INTERNATIONAL JOURNAL OF INTELLIGENT CONTROL AND SYSTEMS
VOL. 13, NO. 3, SEPTEMBER 2008,196-207

Computation and Applications of Industrial Leading Indicators to Business Process Improvement

Wei Peng, Tong Sun, Philip Rose, and Tao Li

Abstract— Within Business Intelligence (*BI*) systems, an industrial Key Performance Indicator (*KPI*) is a measurement of how well the industrial process in the organization performs an operational activity that is critical for the current and future success of that organization [1]. The industrial leading indicators are one type of KPIs that present key drivers of industrial business value, are predictors of future outcomes. Thus leading indicator discovery is critical to success of the industrial objectives. There are some challenges in leading indicator discovery. The traditional approach depending on domain experts' experiences is labor-intensive and error-prone. In addition, because the time shifts between industrial KPIs are vague and often inconstant for variability of business concerns, the correlation between them cannot be correctly calculated using the traditional distance functions. In this paper, we propose a semi-automatic system with an iterative learning process for discovering leading indicators to help trace anomalies and optimize the industrial objectives. Finally two industrial case studies are conducted by applying the proposed methods in the production printing application. The proposed system has two key differentiations and novelties: (1) the semi-automatic framework uses temporal data mining techniques combined with domain knowledge to enable timely access to KPI analysis, and anomaly tracing; and (2) an iterative learning method continuingly uncovers the "root" leading indicators along with the changes of business environment.

***Index Terms*—KPI, Leading Indicator, Business Process, Iterative Learning, Anomaly, Time Series.**

I. INTRODUCTION

Within Business Intelligence (BI) systems (e.g., business dashboard, business performance management, management information system, etc.), industrial Key Performance Indicators (also known as KPI) are one tool used to convey the relative health of the business, or a portion of that business. A KPI is a specific metric (a quantitative, periodic measurement of one or more processes), chosen from all of the collected or possible industrial metrics within a business in such a manner as to convey the most amount of information in a single measurement – the "key" measurement. As such, a KPI is a quantifiable measurement of how well an operational, tactical or strategic activity is performed and progressed in the industrial process within an organization [1]. KPIs must reflect the critical success factors of an organization. Not all metrics are indicators, not all indicators are KPIs either, but all KPIs are indicators and all indicators are metrics, thus defining an ontological hierarchy of measurements.

There are three types of KPI [1]: (1) **Leading indicator**: "a KPI that measures activities that have a significant effect on future performance of industrial objectives" which are causal roots of the outcome (i.e. lagging indicator) they influence, and actionable for the future performance against one or more lagging indicators; (2) **Lagging indicator**: a KPI that measures the output of past activities; and (3) **Diagnostic measure**: a KPI that is neither leading nor lagging, but signals the health of industrial processes or activities. For example[1], "Number of clients that sales people meet with face to face each week" may be a leading indicator of "Sales Revenue" (a lagging indicator or outcome); "Complex repairs completed successfully during the first call or visit" be a leading indicator of "Customer Satisfaction" (a lagging indicator or outcome). "Throughput" should be a diagnostic measure of "Efficiency" of a production workflow. Leading indicators are very powerful metrics in that they possess not only the predictive and insightful causal relationship(s) within the business process(s), but also enable the actionable course for continuing industrial process improvement. Therefore, creating effective leading KPIs is critical to the success of any business organization so that not only it is agile to changes, but also is prepared for changes in advance. However, identifying leading indicators is often hard and tricky requiring months to collect requirements, standardizing definitions and rules, prioritizing metrics, and soliciting feedback, etc. Moreover, because the time shifts between the leading indicators and the corresponding

This work is partially supported by NSF CAREER Award IIS-0546280 and a Xerox UAC (University of Affairs Committee) award.

Wei Peng, Tong Sun and Philip Rose are with the Xerox Innovation Group, Xerox Corporation, Webster, NY 14580 USA (e-mail: {Wei.Peng, Tong.Sun, Philip.Rose}@xerox.com).

Tao Li is with the School of Computer Science, Florida International University, Miami, FL 33199 USA (e-mail: taoli@cs.fiu.edu).

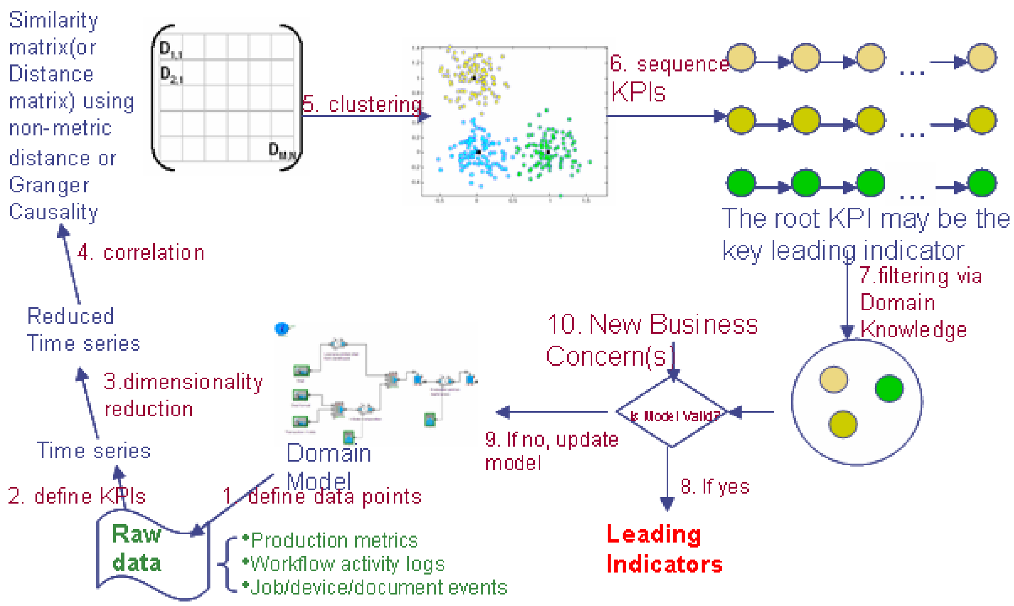

Fig. 1. The Schematic View and Iterative Steps

affected lagging indicators are vague and often inconstant for variability of business concerns, the traditional approach depending on domain experts' experiences is labor-intensive and error-prone. Given that 20% of the work in developing a successful KPI comes from a deploy-observe-adjust cycle [1] – i.e. putting a KPI into a BI system, seeing how it impacts behavior and performance, and then adjusting accordingly – many of the KPIs visible in a typical BI system at any point in time are not delivering full value to the users of these systems. Furthermore, it is much more challenging to identify the more powerful value-driver KPIs (i.e. leading indicators). Finally, many traditional BI systems, such as Business Scorecard [12], Executive Information Systems [13], Dashboard [14,15], deliver primarily pre-defined historic metrics (or lagging indicators) for a long-term strategic or mid-term tactical analysis, and lack the necessary flexibility to support evolving metrics or data collection points over time for real-time operational analysis. There are also some BI systems improving business process quality through managing exceptions or execution quality [16,17].

In this paper, we propose a semi-automatic system with an iterative learning process for analyzing operational metrics, factoring out the key performance indicators (KPIs) and then further discovering leading indicators by applying certain data mining techniques incorporated with the domain knowledge. The system involves domain experts to define the KPIs and validate the leading indicators to avoid "black-box" effects [20]. In order to illustrate how the proposed system and methods can be applied in a specific business process, two case studies have been conducted for typical real-world production print workflows (e.g., transaction printing, and book printing, as examples). The major challenges being investigated and addressed in this paper are:

- How to automate the computation intensive data analysis while incorporating as much domain knowledge as possible?
- How to effectively factor out overlapping or non-critical metrics?
- How to calculate the approximate influencing time shifts between correlated metrics?
- How to measure the significance of causal relationships between the leading indicators and the metrics they predict?
- How to enable the addition of new data collection points if the resulting leading indicators are not satisfactory based on domain knowledge, when business concerns change, or when other alternations occur within the underlying organization?

Compared with the traditional analytic capabilities in most existing BI systems, the proposed system and methods have two key differentiations and innovations: (1) the semi-automatic framework simplifies many traditional labor-intensive and error-prone steps through the application of temporal data mining techniques (e.g., dynamic time warping, Granger Causality, adaptive agglomerative clustering) combined with specific domain knowledge, thus enabling timely access to operational metrics, KPI analysis, and powerful leading indicator discovery; and (2) an iterative learning methodology not only continuingly uncovers the "root" leading indicators, but also enables the flexibility and adaptability for metrics updates and additional data collection points.

The discovered leading relationship also helps to correlate the detected anomalies in business processes. Anomalies indicate some unexpected situations such as operational or business measurement below a certain threshold [23,24]. As we know, anomaly detection and tracing are critical to keep a business process workflow running soundly. Even small faults may cause the whole workflow process unable to work properly. However detecting, understanding, and tracing anomalies in the complex business process are not trivial since the states of business workflow components vary along time, anomalies usually propagate, and the number of anomalies is too small. We propose to trace anomalies and identify the anomaly sources by looking for the correlation between the KPIs where the anomalies present instead of looking for the correlation between the anomalies directly.

The rest of our paper is organized as follows: Section II describes the system scheme and the iterative learning methodology for KPI analysis and leading indicator discovery. Section III introduces Dynamic Time Warping (DTW) and Granger Causality in leading indicator identification. Section IV specifies how to use adaptive agglomerative clustering to construct causal relation hierarchy between KPIs, and find the critical root leading indicators. Section V presents how to detect anomalies and trace them in a business process by using leading indicator discovery techniques. Two case studies that are based on the discrete event simulations of the real-world production printing workflows are illustrated in Section VI. Section VII concludes the paper with summary and future research topics. The preliminary version of this paper is published in [19].

II. SYSTEM SCHEME FOR KPI ANALYSIS LEADING INDICATOR DISCOVERY

Figure 1 illustrates the schematic view of various elements and steps within the proposed system and iterative learning process for KPI analysis and leading indicator discovery. It consists of the following 10 major steps (numbered from 1 to 10 in Figure 1):

1. For any business organization, the underlying domain model, which encapsulates business goal, business process and workflow model, key terminology and relationships, business assumption, etc., drives the definition of data collection points (or data model) that need to be tracked in a BI system. The Raw Data is the basic data entities, which may be culled from recorded events, workflow or machine logs, metrics, etc., and provides a baseline data source for further data-driven analysis.

2. Usually, the domain experts identify and define KPIs and their formulas based on existing raw data, business goals, and personal experience. Some KPIs are generic -- for example, process efficiency and throughput are generic measurements in the manufacturing domain. Some KPIs are customized to a particular business goal or application – for instance, the number of impressions achieved on a press is customized to the printing domain, but bears little meaning in other manufacturing environments. The metrics used in this step are specific sets of time series data from step 1 as identified by the domain expert.

3. In a complex system, a significant number of metrics may be collected. Some metrics contain little information related to certain business goals, and some metrics are overlapping. It is critical to filter out the less important metrics or noises and focus on the small number of metrics in a particular business context that yield the greatest business value. We propose to use unsupervised dimensionality reduction techniques, such as *Principal Component Analysis (PCA)* [2]/*Singular Value Decomposition (SVD)* [3] to filter out the less significant metrics. In addition, an unsupervised dimensionality reduction technique, *Piecewise Aggregate Approximation (PAA)* [4] can help to reduce the time dimensionality of each time series if the computational efficiency is required.

4. In order to discover leading indicators, we explore the correlations among the reduced indicator sets by considering the time-shifts [11]. Traditional metric distance functions, including Euclidean distance and correlation coefficients, are not suitable for detecting correlation between time series. Non-metric distance functions like *Dynamic Time Warping (DTW),* which is widely used in speech recognition, is well suited for leading indicator discovery. DTW applies dynamic programming with time composition and decomposition [5] to discover the best alignment warp with the minimum alignment distortion (distance). This alignment warping path can then be used to compute the time shifts between various metrics, to help determine the time order of these metrics, and to construct the causal hierarchy. In the case of two highly correlated metrics it is often the case that the one appearing before the other is a leading indicator to the other. However, there are two disadvantages of utilizing DTW for leading indicator discovery: (a) the highly correlated time series does not provide theoretic foundation of any causal relationships between the series themselves; and (b) it is difficult to define the threshold alignment warping distance. In order to overcome these drawbacks we next apply *Granger Causality* [6] in order to test causal direction between the time series.

This determines whether one time series variable can help to predict the other. Although Granger Causality is unable to obtain the time shifts, it provides the solid statistical foundation for the leading indicator discovery and establishes the significance scores of causal relationships between metrics. Therefore, we use both DTW and Granger Causality to derive the distance matrix among indicators, which dictate the correlation between any two indicators. Detailed technical content regarding DTW and Granger Causality are provided in Section 3.

5. Any specific metric may have many leading indicators, each of which could in turn lead to additional metrics. This many-to-many directional complexity increases polynomially as the number of indicators increases. Therefore, after the correlation distance matrix is derived in step 4, we propose the use of adaptive agglomerative clustering to obtain the hierarchical relationships in a dendrogram, thus partitioning all of the metrics into clusters. The details behind this proposed clustering algorithm are provided in Section 4.

6. All metrics falling into the same cluster at this point are more highly correlated with each other than across the clusters. Within an individual cluster, the metrics preceding others are assumed to be the root leading indicators for this cluster.

7. After the root leading indicator is discovered from each cluster in step 6, the domain experts will examine them against the underlying domain model and determine whether the resulting leading indicators are actionable and/or meaningfully critical in certain business context.

8. In some cases, the domain experts will discover one or more leading indicators with actionable characteristics, thus designating these indicators as the desired leading KPI. As business processes change, the entire discovery process, from step 1 to step 7, continuingly progresses, such that the resulting leading indicators also change over time.

9. In other cases, the domain experts may decide to delete or devalue some indicators, or decide to add more metrics (which may imply new data collection points to be added) to the domain model. As the underlying domain model is updated, the scheme of raw data collection is evolved via expansion or merging. Meanwhile, the entire discovery process, from step 1 to step 7, continuingly progresses, such that the resulting leading indicators also change over time.

10. Over time the business concerns or context will change, and the domain model is updated accordingly. Similarly, the entire discovery process, from step 1 to step 7, continuingly progresses, such that the resulting leading indicators may evolve over time.

In summary, the above 10 steps illustrate an iterative learning methodology to discover the leading indicators in a business process over time via data mining techniques combined with domain knowledge guidance. The learned leading indicators can also evolve with changing business objectives, and those that are determined to be actionable are implemented as KPIs within the environment so that continuing process improvement is enabled.

III. LEADING INDICATOR IDENTIFICATION METHODS

As we discussed in the last section (step 4 of the iterative process), in order to discover leading indicators, we first need to compute the "time shifts" between the time series indicators, define the threshold alignment warping distance for "high correlation" between the indicators, and then determine whether there is any "causal relationship" from the "high correlative" pairs. We find that DTW and Granger Causality complement each other and serve well for the above two purposes in leading indicator discovery. This section provides detailed descriptions on how these two techniques can be used jointly in uncovering leading indicators from a set of time series indicators.

A. Using Dynamic Time Warping to Discover the Correlation and the Time Order of Indicators

First, we use Z-score normalization [7] to remove the baseline and re-scale the indicators so that the range of all indicators has mean of zero and variance of one. Then we use Dynamic Time Warping (DTW) that is often applied in speech and handwriting recognition [24] to discover the nonlinear alignment similarity between KPIs.

Suppose we have KPI X, which has the value sequence $[X_1, X_2, \dots, X_n]$, and KPI Y, which has the value sequence $[Y_1, Y_2, \dots, Y_m]$ over time. The best warping distance $DTW(X, Y)$ between X and Y is presented as:

$$DTW(X,Y) = D(X_n, Y_m) + min\{DTW(X_{(1,n-1)}, Y_{(1,m-1)}), DTW(X_{(1,n-1)}, Y), DTW(X, Y_{(1,m-1)})\}, \quad (1)$$

where $D(X_n, Y_m)$ is the local distance between the elements X_n and Y_m, and $X_{(1,n-1)}$ and $Y_{(1,m-1)}$ are the subsequences $[X_1, X_2, \dots, X_{n-1}]$ and $[Y_1, Y_2, \dots, Y_{m-1}]$ respectively. A threshold need to be set such that if the best warping distance $DTW(X, Y)$ is below it there should exist a high correlation between X and Y.

For leading indicator discovery based on DTW, we use the alignment warp path to discover the time order in highly correlated indicators and the approximate time shifts between them. The alignment warp path is composed of two arrays which are of the same length, each of which consists of the increasing or decreasing position numbers in an indicator. The elements of the same array number in the two arrays are matched positionally by DTW to achieve the optimal alignment. Let the alignment warp path between KPI X and KPI Y be composed of arrays PX and PY from X and Y respectively. The time shift between X and Y is *abs*(*mode*(PX - PY)), where *abs* is the absolute value function. *Mode*(x) returns the element with the highest frequency in the array x.

- If *mode*(PX - PY) < 0, KPI X precedes Y , and is considered as the leading indicator of Y if X and Y are highly correlated.
- If *mode*(PX - PY) > 0, Y is the leading indicator.
- If *mode*(PX - PY) = 0, X and Y *do not exhibit a leading relationship in either direction*.

For example, given X be [*4, 5, 3, 6, 8*] and Y be [*5, 2, 7, 8, 9*], the alignment warp path PX is [*1, 2, 3, 4, 5, 5*] and PY is [*1, 1, 2, 3, 4, 5*]. The time shift between them is 1, and KPI Y is preceding KPI X if they are regarded to be highly correlated.

B. *Using Granger Causality to Score the Significance of Causal Relationship*

Granger Causality [6] is a measure to determine whether one time series helps to predict another time series. Given lagged values of X and Y from time 1 to t-1, we want to forecast the value of Y at time t. We say that X Granger-cause Y, if the variance of the optimal linear prediction based on lagged X and Y is smaller than if only based on lagged Y. In other words, the addition of lagged X to lagged Y makes better prediction than only lagged Y. Granger Causality usually uses an F-test on the lagged values of X and Y to test whether X provides significant information of the future values of Y.

Let Y_i and X_i be the values of Y and X at time i respectively. The data are described with a bivariate vector regressive model:

$$Y_t = \mu + \sum_{i=1}^{k} \alpha_i X_{t-i} + \sum_{i=1}^{k} \beta_i Y_{t-i} + \epsilon_t, \quad (2)$$

where k is the lag length, α and β are coefficients, and ϵ_t is the error term. The null hypothesis H_0 is $\alpha_1 = \alpha_2 = ... = \alpha_k = 0$. The equation (2) restricted under the null hypothesis is the model:

$$Y_t = \mu' + \sum_{i=1}^{k} \delta_i Y_{t-i} + \eta_t \quad . (3)$$

The residues Res_1 and Res_0 of these two models are $\sum_{t=1}^{n} \epsilon_t^2$ and $\sum_{t=1}^{n} \eta_t^2$, where n is the whole test time. The sum of squares of residues in these two models can be transformed to a modified ratio which is:

$$TS = \frac{(Res_0 - Res_1)/k}{Res_1/(n-2k-1)} \sim F_{k,n-2k-1} \quad . (4)$$

TS, or "test statistic" follows an F-distribution if the null hypothesis is true. The value of test statistic is assigned a significance p-value, which is in the range of [0, 1], by comparing to the corresponding entry in the table of F-test critical value. The smaller the significance score, the higher the possibility to reject the null hypothesis, or to accept the causal relationship. The significance score is directional. If the significance score of X to Y is small enough, while the significance score of Y to X is not, we regard X as the leading indication of Y, and vice versa. If the scores of two directions are very close, neither is presumed to have a leading relationship to the other.

IV. ROOT LEADING INDICATOR IDENTIFICATION METHODS

Any specific metric may have many leading indicators, each of which could in turn lead to additional metrics. This many-to-many directional complexity increases polynomially as the number of indicators increases. Domain experts are unable to determine all of the leading indicators, and decision making is impossible when there are too many leading indicators to be optimized. Therefore complexity must be reduced to provide focus on the critical leading indicators. These critical leading indicators are called root leading indicators and are root causes of one or more key metrics. We propose the use of adaptive agglomerative clustering to obtain the hierarchical relationships in a dendrogram, thus partitioning all of the metrics into clusters and ultimately identifying the "root" leading indicators.

From the results of DTW and Granger Causality, we derive the distance matrix wherein each cell is a value indicating the correlation between the corresponding row metric (represented as K_i) and the column metric (represented as K_j), and the diagonal values are zeros. In case of DTW, the cell value is the alignment warping distance. In case of Granger Causality the significance scores are directional, thus we need a transformation to determine the distance matrix. In the Granger Causality distance matrix, the cell value of row K_i and column K_j is equal to the cell value of row K_j and column K_i, and is equal to the smaller significance score in the scores from K_i to K_j and from j to i. After the distance matrix is derived, agglomerative clustering constructs a dendrogram on top of the metrics. The metrics are then partitioned into clusters such that metrics within the same cluster have a higher correlation between each other as compared to the metrics in other clusters.

By cutting different edges in the aforementioned dendrogram we can obtain a different number of clusters. However, where to cut the edges is uncertain in the traditional agglomerative clustering, and even domain experts cannot determine the number of clusters in a complex set of metrics. Modified agglomerative clustering is able to solve this problem. The cluster number can be determined by *Akaike Information Criterion* (AIC) - essentially the log-likelihood of the model increasingly penalized by the number of parameters. The AIC score of a cluster assignment C_i is defined as:

$$AIC(C_i) = 2L(C_i) - 2K \times m \quad (5)$$

where $L(C_i)$ is the log-likelihood of C_i. $K \times m$ is the number of parameters in the model. K is the number of clusters, and m is the number of coordinates of each metric. We assume that each cluster is following multivariate Gaussian distribution. The log-likelihood $L(C_i)$ is:

$$L(C_i) = \sum_{i=1}^{K}[-\frac{n_j}{2}log(2\pi) - \frac{n_j \cdot m}{2}log(\sigma_j^2) - \frac{n_j - K}{2} + n_j \cdot logn_j - n_j \cdot logn], \quad (6)$$

where n_j is the number of metrics in cluster j, and σ_j is estimated by the average distance between all pairs of metrics in cluster j. The cluster number that obtains the highest AIC score is chosen.

Because the metrics in the same cluster are cohesive based on the non-traditional correlation function, σ_j is calculated using the warping distance matrix from DTW. Moreover, we also project the metrics into the two dimensional space where the traditional distance function preserves the property of the correlation obtained from DTW in the original space. The visualization of metrics in this latent two-dimensional space helps to determine the number of clusters, and offers vivid intuitions of the relationships between metrics [22]. We use *multidimensional scaling* (MDS) to create a space that faithfully captures the observed correlation between entities in this space [8, 9]. In our experiment, the input $n \times n$ pairwise distance matrix of n metrics is transformed to an $n \times 2$ matrix such that every metric is projected into a two dimensional space. In each of the obtained clusters the metrics are sequenced based on the time order and shifts between them. The metrics preceding all other metrics in the cluster are regarded as the root leading indicators of the other metrics in this cluster.

V. TRACING ANOMALIES

We adopt Adaptive Threshold [18] to detect anomalies in each KPI. It uses Exponentially Weighted Moving Average (EWMA) [21] to estimate the recent mean of the time series KPI, and adaptively sets the alarm threshold. If the KPI is below the threshold at a particular time, the alarm signals. Given a KPI w, and the value of w at time i w_i, the alarm signals when: $w_i < \alpha\mu_{i-1}$, where μ_{i-1} is calculated from the past history of w before time i. α in the range (0, 1) is the percentage of mean value below which the alarm prepares to signal. μ_i can be estimated by using EWMA as follows

$$\mu_i = (1-\gamma)w_i + \gamma\mu_{i-1}, \quad (7)$$

where λ is the EWMA factor. The direct use of the above algorithm may yield many false alarms. Usually it is modified such that if there are a certain number of successive violations of the threshold, then the alarm signals.

Before tracing anomalies in workflow, the correlations between anomalies need to be mined. Since the number of anomalies is usually too few to extract any correlation rules, we assume that highly correlated anomalies should lie in highly correlated KPIs. If anomalies show the leading relationship extracted from their KPI analysis process, the anomaly appearing first within a range of a certain time period is the anomaly source. For example, the anomaly a is detected in KPI A at time t_a, the anomaly b is detected in KPI B at time t_b, and if we find the KPI A is the leading indicator of KPI B with the time shift T and t_b-t_a is in the range (0, T+ ε), where ε is a small positive value due to time shift estimation error, then a is the anomaly source of b. If an anomaly c in KPI C is detected at time t_c, C is leading B in time T', and t_c-t_b is in the range (0, T'+ ε), then c is the anomaly source of a and b.

VI. CASE STUDIES

In order to illustrate how the proposed system and methods can be applied in specific business processes, and how the leading indicators help guide decision making, a study was made of two typical real-world workflows in production printing domain: transactional printing and book production. The first case study gives a detail description of applied techniques of leading indicator discovery and anomaly tracing process. The second case study mainly focuses on the iterative process. A domain-expert helped construct two *discrete event* driven domain models based on real world printing applications that in turn drove the data collection points (or metrics), KPI filtering, and leading indicator validation. The *discrete event model* tracks asynchronous discrete incidents (or events) with the system state transitions based on specified time intervals [10], thus delivering the time series data required for these case studies.

A (1) Leading Indicator Discovery in a Transactional Production Printing Workflow Scenario

Figure 2 illustrates a transactional production printing workflow model, consisting of 7 activities or operations (shown in circles) connected via transactions (shown as rectangles). Some of the activities or operations can be