate the dynamics in a modal decomposition:

$$G(s) = \sum_{i=1}^{n_s} \frac{B_i(s)}{A_i(s)}.$$

This is in sharp contrast to traditional methods, where interpretability is mostly lost. Here, we consider methods based on pseudolinear regression:

$$y(t_k, \theta) = \varphi^T(t_k, \theta)\theta + \varepsilon(t_k, \theta).$$

Here, the regressor $\varphi(t_k, \theta)$ contains filtered versions of the input-output data and its derivatives at time-sample t_k. The estimated parameter vector of each mode, θ_i, directly relates to the gains, damping and resonance frequency associated with each modeshape, i.e., c_i, b_i, ζ_i and ω_i, respectively.

To determine the parameters associated with each of the modes $G_i(s) = B_i(s)/A_i(s)$ recursively and in real-time, we compute:

$$L_i(k) = P_i(k-1)\hat{\varphi}_i(k)\left(\lambda(k) + \varphi_i^T(k)P_i(k-1)\hat{\varphi}_i(k)\right)^{-1}$$

$$\theta_i(k) = \theta_i(k-1) + L_i(k)\left(y(k) - \sum_1^{n_s} y_i(k)\right)$$

$$P_i(k) = \frac{1}{\lambda(k)}\left(P_i(k-1) - L_i(k)\varphi_i^T(k)P_i(k-1)\right).$$

This is done at every t_k for every mode i, where y_i is an estimate of the modal contribution to the total output. A forgetting factor $\lambda(k)$ is used to give more weight to recent data while gradually reducing the influence of older data. To address bias in closed-loop estimation due to the correlation between d and u, so-called instrumental variables (IVs) are used.

An overactuated and oversensed flexible beam set-up was considered. Due to an additional actuator-sensor pair, an artificial control loop was created, which allowed to artificially increase or decrease the stiffness and/or damping of the beam; see Figures 10 and 11. The internal damping and stiffness of the beam was manipulated during an experiment between t = 20 s and t = 80 s; the parameter estimates associated with the modes are shown in Figure 12, where the parameters in the left column relate to the first mode, i.e. $B_1 = b_{10}$ and $A_1 = a_{12}s^2 + a_{11}s + 1$. Similarly, the parameters in the second column relate to the second mode with $B_2 = b_{20}$ and $A_2 = a_{22}s^2 + a_{21}s + 1$. The changing second mode was indeed detected and estimated, illustrating the effectiveness of the proposed approach.

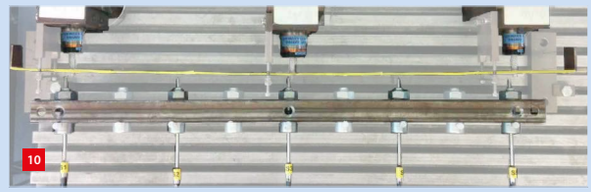

Overactuated and oversensed flexible beam system. The beam is equipped with yellow tape and suspended with wire flexures. It is actuated with voice-coils (top) and the deflection is measured with fibre-optic sensors (bottom).

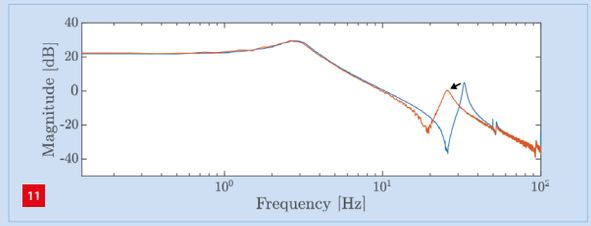

Frozen frequency response functions (FRFs) illustrating the shift of the internal dynamics of the system. This corresponds to a change in the second resonance peak.

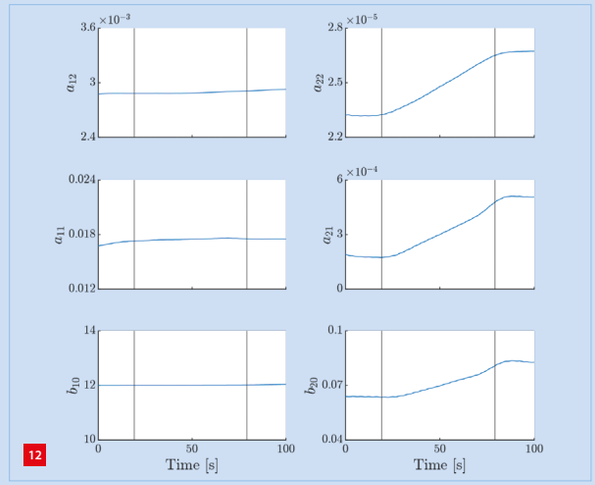

Parameters associated with the first and second mode on the left and right, respectively.

Reconfiguration and self-healing

If a fault is successfully detected, or even better, an estimate of the fault is available, then this information can be used to reconfigure the controller and minimise or even fully mitigate performance loss [24]. This allows to significantly prolong uptime in the event of a fault. To this end, several approaches are envisioned [12] [V1], ranging from virtual sensors/actuators to adaptive modal decoupling. These methods allow for reallocation of the control inputs and mainly exploit redundancy in the number of sensors or actuators. This is illustrated on the prototype wafer stage in the next case study.

Case study: Reconfiguration for a next-generation prototype wafer stage

The overactuated test-rig (OAT) has more actuators than degrees of freedom. This allows redistributing control effort of a (partly) broken actuator over the remaining healthy actuators. Figures 13 and 14 show the response when relying on a faulty actuator. This results in heavy vibrations of the chuck, leading to major performance loss. Next, the response is shown where the control effort of the faulty actuator has been gradually redistributed over the healthy actuators, thereby mitigating performance loss.

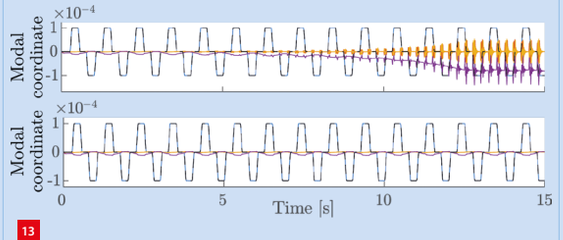

13

*Response of the OAT without (top) and with (bottom) controller reconfiguration. The blue lines indicate the **z**-position and the black dashed lines indicate the setpoint. The red, yellow and purple lines show the tip, tilt and torsional coordinates of the stage, respectively. Without reconfiguration, the chuck starts to vibrate heavily, leading to major performance loss.*

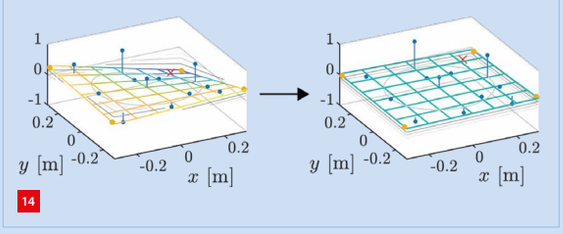

14

Response of the OAT without (left) and with (right) controller reconfiguration. Without reconfiguration, the chuck starts to vibrate heavily. Reconfiguration mitigates this effect.

Conclusion and outlook

Model-based fault-diagnosis systems for high-tech systems are highly promising and have many benefits to offer in terms of predictive maintenance and controller reconfiguration. This is enabled by the readily available models in high-tech systems that can easily be repurposed. The recent theoretical advancements and practical case studies give a broad perspective on how these concepts can drive substantial enhancements in system uptime and hence productivity.

Acknowledgement

This work is supported by *Topconsortia voor Kennis en Innovatie* (TKI) and ASML Research, Veldhoven (NL). Special thanks to the involved students Tjeerd Ickenroth, Mike Mostard, Paul Munns, Stan de Rijk, and Stan Verbeek.

OVERVIEW VIDEO

A short overview of the associated research activities has been published in the form of a video [V1].

[V1] K. Classens, W.P.M.H. Heemels, and T. Oomen, Digital-Twin Enabled Predictive Maintenance for Machines (July 25, 2023). Accessed: Oct. 18, 2023, [online video],

www.youtube.com/watch?v=i3DoI1vxAqA

REFERENCES

[1] T. Zonta, *et al.*, "Predictive maintenance in the Industry 4.0: A systematic literature review", *Comput. Ind. Eng.*, vol. 150, pp. 106889, *2020, doi:10.1016/j.cie.2020.106889*

[2] Y. Ran, *et al.*, "A survey of predictive maintenance: Systems, purposes and approaches", 2019, *arXiv:1912.07383*

[3] R. Isermann, "Model-Based Fault Detection and Diagnosis - Status and Applications," *Annu.* Rev. *Control*, vol 29 (1), pp. 71-85, 2005, *doi:10.1016/j.arcontrol.2004.12.002*

[4] Z. Gao, C. Cecati, and S.X. Ding, "A Survey of Fault Diagnosis and Fault-Tolerant Techniques-Part I: Fault Diagnosis With Model-Based and Signal-Based Approaches", *IEEE Trans. Ind. Electron.*, vol. 62 (6), pp. 3757-3767, 2015, *doi:10.1109/TIE.2015.2417501.*

[5] F. Tao, *et al.*, "Digital Twin in Industry: State-of-the-Art", I*EEE Trans. Industr. Inform.*, vol. 15 (4), pp. 2405-2415, 2019, *doi:10.1109/TII.2018.2873186*

[6] D. Jones, *et al.*, "Characterising the Digital Twin: A systematic literature review", *CIRP J. Manuf. Sci. Technol*, vol. 29, pp. 36-52, 2020, *doi:10.1016/j.cirpj.2020.02.002*

[7] A. Fuller, *et al.*, "Digital Twin: Enabling Technologies, Challenges and Open Research", *IEEE Access*, vol. 8, pp. 108952-108971, 2020, *doi:10.1109/ACCESS.2020.2998358*

[8] A. Rasheed, O. San, and T. Kvamsdal, "Digital Twin: Values, Challenges and Enablers From a Modeling Perspective", *IEEE Access,* vol. 8, pp. 21980-22012, 2020, *doi:10.1109/ACCESS.2020.2970143*

[9] Z. Gao, C. Cecati, and S.X. Ding, "A Survey of Fault Diagnosis and Fault-Tolerant Techniques-Part II: Fault Diagnosis With Knowledge-Based and Hybrid/Active Approaches," *IEEE Trans. Ind. Electron*, vol. 62 (6), pp. 3768-3774, 2015, *doi: 10.1109/TIE.2015.2419013.*

[10] Y. Lei, *et al.*, "Applications of machine learning to machine fault diagnosis: A review and roadmap", *Mech. Syst. Signal. Process.*, vol. 138, pp. 106587, 2020, *doi:10.1016/j.ymssp.2019.106587*

[11] K. Patan, *Robust and Fault-Tolerant Control: Neural-Network-Based Solutions*, Springer, vol. 197, 2019.

[12] K. Classens, W.P.M.H. Heemels, and T. Oomen, "Digital Twins in Mechatronics: From Model-based Control to Predictive Maintenance", *2021 IEEE International Conference on Digital Twins*

and Parallel Intelligence (DTPI), Beijing, China, 2021, pp. 336-339, *doi:10.1109/DTPI52967.2021.9540144*
[13] M. Grieves, *Digital twin: manufacturing excellence through virtual factory replication*, White Paper, 2014.
[14] K. Classens, W.P.M.H. Heemels, and T. Oomen, "Closed-loop Aspects in MIMO Fault Diagnosis with Application to Precision Mechatronics", *2021 IEEE American Control Conference (ACC)*, New Orleans, LA, USA, 2021, pp. 1756-1761, *doi:10.23919/ACC50511.2021.9482785*
[15] K. Classens, *et al.*, "Fault Detection for Precision Mechatronics: Online Estimation of Mechanical Resonances", in *2nd Modeling, Estimation and Control Conference (MECC 2022)*, Jersey City, New Jersey, USA, 2022, *IFAC Papers OnLine*, vol. 55 (37), pp 746-751, *doi:10.1016/j.ifacol.2022.11.271*
[16] R. Isermann, *Fault-diagnosis systems: an introduction from fault detection to fault tolerance*, Springer, 2006, *doi:10.1007/3-540-30368-5*
[17] M. Blanke, *et al.*, *Diagnosis and Fault-Tolerant Control*, Springer, 2006, *doi:10.1007/978-3-540-35653-0*
[18] S.X. Ding, *Model-based Fault Diagnosis Techniques: Design Schemes, Algorithms, and Tools*, Springer, 2013, *doi:10.1007/978-1-4471-4799-2*
[19] A. Varga, *Solving Fault Diagnosis Problems: Linear Synthesis Techniques*, Springer, 2017, *doi:10.1007/978-3-319-51559-5*
[20] K. Classens, W.P.M.H. Heemels, and T. Oomen, "Direct Shaping of Minimum and Maximum Singular Values: An H_-/H_∞ Synthesis Approach for Fault Detection Filters", *Proceedings of the IFAC 22nd Triennial World Congres*, Yokohama, Japan, 2023, *doi:10.48550/arXiv.2305.07258*
[21] N. Liu, and K. Zhou, "Optimal solutions to multi-objective robust fault detection problems", *46th IEEE Conference on Decision and Control*, New Orleans, LA, USA, 2007, pp. 981-988, *doi:10.1109/CDC.2007.4434123*
[22] S.X. Ding, *et al.*, "A unified approach to the optimization of fault detection systems," *Int. J. Adapt. Control Signal Process.*, vol 14 (7), pp. 725-745, 2000, *doi:10.1002/1099-1115(200011)14:7<725::AID-ACS618>3.0.CO;2-Q*
[23] W. Gawronski, *Advanced structural dynamics and active control structures*, Springer, 2004, *doi:10.1007/978-0-387-72133-0*
[24] I. Hwang, *et al.*, "A Survey of Fault Detection, Isolation, and Reconfiguration Methods", *IEEE Trans. Control Syst. Technol.*, vol. 18 (3), pp. 636-653, 2010, *doi:10.1109/TCST.2009.2026285*