Weel & Sandvig
ENERGY AND PROCESS INNOVATION

Reliable Emission monitoring with WS.PEMS

WS.PEMS: Using data reconciliation for predictive emission monitoring systems

Joachim Weel Rosbo
Emission Modeling Engineer
Weel & Sandvig

Peter Jul-Rasmussen
Emission Modeling Engineer
Weel & Sandvig

Agenda

1. **Weel & Sandvig and WS.PEMS**
   Who we are, and what a predictive emission monitoring system does

2. **Data reconciliation**
   Redundancy and observability — why a model can correct a sensor

3. **The gas turbine model**
   From machine to unit components, closed with Stodola and Greitzer

4. **Gross error detection**
   The global test, and four strategies for naming the faulty sensor

5. **Results**
   What the methods catch, what they miss, and what it costs

About Weel & Sandvig

Organisation

- Main office in Lyngby (25 employees)
- Office in Abu Dhabi
- Costumers all over the world
- More than 20 years of experience in emission monitoring & performance optimization

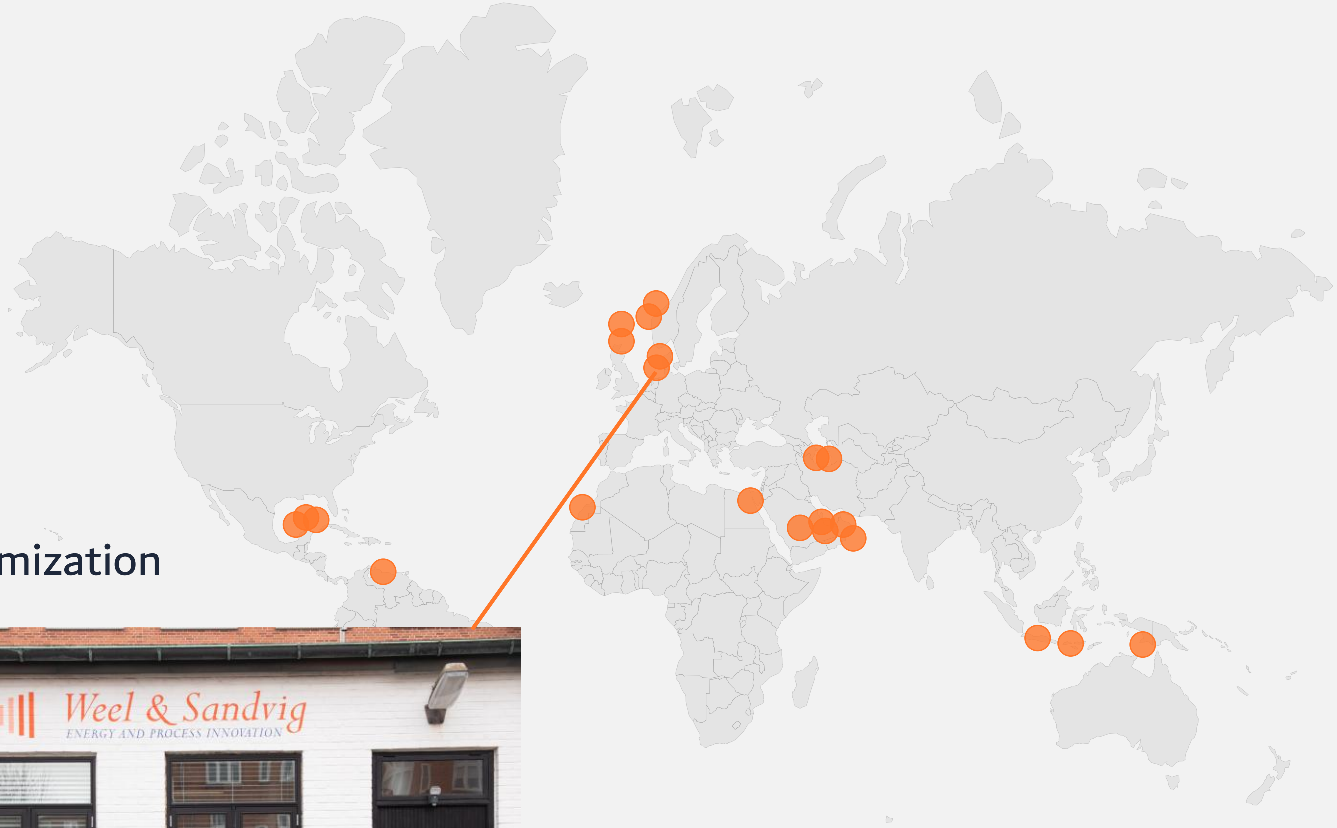

Selected Clients

ADNOC bp OXY

aramco ConocoPhillips

Learn more on **weel-sandvig.com**

Our Key Competences

Fundamental insight

- Combustion processes
- Energy intensive industrial processes

Process modelling and simulation

- First Principle Models
- Online monitoring systems

Stack emission measurements

- Planning
- Execution


CONTENTS 0.2. VARIABL

The mass burn rate is calculated assuming that reaction R1 accounts for all methane combustion. And that all heat release is fr the reaction enthalpy. The mass burned is calculated by dividing the amount of heat released by the reaction enthalpy multiplied over the mass per mole burned.

0.2.2 Massflow

The mass flow through the inlet and the outlet is calculated by using the flowing relations, where P_{in} is the pressure in the intake P_{ex} is the pressure in the exhaust. R_m is the gas constant divided with the molar mass of the gas.

$$r_c = \left(\frac{\gamma+1}{2}\right)^{\frac{\gamma}{\gamma-1}}$$

An empirical correlation is used to model the flow inside the cylinder.

$$w = C_1 S_p + C_2 V_d \frac{T_r}{P_r V_r}\left(P - P_r\left(\frac{V_r}{V}\right)^{\gamma}\right)$$

$$C_1 = \begin{cases} 6.18 & \forall\theta \in [\theta_{EO}, 720] \vee [0, \theta_{IC}] \\ 2.28 & \forall\theta \in]\theta_{IC}, \theta_{EO}[\end{cases}$$

$$C_2 = \begin{cases} 0.00324 & \forall\theta \in [0, 180] \vee [\theta_{PC}, 540] \\ 0 & \forall\theta \in]180, \theta_{PC}[\vee]540, 720] \end{cases}$$

An empirical formula for an internal combustion engine calculates the heat transfer coefficient. Where the pressure is instead in Pa and the unit of the coefficient is $\frac{W}{m^2K}$

$$h = 3.26 \cdot 10^{-3.4} \cdot B^{-0.2} P^{0.8} T^{-0.55} w^{0.8}$$

The heat transfer to the wall can then be calculated as seen in equation (32), where it is assumed that the wall temperature T_w constant.

$$\dot{Q}_W = A_w h (T - T_w)$$

The area where a heat transfer occurs is calculated from the volume.

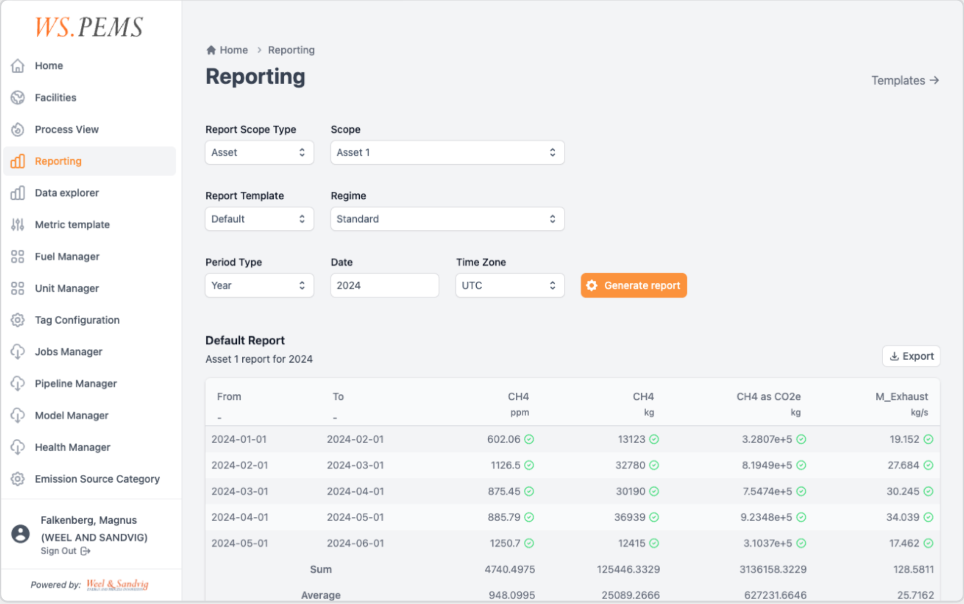

Introduction to WS.PEMS

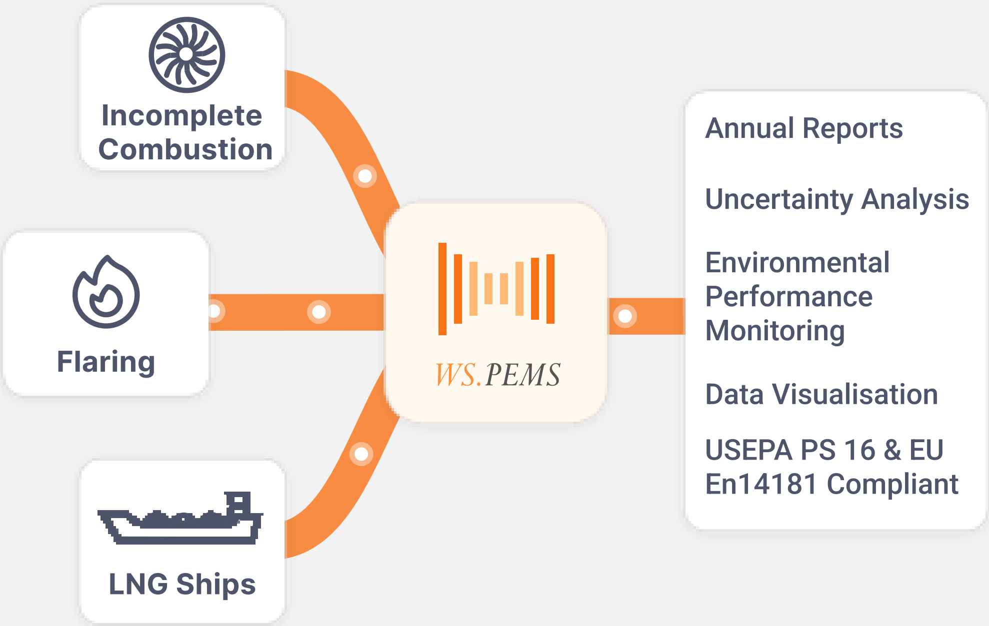

Introduction to WS.PEMS

Emission monitoring methodologies

- Emission factors (Cheap and inaccurate)
- Online gas analysers (CEMS) (Accurate but costly)
- Predictive Emission Monitoring Systems (PEMS)

PEMS – A Cost Effective Alternative to CEMS

- Same compliance and regulation requirement
- Better insight

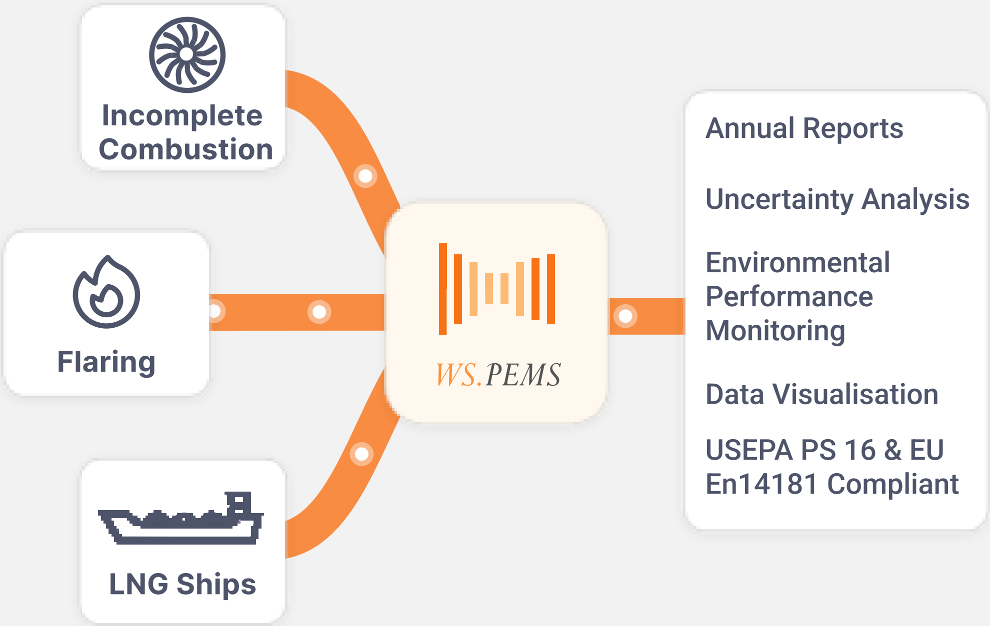

Industrial turbine example

Two-shaft industrial gas turbine

- Shaft power ≈ 25 MW
- Heat input ≈ 75 MW (1.50 kg/s fuel, LHV ≈ 50 MJ/kg)
- Thermal efficiency ≈ 33 %
- Gas generator ≈ 9 580 rpm, power turbine ≈ 3 591 rpm
- Compressor discharge ≈ 20.5 bar, 740 K
- Turbine exit ≈ 1 098 K

PEMS requirements:

- Monitoring of pollutants and climate emissions (NO_x, CO, SO_2, CO_2, CH_4)
- 1-minute datapoint frequency
- No CFD!

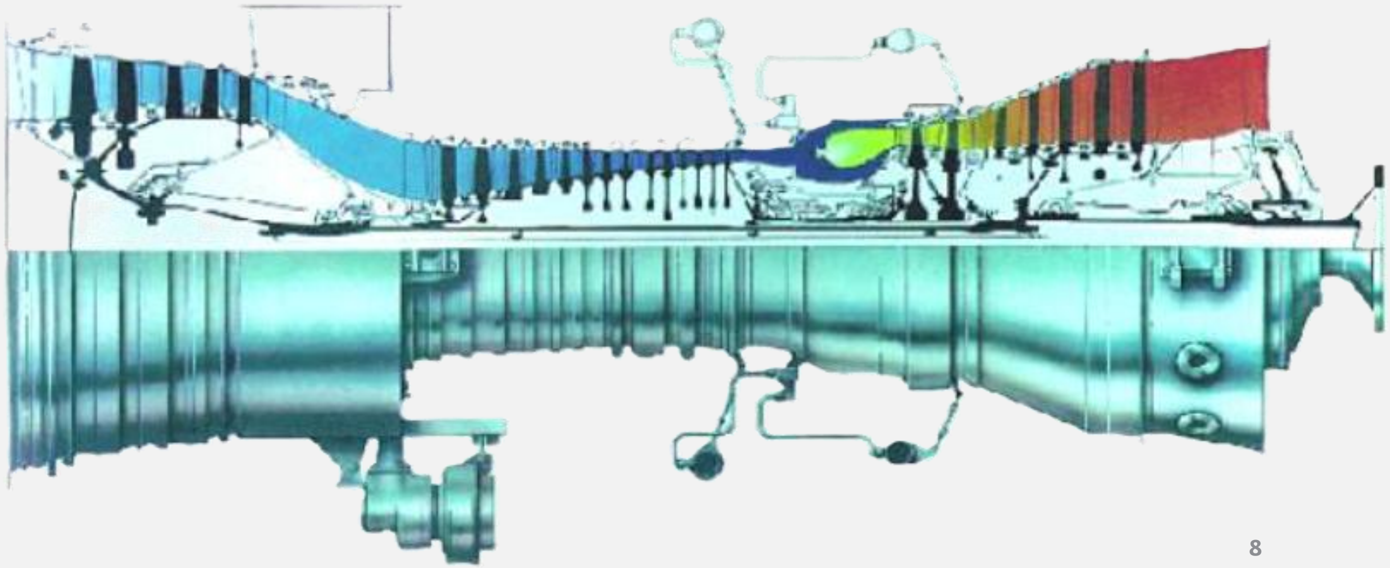

From machine to unit-component model

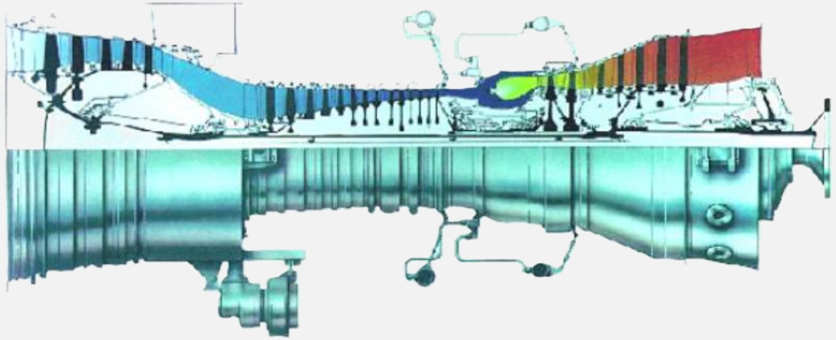

The machine becomes a network of unit components

- Inlet → compressor → combustor → turbines → exhaust

Each unit contributes balance equations

- Mass, energy and pressure relations per component
- 32 balance residuals in total for this model

Process data (14 measurements)

- Ambient T_0, P_0, RH_0 and inlet ΔP
- Compressor T_2, T_3, P_3
- Turbine T_{45} / P_{45} and exhaust T_5
- Fuel flow and both shaft speeds

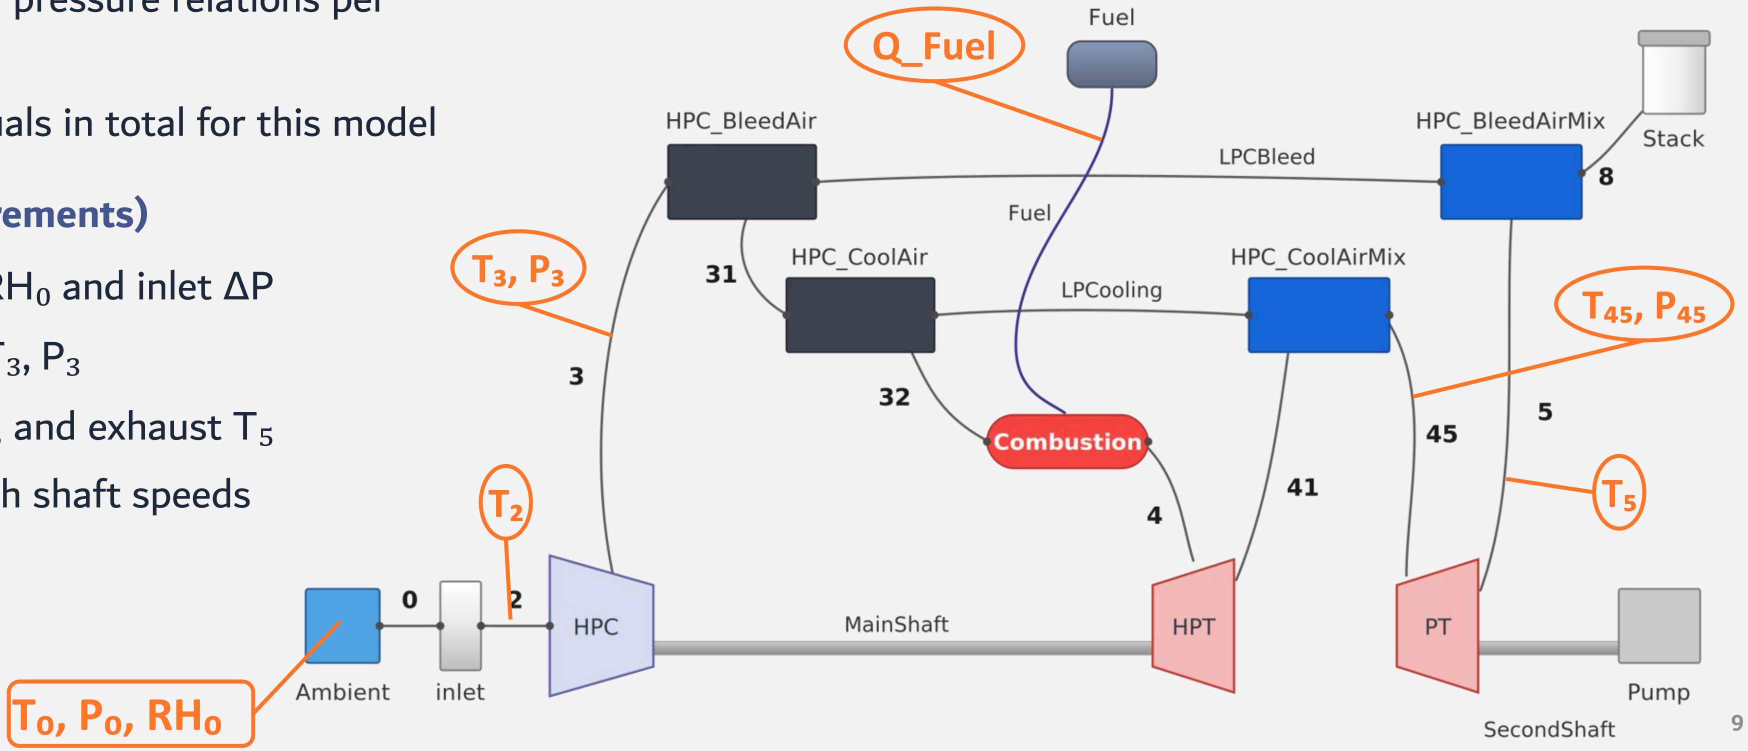

Process data and measurement uncertainty

Real operating data from the unit

- Steady load, sampled every minute

σ estimated directly from the signal scatter

- These σ set the weights in the objective
- They also define what counts as a gross error

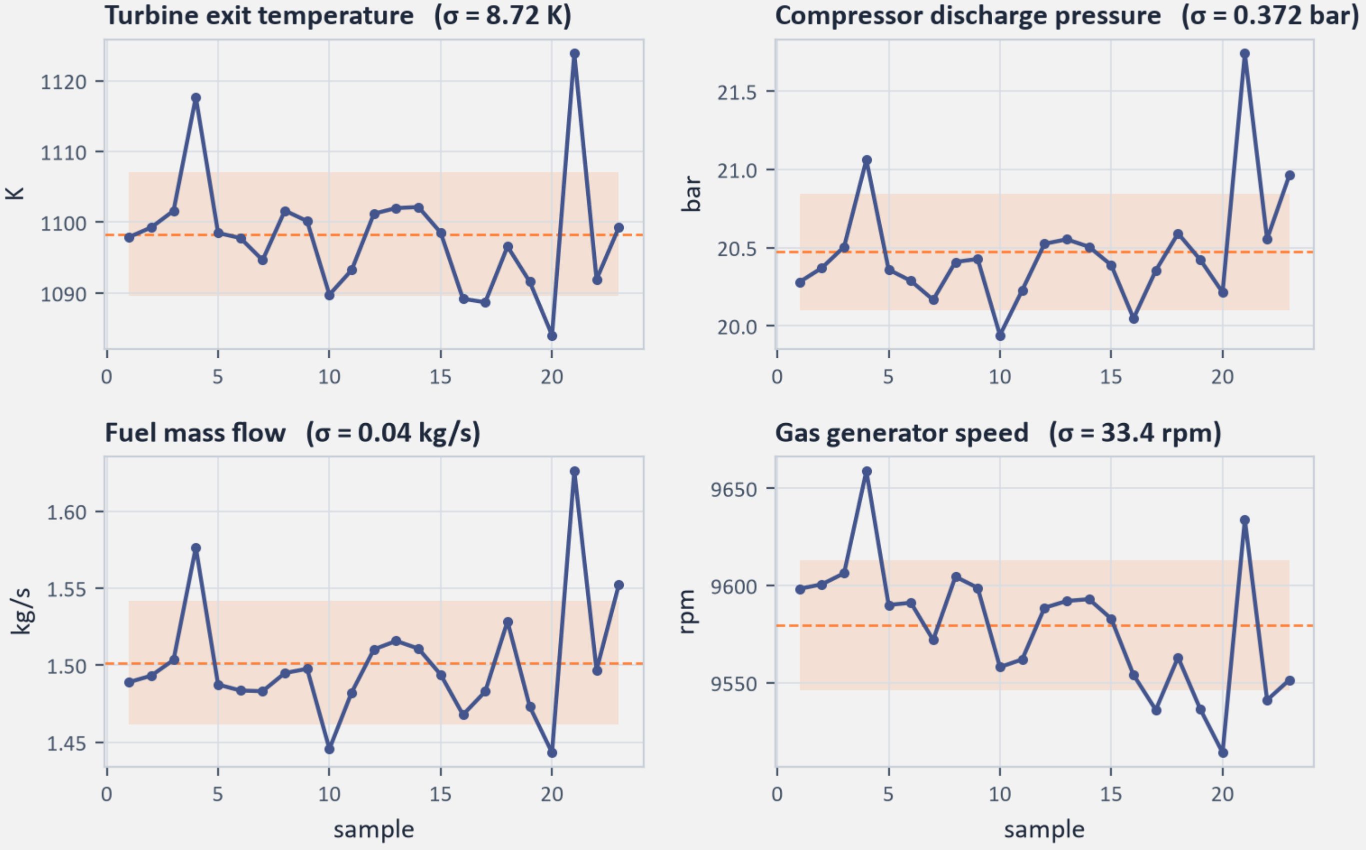

Measurement	Mean	σ
T_{45} [K]	1098	8.7
T_5 [K]	812.7	5.4
T_3 [K]	740.0	4.4
P_3 [bar]	20.47	0.37
P_{45} [bar]	4.38	0.077
Fuel [kg/s]	1.502	0.040

Data reconciliation

Plant sensors never agree exactly

- Balances are violated by measurement noise

Reconciliation finds the state that best fits every measurement

- Weighted by each sensor's own uncertainty σ_i

Redundancy is what makes it possible

- ν = (independent constraints) – (unmeasured variables)
- $\nu > 0$ means surplus information: errors become visible
- Only redundant measurements can be corrected or checked

$$\min_x \sum_i \left(\frac{y_i - x_i}{\sigma_i}\right)^2$$

$$\text{subject to } \; h(x) = 0$$

measurements $y \in \mathbb{R}^m$, unmeasured $u \in \mathbb{R}^p$

constraints $h(x, u) = 0, \quad h : \mathbb{R}^n \to \mathbb{R}^q$

degree of redundancy $\nu = q - p$

$\nu > 0 \;\Rightarrow\;$ surplus information $\;\Rightarrow\;$ errors detectable

Increasing observability

With only mass and energy balances the turbine is sometimes under-determined

- Standard instrumentation does not measure enough of the internal state

Generalised data reconciliation adds unit characteristics

- Stodola's ellipse for each turbine stage
- Greitzer-style map for the compressor

Added as soft constraints, not hard equalities

- Each characteristic residual carries its own variance σ_j
- Raises the degree of redundancy to $\nu = 6$ for this machine

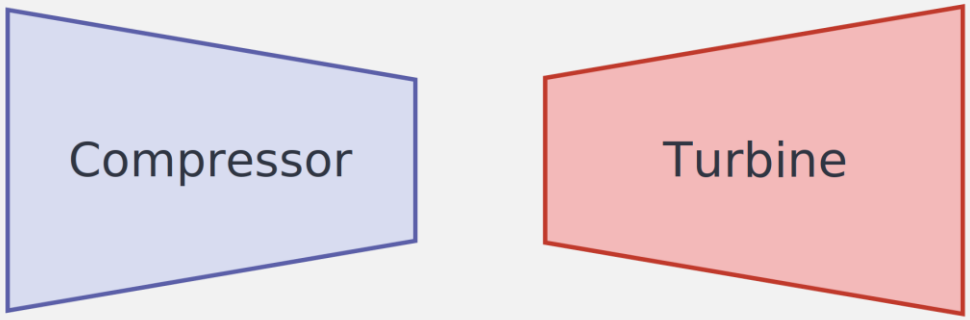

$$\min_x \sum_i \left(\frac{y_i - x_i}{\sigma_i}\right)^2 + \sum_j \left(\frac{h_j(x)}{\sigma_j}\right)^2$$

$$\text{Stodola}: \quad \dot{m}\sqrt{T_{in}}/p_{in} = f(p_{out}/p_{in})$$

$$\text{Greitzer}: \quad \pi_c = f(\dot{m}\sqrt{T}/p,\ N/\sqrt{T})$$

Coupled compressor–turbine matching

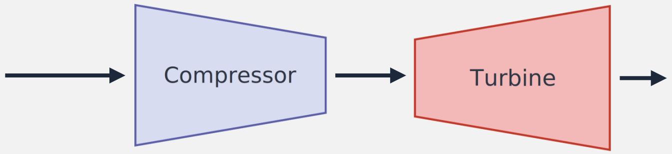

- Compressor pressure ratio and speed -> Mass flow
- Turbine pressure ratio and mass flow -> Turbine inlet temperature
- Connected to the fuel flow

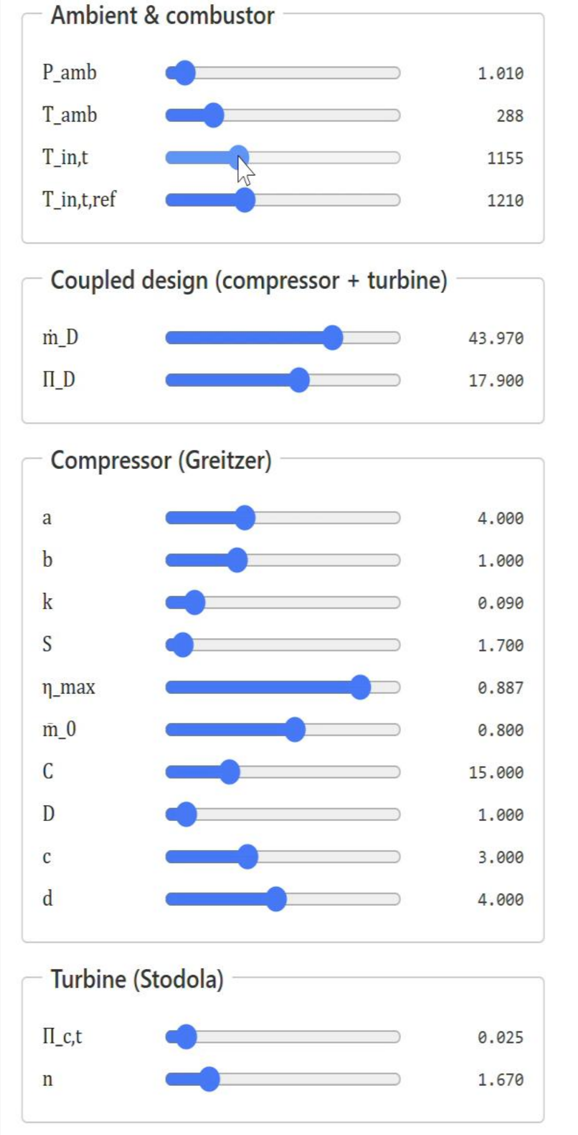

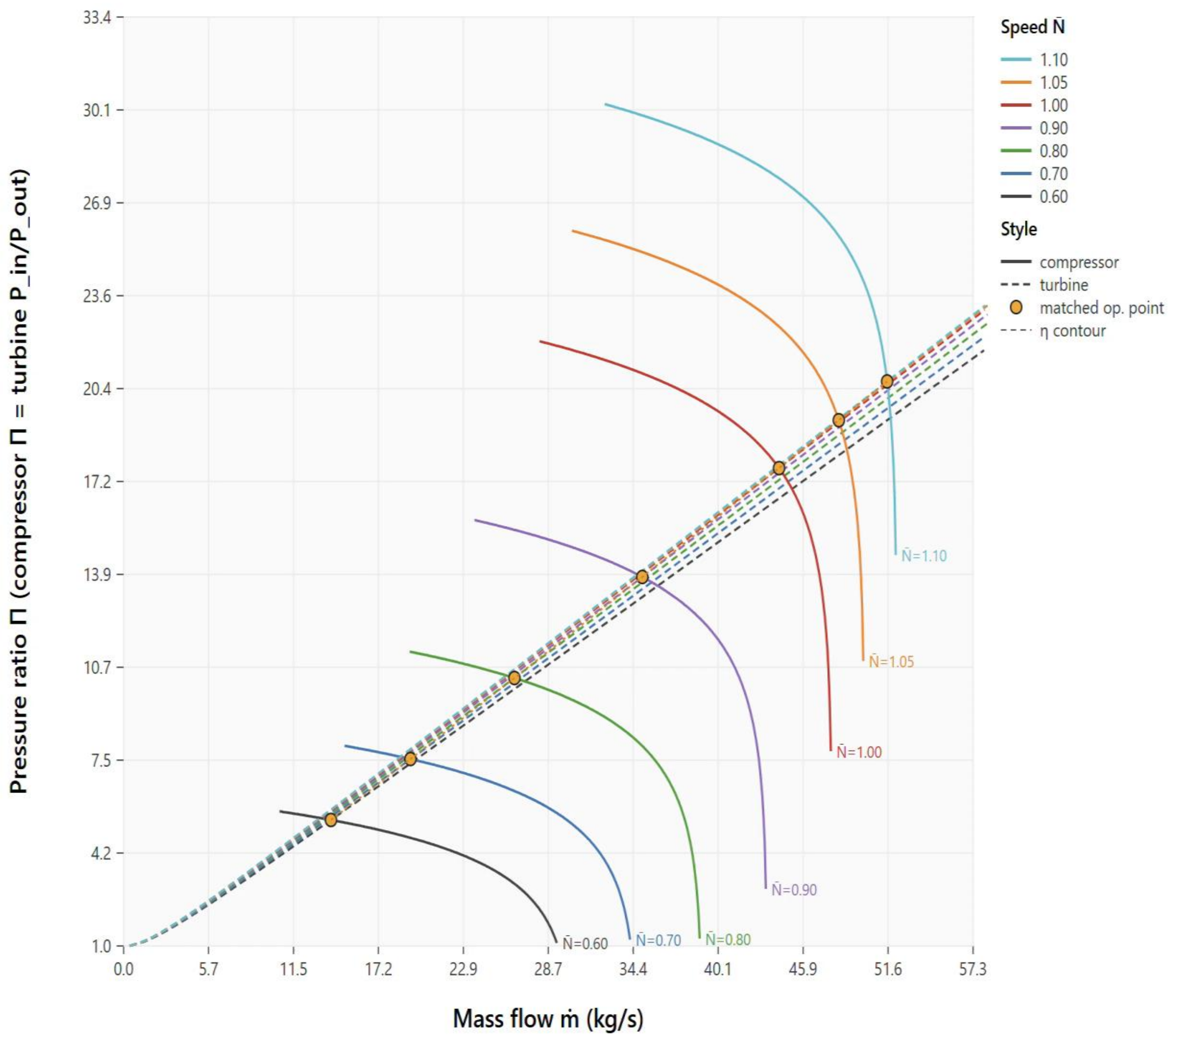

Weel & Sandvig
ENERGY AND PROCESS INNOVATION

Reconciliation result

14 measurements are reconciled simultaneously

- Solved with sequential least-squares programming (SLSQP)
- ≈ 120 ms per reconciliation — comfortably real-time at 1-minute sampling

At a healthy operating point the model reconciles cleanly

- Test: $GT = \min \sum_{i \in n_m} \left(\frac{y_i - \hat{x}_i}{\sigma_i}\right)^2 < \chi^2{}_{1-\alpha}(\nu)$
- Global test statistic $\chi^2 = 8.54$
- Threshold $\chi^2_{0.95}(6) = 12.59$ ⟶ passes
- Degree of redundancy $\nu = 6$

Gross error detection

A biased sensor is spread across every reconciled value

Global test: flag when χ^2 exceeds $\chi^2_{0.95}(\nu)$

- Says something is wrong — not which sensor

M1 — measurement test

- Largest normalised error is the prime suspect. One reconciliation, cheapest

M2 — exhaustive masking

- Mask each candidate, re-reconcile, keep whichever gives the lowest χ^2. Most reliable, most expensive

M3 — early termination

- Same loop, but stops as soon as a removal passes the test. Faster, slightly less precise

M4 — MT pre-ranked masking

- Rank candidates by M1 first, then mask in that order — the stop fires sooner

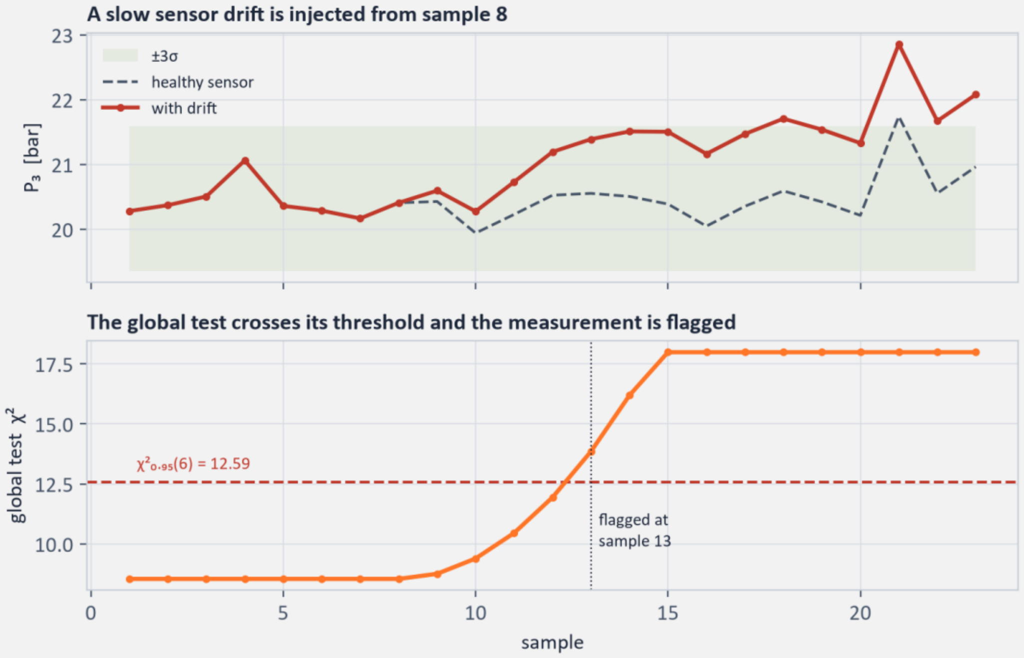

linearise the constraints at the solution : $A = \partial h/\partial x|_{\hat{x}}$

$$a = y - \hat{x} = \Sigma A^T (A \Sigma A^T)^{-1} A\, y$$

$$\mathrm{Cov}(a) = \Sigma A^T (A \Sigma A^T)^{-1} A\, \Sigma \quad \Rightarrow \quad \mathrm{Var}(a_i) = r_i\, \sigma_i^2$$

$$z_i = \frac{|a_i|}{\sigma_i \sqrt{r_i}} \sim \mathcal{N}(0,1) \qquad \text{suspect} = \arg\max_i z_i$$

Test methodology

Controlled perturbation study on real data

- Each of the 14 reconciled measurements biased in turn
- Varies 40σ both directions
- 560 trials; 524 completed

	Flagged	**Not flagged**
Measurement with error	True Positive = TP	False Negative = FN
Measurement without error	False Positive = FP	True Negative = TN

$$\text{Error rate} = \frac{FP + FN}{TP + FP + FN + TN}$$

$$\text{FP ratio} = \frac{FP}{FP + TN}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

14 measurements × 5 magnitudes × 2 signs = 140 trials per method

Weel & Sandvig
ENERGY AND PROCESS INNOVATION

Global-test valley — compressor discharge P_3

A steep valley — all three slides share a ±40σ axis

- Minimum at the unbiased operating point
- Redundancy r ≈ 1.0: nearly the whole bias reaches the test

The most detectable measurement on the machine

- A bias of about 2σ already trips the test

Green markers: the method blamed the right sensor

- Red markers: it blamed a different one

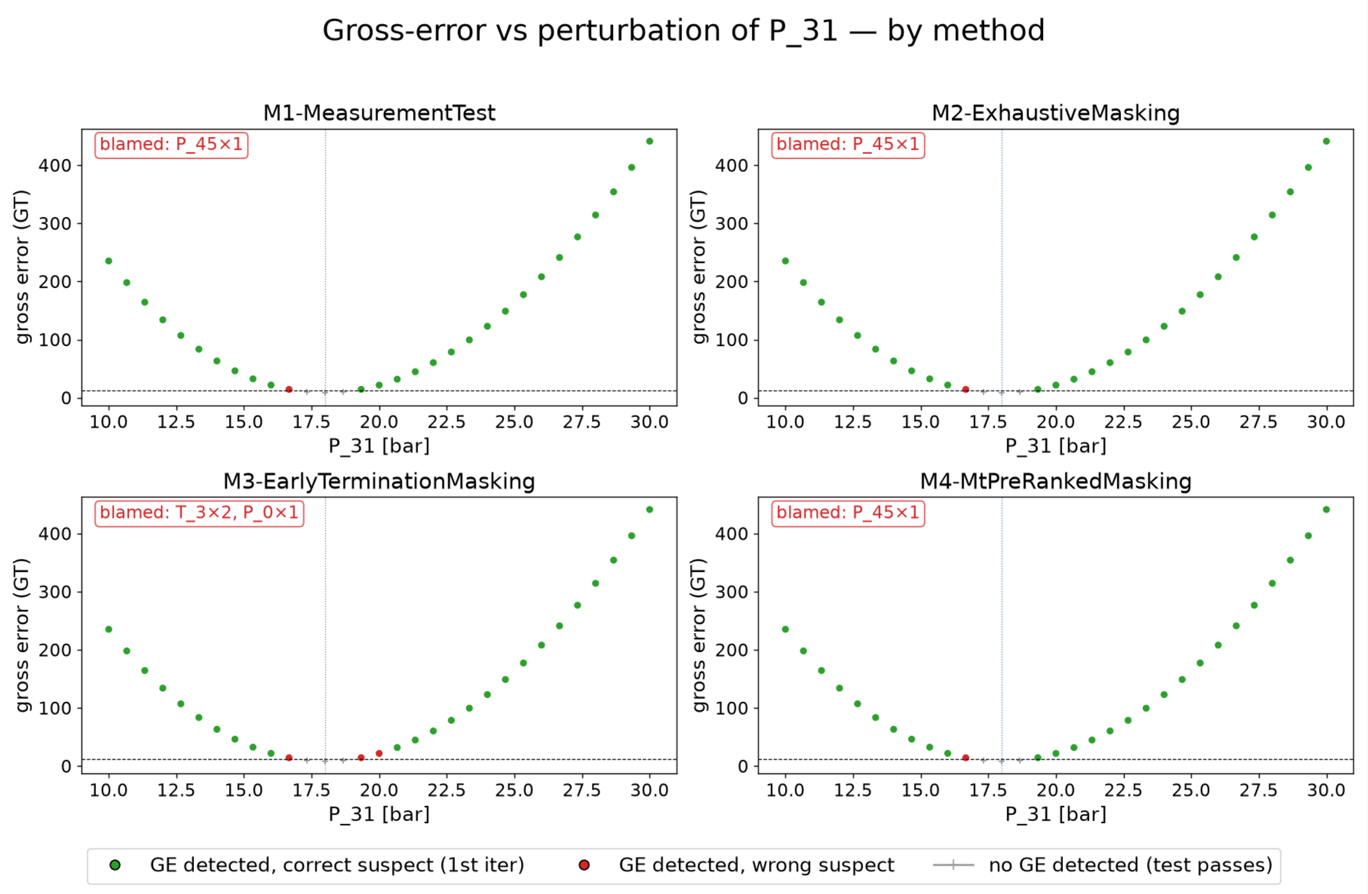

Global-test valley — fuel flow

A wide, shallow valley — note the axis spans ±40σ

- About 11σ before the statistic crosses the threshold

A blind spot that matters

- Fuel flow dominates the emission calculation
- Yet it is the measurement we can police least

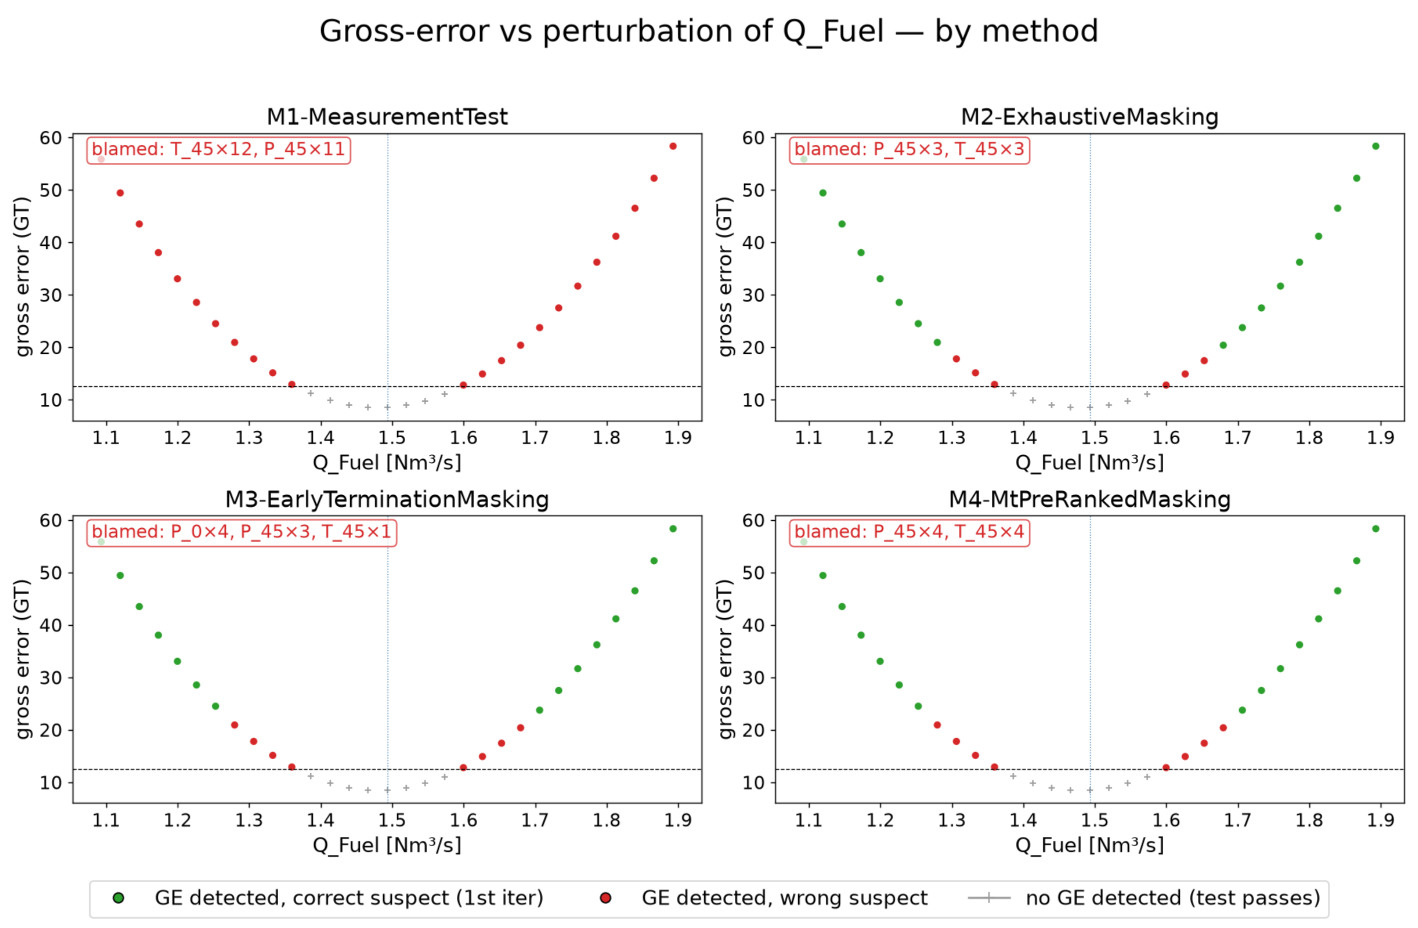

Global-test valley — compressor discharge temperature T_3

An intermediate case

- Redundancy r ≈ 0.33
- Detected from roughly 3.5σ

Attribution is harder than detection

- The test fires, but the wrong sensor is often blamed for M1

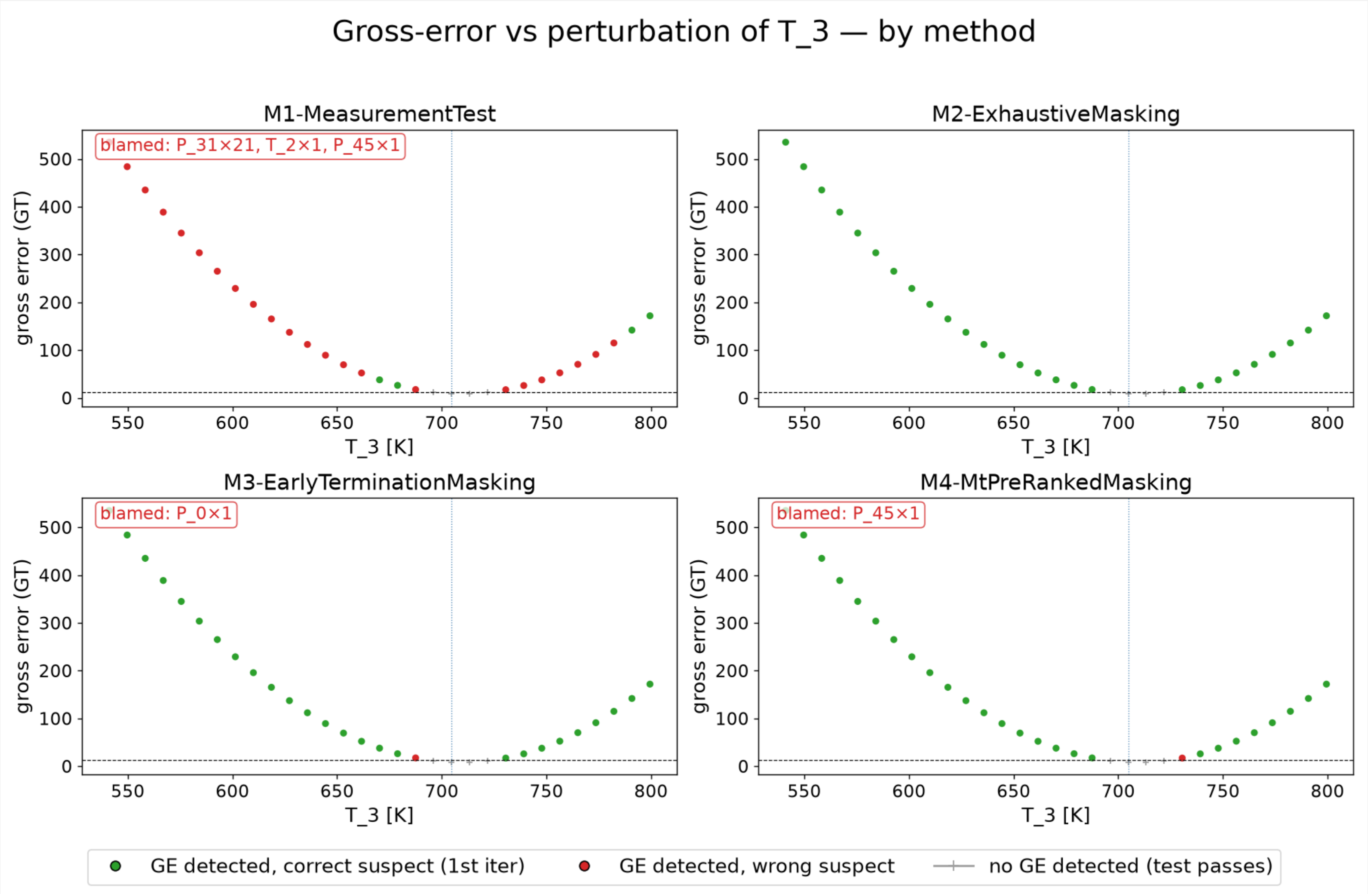

Results — faulty sensor identification

M2 identifies the faulty sensor most reliably

- Precision 0.800 against 0.686 for M1
- But the most expensive — 250 ms per reconciliation

M4 is the practical compromise

- Within 4 % of M2's precision at 37 % of the cost

False positives stay low throughout

- FP ratio ≤ 0.6 % for every method

Metric	M1	M2	M3	M4
Error rate	0.064	0.060	0.063	0.061
FP ratio	0.006	0.004	0.006	0.005
Precision	0.686	0.800	0.714	0.771
Mean time [ms]	80	250	139	93

$$\text{Error rate} = \frac{FP + FN}{TP + FP + FN + TN}$$

$$\text{FP ratio} = \frac{FP}{FP + TN}$$

$$\text{Precision} = \frac{TP}{TP + FP}$$

20

What limits detection

Identification improves with bias magnitude

- Below 2σ no method can succeed

Detectability, not the method, sets the floor

- Only redundant measurements can be checked at all

Fuel flow is the blind spot

- $r_i = 0.03$ — needs ≈ 11σ before the test responds
- Pressures and temperatures are caught at 2–3σ

How r_i is obtained

$$r_i = [\Sigma A^T (A \Sigma A^T)^{-1} A]_{ii} = \frac{\mathrm{Var}(a_i)}{\sigma_i^2}, \qquad 0 \leq r_i \leq 1$$

Identification rate rises with bias magnitude

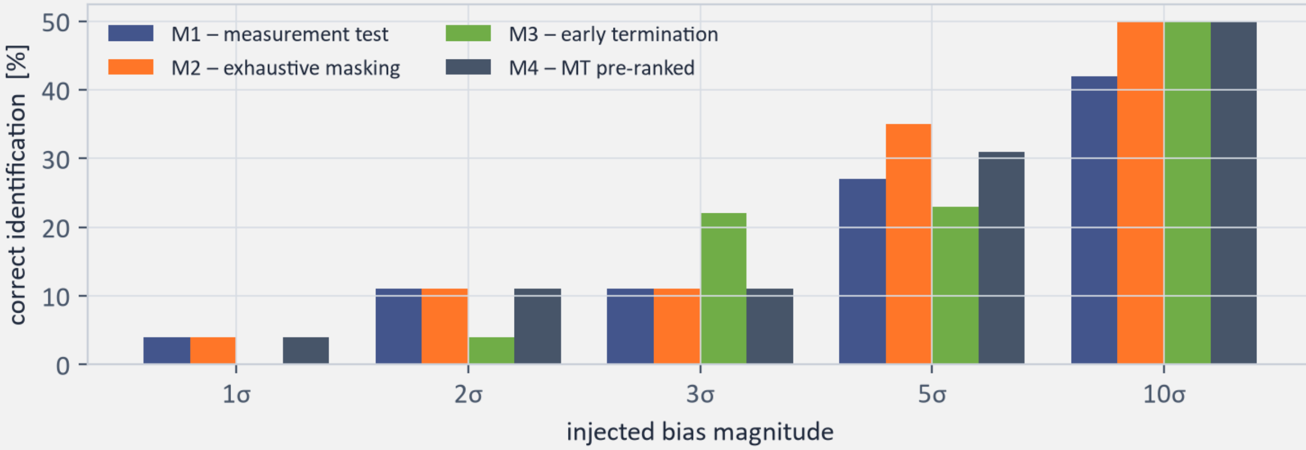

Only measurements with redundancy can be policed

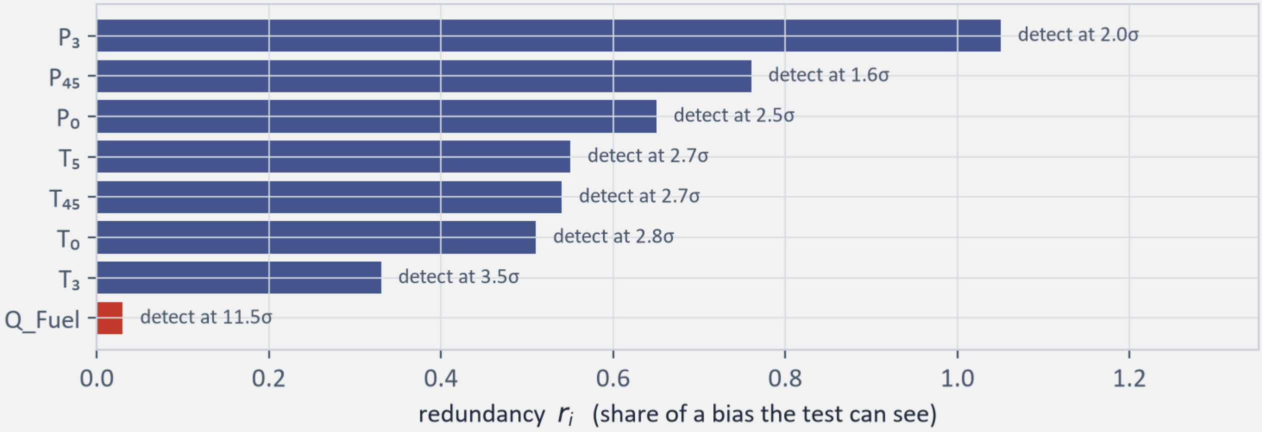

Weel & Sandvig
ENERGY AND PROCESS INNOVATION

Results — leaving the faulty sensor in

A biased P_3 shifts combustor state (top) and emissions (bottom)

- Masking restores T_4, M_4, flame temperature, NO_x and CO — flat green lines
- CH_4 is the exception: it still follows the bias

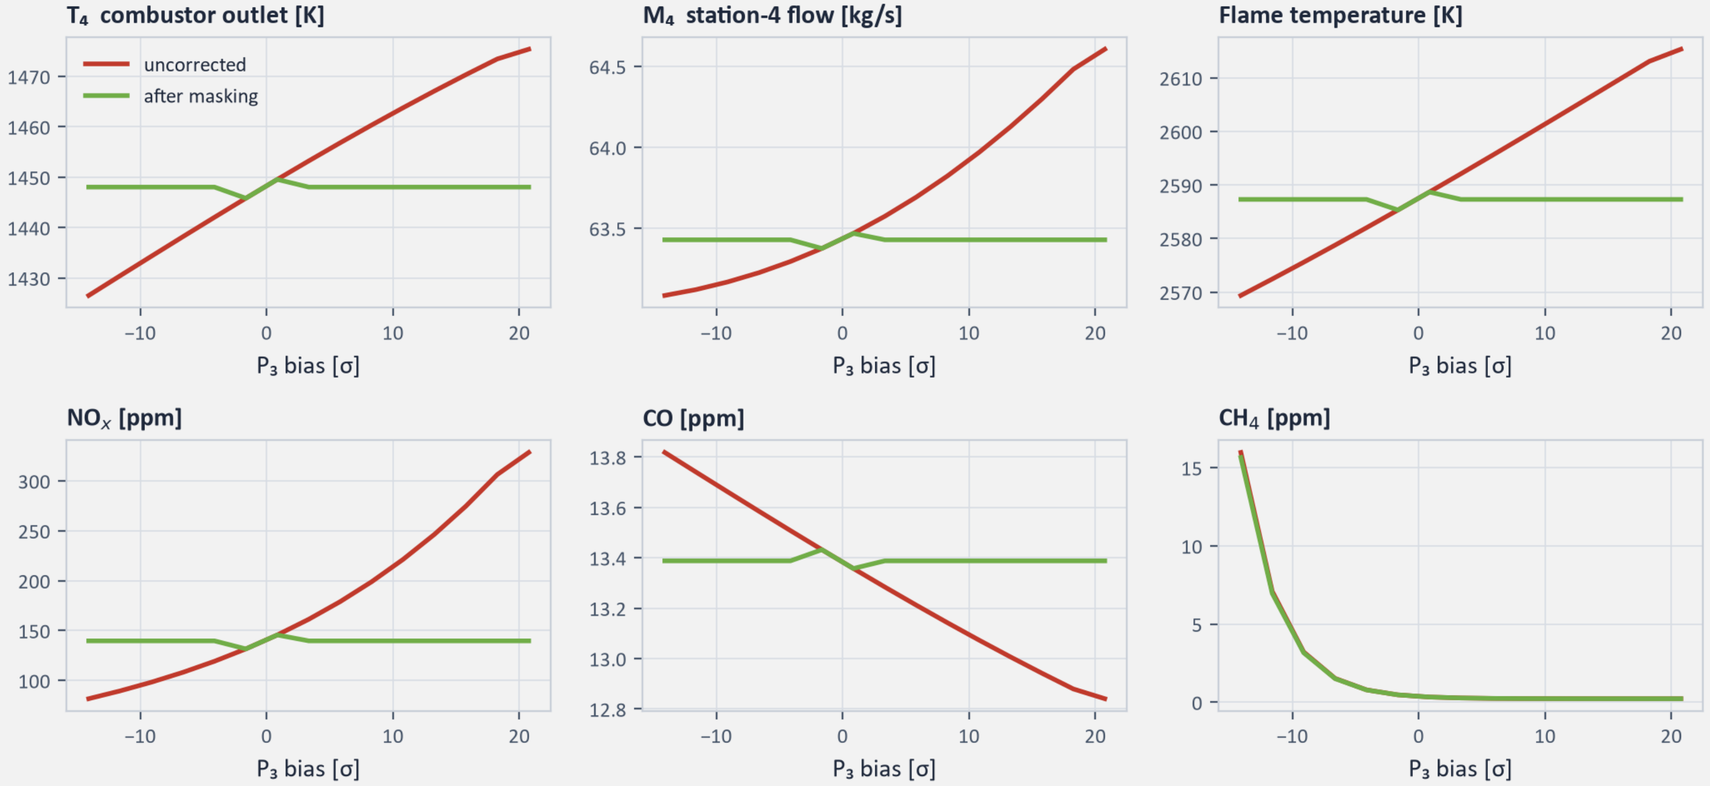

Conclusion

Reconciliation gives a coherent state and validates sensors at the same time

Characteristic equations restore observability

- Stodola and Greitzer relations take the model from under-determined to $\nu = 6$

The χ^2 global test detects gross errors; masking identifies them

- M2 is the most reliable, M4 the best accuracy-per-cost

Detectability is set by redundancy, not by the algorithm

- Pressures and temperatures are caught at 2–3σ; fuel flow needs ≈ 11σ
- More redundancy where it matters would improve detection further

Thank you for your attention

- You can follow us on linkedin