RC25563 (WAT1510-012) October 6, 2015
Computer Science

IBM Research Report

Cross Industry Analytics Solution Library for Resource and Operations Management

Jayant Kalagnanam, Young Min Lee, Tsuyoshi Ide
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598 USA

CROSS INDUSTRY ANALYTICS SOLUTION LIBRARY FOR RESOURCE AND OPERATIONS MANAGEMENT

Jayant Kalagnanam, Ph.D. Distinguished Research Staff Member,
Young Min Lee, Ph.D. Research Staff Member
Tsuyoshi Ide, Ph.D. Senior Technical Staff Member
contact: jayant@us.ibm.com

IBM Research
IBM T.J. Watson Research Center
110 Kitchawan Road
Yorkown Heights, NY 10598, USA

ABSTRACT

In this paper, we describe a framework of analytics solution library for solving variety of industrial problems in resource and operations management. This Big Data enabled analytics solutions library is targeted for manufacturing industries including Electronics, Semiconductor, Automotive, Chemical & Petroleum, Oil & Gas industries, Energy & Utility, and Mining & Metals industries. The analytic library contains families of state of art analytics including predictive asset management, predictive failure analysis of process and equipment, process and equipment monitoring, analysis and optimization, predictive environmental analytics system and spatio-temporal analytics for safety and operational effectiveness. The analytics is powered by variety of advanced algorithms, is cataloged based on a comprehensive ontology, and guides common users without specialized training in mathematics with application specific workflows and a pre-engineered application programming interface (API), data structures and Graphical User Interface (GUI) widgets for visualizing input data and output solution. The analytics solutionlibrary can ingest remote sensing data from the Internet of Things (IoT) 3.0 foundations and is designed to consume variety and vast amount (i.e., Big Data) of structured and unstructured data including video, text and real time sensor data. The analytics services library bridges the gap between analytics algorithms and real business problems, and promotes the convergence of Information Technology (IT) and Operations Technology (OT) and promotes a vision of Industry 4.0 through effective and efficient utilization of advanced analytics provisioned through an increasingly sensored physical world.

1 INTRODUCTION

A new era of computing is here. The new era of computing is the era of cognitive computing that understands the world in the way that humans do through sensing, learning, and adapting behavior automatically based on new knowledge. The first computing era started around the 1900s and was made up of tabulating machines and the second era started around 1950s and was made up of programmable computers. Cognitive computing represent a whole new approach to solving complex problems with large amount data, machine learning and analytics that goes beyond just computing. There are enormous growth of Big Data, which is being fueled by the confluence of the internet of things (IoT), social, mobile and cloud. Industrial operations, processes and equipment are now also generating enormous amount of data. Currently, one trillion connected objects and devices on the planet are generating 2.5 billion gigabytes of data every day, and the rate is growing. The size of the machine generated data globally is about 8.5 ZB (zettabytes, or trillions of gigabytes), and is expected to grow to exponentially to 40 ZB in 5 years (2020) [1]. By 2020, machine to machine data will dominate constituting 40% of all the data. The data generated from sensors and devices are becoming bigger than the social media data, VoIP and enterprise data.

The data are being generated from variety of places and objects including manufacturing shop floor, industrial equipment, sensors from factories and homes, social media, smart phone, satellite images, weather laboratory, home appliances and even from automobiles that we drive. The types of the data are also very diverse. Data can be structures, un-structured, video, audio, text and so on.

Data has become the world's new natural resource. Industry is being revolutionized with the data. By utilizing the data properly, manufacturing firms can improve product quality, reduce production costs, improve safety in manufacturing plants, improve energy efficiency of production, and improve the profit margin. Data is now the basis of competitive advantage. The data can also be used improve the quality of human lives by promoting smarter buildings, smarter cities and smarter planet. However, collection of the Big Data itself does not generate any benefits. It is important to interconnect all the data sources, and to collect the big data; however, without appropriate analytics to process the data and transform the data into useful decision supporting information, the Big Data can be just a big waste. Only when the data is properly analyzed and utilized, the data can help industries and society. Analytics in this new era of computing is different from conventional analytics. The traditional analytics can no longer be able to handle the volume, variety, velocity and veracity of Big Data. What we need in the new era of computing is advanced analytics that can be easily built, configured and deployed, and fully integrated with IoT and Big Data. The analytics in the new era of computing is complex aggregation of various algorithms and techniques including advanced mathematics, statistics, data mining, machine learning, cognitive computing and optimization.

One important technical domain in industry that can take advantage of the analytics in the new era of computing is *Resource and Operations Management*, which is a collection of key problems that deals with assets, resources, operations and processes in manufacturing industries. The ability to integrate with IoT, and store, process and analyze Big Data will improve various aspects of resource and operations management problems, and bring competitive advantage to industries.

By effectively and efficiently utilizing Big Data and IoT with analytics, the opportunities to improve the various industrial processes and operations can be discovered. However, it is quite challenging to develop such analytics because it typically requires deep expertise across data mining, machine learning, statistics and optimization as well as domain knowledge in order to develop practical solutions. Substantial resources and time are required to develop such analytics. Therefore, not many enterprises can afford to develop and utilize such analytics. It is often difficult for common users without specialized training in mathematical modeling or statistics to identify an appropriate analytics for solving a specific problem, develop and run the analytics in appropriate steps and manners to produce desired results. Increasingly along with the complexity of the data itself, the complexity of choosing the right analytics algorithms is emerging as another challenge in the IoT era.

In this work, we are developing a solution platform that provides an analytics services library, called SROM (Smarter Resource and Operations Management), with a hierarchy of analytics that can solve variety of problems in the area of resource and operations management. The platform enables the common users without specialized training in mathematical modeling to easily identify, develop and use the analytics for solving complex industrial problems in resource and operations management with much less effort and time.

2 CAPABILITIES OF SROM

Smarter Resource and Operations Management (SROM) is a next generation analytics solution framework that can fill the gap between real business and analytics in the areas of resource and operations management. Most of the existing analytics solutions for industrial problem lack realistic use cases and industry specific solutions. A significant gap exists between algorithms and real industrial problems. SROM integrates business solution library that is provisioned by rich, state-of-the-art algorithms. Analytics platforms for enterprise systems are beginning to evolve from pure analytics' function engines to business/enterprise function oriented solutions that are purpose built for specific industries. The first generation analytics typically consists of bare statistics function and machine learning library, and the examples of these analytics include products such as S, R, Weka, etc. The second generation analytics are typically machine learning library armed with composable GUI enterprise system, and the examples of these are SAS, SPSS, etc. The third generation analytics are cloud-based, solution-oriented, function based-APIs, and SROM is one of them. SROM is designed to address the rapidly evolving solution needs for asset heavy industries.

SROM is an ontology based services solution library, and it is enabled to be integrated with Big Data and IoT. SROM is designed for cross-industry reuse, and targeted for manufacturing industries including Electronics, Semiconductor, Automotive, Chemical & Petroleum, Oil & Gas industries, Energy & Utility, Mining & Metals industries. The analytics are powered by variety of advanced algorithms, are cataloged based on comprehensive ontology, and guide common users though application specific workflow with pre-engineered application programming interface (API), data structures and GUI widgets for visualizing input data and output solution.

The analytic library includes, but not limited to, families of state of art analytics such as predictive asset management, predictive project health, process monitoring and optimization, and spatio-temporal analytics for safety and operational effectiveness as can be seen in Figure 1. SROM is a cloud-enabled services, and it is integrated with data services, which have connected to two types of data. One type of data is IoT data, which is dynamically collected from sensors and equipment in industrial operations. This data can include data from SCADA (supervisory control and data acquisition) systems, DCS (Distributed Control System), MES (Manufacturing Execution System) and other OT system as well as IT system etc. The other type of data is systems of record such as ERP (enterprise resource planning) system and work order management system and asset management system and so on. SROM analytics consume both types of data. Depending on the particular problem of interest, the data size can be very large (i.e., Big Data), and the data can be structured and unstructured data including video, text and real time sensor data. SROM analytics services layer in the middle of Figure 1 can be packaged into specific industry specific solution pack, for example, an analytic library packaged for automotive industry, chemical/ petroleum, electronics or energy/utility etc.

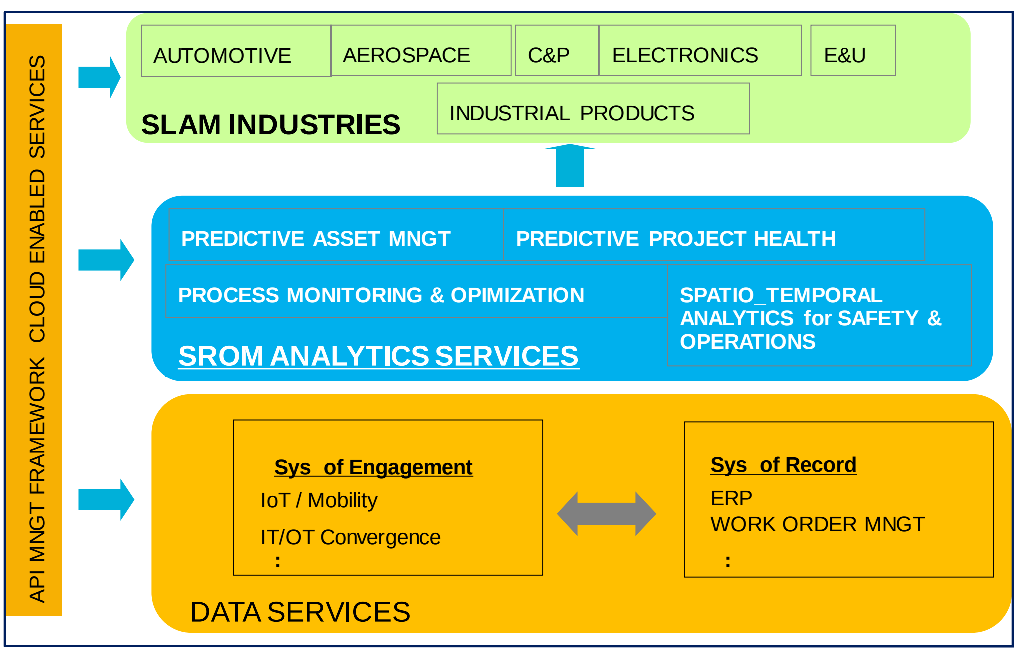

Figure 1: An Overview of SROM

There are many types of analytics that can solve various problems in industry in the areas of resource and operations management, but it is often difficult for common users to identify appropriate analytics and algorithm for solving a specific problem, and also to try out example applications or use cases that are already built for other firms in the same industries. The analytics in SROM is catalogued based on ontology of analytics. The ontology enables the common users without specialized training in mathematical modeling to understand what are typical analytics that are useful in different industries, how various analytics are related to one another, identify, and ultimately develop a specific analytics for solving complex industrial problems in resource and operations management with much less effort and time. A simplified view of the overall structure of SROM ontology is shown in Figure 2.

As shown in Figure 2, there are four levels of elements in SROM ontology in the area of resource and operations management. The top level is the family of analytics and they includes: *Predictive Maintenance*, *Predictive Failure Analysis*, *Process and Equipment Analysis* and *Process Monitoring and Optimization*.

Predictive Maintenance Analytics is used to determine the condition of various assets and equipment in processes and operations and predict when maintenance should be performed. The maintenance work can be inspection, repair or replacement of equipment and parts. This analytics are typically formulated as optimization model such as Mixed Integer Linear Programming (MILP) with a goal of minimizing objective function of cost of maintenance and failure and with constraints of resources, spare parts, budget and operations. Predictive maintenance analytics help to reduce maintenance cost, increase lifetime of equipment and increase the quality of products and improve the safety of manufacturing plant. A longer term planning activity of Predictive Maintenance is Maintenance Planning, which determines a maintenance plan typically in weekly or monthly time bucket with a planning horizon of one or more years considering failure risk of asset. A shorter term scheduling activity of Predictive Maintenance is Maintenance Scheduling, and it determines scheduling of maintenance crew, resources and routing typically in hourly or sub-hourly time resolution for a time horizon of one day or one week.

Predictive Failure Analysis is used to predict future failure of assets and equipment so that corrective actions are taken to avoid the failure. Failure of assets and equipment in production can cause significant negative impact on product quality, production cost and safety. There are two types of predictive failure analysis. One is Failure Pattern Analysis, which identifies conditions (also referred as patterns) that are precursors to a failure. The other is Failure Risk Analysis, which estimates of the life expectancy of an asset based on historical failure data and operational conditions.

Process and Equipment Analysis is used to analyze the stream of sensor data generated from the process and equipment and identify anomalies or faults. The analysis is also called Fault Detection and Diagnosis (FDD). A good instance of this analysis is Anomaly Detection, which identifies asset behavior that is outside of a defined norm that constitutes normal behavior. Anomaly detection first characterizes the normal behavior using historical measurements of the asset behavior under different conditions, and then identifies the excursions outside these norms as potential anomalies, and generates alerts.

Process Monitoring and Optimization refers to a group of analytics that continuously monitors process and equipment, and optimizes the controllable variable settings so that the processes or operations run in most cost-efficient, energy-efficient and safe manners with maximized throughput. A predictive model for the processes or operations, either in physics model or data-driven model or combination of both needs to be built for this analysis. And then an optimization model is built to compute the optimal control variables values using on the predictions from the predictive model.

For each analytics family, there are multiple instances of analytics (second level). For example, the *Predictive Maintenance* has *Maintenance Planning* and *Maintenance Scheduling* analytics among others. And, *Predictive Failure Analysis* family contains the analytics of *Failure Pattern Analysis* and *Failure Risk Analysis*. Each analytics can be solved by multiple algorithms (level 3). For example, *Failure Pattern Analysis* can be solved by three different algorithms, i.e., *Decision Tree*, *Support Vector Machine* and *Random Forest*. Then at the last level (level 4), sample applications of use cases are included. For example, for *Graphical Method* of *Anomaly Detection* analytics has use cases of *Mining Machinery Anomaly Detection* and *Offshore Oil Platform Anomaly Detection*. The more detailed hierarchies of *Predictive*

Maintenance and *Predictive Failure Analysis* are shown in Figure 3. Note that the particular ontology shown here is still being developed and are not complete. Continuous work is being done to streamline and improve the ontology as we deal with more industrial problems. Some instances of the analytics, algorithms and use cases are described in the section 6 below.

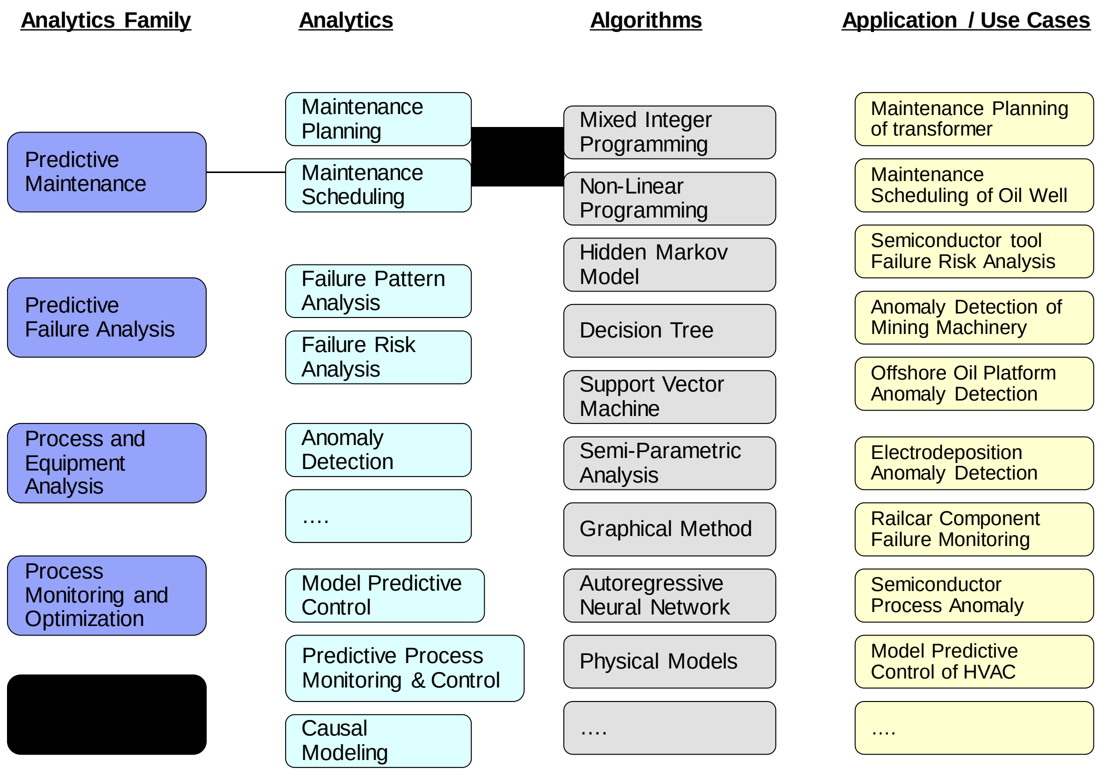

Figure 2: An Overview of SROM Ontology

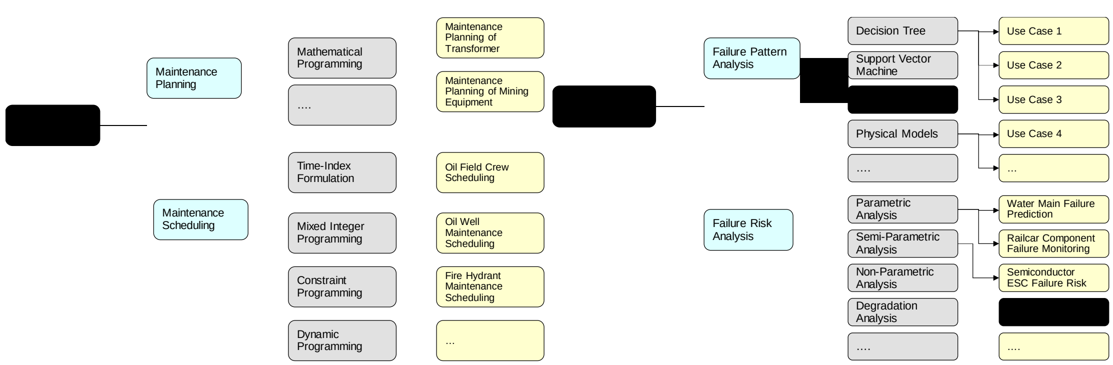

Figure 3: Ontology of Predictive Maintenance and Predictive Failure Analysis

Each analytics uses specific input data and produces certain output data as solutions. In addition, the solution of one analytics can be used as an input to another analytics. An example of the relationship among analytics and corresponding data is shown in Figure 4. In this example, data of failure history, asset attributes, environmental data and operational data are used for *Failure Pattern Analysis* model and *Failure*

Risk Analysis model. The output from the *Failure Risk Analysis* model, hazard function along with cost factors of maintenance and failure, reliability estimation and operational constraints such as budget and resource of maintenance, are used as input for *Maintenance Planning* optimization model. The maintenance plan and anomaly detected can be used as input to *Maintenance Scheduling* model.

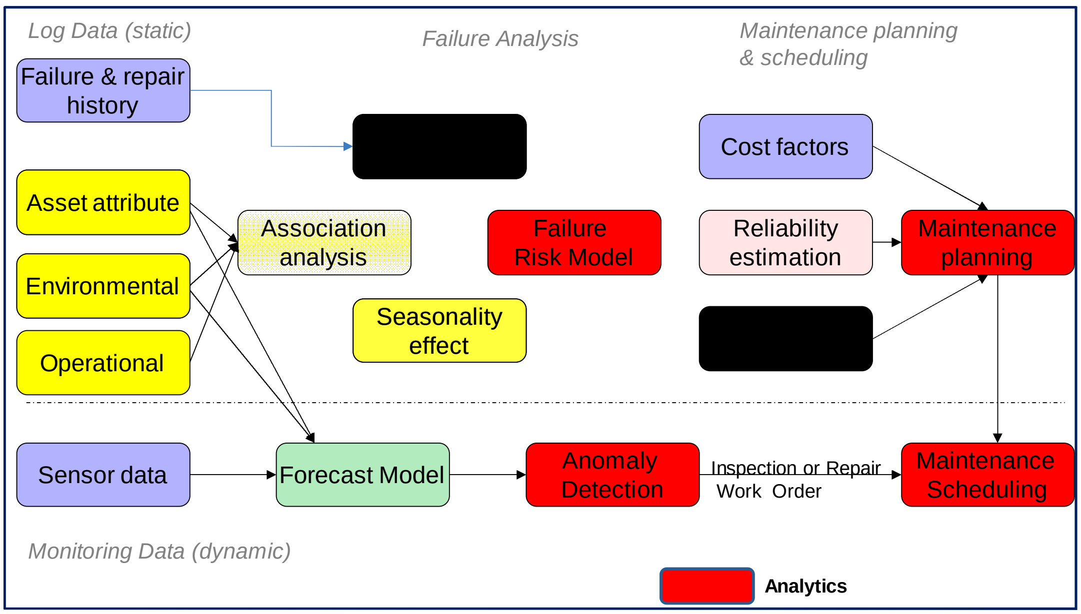

Figure 4: Analytics and Data Relationship

The ontology of analytics is implemented as a hierarchical menu for easy navigations in a web-based GUI as shown in the left side of Figure 5.

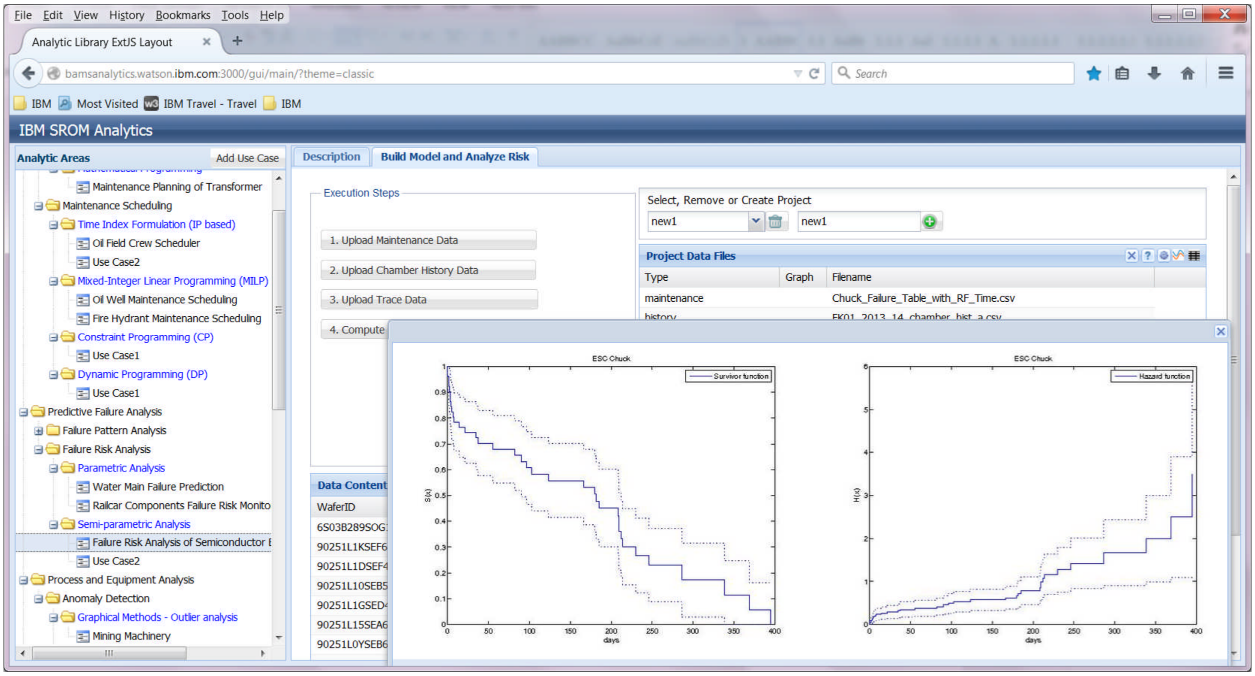

Figure 5: An Example for the Semantically defined SROM Analytics Solution

3 ARCHITECTURE OF SROM

The core of the SROM is the advanced analytics library that consists of analytics families such as *Predictive Asset Management*, *Predictive Project Analytics* and *Process Monitoring/Optimization* etc. as shown in Figure 6. A subset of this library can be packaged for a specific industry as an industry pack. For example, *industry pack – C&P* is a subset of SROM library for chemical and petroleum industry and *industry pack – E&U* is for energy and utility industry. The analytics library is semantic and ontology based so that common users can easily identify a specific analytics for a specific problem to solve by navigating through the hierarchy of the ontology.

The SROM analytics are supported by a computational layer, where computational engines are located. For the analytics that involve optimization and decision management, an optimization engine such as ILOG^{TM} Cplex, an optimization platform such as ILOG Decision Optimization Center (DOC) and decision management tool such as SPSS^{TM} DM can be utilized. For the analytics that involve statistical analyses, statistical engines such as SPSS Modeler, R, Python or Matlab can be accessed. For the analytics that involve business analyses, a business intelligence tool such as COGNOS^{TM} BI can be accessed. For the analytics that require visualization and discovery, Watson Explorer can be accessed for example.

The computational engines consume variety and vast amount (i.e., Big Data) of structured and unstructured data including video, text and real time sensor data, and this is supported by data storage layer. The data storage layer includes data repository such as HDFS, HBase, Hive, Spark for archiving large volumes of operational data and sensor data. It also includes database system such as Cassandra, DB2, Informix, PureData Analytics with pre-built data schema for storing quality, machine and prod data, configuration. For IoT application, real-time event processing capability such as InfoS phere Streams can also be utilized.

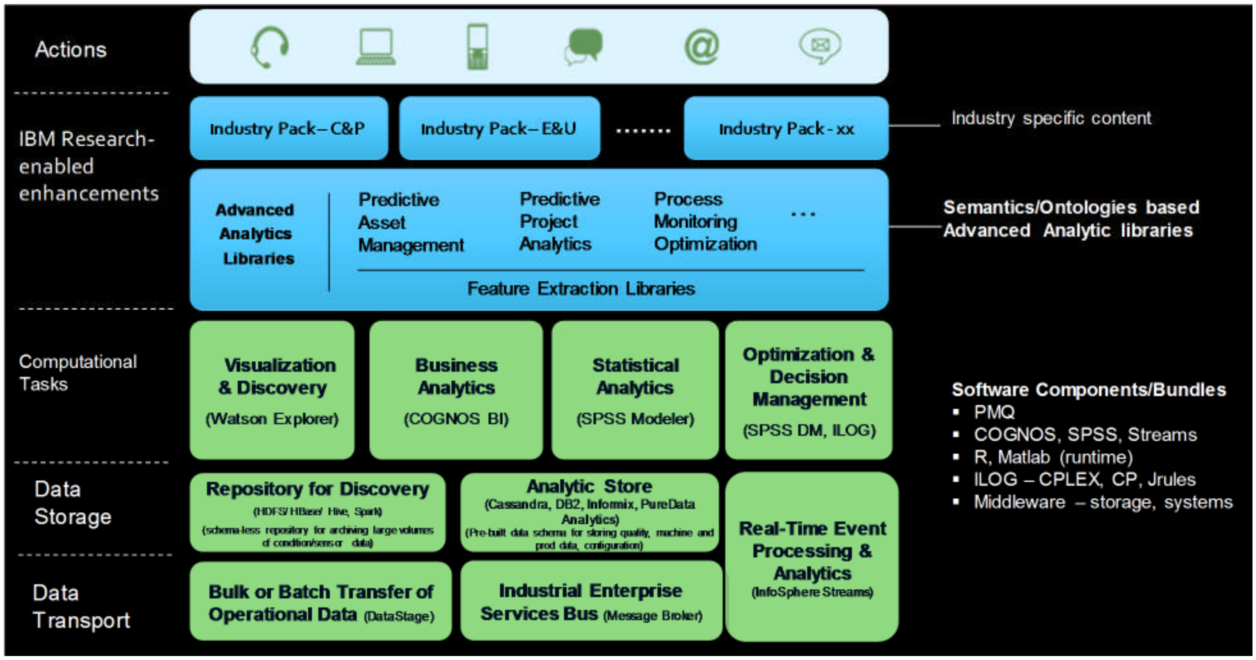

Figure 6: An Architecture for the Semantically defined SROM Analytics Solution

The data storage layer is supported by data transport layer which enables bulk or batch transfer of operational data for data stage, and industrial enterprise services bus, i.e., message broker.

SROM is a cloud-based solution, and the solution computed by the SROM analytics can be communicated to users through computer screen, mobile devices, email, social media and voice etc. in the action layer.

4 SROM WITH IOT AND BIG DATA

The Internet of Things (IoT) is a technological revolution in the future of computing and communication that is based on the concept of anytime, anyplace connectivity for anything [5]. IoT represents a network of Internet-enabled objects and devices, such as factory equipment, consumer electronics, home appliances, sensors, meters and smart phones. We are already using IoT solutions in practical ways in our daily lives. For example, we can monitor our home security, lights, heaters, air conditioners and status of other electric appliances from our smartphones. In manufacturing environment, even more data are generated in faster speed from various SCADA (supervisory control and data acquisition) systems, DCS (Distributed Control System), MES (Manufacturing Execution System) and other OT system as well as IT system etc. IoT affects almost all industries including automotive, chemical and petroleum, mining and metal and electronic.

Big Data is a phenomenon that is characterized by rapid accumulation of data from numerous sources. The data are generated and collected so quickly in such a way that it is flooding industries and society. The data are valuable, new raw materials in industries. We can utilize it to produce new services and products, and to improve the existing services and products. Therefore, Big Data represents both a challenge and an opportunity. The challenge is related to how this volume of data is harnessed, and the opportunity is related to how the effectiveness of industries and society is enhanced by properly analyzing and utilizing this information. It is now commonplace to distinguish Big Data solutions from conventional IT solutions by considering the following four dimensions [3]:

- Volume. Big data solutions must manage and process larger amounts of data.
- Velocity. Big data solutions must process more rapidly arriving data.
- Variety. Big data solutions must deal with more kinds of data, both structured and unstructured.
- Veracity. Big data solutions must validate the correctness of the large amount of rapidly arriving data.

Big data is being generated by everything around us at all times. Every digital process and social media exchange produces it. Systems, sensors and mobile devices transmit it. Big data is arriving from multiple sources at an alarming velocity, volume, variety and veracity.

Big data and the IoT are computing paradigms that, together, fundamentally change the nature of how we work, play, and interact with our environment. Where Big Data is all about volume, velocity, verity, and veracity, the IoT is about using that data in meaningful ways to improve productivity and quality of life [6]. However, IoT and Big Data by themselves are not useful. Only when the data are properly analyzed and converted to useful information, benefits of Big Data can be realized. Analytics are the key for extracting meaningful value from Big Data. Big Data is changing the way people within organizations work together. It is creating a culture in which business and IT leaders must join forces to realize value from all data. Insights from big data can enable all employees to make better decisions—deepening customer engagement, improving product quality, optimizing operations, preventing threats and fraud, and capitalizing on new sources of revenue [2]. But all these are possible only through effective use of analytics. Analytics is the collection of advanced mathematics, statistics, data mining, machine learning and cognitive computing. The intent of the Cross Industry Analytics Solution Library, e.g., SROM is to rapidly distribute the analytics capability throughout industries.

SROM (Smarter Resource and Operations Management) is a library and a development platform that contains a set of analytics that can solve variety of problems in the area of resource and operations management. The platform enables the common users without specialized training in mathematical modeling to easily identify, develop and use the analytics for solving complex industrial problems in resource and operations management with much less effort and time.

The description of SROM's platform with IoT and Bid Data are shown in Figure 7. The source of data for SROM includes smart devices such as smart phones and tablets, IT and OT systems from manufacturing plants such SCADA system, and other external sources for weather data, geographical data and traffic data etc. as shown in the bottom of Figure 7. These IoT data are serviced by IoT operation service layer, which includes IoT Data Infrastructure Services, IoT Services Composition and Lifecycle Services, IoT Device Connectivity/Management Services, IoT Platform Operation Analytics Services and IoT Security Intelligence Services. The data collected from IoT layer are aggregated and distributed in the Big Data layer, which include Hadoop Distributed File System (HDFS), Spark and MapReduce etc. The BigData layers is accessed by SROM analytics layers, which include the families of various analytics in resource and operations management. SROM analytic services are packaged into various industry solution pack such as mining and metals, chemical and petroleum, and electronics etc.

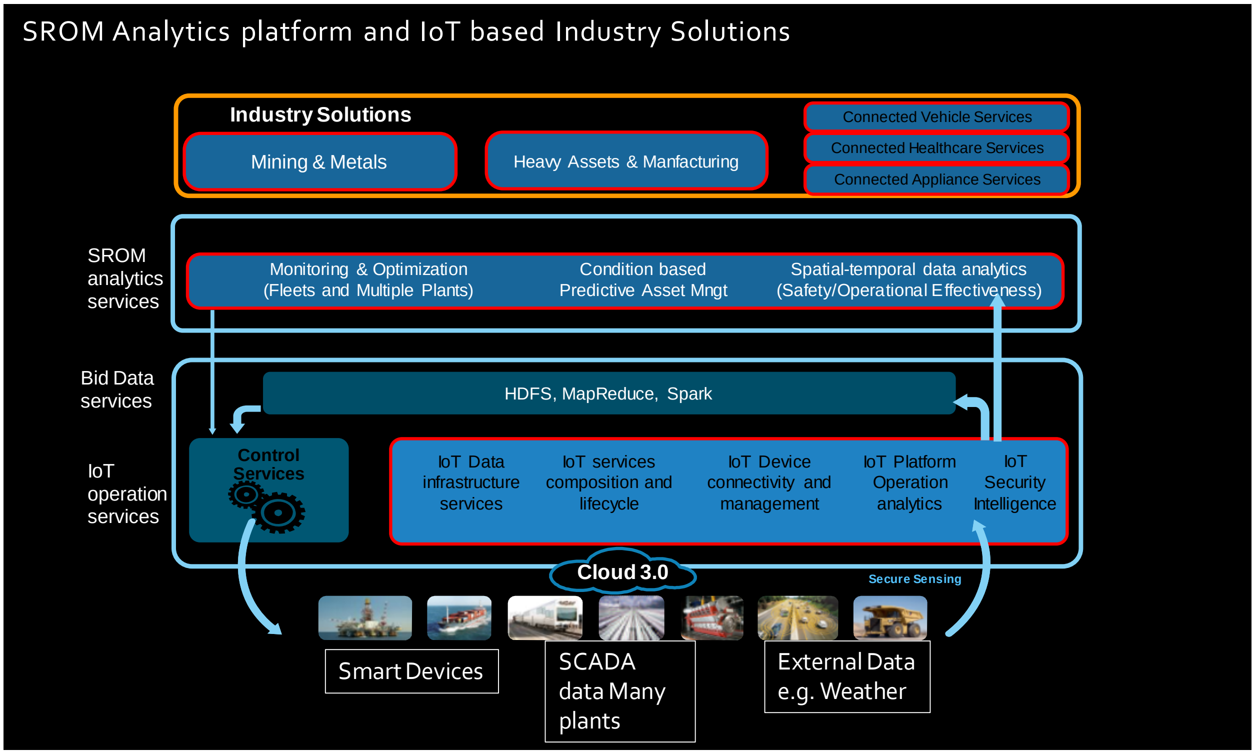

Figure 7: SROM Analytics platform and IoT/Big Data based Industry Solutions

5 DEPLOYMENT OF SROM

The analytics in SROM require substantially configured computing hardware and software including data interface programs, database, data integration and transformation, data warehouse, statistical engine, optimization engine, and visualization. Also depending upon the size of the data and amount of analysis there is a need to scale up and scale down the install base. Typical users may not afford such resources individually. The cloud is a secure, flexible and scalable way of delivering such analytic services to many users by sharing all the resources in a cost effective and secure way among many users. Running this kind of analytics service on cloud for multiple clients using the shared computing resource is more cost-efficient than providing the same service to individual clients with a traditional computing medium without the cloud. Depending on the number and roles of the users, an appropriate level of resource can be easily configured virtually on the cloud, and the level of access to resources can be adjusted on demand. Using cloud, all the analytics can be readily accessible by the users through web browsers without installing any special hardware, software, data interface and mathematical tools.

SROM is designed to be deployed as Analytics as a Service (AaaS) model. SROM is hosted as a service and is provided to customers across the Internet. Deployment of SROM on the cloud eliminates the need to install and run the analytics on the customer's own computers along with alleviating any burden to the customer with regard to software maintenance, ongoing operation and support. SROM cloud deployment is designed to provide rapid access to security-rich, enterprise-class virtual server environments, well suited for dynamic workloads. SROM cloud offers the capabilities to control access and configure security.

6 DESCRIPTIONS OF ANALYTICS IN SROM

A few selected analytics are described here. SROM is still under development and additional analytics and use cases are being added.

6.1 Maintenance Planning

Maintenance planning is targeted at providing a cost minimizing maintenance plan over the lifetime of an asset. This analysis considers the cost of repair, downtime and replacement to provide a repair/replace plan. The planning horizon is the lifetime of the asset and hence for high capital assets it could be decades. The maintenance planning is an optimization exercise that prescribes the optimal time for maintenance, repair and replacement. The planning exercise requires as input the failure risk analysis (i.e. the survival function) which provides the life expectancy of an asset based on its history and conditions of use.

Use Case: Maintenance Planning of Transformer

This optimization application computes an optimal preventive maintenance/replacement schedule for a set of assets under given constraints such as budget and labour. The input data for this models includes asset information, hazard function data and constraints such as budget and personnel availability for maintenance. This use case computes the optimal maintenance planning of transformers of an utility company. The optimal problem is formulated and solved as Non-Linear Programming (NLP) and Mixed Integer Linear Programming (MILP).

Figure 8: A Sample Output of a Maintenance Planning Optimization Analytics

6.2 Maintenance Scheduling

Maintenance scheduling is an operational scheduling effort with a horizon of weeks or months. Consider a plant with a portfolio of assets. Each asset has a maintenance plan over its life time. Now for any given month (or week) each asset can be examined to see if there is a prescribed inspection/repair work order. If we collect all the activities that need to be performed within the month then we have a problem of scheduling the work orders keeping in mind constraints re crews, vehicles, spare parts inventory and tools. This scheduling might also require considerations of the different locations and travel distances. The scheduling exercise is a prescriptive analysis that takes as input the work orders, the required resources, the time windows when the work order needs to be performed and generates a cost minimizing schedule for maintenance.

Use Case: Maintenance Scheduling of Oil Wells
This application computes optimized scheduling and routing of maintenance resources (e.g. operators, equipment, material etc.) for maintenance and inspection of oil wells so that maintenance requests can be fulfilled at minimal cost. The tool allows users to route and schedule people and resources to attend geographically sparse tasks. The user inputs the locations and durations of tasks, which may be different depending on the types of technicians and equipment which are used to complete the task. The tool then uses Mixed Integer Linear Programming (MILP) techniques to search for a solution which assigns as many tasks as possible within a given time frame, whilst minimizing travel distance and total time spent.

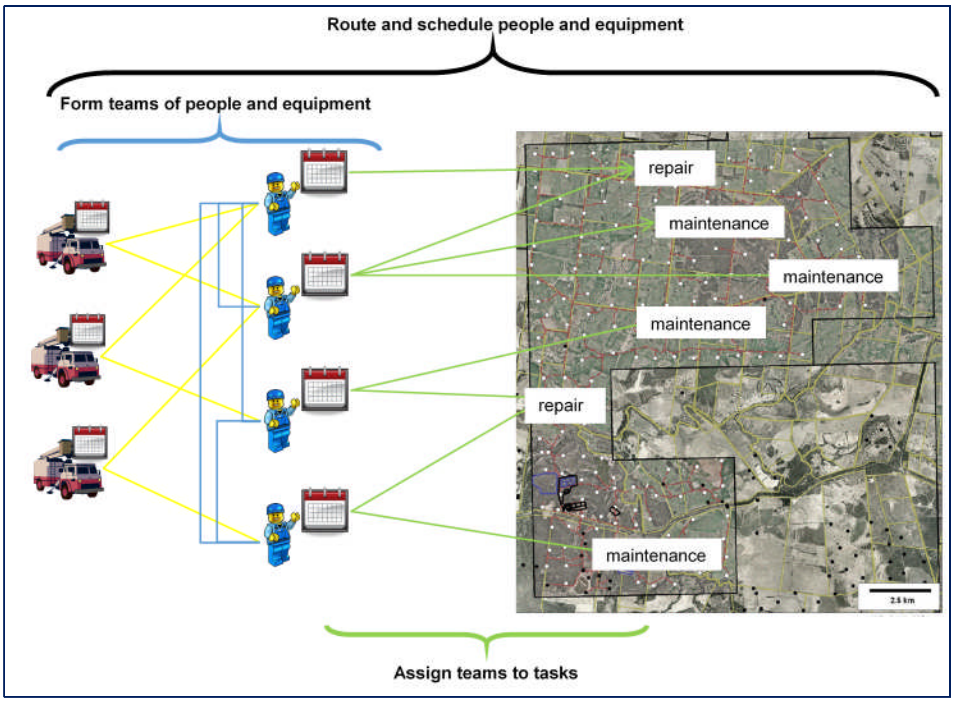

Figure 9: Overview of a Maintenance Scheduling Optimization Analytics

6.3 Failure Risk Analysis

Failure Risk Analysis estimates the life expectancy of an asset based on historical failure data for assets of that class. The life expectancy of an asset depends on its life history, which is characterized by its own attributes and the environmental and operational stress factors. A typical failure risk model provides the survival rate as a function of the various environmental, operational factors and its own attributes. The survival rate provides the probability that asset will survive a given duration into the future based on its own history. Survival analysis consists of parametric, semiparametric, and nonparametric methods. It is used to estimate the most commonly used measures in survival studies, survivor and hazard functions, compare them for different groups, and assess the relationship of predictor variables to survival time.

Use Case: Failure Risk Analysis of Semiconductor Electrostatic Chuck (ESC)

This application predicts the life expectancy of Electrostatic Chuck (ESC) in semiconductor manufacturing process. Based on historical maintenance data, operational history and sensor data from the manufacturing equipment, this tool predicts survivor/failure function, and identify what/how process variables impact the chuck failure, and identify the conditions leading to ESC maintenance or to ESC replacement. The risk analysis is then used to compute the optimal maintenance and replacement plan that minimize the cost of maintenance and failure. This analytics is based on Cox proportional hazards regression, which is a semiparametric method for adjusting survival rate estimates to quantify the effect of predictor variables. The method represents the effects of explanatory variables as a multiplier of a common baseline hazard function. The hazard function is the nonparametric part of the Cox proportional hazards regression function, whereas the impact of the predictor variables is a log linear regression.

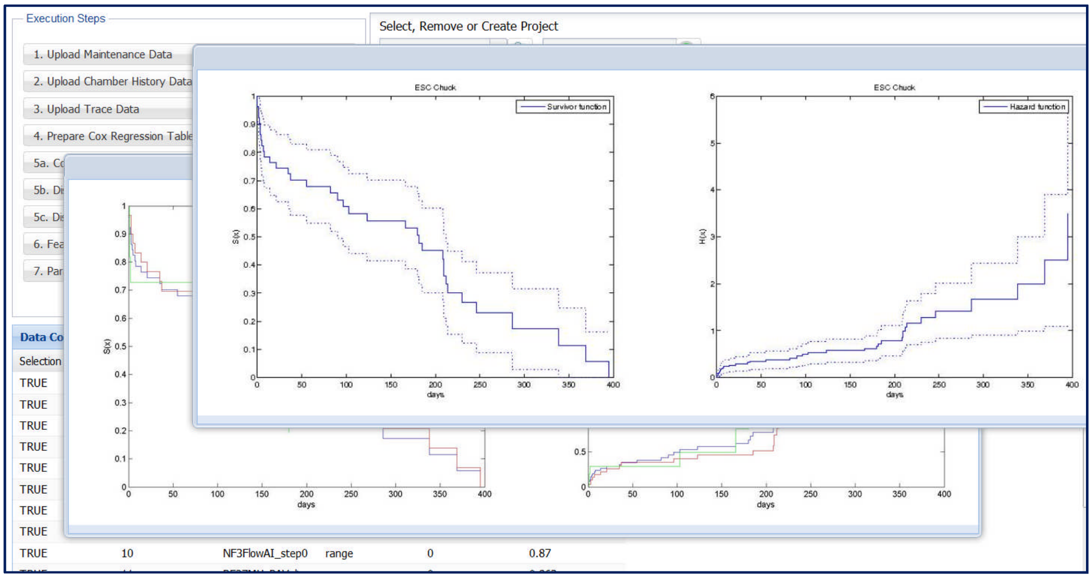

Figure 10: A Sample Output from a Failure Risk Analytics

6.4 Anomaly Detection

Anomaly Detection is the task to detect behaviors that look differently from normal operations. Typically anomaly detection needs three components: (1) A normal state model. In modern analytics, a probability distribution is used to model the normal state for practical application. (2) A definition of anomaly score, which is generally defined as a function of the probability distribution of the normal state. (3) A threshold to the anomaly score to decide whether or not the current observation or state is anomalous. Among several subtasks that can be defined based on anomaly detection, change detection is of particular practical importance for IoT sensors. In change detection, the sliding window techniques is typically used to define the past and the current states, and a score representing the excursion of the current state from the past is computed.

Use Case: Mining Machinery Anomaly Detection

The application detect anomalies from expensive mining machineries such as shearers in underground coal mining. In this mining operations, sensor data obtained from mining machineries are highly dynamic and traditional outlier detection approaches do not work. Due to the intractable nature, the monitoring solution had suffered from too many false alerts. To tackle the challenge, this use case leverages the technique of sparse structure learning and formalizes the anomaly detection task as quantification of structural changes of inter-variable dependencies, as represented by Figure 11. Here the highest anomaly score is attributed to a major structural change with respect to the 24th variable, whose dependency with 40th variable does not exist any more

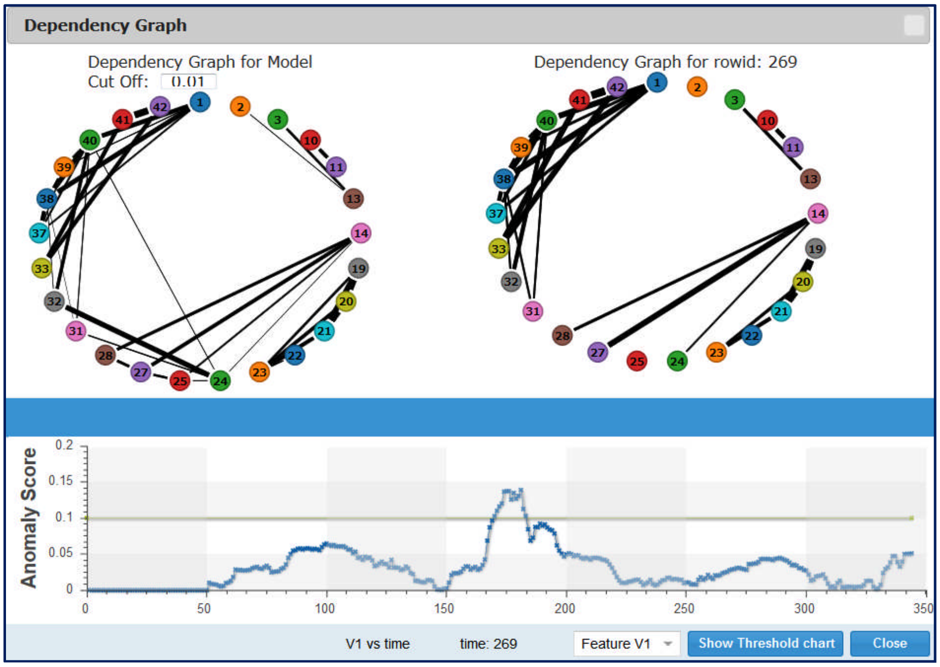

Figure 11: A Sample Output from a Failure Risk Analytics

6.5 Predictive Statistical Quality/Process Control

Predictive Statistical Quality/Process Control is a data mining approach for developing a model that predicts the quality of process outputs based on various process variables and using the prediction to enhance feedback control. Typical methods include ridge regress, lasso regression, customized prediction model, EWMA, causal modeling, neural networks and hierarchical modeling.

<u>Use Case: Virtual Metrology of Semiconductor Manufacturing</u>
This application is a model based prediction capability of process outcome when there is no physical measurement of that outcome. The underlying models are learned from histories of the actual physical outcomes and process trace data. The tool helps in detecting faulty product early and improving process control, and reduce physical measurements for process monitoring and control. This use case is for prediction of plasma oxide deposition rate of a semiconductor fabrication process.

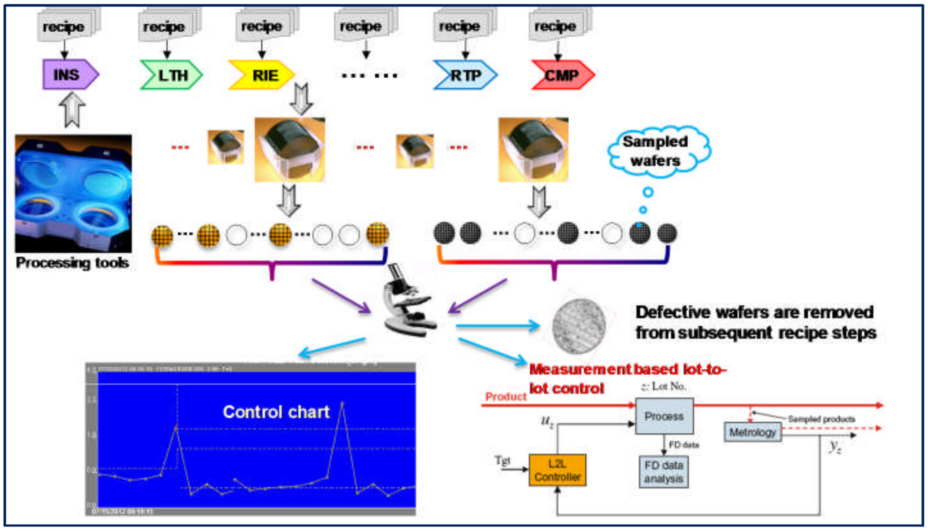

Figure 12: An Illustration of a Predictive Statistical Quality/Process Control Analytics

6.6 Model Predictive Control

Model Predictive Control (MPC) is an optimization method of process control that are typically used continuous processes such as chemical and petroleum plants and power systems. MPC uses dynamic models of the process, typically physics-based models or data-driven model that are calibrated using dynamic sensor data. The optimal control problem is typically solved as a Nonlinear Programming (NLP). MPC predicts future events and computes the optimal control profile accordingly.

<u>Use Case: Optimal Control of HVAC System</u>

This is a model predictive control (MPC) that optimally determines control profiles of the HVAC (Heating Ventilation and Air Conditioning) system as demand response. The thermal behavior of the building zone is modelled by a Nonlinear Autoregressive Neural Network (NARNET) and the optimal control problem is formulated as a Mixed-Integer Non-Linear Programming (MINLP) model. The optimal control model minimizes the total costs of powering HVAC system and corresponding GHG emission considering dynamic demand response signal, on-site energy storage system and on-site energy generation system while satisfying thermal comfort of building occupants within the physical limitation of HVAC equipment and physical limitation of on-site energy storage system and on-site energy generators.

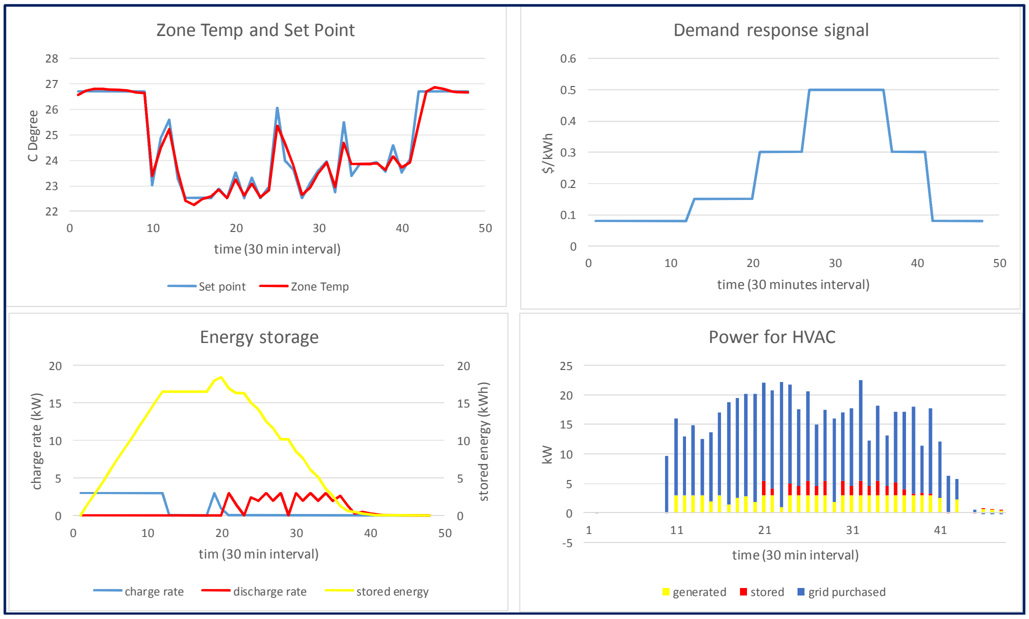

Figure 13: A Sample Output from a Model Predictive Control Analytics

6.7 Causal Modeling

Causal Modeling is a quantitative representation of cause and effect of real-world dynamics. A causal model describes the causal and other relationships, among a set of variables. Causal models incorporate the idea of multiple causality, that is, there can be more than one cause for any particular effect.

Use Case: Causal Analysis of IT Operations

This tool monitors, detects and diagnoses anomalies in computer networks and systems. With the prevalence of IT data monitoring techniques and the surge of IT operations data available for analysis, there are increasing expectations on the many benefits of analytics based operations management, such as increasing operational efficiencies and understanding problems that have occurred. This analytics can identify problems before they interrupt services.. Additionally, it is hoped to be able to assist not only in identifying anomalous behaviors that deviate from the baseline, but in providing root cause analysis, and hence proactive failure prevention

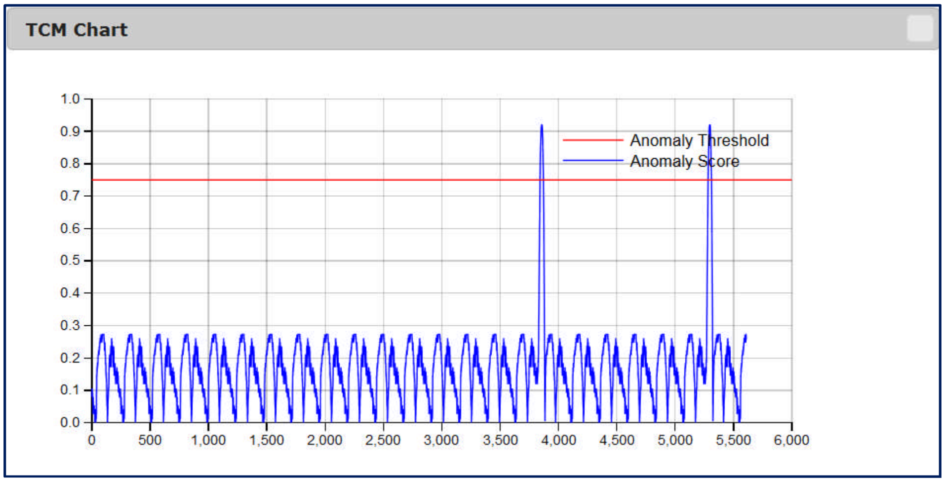

Figure 14: A Sample Output from a Causal Modeling Analytics

7 CONCLUSION

Analytics is the engine for converting the world's new natural resource, data, into good products and services in industries. Without analytics, the large amount of data we collect is simply data, i.e. underutilized natural resources. By utilizing the data properly, manufacturing firms can improve product quality, reduce production costs, improve safety in manufacturing plants, improve energy efficiency of production, and improve the profit margin. Analytics can bring competitive advantage. The data can also be used to improve the quality of human lives by promoting smarter buildings, smarter cities and smarter planet. The analytics in the new era of computing is complex aggregation of various algorithms and techniques including advanced mathematics, statistics, data mining, machine learning, cognitive computing and optimization. Therefore, it takes a lot of expertise, effort and time to develop and use analytics in industries. There is an urgent need to develop and make available analytic solutions and services libraries in industry that can help firms utilize the analytics easily without significant investment in resources, expertise and time. The cross industry analytics solution & services library for resource and operations management we introduce here is intended to spark the engine.

REFERENCES

[1] IDC's sixth annual study of the digital universe, 2012. http://www.emc.com/leadership/digital-universe/2012iview/big-data-2020.htm
[2] Big Data, Data savvy. Insight driven. http://www.ibm.com/big-data/us/en/
[3] IBM Redbook, Performance and Capacity Implications for Big Data, http://www.redbooks.ibm.com/redpapers/pdfs/redp5070.pdf
[4] IBM Redbook on The Interconnecting of Everything, http://www.redbooks.ibm.com/redpapers/pdfs/redp4975.pdf99
[5] ITU Internet Reports 2005: The Internet of Things: Executive Summary: http://www.itu.int/osg/spu/publications/internetofthings/InternetofThings_summary.pdf
[6] Big Data at the Speed of Business: What is big data: http://www.ibm.com/software/data/bigdata.