PLANT-WIDE FAULT AND DISTURBANCE SCREENING USING COMBINED TRANSFER ENTROPY AND EIGENVECTOR CENTRALITY ANALYSIS

A PREPRINT

Simon J. Streicher
Department of Chemical Engineering
University of Pretoria
South Africa

Carl Sandrock
Department of Chemical Engineering
University of Pretoria
South Africa
carl.sandrock@up.ac.za

April 9, 2019

ABSTRACT

Finding the source of a disturbance or fault in complex systems such as industrial chemical processing plants can be a difficult task and consume a significant number of engineering hours. In many cases, a systematic elimination procedure is considered to be the only feasible approach but can cause undesired process upsets. Practitioners desire robust alternative approaches.

This paper presents an unsupervised, data-driven method for ranking process elements according to the magnitude and novelty of their influence. Partial bivariate transfer entropy estimation is used to infer a weighted directed graph of process elements. Eigenvector centrality is applied to rank network nodes according to their overall effect. As the ranking of process elements rely on emerging properties that depends on the aggregate of many connections, the results are robust to errors in the estimation of individual edge properties and the inclusion of indirect connections that do not represent the true causal structure of the process.

A monitoring chart of continuously calculated process element importance scores over multiple overlapping time regions can assist with incipient fault detection. Ranking results combined with visual inspection of information transfer networks is also useful for root cause analysis of known faults and disturbances. A software implementation of the proposed method is available.

Keywords advanced process monitoring · transfer entropy · eigenvector centrality · fault detection · data-driven

1 Introduction

Plant-wide fault detection and diagnosis (FDD) and advanced process monitoring (APM) is an active topic in control engineering research, and has been for many decades [1, 2]. However, many proposed methods require either a database of known faults or a benchmark of nominal, fault-free operation. These requirements are a significant barrier to entry for widespread industrial adoption which requires methods to be robust and easy to implement and maintain.

An optimal APM system will assess all process states in near real-time, drawing attention to the most critical behaviour, provide insight into the causal flow of events and offer actionable advisories such that operating personnel can take optimal corrective remedial action. Ranking faults and poor control loop performance based on overall impact and importance is a useful triaging strategy for over-extended control departments [3, 4, 5]. A good ranking tool will rank faults and disturbances such that those affecting the most important variables in the system is assigned higher priority than conditions which may be more severe but affecting less important variables. The flagged areas should be as close to the real source as sensor placement allows, even if significantly amplified by other processes or control loops.

2 Research framework

The time-series signals captured from periodic sampling of measurement devices from a chemical process under automatic closed-loop control can be expected to show the behaviour of a continuous, dynamic multivariate system that is complex, non-linear and non-stationary. Systems with these properties are relatively challenging to analyse compared to discrete, stationary, univariate and linear systems, for which the theory is well developed.

Chemical processes are inherently deterministic. Every single measurement reflected by every probe and each floating point value stored in a process historian is ultimately the result of a deterministic physical system that follows a set of heavily convoluted but individually simple laws. Fundamental processes that transfer heat and mass between elements and chemical reactions that alter the composition of streams do not generate original information. If the entropy of states exciting a node is more than the incoming entropy, the additional information must have originated from a disturbance or fault. This property is not true for some other complex systems frequently studied, such as a stock market. The proposed method for FDD presented here relies on the premise that uncertainty introduced into a system that impacts significant variables indicates the severity and location of a fault or disturbance, and that a combination of information and graph theory can trace this uncertainty back to its original source.

3 Prior art and context of proposed methods

3.1 Use of directed graphs in FDD

A directed graph can be represented in matrix form as an adjacency matrix,

$$A_{ij} = \begin{cases} 1, & \text{if there is a connection from node } j \text{ directed towards node } i \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

as indicated in Figure 1. This work follows the convention of sources along columns and sinks along rows. In a weighted digraph, the Boolean entries of the adjacency matrix are replaced by the appropriate weights.

Figure 1: Example directed graph with associated adjacency matrix

The first known use of a signed directed graph (SDG) for modelling chemical processes dates almost 40 years back [6], and was later extended to include weights based on sensitivity coefficients [7]. Data-driven methods and heuristics captured from operational experience can be used for their development, and a complete quantitative description is not required [8]. SDG models provide a representation that captures cause and effect relationships which can be easily applied in diagnostic reasoning, making them a popular component of several proposed approaches to plant-wide FDD [9, 10].

3.2 Transfer entropy based causal inference

The idea of inferring an information transfer network (ITN) using the transfer entropy (TE) functional for FDD applications has gathered significant traction among several research groups over the last decade. TE relates to other information theoretic measures as follows: The average *entropy* of a signal is a measure of its uncertainty; observations of uncertain states carry more information than samples from predictable signals. *Conditional entropy* gives the remaining entropy after accounting for uncertainty introduced due to the effect of other processes. *Mutual information* measures entropy not unique to a specific signal but shared with other signals evaluated, while *conditional mutual information* is the remaining entropy shared among a set of signals after accounting for the effect of other processes. *Transfer entropy* is the entropy shared between a source and target signal after conditioning on the target signal's past. It can be viewed as a measure of conditional mutual information in time. Finally, *total TE* is the entropy *uniquely* shared

between a source and target signal measured by conditioning on every other signal in the system under investigation in addition to the target's past [11, 12].

Most proposed methods involving TE in the context of FDD try to develop an ITN for the express purpose of reasoning based on its structure [13, 14, 10]. Significant effort has been made to avoid spurious or indirect connections from being included [15] and to reduce the computational load by symbolic [16] and discrete approximations of signals based on alarm limits [17, 18]. In contrast, the method proposed here relies on emerging properties that depend on the aggregate of many connections, and as such is more robust to errors in the estimation of individual causal links. In fact, including indirect connections between source and target nodes will boost the importance of influential nodes, making the detection of changes in the plant's causal structure easier.

Even simplified models of a typical chemical plant are likely to involve several hundreds of states. Rigorous data-driven information decomposition for a general non-linear process becomes computationally intractable past a few states, while analytical solutions for multivariate Gaussian processes exist [19]. As a second reward to full multivariate information decomposition, *complete* bivariate TE analysis that accounts for the effects of all potential sources can provide a sufficient approximation of an ITN to allow for causal reasoning [20, 21]. Estimating complete TE still suffers from the curse of dimensionality, but graph theory based simplifications [20] and greedy-algorithms [22] proved successful in truncating the problem to a manageable number of dimensions.

Simplifying things further, a *partial* bivariate information transfer metric combined with graph centrality measures can indicate the most important sources of original information, without necessarily providing an accurate representation of the system's unique and direct internal information pathways. As the intended goal of the proposed method is to be a useful plant-wide fault triaging indicator and not provide causal reasoning support *per se*, this approach that is two steps (partial and bivariate) away from the near-unattainable theoretical optimum of full information decomposition provides an appropriate balance between rigour and computational complexity.

3.3 Influence based ranking of process elements

Farenzena [3] suggested prioritising control loop maintenance according to a variability matrix that indicates how improving the performance of a specific loop to that of a minimum variance controller is expected to influence the variability in other control loops. In many cases, increased performance in one section of a plant will offset disturbances to interacting units, and solving the ranking problem will give the best compromise while being insensitive to model mismatch.

The idea of using a variant of Google's PageRank-type algorithm for analysing connectivity between control loops was introduced in 2009 and dubbed LoopRank [4]. Partial correlation was used to determine a weighted directed graph. Recently a combination of transfer entropy based causal inference and eigenvector network centrality with additional post-processing steps was successfully applied to detect epilepsy seizure onset zones from neurological data [23]. The nature of this problem is similar to the task of detecting the source of an incipient disturbance in a chemical plant.

4 Information transfer estimation

4.1 Formal definitions

Transfer entropy was introduced almost two decades ago as a dynamic, non-linear and non-parametric measure for measuring information flow [24]:

$$T_{Y \to X} = \sum_{\left(x_{i+h}, \mathbf{x}_i^k, \mathbf{y}_i^l\right)} p\left(x_{i+1}, \mathbf{x}_i^k, \mathbf{y}_i^l\right) \log \frac{p\left(x_{i+1} | \mathbf{x}_i^k, \mathbf{y}_i^l\right)}{p\left(x_{i+1} | \mathbf{x}_i^k\right)} \tag{2}$$

where Y and X are the source and target variables with embedding dimensions of l and k respectively. Using the natural logarithm will return values in the units of nats while using base 2 will return bits.

The need to optimise the measure over a range of potential lags was identified when TE was applied to process engineering signals [13], resulting in the following formulation [25] that includes a lag parameter h:

$$T_{Y \to X} = \sum_{\left(x_{i+h}, \mathbf{x}_{i+h-1}^k, \mathbf{y}_i^l\right)} p\left(x_{i+h}, \mathbf{x}_{i+h-1}^k, \mathbf{y}_i^l\right) \log \frac{p\left(x_{i+h} | \mathbf{x}_{i+h-1}^k, \mathbf{y}_i^l\right)}{p\left(x_{i+h} | \mathbf{x}_{i+h-1}^k\right)}. \tag{3}$$

As calculation will involve a finite number of observations, the estimator can be re-formulating as a global average of local transfer entropies [26], eliminating one probability density function (PDF) estimation and significantly reducing

3

the number of states that need to be summed:

$$T_{Y \to X} = \frac{1}{N} \sum_{i=1}^{N} \log \frac{p\left(x_{i+h}|\mathbf{x}_{i+h-1}^k, \mathbf{y}_i^l\right)}{p\left(x_{i+h}|\mathbf{x}_{i+h-1}^k\right)}. \tag{4}$$

Some applications in prior art used the difference between forward and backwards transfer entropy estimates to maximise the potential benefit obtained from the measure's asymmetrical properties [13]. The term directional as opposed to simple TE is used to refer to this formulation in the scope of this work:

$$t_{Y \to X} = T_{Y \to X} - T_{X \to Y}. \tag{5}$$

4.2 Estimators

Since TE is a functional of multi-dimensional and sometimes difficult to estimate PDFs, the method used for PDF approximation can have a significant influence on its accuracy. Discrete signals with countably limited options allow for easy computation of precise PDFs. However, for continuous variables not modelled according to a specific statistical distribution with an analytical solution for entropy, the bandwidth and/or other parameters of the estimator will influence results. The application of parameter-free zero information calculation to TE provides a robust estimator applicable to non-stationary processes but tends to be overly conservative with a high false negative rate [21].

The simplest approach to estimating TE on continuous signals is to assume that the signal distributions are Gaussian, resulting in a simple closed-form expression. A mean and covariance function completely specifies Gaussian processes and can be used to describe time series data from various sources with reasonable success. However, approaching a non-Gaussian distribution as a Gaussian can result in significant entropy estimation errors [27]. Numerous alternatives to Gaussian kernels are available, and are typically hand-picked according to the specific application and nature of the data [28]. There are no automated routines for estimating the best kernel choice for non-Gaussian distributions, making the use of individually tuned kernels impractical as FDD methods should be applicable to a variety and mix of signal distributions.

The Java Information Dynamics Toolkit (JIDT) contains routines for box (or uniform) kernel estimation together with dynamic correlation exclusion while optimising box counting for computing differential TE [26]. Box kernels make no assumptions about the nature of distributions and are therefore robust to a wide variety of datasets from different sources at the cost of requiring more samples to achieve comparable accuracy as a shaped kernel that does correctly represent the distribution. Box kernels are discontinuous with a square shape and are similar to histograms in univariate cases, producing non-smooth results. The measure centres a box over each data point, and the overlap of boxes determine the density estimate at each point. The box size also referred to as the bandwidth, is the sole parameter of this estimator and affects the level of fidelity.

The k - nearest neighbour (k-NN) based Kraskov-Stügbauer-Grassberger (KSG) estimator [29] is considered the current best-of-breed approach for empirical (conditional) mutual information [30] and is recommended for TE estimation [31]. It estimates mutual information without explicitly estimating a PDF as an intermediate step or assuming a specific distribution. The number of neighbours is the sole parameter and the measure is pleasantly insensitive to it. A default value of $k = 4$ is recommended. This measure is also available as part of JIDT.

The use of both the box-kernel and KSG estimators was evaluated on test cases, and the estimated magnitudes differed significantly, with KSG estimates results exceeding significance thresholds by a larger margin. Additionally, the KSG directional results are very similar to that of simple estimates, indicating that this estimator correctly assigns a near-zero value to the inverse relationship.

4.3 Embeddings

In the most naive implementations, the embedding dimensions l and k are both set to unity. Ideally, to separate information storage from transfer, the embedding dimensions of the target signal k needs to be optimised such that the self-information between x_{i+1} and $\mathbf{x}_i^k$ is maximised [31]. Failure to do so will lead to an over-estimation of TE. The target embedding dimension k can be optimised to account for cases where multiple past values are causal to the target. An embedding delay τ can be used to better empirically capture the state from a finite sample size [31]:

$$\mathbf{x}_i^k = \left\{ x_{i-(k-1)\tau}, \ldots, x_{i-\tau}, x_i \right\}. \tag{6}$$

The JIDT package provides an option to perform an automated search for the optimal embedding dimensions and delays for the source and target variables when the KSG estimator is used. The use of auto-embedding was evaluated on test cases and significantly reduces the absolute magnitude of results obtained. Sub-harmonic peaks were also removed

4

and the optimising delay was preceded by a plateau instead of forming a well-defined peak. If any auto-embedding parameter resolves at an edge of the allowed search space it should be widened.

4.4 Parameter selection

4.4.1 Sample size

In this context sample size refers to the number of data points involved in a single bivariate partial TE estimation. The main focus in selecting this parameter should be on identifying the smallest number that will allow for accurate and stable estimation of the conditional PDFs involved. Computation time scales exponentially with sample size depending on the number of embedding dimensions involved.

Small sample size leads to increased variability and varying absolute values in the case of box-kernel and KSG estimators respectively. The complexity and stability of process signals will dictate the minimum number of samples needed in order to estimate a representative PDF. Highly non-linear or directional (large condition number as obtained by Singular Value Decomposition (SVD) analysis [32]) processes will result in a more complex relationship between the signals and therefore a more complicated conditional PDF in need of approximation. In exceptionally difficult cases the practitioner might need to evaluate estimator convergence with regards to varying sample size to make an appropriate selection of this parameter. In test cases estimated values converge for sample sizes above 2 000. This recommendation can also be found in prior art [13]. It might be worthwhile to investigate the effect of sample size on the specific process studied and compromise accordingly.

The desire to analyse an extensive period to capture an effect that occurred at an uncertain time should never motivate sample size selection. Instead, performing multiple time region (MTR) analysis is recommended as this will be significantly faster, scaling linearly with total samples covered, and will also more accurately pinpoint the exact time at which a disturbance or fault entered the system [33].

4.4.2 Sampling period

The product of the sample size and sampling period define the time span covered by a single estimation. It is highly unlikely that time series data captured at a resolution faster than one second is available in a plant historian and the suggested sample size of 2 000 points translates into a minimum time span of just over half an hour covered. The sampling period is the primary parameter that should be varied to adjust the time span as sample size will always be constrained to a few thousand data points due to its impact on computational complexity. However, the selected period should always satisfy the Nyquist sampling theorem for the frequency region of interest [32]. However, as the time constants of some processes, such as distillation columns, can be several hours, it might be useful to perform analyses on data sets sampled at lower frequencies.

A fast Fourier transform (FFT) analysis can be used to assist with identifying time constants of interest. If more than one frequency region is active, it is likely that this is due to the presence of two distinct disturbances. Band-gap filtering the time series data and analysing each frequency region in isolation is recommended. Disturbances with faster dynamics are usually more troublesome than slower ones. However, as slow disturbances that go unchecked can result in oscillations of growing magnitude leading to the plant becoming unmanageable after a few days, it might be worthwhile to schedule analyses targeting slow dynamics periodically.

4.4.3 Delay optimisation range

Ideally, a large range of delays should be analysed at high resolution, but due to processing time constraints, compromises need to be made. The dead-time of processes can vary significantly. The throughput of a typical petrochemical plant is rather high compared to the hold-up volumes, and therefore the maximum time delays are expected to be relatively small. However, this is not necessarily the case for mineral and pharmaceutical plants. The initial delay optimisation search space covers should cover at least five minutes. If many directed variable pair weights maximise at the search space borders, the analyst should increase the delay range.

The optimising time delay is a property of the processes involved and not likely to change significantly under the influence of different faults and disturbances. Over time, the ideal delay range that optimises each particular directed variable pair for a specific system will become known. The computational load can be reduced by detecting an appropriate time delay range for each directed variable pair and narrowing the search space accordingly. Using mutual information instead of TE to determine the optimal delay range is another successful strategy for reducing computational load [18].

4.5 Significance testing

TE estimation methods are likely to return non-zero values for any set of random signals. As the proposed scheme for fault and disturbance screening is fully data-driven, unsupervised and model-free, no first principle or model-based knowledge can be used to force any meaningful structure to the results obtained. Significance testing is necessary to confirm whether an estimated TE value is due to real interaction.

4.5.1 Magnitude

A significance threshold on the magnitude of a bivariate TE estimate is typically derived by studying the distribution of TE estimates obtained on a surrogate set of data generated such that the temporal ordering between the original data vectors is broken while maintaining as much of the other properties of the signal as possible. Some researchers simply random shuffled the source data vector in time [34, 35] while others advocated the use of the iterative amplitude adjusted Fourier transform (iAAFT) method [36] that strives to retain amplitude both the amplitude distribution and power spectrum [13].

Numerous researchers applying TE to process data derive a significance threshold based on a certain number of standard deviations away from the mean of values calculated on the surrogate set [13, 18]. However, this practice is discouraged for non-linear metrics and non-parametric rank-order approaches are recommended instead [36]. According to a rank-order test, a residual probability of a false rejection α corresponds to a significance level of $(1 - \alpha) \times 100\%$. A test for TE significance in a specific direction of time is one-sided for which $M = 1/\alpha - 1$ surrogate sequences are generated, giving a total of $1/\alpha$ data sets including the data set itself. The probability that the original data results in the largest causality measure value by coincidence is, therefore, α. For a significance level of 95%, the test thus generates and calculates the TE of 19 surrogate data sets and requires that the real data returns results higher than all the surrogates to pass.

4.5.2 Directionality

Although TE is an asymmetric measure and therefore implicitly tests significance in a specific direction, automated optimisation over a range of delays complicates matters as an affected variable might appear to influence their causal variables if the affected variable value is lagged far enough behind the causal variables.

An optional additional directionality significance test is applied to help reduce spurious connections detected due to this effect. The test involves calculating the TE over a range of delays shifting the source target forward and backward, recording the maximum values and their associated delays in each direction: $\left[\psi^{\text{forward}}; \delta^{\text{forward}} \right]$ and $\left[\psi^{\text{backward}}; \delta^{\text{backward}} \right]$ respectively. The directionality test passes on condition that $\psi^{\text{backward}} > \psi^{\text{forward}}$ or $\delta^{\text{backward}} < \delta^{\text{forward}}$. Note that application of this test will suppress the detection of true bi-directional interactions.

4.6 Scaling

Ideally the value of entropy measured on any signal should have a consistent practical interpretation across variables regardless of the engineering units used to record the measurements. However, the estimated entropy of a continuous signal using kernel methods is dependent on the relationship between its variance and the kernel bandwidth used to estimate the underlying PDF. To ensure that the magnitude of estimated TE between variable pairs has a consistent meaning throughout an analysis, the scaling of data before estimation of TE is critical when using kernel estimators.

Even though standardisation (subtraction of the mean and division by the standard deviation) is a commonly applied scaling strategy [13, 37] its use removes all information about how much the signal varies relative to its desired operating range and thereby introduces the requirement that sufficient and representative excitation of the process in all directions should be present in the sample data. If this is not the case, division by the small standard deviation will blow a small amount of noise on an otherwise quiet signal out of proportion, misrepresenting the relative real information content carried by the signal. For example, we do not wish to flag a disturbance that causes a level (an unimportant inventory variable) to fluctuate between 48 and 52% while the desired control limits are 30-80% as critically important. However, a fault that causes causing an analyser composition reading (an economic variable) to fluctuate between 0.85 and 0.86 while the minimum requirement is 0.84 should be regarded as highly significant.

Strong cases for the necessity of getting a scaling independent representation of data to make reliable conclusions about relative interaction strengths among process elements using weighted graphs have been made elsewhere [37]. A suitable scaling strategy involves subtracting the nominal operating value and division with the maximum expected or allowed change [32, 37]. If the nominal operational point does not lie in the middle of the allowed or desired range, the most conservative approach is to use the biggest difference to scale disturbances and reference (set-point) values, while the

smallest difference is used to scale allowed input changes and control errors. The result is that all values have less than unity magnitude if they remain within the predetermined limits informed by process knowledge, and similar magnitudes indicate a deviation of equal severity.

Both standardisation and process limit base scaling was evaluated on test cases. The KSG estimator has an adaptive bandwidth and is not sensitive to the scaling approach used while a significant improvement in ranking results is observed when using the kernel estimator which depends on a fixed bandwidth parameter.

5 Ranking nodes using graph centrality

5.1 Eigenvector centrality

Eigenvector centrality measures a node's influence in a weighted directed graph as the sum of importance scores nodes with an edge incident to it multiplied by the edge weights,

$$C_{\text{eig}}(k) = \sum_{j \in \mathbf{L}_k} w_{kj} x_j, \tag{7}$$

where $C_{\text{eig}}(k)$ is the importance of node k, $\mathbf{L}_k$ is the set of nodes with connections to x_k and w_{kj} are entries of the edge weight matrix $\mathbf{W}$. In the context of this application the edge weight matrix $\mathbf{W}$ must be column-stochastic (all columns should sum to one) with real positive entries representing a measure of connection strength between nodes. The problem can be rewritten in the format of a standard (right) eigenvalue problem

$$\mathbf{W}\mathbf{x} = \lambda \mathbf{x}. \tag{8}$$

Although there might be many eigenvalues λ with different eigenvectors $\mathbf{x}$ which satisfy this equation, the eigenvector with all positive entries and associated with an eigenvalue of unity, $\lambda = 1$, contains the relative importance scores. This eigenvector corresponds to the stationary probability vector defined on a stochastic matrix and the Perron-Frobenius theorem guarantees its existence and uniqueness.

5.2 Markov chain interpretation

The eigenvalue centrality measure can be interpreted as the results of a Markov chain operator on the graph that assigns node importance relative to the distribution of time a random walk starting on any node will spend on the various nodes. The proof relates to the idea that the long-term probability of being in a state is independent of the initial state. A Markov chain problem formulation usually consists of a transistion probability matrix (TPM) $\mathbf{G}$ (also called a substitution matrix) together with a reset probability matrix (RPM) $\mathbf{Q}$ weighed by a mixing factor m

$$\mathbf{M} = m\mathbf{G} + (1 - m)\mathbf{Q}. \tag{9}$$

Entries in the TPM indicate the probability that an actor currently residing on a specific node will move to a neighbouring node – in the appropriate direction for the case of directed graphs. The RPM represents the probability that an actor will move randomly to any node on the network. In basic algorithms, this probability distribution is usually considered to be equal among all nodes. Since the sum of all probabilities for a given event should be unity, both the TPM and RPM have stochastic column vectors containing non-negative real numbers.

For the Markov chain operator interpretation to hold, it is essential that there be no dangling nodes, that is, columns in $\mathbf{M}$ which have no non-zero entries (whose sum is, therefore, zero). In the context of this work dangling nodes are those that do not influence any other, a situation that is very likely to occur. A straightforward means of dealing with dangling nodes is to add very weak connectivity between all nodes in the network. Weighing a TPM $\mathbf{G}$ with an equal probability RPM $\mathbf{Q}$ with a m very close to unity implements this solution.

5.3 Direction of analysis

There are two different directions in which the eigenvector centrality can be used to assess importance. Some applications, such as the ranking of web pages, attribute importance based on the quantity and significance of *inbound* connections to a specific node, or references back to a particular page. In the context of determining which node has the most considerable extent of causal influence on other nodes, it is desired to attribute importance based on the quantity and significance of *outgoing* connections from a specific node.

The convention used to define entries in the adjacency or weight matrices according to row or columns and whether the right or left eigenvector problem is solved determines if the ranking results refer to the inbound or outbound metric. In this work, eigenvector centrality always refers to the results obtained by calculating the right eigenvector with the adjacency matrix arranged source (causal) nodes along the columns and destination (affected) nodes along the rows. Under this arrangement it is necessary to use the transposed weighted digraph adjacency matrix $\mathbf{M}$ as the ranking weight matrix $\mathbf{W}$ in 8 to get the desired result of ranking based on outgoing connections

$$\mathbf{W} = \mathbf{M}^\mathsf{T}. \tag{10}$$

5.4 Introducing bias

Identification of the most critical faults might be enhanced by incorporating individual controller performance key performance indicator (KPI)s and stream importance measures, such as economic value, as meta-data in the node ranking problem. Adjusting the relative reset probabilities is an ideal means of biasing the network ranking towards specific nodes.

A relative reset bias vector $\mathbf{q}$ needs to be provided and normalised such that it sums to unity. The relative, not absolute, values of the entries is of importance. The reset probability matrix $\mathbf{Q}$ is then defined as the $n \times n$ matrix with each row equal to the normalised reset bias vector $\mathbf{q}$:

$$\mathbf{Q} = [\ \mathbf{q}, \mathbf{q}, \cdots\ ]^\mathsf{T}. \tag{11}$$

The transition and reset probability mixing weight m can be considered a tuning factor that will determine how significant the edge weightings are relative to the node biases in determining the final ranking score.

5.5 Weight normalisation

Edge weights should be normalised to maintain a consistent interpretation of the ranking results independent of the units used to record connection strength and to provide a scale with regards to the (biased) reset probability weights. Additionally, for $\mathbf{W}$ to be column stochastic the rows of $\mathbf{M}$ must sum to unity. To preserve the meaning of reset probability, $\mathbf{G}$ needs to be row-normalised before linearly combining with $\mathbf{Q}$ to form $\mathbf{M}$. Row normalisation involves dividing each element of a matrix by the sum of the elements in the same row

$$X_{ij_{\text{norm}}} = \left(\frac{X_{ij}}{\sum\limits_{i} X_{ij}} \right). \tag{12}$$

If dangling nodes are present in the network, some rows of $\mathbf{G}$ will sum to zero instead of unity. These rows will not result in a row-normalised weight matrix $\mathbf{M}$ after linear combination with $\mathbf{Q}$ and a final row-normalisation step is needed to ensure that $\mathbf{W}$ is column-stochastic. unities for future research, see Section **??**.

5.6 Katz centrality

Katz centrality measures the total number of walks (sequence of edges and nodes connecting two nodes) between two nodes in the network. In effect, it is similar to eigenvector centrality with the additional accounting for higher order connections weighed by an attenuation factor and can assist with boosting the difference in ranking scores between nodes.

The formal definition of Katz centrality is

$$C_{\text{Katz}}(k) = \sum_{p=1}^{\infty} \sum_{j=1}^{N} \alpha^p (W^p)_{kj}, \tag{13}$$

where $C_{\text{Katz}}(k)$ is the importance of node k according to the Katz centrality measure, $\mathbf{W}$ is the weight matrix and α is as an attenuation factor that cannot exceed a value of $1/\lambda$ where λ is the largest eigenvalue of $\mathbf{W}$. Unlike eigenvector centrality, it is possible to compute Katz centrality on directed acyclic graphs (DAGs), which removes the need for adjusting the adjacency matrix with a fully connected reset matrix.

The principal eigenvector associated with the largest eigenvalue of $\mathbf{W}$ is returned in the limit that α approaches $1/\lambda$ from below. This is identical to the result calculated according to eigenvector centrality when an acyclic graph is combined with a uniform reset bias vector with m approaching zero from the right. The inclusion of weak indirect connections in ITNs by partial bivariate TE estimators is comparable to calculating a ranking based on Katz centrality.

6 Recommended procedures

Efficient procedures for applying the presented method will depend on the fault or disturbance's nature and how much the analyst already knows about it. Applications can range from an original assessment of the key influences in a plant to confirmation of already existing hypotheses concerning propagation pathways. The intended use of the presented method is as a triaging aid for use by process experts and is not refined enough for drawing blind conclusions, but should be integrated into a larger FDD workflow.

It is difficult to provide performance guarantees under a range of different conditions for unsupervised methods. Instead, sanity checks are required to eliminate spurious results, and visual graph interpretation is occasionally necessary to locate the actual source of information flow.

6.1 Variable selection

Not all data sources associated with a specific section of a plant will contribute useful information to the analysis, and it is recommended that they are filtered to reduce the scope to no more than 50 elements. If it is desired to analyse a specific disturbance event or fault, care should be taken to ensure that it is included in the period covered by the actual calculations. For complex processes a high-level analysis of the most important economic variables can be used to inform more detailed, focused analysis of flagged process sections.

Unchanged set-points and other constant data vectors will not play a role in any causal relationship, and the analyst can safely eliminate these from the data set without any loss in fidelity. The analyst may also eliminate redundant measurements upon confirming that no sensor faults took place during the period under investigation. Measurements at the very end of the process stream not connected via any controllers or other feedback mechanisms only have to be included if they are affected.

6.2 Multiple time region analysis

A single snapshot of network ranking scores might be difficult to interpret without comparison to nominal fault-free operation. Monitoring the relative changes in node rankings over time might provide more insight. A plot of node importance calculated for multiple highly overlapping rolling windows of time provides a single comprehensive overview of changes in the process' causal structure and is considered the primary output of the proposed method for the purpose of incipient fault detection.

6.3 Visual inspection

Visual inspection of the ITN while continually referring to a process flow diagram (PFD) or basic description of the process, such as is commonly found in the operator training manual, is arguably the most useful means of interpreting the results and forming hypotheses about the source of faults and disturbances.

Colour grading nodes based on their centrality scores make it easy to visually process their relative locations and the areas they influence. However, if the indicated nodes are unexpected or no clear link between them and known faults is apparent this information is of limited use in isolation. If the most important edges are highlighted it becomes significantly easier to understand why some nodes have high centrality scores. The associated software saves the ITN in GML format. Cytoscape [38] is an open source and user-friendly tool for viewing and analysing large networks and supports applying the analysis steps mentioned above.

An appropriate layout enables visual interpretation of a network's structure. In the context of tracing an effect to its source it is desirable to arrange nodes from most independent to dependent. A hierarchical network layout serves this purpose. Hierarchical layouts separate the nodes into clusters influenced by different original information source nodes, making the method applicable to situations that involve multiple distinct faults. The yFiles hierarchic network layout algorithm provides the most useful results amongst the options available in Cytoscape.

A hierarchical layout cannot be computed on a fully connected graph as produced if the recommended eigenvector centrality method is applied. Deleting all edge weights below a certain percentile or threshold is usually necessary for visualisation purposes.

7 Results

7.1 Simple mixing process

The simple mixing process presented in Figure 2 exhibits many of the aspects we wish to study, including non-linearity, interacting control and integrating processes. To assess the efficiency of different methods, random noise is introduced in F_B from $t = 0$ until $t = 20$, and the composition set-point is perturbed thereafter until the end of a 40-hour simulation. Figure 3 provides high density time series plots of the standardised and process limit scaled simulation results respectively. Note the attenuation of outlet flow rate and height signals together with amplification of the composition signal in the latter.
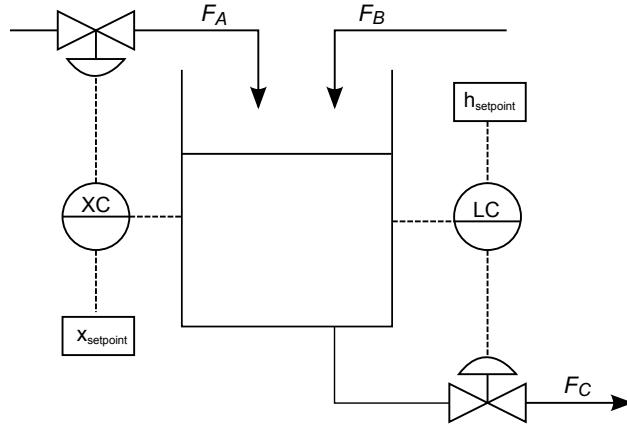
Figure 2: Level and composition problem process flow diagram

(a) standardised   (b) process limit scaled
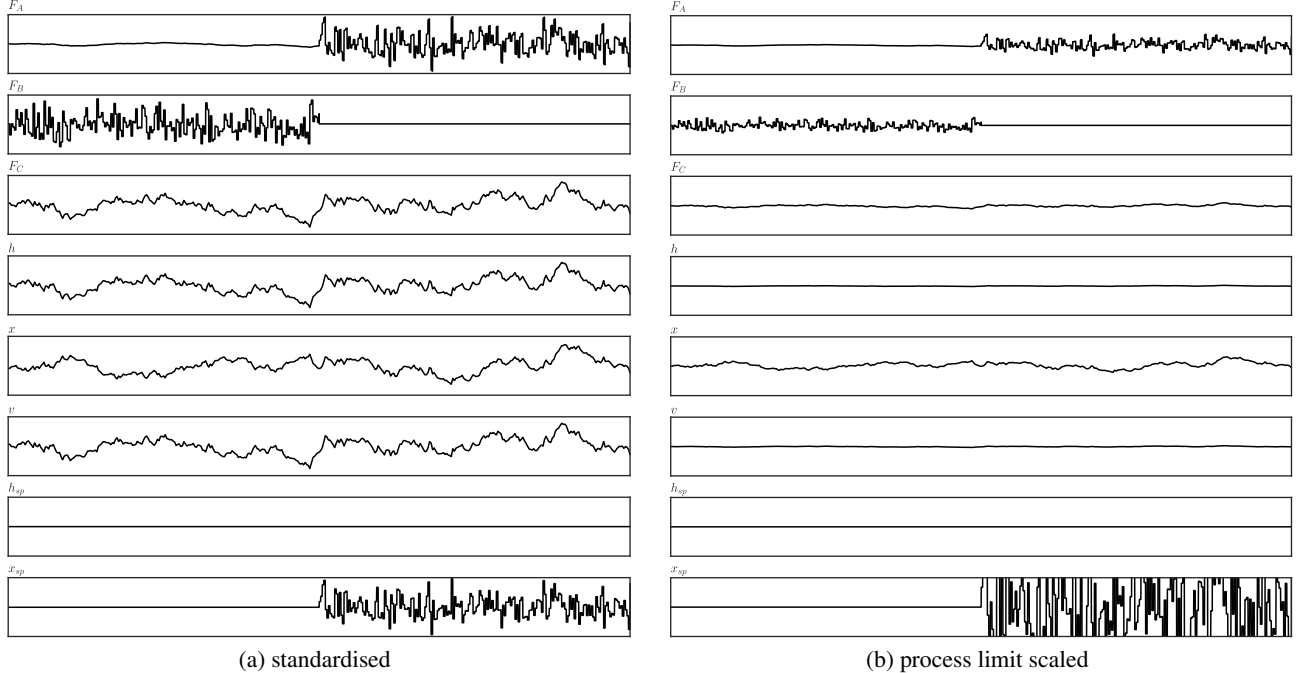
Figure 3: High density time series plot of disturbances

Figures 4 plots node importance scores over multiple time regions obtained without the use of significance testing. Note how process limit scaling enables the kernel method to rank x_{sp} higher than F_A, but made rankings less distinct when the KSG method is used without embedding. The use of directional TE and significance testing did not result in better MTR plots.

(a) kernel estimator; standardised

(b) kernel estimator; process limit scaled

(c) KSG estimator; standardised

(d) KSG estimator; process limit scaled

(e) KSG estimator; standardised; embedded

(f) KSG estimator; process limit scaled; embedded
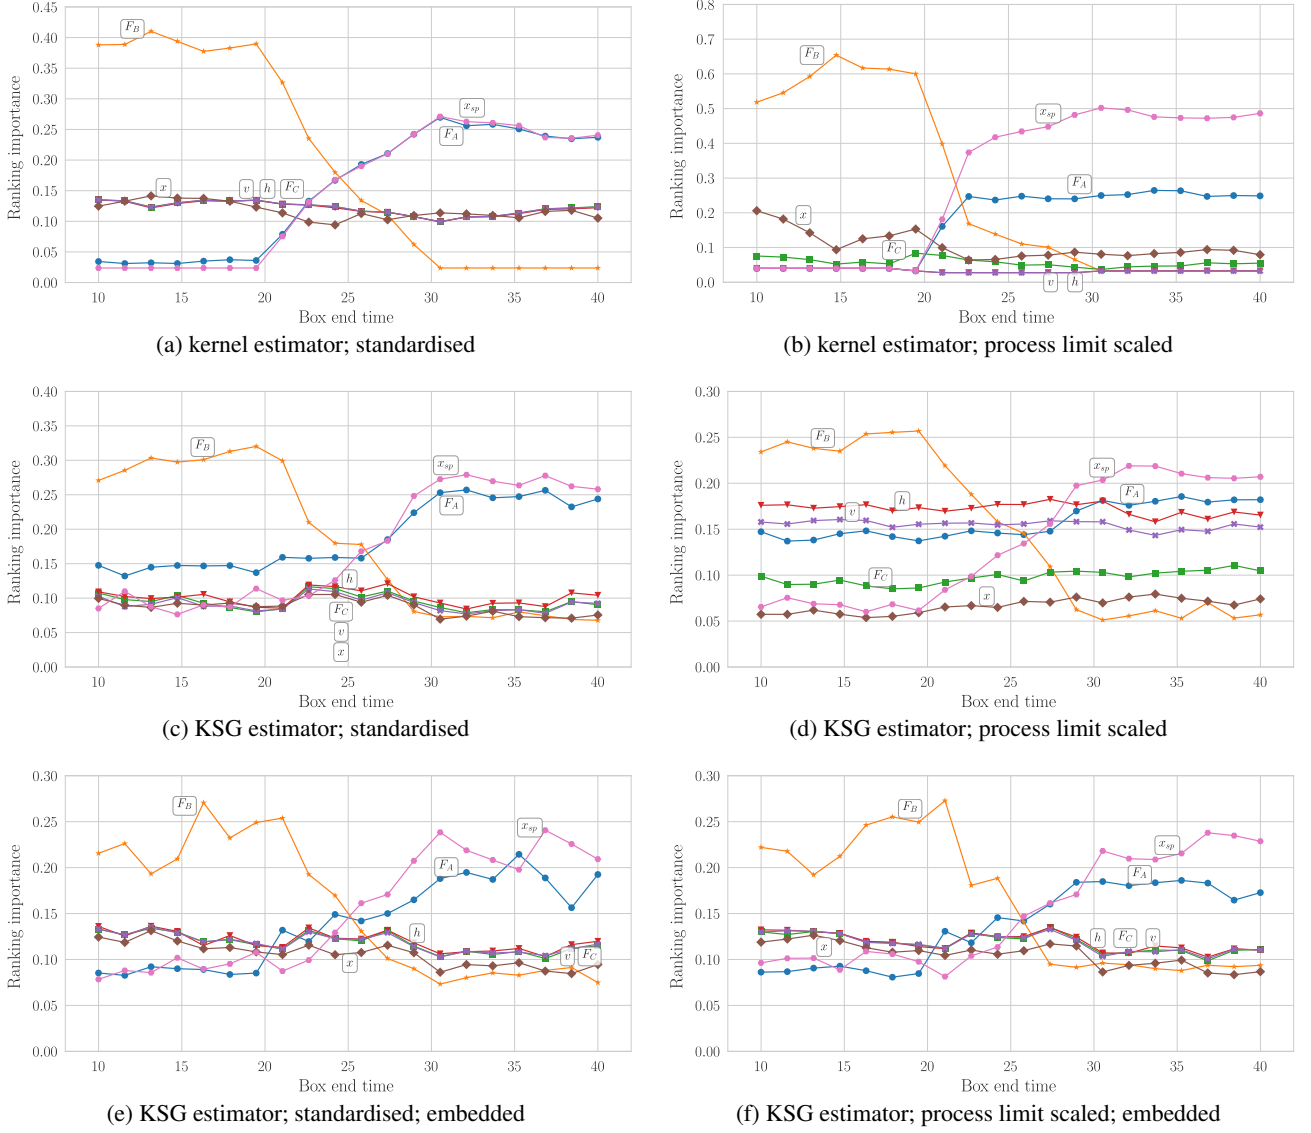
Figure 4: Multiple time region node rankings for mixing process [simple estimates; no significance testing]

Figure 8 shows ITNs for the last time window investigated where the composition set-point disturbance has been active in isolation for a long time. Significance testing was applied and the graphs have been further simplified by deleting the bottom half of edge weights and resolving higher-order connections up to the fifth degree of separation. Note that the two constant signals of F_B and h passed a significance test in a number of cases involving the KSG estimator. This is due to the KSG estimator adding a small amount of noise to signals in order to ensure differentiation between points, a requirement for its convergence. We therefore conclude that it is important to remove stationary signals before analysis with the KSG estimator. Another obvious mistake is that although x_{sp} is an independent source of information, some methods indicate incoming connections. This is due to the delay optimisation lagging the real affected variable behind the source, and applying the bi-directional significance tests described in Section 4.5.2 might eliminate these errors. In general, the ITNs produced by the KSG estimator are a better resemblance of reality, and process limit scaling allows for better detection of propagation paths resulting in a more hierarchical graph (compare Figure 8 (g) and (h), for example).

The significant contrast between the accuracy of individual ITNs and the MTR ranking results clearly demonstrate the significant benefit of performing node centrality analysis for the purpose of fault detection. Ranking does not only condense information but also filters out the effect of individual erroneous connection by aggregating over all the results.

11

A similar result is seen in boosting models for regression and classification, where an aggregate of mediocre models manages to provide results that are better than any individual model.

## 7.2 Tennessee Eastman Challenge Problem

The Tennessee Eastman challenge problem [39] is a popular FDD method benchmark [18] simulation that comes with 20 pre-defined disturbances that can be activated. Some base layer control strategies have been suggested in literature; the scheme presented by Ricker [40] is used as the basis for all tests presented here.

To test the MTR approach and ability to handle multiple simultaneous effects two of the random variation disturbances were activated for specific periods of time with an overlapping window. A random variation in the reactor cooling water (CW) inlet temperature was introduced at a simulation time of $T = 37$ hours and deactivated at $T = 145$, and random variations in the composition of feed stream E was introduced at $T = 109$ and deactivated at $T = 217$. 252 hours of operation were simulated at an interval of $5\text{E-}4$ hours, capturing a total of 504 000 data points for each variable.

Figure 6 presents standardised high density time series plots of all relevant signals generated in the simulation. Note the overlapping disturbances by comparing XMV 10 with XMV 8, for example.
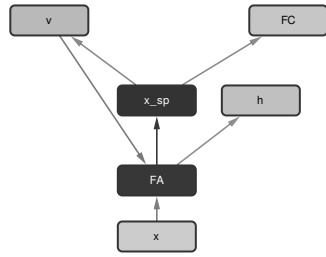
(a) simple; kernel estimator; standardised
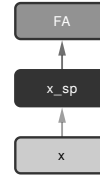
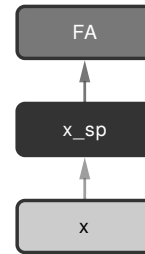(b) simple; kernel estimator; process limit scaled
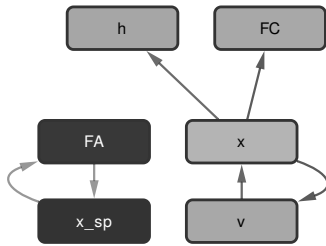
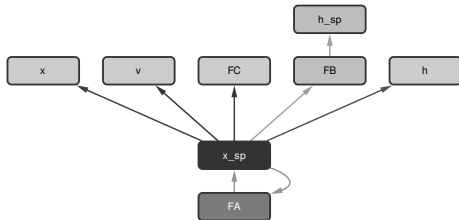(c) simple; KSG estimator; process limit scaled
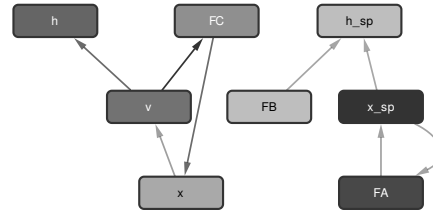
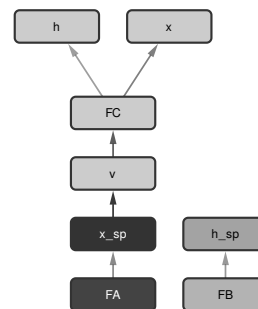(d) simple; KSG estimator; process limit scaled
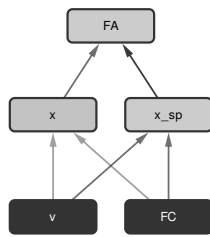
(e) directional; kernel estimator; standardised

(f) directional; kernel estimator; process limit scaled

(g) directional; KSG estimator; standardised

(h) directional; KSG estimator; process limit scaled

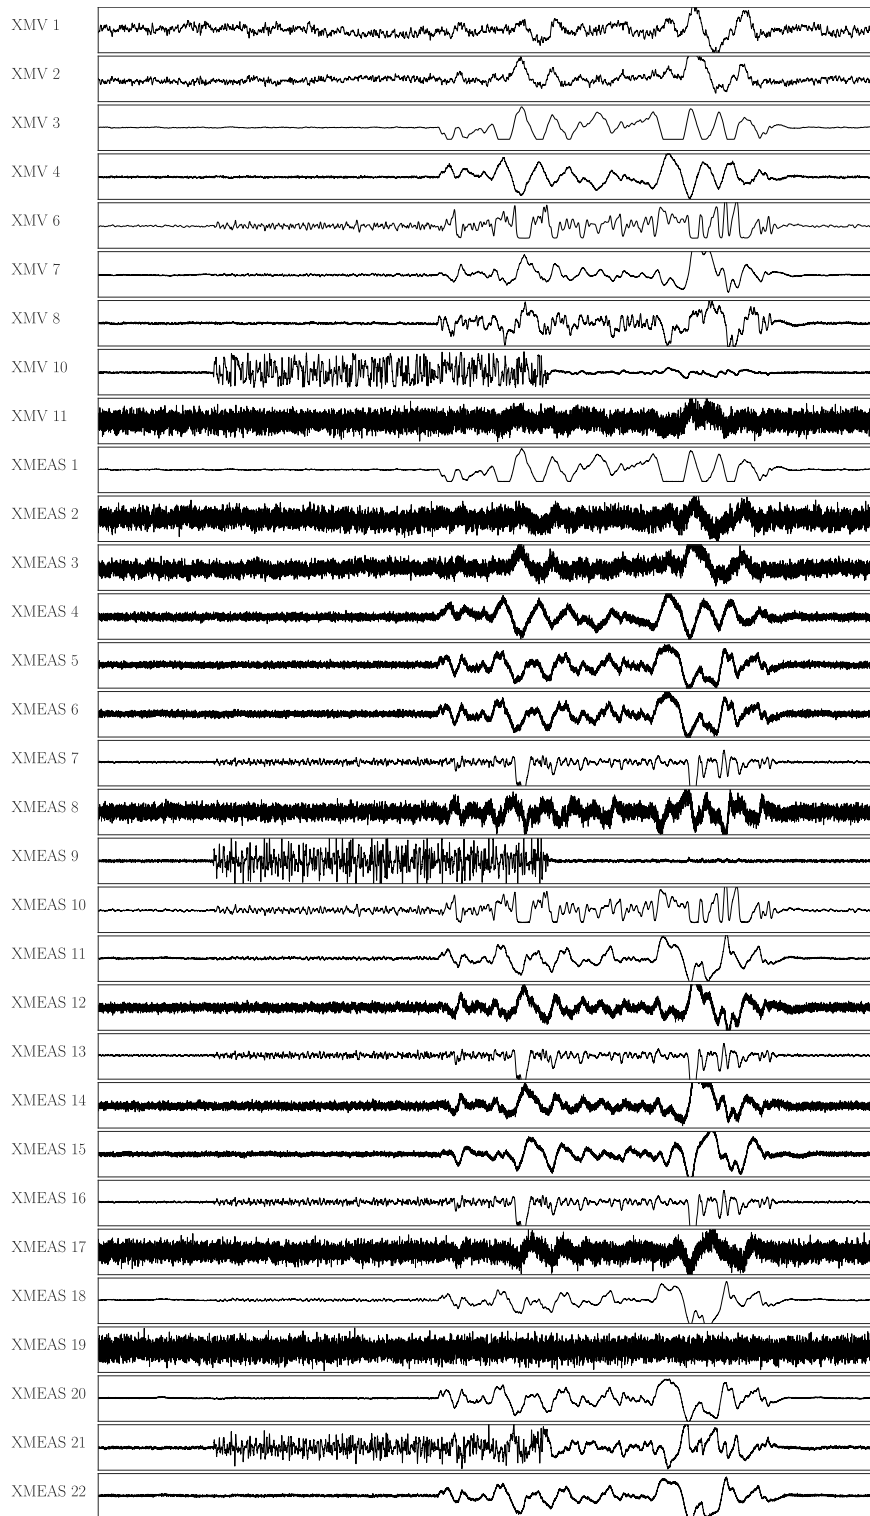Figure 5: Reduced edge ITNs for x_{sp} disturbance [significance tested]

Figure 6: High density time series plot of disturbances [standardised]

As discussed in Section 4.4, a balance between sampling rate and sample size must be achieved with respect to the process time constants of interest. The original data was subsampled by a factor of 30 giving a final sampling interval of 0.015 hours or 54 seconds. At this resolution, if the suggested number of samples of 2 000 is used, each analysis period covers 30 hours which should be enough to capture the slow dynamics of this process.

In order to achieve a 75% overlap between consecutive time windows, 28 periods covering 36 hours each was evenly over the 252 hours of simulated operation. The 36 hours covered resolves to 2 400 samples per time window. It was decided to search for the optimal delay over 100 samples or 1.5 hours in each direction, resulting in a final sample size of 2 200.

The reactor cooling water inlet temperature disturbance was active in isolation during bins 2-10, while the feed composition disturbance was active in isolation for 19-27, with both disturbances present in bins 11-18. To demonstrate the nature of results obtained, the two disturbances are investigated in isolation as and in an overlapping time window. It was selected to use the bins where each disturbance has been active for the longest in isolation over the full duration of the bin, that is, bins 10 and 23 for the reactor cooling water and feed composition disturbances respectively.

To investigate the effects of a reactor cooling water inlet temperature disturbance in isolation, an analysis of the results obtained for bin 10 is presented. At the start of this bin, the disturbance was already active for 35 hours and can be expected to have propagated through process units with long time constants. A list relative tag rankings are presented in Table 7.2 and the tag closest to the actual source of the disturbance is given the highest importance score by a significant margin, followed by the reactor temperature cooling water flow, which is connected to the reactor temperature via a control loop as can be seen in the schematic Figure 7.

By looking at the network arranged in a hierarchical layout, a very clear picture of the situation is obtained. The reactor cooling water outlet temperature is not receiving any incoming connections indicating that it is an independent source of information flow in the system. A clear propagation path from the cooling water outlet temperature to reactor temperature and the cooling water flow manipulated variable (MV) is indicated. Note all the knock-on effects of the cooling water disturbance indicated by the KSG estimator, including a number of feedback loops.

Table 1: Transfer entropy ranking results for CW inlet temperature disturbance [kernel estimator; directional weights; forward delays; significance tested]

Rank	Tag name	Tag description	Tag type	Relative score
1	XMEAS 21	**Reactor CW outlet temperature**	**PV**	
2	XMEAS 9	Reactor temperature	CV	
3	XMV 10	Reactor CW flow	MV	
4	XMEAS 19	Stripper steam flow	PV	
5	XMEAS 17	Stripper underflow	CV	
6	XMEAS 16	Stripper pressure	PV	
7	XMEAS 7	Reactor pressure	CV	
8	XMV 6	Purge valve	MV	
9	XMEAS 13	Product separator pressure	PV	
10	XMV 2	E feed flow	MV	

To investigate the effects of a feed composition disturbance in isolation, a list of the top tags for bin 23 is presented in Table 7.2. The actual source of the disturbance is not immediately apparent. It is clear that the control loops that get their set points from the composition controller are very active. However, note that the stripper temperature and compressor work are some of the measurements closest to the disturbance, while most of the other high-ranking nodes point to a disturbance in the composition.

Node centrality scores might be high for units that serve as distribution points for disturbances, even if they are not powerful sources of disruption themselves. Stripper steam flow (XMEAS 19) was weakly influenced and influences almost all other nodes which result in an unfortunately high place in the ranking scores.

15

Table 2: Transfer entropy ranking results for feed composition disturbance [kernel estimator; directional weights; forward delays; significance tested]

Rank	Tag name	Tag description	Tag type	Relative score
1	**XMEAS 22**	**Separator CW outlet temperature**	**PV**	
2	XMEAS 19	Stripper steam flow	PV	
3	XMEAS 20	Compressor work	PV	
4	XMEAS 13	Product separator pressure	PV	
5	XMV 1	D feed flow	MV	
6	XMEAS 21	Reactor CW outlet temperature	PV	
7	XMV 7	Separator pot liquid flow	MV	
8	XMEAS 16	Stripper pressure	PV	
9	XMV 3	A feed flow	MV	
10	XMEAS 11	Product separator temperature	CV	

Figure 9 shows the MTR monitoring plot generated over a range spanning both the cooling water and stripper feed composition disturbances, with the most important nodes detected in bins 10 and 23 above indicated.

Note that bin 20 seems to relate more significant disturbance in stripper related variables compared to bin 23 which was analysed in detail. This is in line with the time series plots in Figure 6, and serves to illustrate the potential importance of continuous ranking calculations in facilitating a proper understanding of how a disturbance evolves through a process over time, ultimately allowing for more accurate diagnosis.
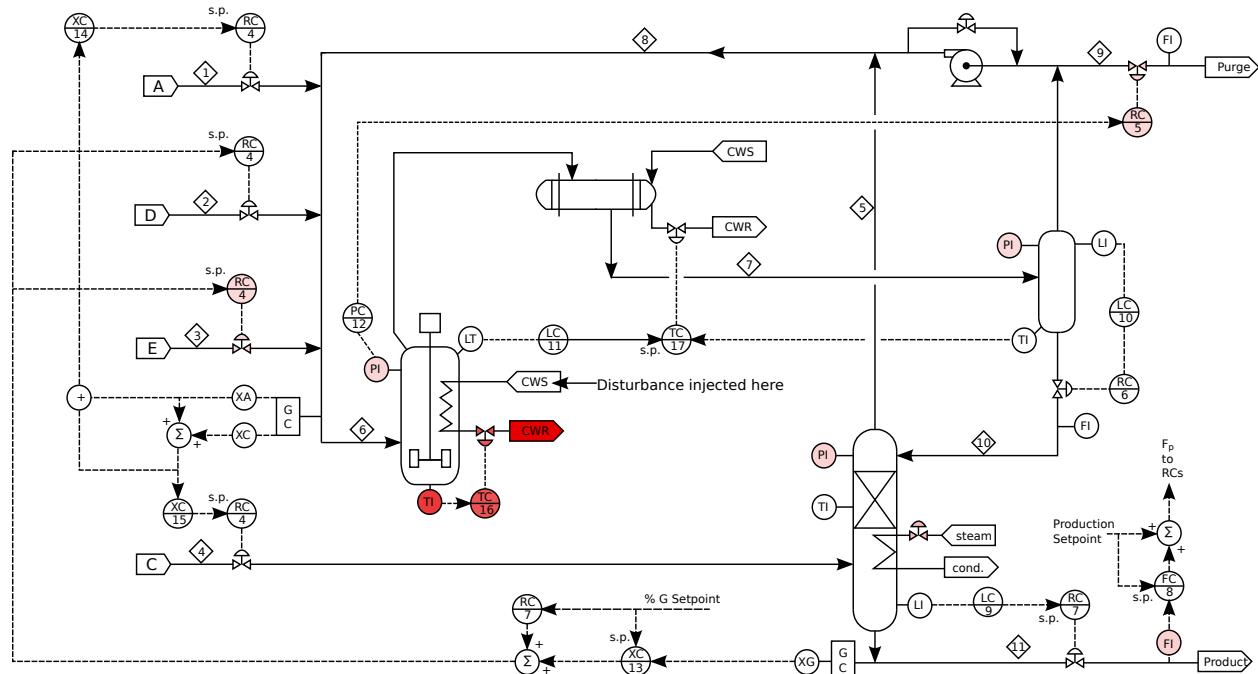

Figure 7: Tennessee Eastman schematic with cooling water inlet temperature disturbance effects highlighted

8 Discussion

The presented techniques show promise in assisting a process expert to identify the source of a disturbance or fault significantly faster than by following more classical workflows, usually limited to studying time series plots or individual loop KPIs. The methods have proven helpful in two non-trivial industrial fault diagnosis problems encountered in the petrochemical industry and was independently evaluated on test cases [41].

Techniques such as auto-embedding and significance testing designed to infer an ITN that is more representative of a true causal structure does not assist with obtaining better rankings and may even have a detrimental effect. When
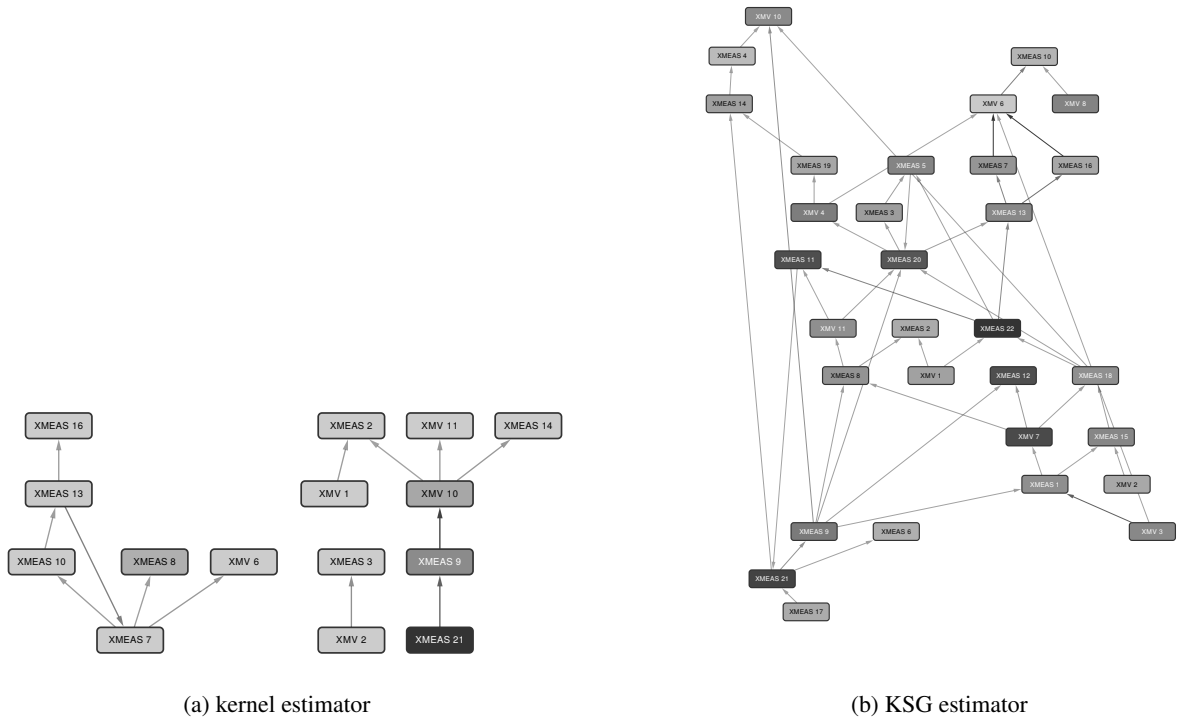
(a) kernel estimator

(b) KSG estimator

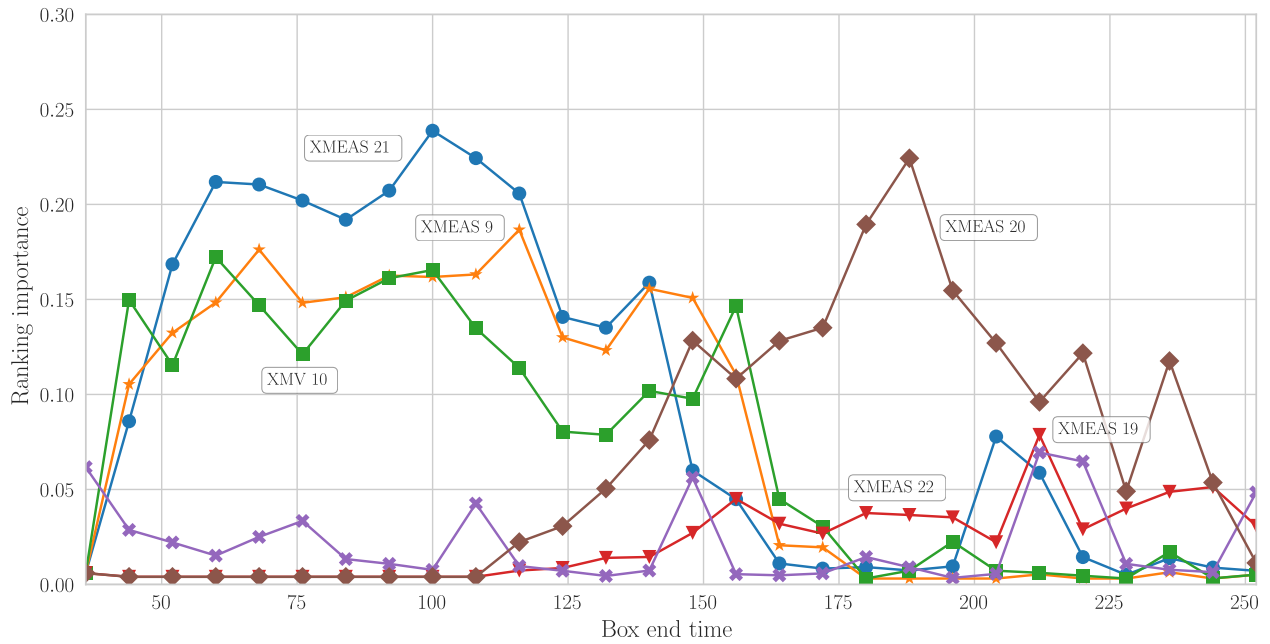Figure 8: Reduced edge ITNs for CW disturbance [simple; bidirectional delays; significance tested]

Figure 9: Multiple time region absolute ranking scores for periods covering both reactor CW inlet temperature and stripper feed composition disturbance in Tennessee Eastman process [kernel estimator; directional weights; process limit scaled; forward delays; significance tested]

process knowledge based scaling limits are available, the simple kernel transfer entropy method should be used to infer the ITN, otherwise directional naive KSG estimation should be used. There is little evidence that embedding improves the ranking results obtained with the KSG estimator. Significance testing is not necessary for getting good MTR node ranking plots, and may be safely skipped for fault detection, but should be enabled when the underlying ITNs are to be used for fault diagnosis.

As the results come without robustness guarantees or confidence intervals and may require some effort in interpretation, they are best presented as advisories to panel operators and should not be the basis of automated action pending further development. The concept of identifying the "true source" of a disturbance or fault in a process area that involves multiple interacting elements (especially bi-directional couplings) is ambiguous, since adjusting any number of the elements involved could be likely to improve rejection of the disturbance. If a "fault" due to a poorly tuned controller (as opposed, for example, to a specific faulty valve), then it is likely that re-tuning most of the other controller elements involved can compensate for it. There are likely to be better and worse solutions, especially with regards to the directionality and potential ill-conditioning of the resulting system. Finding acceptable tuning parameters might require the use of specialised algorithms combined with expert knowledge and trial-and-error.

As far as could be determined, the following elements is not discussed in prior art related to TE based FDD on chemical processing plants:

Scaling according to control limits
> Results indicate that scaling according to control limits provides significant benefits in certain situations.

Bi-directional delay analysis with associated significance testing logic
> Directional significance testing assists in determining the true direction of interaction between coupled variables. This is necessary for an unsupervised optimising search over a range of delays to be robust.

Use of KSG estimator and auto-embedding
> The implementation of these techniques in open source code is relatively new and their application in literature mostly limited to neuro- and geoscience. The KSG estimator exhibits significant advantages in the absence of scaling according to control limits and is also better at accurately detecting long chains of interaction.

9 Recommended areas for future research

9.1 Incorporating meta-data of process elements

All process states are not equally important from economic or safety perspectives. Strict control of inventory variables is not required and might even be detrimental to disturbance rejection. Considering additional information about process elements during the calculation of node importance scores will allow for increased flexibility and relevancy in results.

Biasing the eigenvector network centrality measure with individual weights is discussed in Section 5.4, but this might not be sufficient to properly handle multiple sources of potentially conflicting information such as the relative value of material or energy streams, control performance KPIs, flags indicating whether a loop is in manual or an element of a advanced process control (APC) strategy, etc. The use of ranking algorithms that can optimise a complex objective function [42] is needed.

9.2 Cluster ranking

The probability of indicating a wrong element (false positive) and missing a real source (false negative) is significant when attempting to pinpoint a specific element as the singular source of a fault or disturbance. A tool tasked with screening faults and disturbances on a plant-wide scale will do well by indicating a cluster of likely elements that can be further analysed with more rigorous information decomposition based techniques. Chemical plants and their integrated control systems generally behave as dynamic multi-agent processes, and calculating scores for interconnected groups might make analysis more accurate but less precise.

Efficient methods for organising large networks into hierarchical structures with intergroup connections are available [43]. In addition, methods that differentiate between inter- and intra-cluster edges, such as the Weighted Inter-Cluster Edge Ranking (WICER) algorithm [44] might be useful to consider for this application.

9.3 Multivariate transfer entropy

Multivariate transfer entropy can measure directed causal effects between groups of variables and reveal complex synergies [35]. This approach might be useful in the process engineering context on processes that exhibit strong

directionality where the impact of relative movement between two or more variables can be more important than the movement of a single variable.

10   Implementation

A Python implementation of this method is accessible at `https://github.com/SimonStreicher/FaultMap.git` It relies on the Java Information Dynamics Toolkit (JIDT) [26] and NetworkX package [45] for calculation of information-theoretic and network centrality measures, respectively.

References

[1] Nina F. Thornhill and Alexander Horch. Advances and new directions in plant-wide controller performance assessment. *Proceedings of ADCHEM-2006*, 2006.

[2] Venkat Venkatasubramanian, Raghunathan Rengaswamy, and Surya N Ka. A review of process fault detection and diagnosis Part III : Process history based methods. *Computers & Chemical Engineering*, 27(3), 2003.

[3] Marcelo Farenzena, JO Trierweiler, and SL Shah. Variability matrix: A novel tool to prioritize loop maintenance. *Proceedings of ADCHEM-2009*, 2009.

[4] Marcelo Farenzena and JO Trierweiler. LoopRank: a novel tool to evaluate loop connectivity. *Proceedings of ADCHEM-2009*, 2009.

[5] Anisur Rahman and M.A.A. Shoukat Choudhury. Detection of control loop interactions and prioritization of control loop maintenance. *Control Engineering Practice*, 19(7):723–731, jul 2011.

[6] M Iri, K Aoki, E O'Shima, and H Matsuyama. An algorithm for diagnosis of system failures in the chemical process. *Computers & Chemical Engineering*, 3:489–493, 1979.

[7] R. F. Vianna and C McGreavy. Qualitative modelling of chemical processes - a weighted digraph (WDG) approach. *Computers and Chemical Engineering*, 19(SUPPL. 1):375–380, 1995.

[8] Mano Ram Maurya, Raghunathan Rengaswamy, and Venkat Venkatasubramanian. A signed directed graph-based systematic framework for steady-state malfunction diagnosis inside control loops. *Chemical Engineering Science*, 61(6):1790–1810, mar 2006.

[9] Fan Yang, Sirish Lalji Shah, and Deyun Xiao. Signed directed graph-based hierarchical modelling and fault propagation analysis for large-scale systems. *IET Control Theory & Applications*, 7(4):537–550, 2013.

[10] F. Yang, P. Duan, S.L. Shah, and T. Chen. *Capturing Connectivity and Causality in Complex Industrial Processes*. Springer, 2014.

[11] Thomas M Cover and Joy A Thomas. *Elements of Information Theory*. John Wiley & Sons, Inc., 2nd edition, 2006.

[12] Luca Faes, Alberto Porta, Giandomenico Nollo, and Michal Javorka. Information Decomposition in Multivariate Systems: Definitions, Implementation and Application to Cardiovascular Networks. *Entropy*, 19(1):5, 2016.

[13] Margret Bauer. *Data-driven methods for process analysis*. PhD thesis, University College London, 2005.

[14] R. Landman, J. Kortela, Q. Sun, and S. L. Jämsä-Jounela. Fault propagation analysis of oscillations in control loops using data-driven causality and plant connectivity. *Computers and Chemical Engineering*, 71:446–456, 2014.

[15] Ping Duan, Fan Yang, Tongwen Chen, and Sirish L. Shah. Direct causality detection via the transfer entropy approach. *IEEE Transactions on Control Systems Technology*, 21(6):2052–2066, 2013.

[16] Matthäus Staniek and Klaus Lehnertz. Symbolic transfer entropy. *Physical Review Letters*, 100(15):1–4, 2008.

[17] Weijun Yu and Fan Yang. Detection of causality between process variables based on industrial alarm data using transfer entropy. *Entropy*, 17(8):5868–5887, 2015.

[18] Jianjun Su, Dezheng Wang, Yinong Zhang, Fan Yang, Yan Zhao, and Xiangkun Pang. Capturing causality for fault diagnosis based on multi-valued alarm series using transfer entropy. *Entropy*, 19(12):5–8, 2017.

[19] Luca Faes, Daniele Marinazzo, and Sebastiano Stramaglia. Multiscale information decomposition: Exact computation for multivariate Gaussian processes. *Entropy*, 19(8), 2017.

[20] Jakob Runge, Jobst Heitzig, Vladimir Petoukhov, and Jürgen Kurths. Escaping the curse of dimensionality in estimating multivariate transfer entropy. *Physical Review Letters*, 108(25):1–5, 2012.

[21] Ping Duan, Fan Yang, Sirish L Shah, and Tongwen Chen. Transfer Zero-Entropy and Its Application for Capturing Cause and Effect Relationship Between Variables. *IEEE Transactions on Control Systems Technology*, 23(3):855–867, 2015.

[22] Joseph T. Lizier and Mikail Rubinov. Multivariate construction of effective computational networks from observational data. *Max Planck Institute: Preprint*, 2012.

[23] Yonathan Murin, Jeremy Kim, Josef Parvizi, and Andrea Goldsmith. SozRank: A new approach for localizing the epileptic seizure onset zone. *PLoS Computational Biology*, 14(1):1–26, 2018.

[24] Thomas Schreiber. Measuring Information Transfer. *Physical Review Letters*, 85(2):461–464, 2000.

[25] Yidan Shu and Jinsong Zhao. Data-driven causal inference based on a modified transfer entropy. *Computers & Chemical Engineering*, 57:173–180, oct 2013.

[26] Joseph T. Lizier, Mikhail Prokopenko, and Albert Y. Zomaya. Local information transfer as a spatiotemporal filter for complex systems. *Physical Review E*, 77(2):026110, feb 2008.

[27] K Hlaváčková-Schindler, M Paluš, M Vejmelka, and J Bhattacharya. Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports*, 441(1):1–46, mar 2007.

[28] David Kristjanson Duvenaud. *Automatic Model Construction with Gaussian Processes*. PhD thesis, University of Cambridge, 2014.

[29] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Physical Review E*, 69(6):066138, jun 2004.

[30] Michael Wibral, R Vicente, and Jt Lizier. *Directed Information Measures in Neuroscience*. Springer, 2014.

[31] Joseph T Lizier. JIDT : An information-theoretic toolkit for studying the dynamics of complex systems. Technical report, 2014.

[32] Sigurd Skogestad and Ian Postlethwaite. *Multivariable Feedback Control*. Wiley, 2005.

[33] Chunhui Zhao, Youxian Sun, and Furong Gao. A multiple-time-region (MTR)-based fault subspace decomposition and reconstruction modeling strategy for online fault diagnosis. *Industrial & Engineering Chemistry Research*, 51:11207–11217, 2012.

[34] R. Marschinski and H. Kantz. Analysing the information flow between financial time series. *The European Physical Journal B*, 30(2):275–281, nov 2002.

[35] Joseph T Lizier, Jakob Heinzle, Annette Horstmann, John-Dylan Haynes, and Mikhail Prokopenko. Multi-variate information-theoretic measures reveal directed information structure and task relevant changes in fMRI connectivity. *Journal of computational neuroscience*, 30(1):85–107, feb 2011.

[36] Thomas Schreiber and Andreas Schmitz. Surrogate time series. *Physica D: Nonlinear Phenomena*, 142(3-4):346–382, aug 2000.

[37] Miguel Castaño Arranz and Wolfgang Birk. New methods for interaction analysis of complex processes using weighted graphs. *Journal of Process Control*, 22(1):280–295, jan 2012.

[38] Paul Shannon, Andrew Markiel, Owen Ozier, Nitin S Baliga, Jonathan T Wang, Daniel Ramage, Nada Amin, Benno Schwikowski, and Trey Ideker. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome research*, 13(11):2498–504, nov 2003.

[39] JJ Downs and EF Vogel. A plant-wide industrial process control problem. *Computers & Chemical Engineering*, 17(3):245–255, 1993.

[40] N Lawrence Ricker. Decentralized control of the Tennessee Eastman Challenge Process. *Journal of Process Control*, 6(4):205–221, jan 1996.

[41] Errol Zalmijn and David Sigtermans. Identifying influential nodes in complex networks. In *Coupling and Causality in Complex Systems*, Cologne, 2017.

[42] Bin Gao, TY Liu, W Wei, Taifeng Wang, and H Li. Semi-supervised ranking on very large graphs with rich metadata. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, number 49, pages 96–104, 2011.

[43] Daisuke Tsubakino and Shinji Hara. Eigenvector-based intergroup connection of low rank for hierarchical multi-agent dynamical systems. *Systems & Control Letters*, 61(2):354–361, feb 2012.

[44] D. Padmanabhan, P. Desikan, J. Srivastava, and K. Riaz. WICER: A Weighted Inter-Cluster Edge Ranking for Clustered Graphs. *The 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI'05)*, pages 522–528, 2005.

[45] Aric A Hagberg, Los Alamos National, and Los Alamos. Exploring Network Structure, Dynamics, and Function using NetworkX. In *7th Python in Science Conference (SciPy 2008)*, pages 11–15, 2008.