Analyzing Analytics

Rajesh Bordawekar
IBM Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598
bordaw@us.ibm.com

Bob Blainey
IBM Toronto Software Lab
8200 Warden Avenue
Markham, Ontario L6G 1C7
blainey@ca.ibm.com

Chidanand Apte
IBM Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598
apte@us.ibm.com

ABSTRACT

Many organizations today are faced with the challenge of processing and distilling information from huge and growing collections of data. Such organizations are increasingly deploying sophisticated mathematical algorithms to model the behavior of their business processes to discover correlations in the data, to predict trends and ultimately drive decisions to optimize their operations. These techniques, are known collectively as *analytics*, and draw upon multiple disciplines, including statistics, quantitative analysis, data mining, and machine learning.

In this survey paper, we identify some of the key techniques employed in analytics both to serve as an introduction for the non-specialist and to explore the opportunity for greater optimizations for parallelization and acceleration using commodity and specialized multi-core processors. We are interested in isolating and documenting repeated patterns in analytical algorithms, data structures and data types, and in understanding how these could be most effectively mapped onto parallel infrastructure. To this end, we focus on analytical models that can be executed using different algorithms. For most major model types, we study implementations of key algorithms to determine common computational and runtime patterns. We then use this information to characterize and recommend suitable parallelization strategies for these algorithms, specifically when used in data management workloads.

1. ANALYTICS AT YOUR SERVICE

From streaming news updates on smart-phones, to instant messages on micro-blogging sites, to posts on social network sites, we are all being overwhelmed by massive amounts of data [35, 29]. Access to such a large amount of diverse data can be of tremendous value if useful *information* can be extracted and applied rapidly and accurately to a problem at hand. For instance, we could contact all of our *nearby* friends for a dinner at a local *mutually agreeable* and *well-reviewed* restaurant that has advertised *discounts* and *table availability* for *that night*; but finding and organizing all that information in a short period of time is very challenging. Similar opportunities exist for businesses and governments, but the volume, variety and velocity of data can be far greater. This process of identifying, extracting, processing, and integrating *information* from *raw data*, and then applying it to solve a problem is broadly referred to as *analytics*.

Table 1 presents a sample of analytic applications from different domains, along with their functional characteristics. As this table illustrates, many services that we take for granted and use extensively in everyday life would not be possible without analytics. For example, social networking applications such as Facebook, Twitter, and LinkedIn encode social relationships as graphs and use graph algorithms to identify hidden patterns (e.g., finding common friends). Other popular applications like Google Maps, Yelp or FourSquare combine location and social relationship information to answer complex spatial queries (e.g., finding the nearest restaurant of a particular cuisine that your friends like). Usage of analytics has substantially improved the capabilities and performance of gaming systems as demonstrated by the recent win of IBM's Watson/DeepQA intelligent question-answer system over human participants in the Jeopardy challenge [31]. The declining cost of computing and storage and the availability of such infrastructure in *cloud* environments has enabled organizations of any size to deploy advanced analytics and to package those analytic applications for broad usage by consumers.

While consumer analytical solutions may help us all to better organize or enrich our personal lives, the analytic process is also becoming a critical capability and competitive differentiator for modern businesses, governments and other organizations. In the current environment, organizations need to make on-time, informed decisions to succeed. Given the globalized economy, many businesses have sup-

Application (Domain)	Principal Goals	Key Functional Characteristics
Netflix and Pandora [3, 19] (Consumer)	Video and music recommendation	Analyzing structured and unstructured data, Personalized recommendations
Yelp and FourSquare (Consumer)	Integrated geographical analytics	Spatial queries/ranking, Streaming and persistent data
DeepQA (Watson) [12] (HealthCare/Consumer)	Intelligent question-answer (Q/A) System	Real-time natural language, Unstructured data processing, Artificial intelligence techniques for result ranking
Telecom churn analysis [28] (Telecom)	Analysis of call-data records	Graph modeling of call records, Large graph dataset, Connected component identification
Fraud analytics (Insurance/HealthCare)	Detection of suspicious behavior	Identification of abnormal patterns Real-time data analysis over streaming and persistent data
Cognos consumer insight [30] Twitter sentiment [13] (Hospitality)	Sentiment/Trend analysis	Processing large corpus of text documents, Extraction transformation, Text indexing, Entity extraction
UPS [1], Airline scheduling [20] (Logistics)	Transportation routing	Mathematical programming solutions for transportation
Integrated supply Chain (Resource Planning)	Maximize end-to-end efficiencies	Mathematical solutions for optimizing under multiple constraints
Salesforce.com (Marketing)	Customer data analytics	Reporting, Text search, Multi-tenant support, Automated price determination, Recommendation
Quantitative Trading (Finance)	Identify trading opportunities	Identification of patterns in high-speed data Statistical modeling of financial instruments
Moody's, Fitch, and S&P [7, 6] (Finance)	Financial credit rating	Statistical analysis of large historical data
Amazon retail analysis (Retail)	End-to-end retail management	Analysis over large persistent and transactional data, Integration with logistics and customer information
Energy trading (Energy)	Determining prices	Processing large time-series data, Integrated stochastic models for generation, storage and transmission
Splunk [32] (Enterprise)	System management analysis	Text analysis of system logs, Large data sets
Flickr and Twitter, Facebook and Linkedin (Consumer/Enterprise)	Social network analysis	Graph modeling of relations, Massive graph datasets, Graph analytics, Multi-media annotations and indexing
Voice of customer analytics [4] (Enterprise)	Analyzing customer voice records	Natural language processing, Text entity extraction
Workforce Analytics (Enterprise)	Intelligent staffing	Human resource matching, Intelligent work assignments
Genomics (Medical/BioInformatics)	Genome analysis, sequencing, and assembly	Analysis of large text sequences, Processing of large graphs
fMRI analysis (Medical)	Analyzing synaptic activities	Graph modeling, Graph analytics
Facial recognition [33] (Government)	Biometric classification	Analysis and matching of 2-/3-D images, Large data sets
Predictive policing [21] (Government)	Crime prediction	Spatial and temporal analytics of iamges and streams

Table 1: Well-known analytics solutions and their key characteristics

ply chains and customers that span multiple continents. In the public sector, citizens are demanding more access to services and information than ever before. Huge improvements in communication infrastructure have resulted in wide-spread use of online commerce and a boom in smart, connected mobile devices. More and more organizations are run around the clock, across multiple geographies and time zones and those organizations are being *instrumented* to an unprecedented degree. This has resulted in a deluge of data that can be studied to harvest valuable information and make better decisions. In many cases, these large volumes of data must be processed rapidly in order to make timely decisions. Consequently, many organizations have employed analytics to help them decide what kind of data they should collect, how this data should be analyzed to glean key information, and how this information should be used for achieving their organizational goals. Examples of such techniques can be found in almost any sector of the economy, including financial services [7, 6], government [33, 14], healthcare, retail [28, 24], manufacturing, logistics [1, 20], hospitality, and eCommerce [8, 9].

1.1 Characterizing Analytics Workloads

The distinguishing feature of an analytics application is the use of mathematical formulations for modeling and processing the raw data, and for applying the extracted information [34]. These techniques include statistical approaches, numerical linear algebraic methods, graph algorithms, relational operators, and string algorithms. In practice, an analytics application uses multiple formulations, each with unique functional and runtime characteristics (Table 1). Further, depending on the functional and runtime constraints, the same application can use different algorithms. While many of the applications process a large volume of data, the type of data processed varies considerably. Internet search engines process unstructured text documents as input, while retail analytics operate on structured data stored in relational databases. Some applications such as Google Maps, Yelp, or Netflix use both structured and unstructured data. The velocity of data also differs substantially across analytics applications. Search engines process read-only historical data whereas retail analytics process both historical and transactional data. Other applications, such as the monitoring of medical instruments, work exclusively on real-time or streaming data. Depending on the mathematical formulation, the volume and velocity of data and the expected I/O access patterns, the data structures and algorithms used by analytical applications vary considerably. These data structures include vectors, matrices, graphs, trees, relational tables, lists, hash-based structures, and binary objects. They can be further tuned to support in-memory, out-of-core, or streaming execution of the associated algorithm. Thus, analytics applications are characterized by diverse requirements, but share a common focus on the application of advanced mathematical modelling, typically on large data sets.

1.2 Systems Implications

Although analytics applications have come of age, they have not yet received significant attention from the data management and systems communities. It is important to understand systems implications of the analytics applications, not only because of their diverse and demanding requirements, but also, because systems architecture is currently undergoing a series of disruptive changes. Wide-spread use of technologies such as multi-core processors, specialized co-processors or accelerators, flash memory-based solid state drives (SSDs), and high-speed networks has created new optimization opportunities. More advanced technologies such as phase-change memory are on the horizon and could be game-changers in the way data is stored and analyzed.

In spite of these trends, currently there is limited usage of such technologies in the analytics domain. Even in the current implementations, it is often difficult for analytics solution developers to fine-tune system parameters, both in hardware and software, to address specific performance problems. Naive usage of modern technologies often leads to unbalanced solutions that further increase optimization complexity. Thus, to ensure effective utilization of system resources: CPU, memory, networking, and storage, it is necessary to evaluate analytics workloads in a holistic manner.

1.3 Our Study

In this paper, we aim to understand the application of modern systems technologies to optimizing analytics workloads by exploring the interplay between overall system design, core algorithms, software (e.g., compilers, operating system), and hardware (e.g., networking, storage, and processors).
Specifically, we are interested in isolating repeated patterns in analytical applications, algorithms, data structures, and data types, and using them to make informed decisions on systems design. Over the past two years, we have been examining the functional flow of a variety of analytical workloads across multiple domains (Table 1), and as a result of this exer-

cise, we have identified a set of commonly-used analytical models, called *analytics exemplars* [5]. We believe that these exemplars represent the essence of analytical workloads and can be used as a toolkit for performing exploratory systems design for the analytics domain. We use these exemplars to illustrate that analytics applications benefit greatly from holistically co-designed software and hardware solutions and demonstrate this approach using the Netezza [11] appliance as an example. In spite of the recent efforts in integrating analytics components into database systems, a lot of work still needs to be done [25, 15, 11], in particular, for accelerating analytics workloads within the context of database systems. We hope this study acts as a *call to action* for researchers to focus future data management and systems research on analytics.

2. ANATOMY OF ANALYTICS WORKLOADS

To motivate the study of analytics workloads, we first describe in detail a recent noteworthy analytics application: the Watson intelligent question/answer (Q/A) system [12].

2.1 The Watson DeepQA System

Watson is a computer system developed to play the Jeopardy! game-show against human participants [31]. Waston's goals are to correctly interpret the input natural language questions, accurately predict answers to the input questions and finally, intelligently choose the input topics and the wager amounts to maximize the gains. Watson is designed as an open-domain Q/A system using the DeepQA system, a probabilistic evidence-based software architecture whose core computational principle is to assume and pursue multiple interpretations of the input question, to generate many plausible answers or hypotheses and to collect and evaluate many different competing evidence paths that might support or refute those hypotheses through a broad search of large volumes of content.

This process is accomplished using multiple stages: the first, question analysis and decomposition stage parses the input question and analyzes it to detect any semantic entities like names or dates. The analysis also identifies any relations in the question using pattern-based or statistical approaches. Next, using this information, a keyword-based primary search is performed over a varied set of sources, such as natural language documents, relational databases and knowledge bases, and a set of supporting passages (initial evidence) is identified. This is followed by the candidate (hypothesis) generation phase which

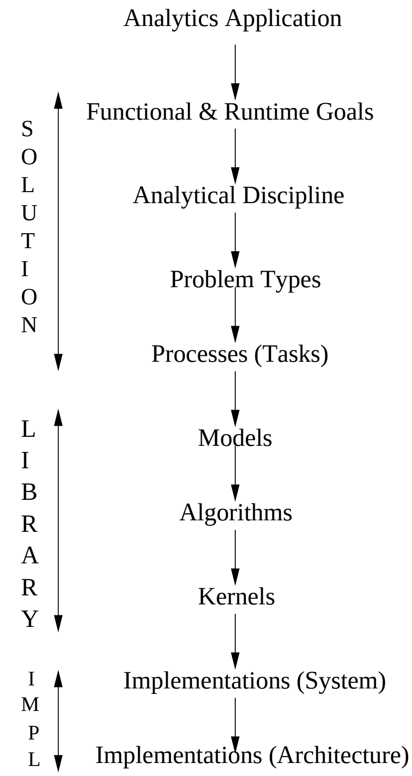

Figure 1: Simplified functional flow of business analytics applications

uses rule-based heuristics to select a set of candidates that are likely to be the answers to the input question. The next step, Hypothesis and Evidence Scoring, for each evidence-hypothesis pair, applies different algorithms that dissect and analyze the evidence along different *dimensions of evidence* such as time, geography, popularity, passage support, and source reliability. The end result of this stage is a ranked list of candidate answers, each with a confidence score indicating the degree to which the answer is believed correct, along with links back to the evidence. Finally, these evidence features are combined and weighted by a logistic regression to produce the final confidence score that determines the successful candidate (i.e. the correct answer). In addition to finding correct answers, Watson needs to master the strategies to select the clues to it's advantage and bet the appropriate amount for any given situation. The DeepQA system models different scenarios of the Jeopardy! game using different simulation approaches (e.g., Monte Carlo techniques) and uses the acquired insights to maximize Watson's winning chances by guiding topic selection, answering decisions and wager selections.

2.2 Functional Flow of Analytics Applications

The Watson system displays many traits that are common across analytics applications. They all have one or more functional goals. These goals are accomplished by one or more multi-stage processes, where each stage is an independent analytical component. To study the complex interactions between these components, it is useful to examine the functional flow of an analytics application from the customer usage to implementation stages. As Figure 1 illustrates, execution of an analytics application can be partitioned into three main phases: (1) solution, (2) library, and (3) implementation.

2.2.1 The Solution Phase

The solution phase is end-user focused and customized to to satisfy user's functional goals, which can be one of the following: prediction, prescription, reporting, recommendation, quantitative analysis, simulation, pattern matching, or alerting[1]. For example, Watson's key functional goals are: *pattern matching* for input question analysis, *prediction* for choosing answers, and *simulation* for wager and clue selection. Usually, any functional goal needs to be achieved under certain runtime constraints, e.g., calculations to be completed within a fixed time period, processing very large datasets or large volumes of data over streams, supporting batch or ad-hoc queries, or supporting a large number of concurrent users. For example, for a given clue, the Watson system is expected to find an answer before any of the human participants in the quiz. To achieve the functional and runtime goals of an application, the analytical solution leverages well-known analytical disciplines such as machine learning, data mining, statistics, business intelligence, and numerical analysis. Specifically, for a given analytical problem, the solution chooses appropriate problem types from these disciplines to build processes. Examples of analytic problem types include supervised and unsupervised learning, optimization, structured and unstructured data analysis, inferential and descriptive statistics, and modeling and simulation.

Table 2 presents a set of analytics applications along with their functional goals and the analytic problem types used to achieve these goals. As illustrated in Table 2, in many cases, a functional goal can be achieved by using more than one problem types. The choice of the problem type to be used depends on many factors that include runtime constraints, underlying software and hardware infrastructure, etc. For example, customer churn analysis is a technique for predicting the customers that are most likely to leave the current service provider (retail, telecom or financial) for a competitor. This analysis can use one of the three problem types: inferential statistics, supervised learning or unstructured data analysis. One approach models individual customer's behavior using various parameters such as duration of service, user transaction history, etc. These parameters are then fed either to a statistical model such as regression or to a supervised learning model such as a decision tree, to predict if a customer is likely to defect [22]. The second approach, models behavior of a customer based on her interactions with other customers. This strategy is commonly used in the telecom sector, where customer calling patterns are used to model subscriber relationships as a graph. This unstructured graph can then be analyzed to identify subscriber groups and their influential leaders: usually the active and well-connected subscribers. These leaders can then be targeted for marketing campaigns to reduce defection in the members of her group [23].

[1]We have expanded the classification proposed by Davenport et al. [8, 9].

2.2.2 The Library Phase

The library component is usually designed to be portable and broadly applicable across multiple analytic solutions (e.g., the DeepQA runtime that powers the Watson system). A library usually provides implementations of specific models of the common problem types. For example, an unsupervised learning problem can be solved using one of many models including associative mining, classification, or clustering [16]. Each model can, in turn, use one or more algorithms for its implementation. For instance, the associative rule mining model can be implemented using the different associative rule mining or decision tree algorithms. Similarly, classification can be implemented using nearest-neighbor, neural network, or naive Bayes algorithms. It should be noted that in practice, the separation between models and algorithms is not strict and many times, an algorithm can be used for supporting more than one models. For instance, neural networks can be used for clustering or classification.

2.2.3 The Implementation Phase

Finally, depending on how the problem is formulated, each algorithm uses specific data structures and kernels. For example, many algorithms formulate the problem using dense or sparse matrices and invoke kernels like matrix-matrix and matrix-vector multiplication, matrix factorization, and linear system solvers. These kernels are sometimes

Analytical applications	Functional goals	Problem types
Supply chain management, Product scheduling, Logistics, Routing, Workforce management	Prescription	Optimization
Revenue prediction, Disease spread prediction, Semiconductor yield analysis, Predictive policing	Prediction	Unsupervised/Supervised learning Descriptive/Inferential statistics
Retail sales analysis, Financial reporting, Budgeting, System management analysis, Social network analysis	Reporting	Structured/Unstructured data analysis
VLSI sensitivity analysis, Insurance risk modeling, Credit risk analysis, Physics/Biology simulations, Games	Simulation	Modeling and simulation Descriptive/Inferential statistics
Topic/Sentiment analysis, Computational chemistry, Document management, Searching, Bio-informatics	Pattern matching	Structured/Unstructured data analysis Unsupervised/Supervised learning
Cross-sale analysis, Customer retention, Music/Video, Restaurant recommendation, Intrusion detection	Recommendation	Unsupervised/Supervised Learning Structured/Unstructured data analysis
Web-traffic analysis, Fraud detection, Geological Sensor networks, Geographical analytics (Maps)	Alerting	Descriptive/Inferential statistics Unsupervised/Supervised learning
Customer relationship analysis, Weather forecasting Econometrics, Computational finance	Quantitative analysis	Descriptive/Inferential statistics Unsupervised/Supervised learning

Table 2: Examples of analytics applications, associated functional goals, and analytical problem types

optimized for the underlying system architecture, in form of libraries such as IBM ESSL [17] or Intel MKL [18]. Any kernel implementation can be characterized according to how it manages parallel execution, if at all, and how it manages data and maps it to the system memory and I/O architecture. Many parallel kernels can use shared or distributed memory parallelism. In particular, if the algorithm is embarrassingly parallel, requires large data, and the kernel is executing on a distributed system, it can often use the MapReduce approach [10]. At the lowest level, the kernel implementation can often exploit hardware-specific features such as short-vector data parallelism (SIMD) or task parallelism on multi-core CPUs, massive data parallelism on GPUs, and application-specific parallelism using Field Programmable Gate Arrays (FPGAs).

3. ANALYTICS EXEMPLARS

Given the wide variety of algorithmic and system alternatives for executing analytics applications, it is difficult for solution developers to make the right choices to address specific performance issues. To alleviate this problem, we have analyzed the functional flow (Figure 1) of a wide set of key applications across multiple analytics domains and have isolated repeated patterns in analytical applications, algorithms, data structures, and data types. We have been using this information to optimize analytic applications and libraries for modern systems and in some cases, specialize our processor and system designs to better suit analytic applications.

Towards this goal, we have identified a set of widely-used analytical models that capture the most important computation and data access patterns of the analytics applications that we have studied [5, 27]. These models, referred to as *Analytics Exemplars*, cover the prevalent analytical problem types and each exemplar can be used to address one or more functional goals. Table 3 presents the list of thirteen exemplars, along with target functional goals and key algorithms used for implementing these exemplars.

3.1 Key Algorithms

As Table 3 illustrates, each exemplar can be implemented by one or more distinct algorithms. Some of the algorithms can be used for implementing more than one exemplars, e.g., the Naive Bayes algorithm can be used in text analytics and for general clustering purposes. Each algorithm, depending on the runtime constraints, i.e., whether the application data can fit into main memory or not, can use a variety of algorithmic kernels (Figure 1). For more details on the algorithms and their implementations, the reader is referred to [5, 36].

3.2 Computational Patterns

Table 4 presents a summary of computational patterns, key data types, data structures and functions used by algorithms for each exemplar. As Table 4 illustrates, while different exemplars demonstrate distinct computational and runtime characteristics, they also exhibit key similirities. Broadly, the analytic exemplars can be classified into two classes: the first class exploits linear-algebraic formulations and the second uses non-numeric data structures (e.g., hash tables, trees, bit-vectors, etc.). Exemplars belonging to the first class, e.g., Math-

Model Exemplar (Problem type)	Functional goals	Key algorithms
Regression analysis (Inferential statistics)	Prediction Quantitative Analysis	Linear, Non-linear, Logistic Probit regression
Clustering (Supervised learning)	Recommendation, Prediction, Reporting	K-Means and Hierarchical clustering EM Clustering, Naive Bayes
Nearest-neighbor search (Unsupervised learning)	Prediction, Recommendation	K-d, Ball, and Metric trees, Approx. Nearest-neighbor Locality-sensitive Hashing, Kohonen networks
Association rule mining (Unsupervised learning)	Recommendation	Apriori, Partition, FP-Growth, Eclat and MaxClique, Decision trees
Neural networks (Supervised learning)	Prediction Pattern matching	Single- and Multi-level perceptrons, RBF, Recurrent, and Kohonen networks
Support Vector Machines (Supervised learning)	Prediction Pattern matching	SVMs with Linear, Polynomial, RBF, Sigmoid, and String kernels
Decision tree learning (Supervised learning)	Prediction Recommendation	ID3/C4.5, CART, CHAID, QUEST
Time series processing (Data analysis)	Pattern matching, Alerting	Trend, Seasonality, Spectral analysis, ARIMA, Exponential smoothing
Text analytics (Data analysis)	Pattern matching Reporting	Naive Bayes classifier, Latent semantic analysis, String-kernel SVMs, Non-negative matrix factorization
Monte Carlo methods (Modeling and simulation)	Simulation Quantitative analysis	Markov-chain, Quasi-Monte Carlo methods
Mathematical programming (Optimization)	Prescription Quantitative analysis	Primal-dual interior point, Branch & Bound methods, Traveling salesman, A* algorithm, Quadratic programming
On-line analytical processing (Structured data analysis)	Reporting Prediction	Group-By, Slice_and_Dice, Pivoting, Rollup and Drill-down, Cube
Graph analytics (Unstructured data analysis)	Pattern matching Recommendation	Eigenvector Centrality, Routing, Coloring, Searching and flow algorithms, Clique and motif finding

Table 3: Analytics exemplar models, along with problem types and key application domains

ematical Programming, Regression Analysis, and Neural Networks, operate primarily on matrices and vectors. Matrices are either sparse or dense, and are used in various linear algebraic kernels like the matrix multiplication, inversion, transpose, and factorization. The second class, which includes clustering, nearest-neighbor search, associative rule mining, decision tree learning, use data structures like hash-tables, queues, graphs, and trees, and operate on them using set-oriented, probabilistic, graph-traversal, or dynamic programming algorithms. Exemplars like mathematical programming, text analytics, and graph analytics can use either of these approaches. The analytic exemplars use a variety of types, such as integers, strings, bit-vector, and single and double precision floats, to represent the application data. This information is then processed using different functions that compare, transform, and modify input data. Examples of common analytic functions include various distance functions (e.g., Euclidian), kernel functions (e.g., Linear, Sigmoid), aggregation functions (e.g., Sum), and Smoothing functions (e.g., correlation). These functions, in turn, make use of intrinsic library functions such as `log`, `sine` or `sqrt`.

3.3 Runtime Characteristics

Table 5 summarizes the runtime characteristics of the analytics exemplars. The key distinguishing feature of analytics applications is that they usually process input data in read-only mode. The input data can be scalar, structured with one or more dimensions, or unstructured, and is usually read from files, streams or relational tables in the binary or text format. In most cases, the input data is large, which requires analytics applications to store and process data from disk. Notable exceptions to this pattern are Monte Carlo Methods and Mathematical Programming, which are inherently in-memory as they operate on small input data. The results of analysis are usually smaller than the input data. Only two exemplars, association rule mining and on-line analytical processing (OLAP) generate larger output. Finally, analytics applications can involve one or more stages (real-time execution can be considered to have only one stage), where each stage invokes the corresponding algorithm in an iterative or non-iterative manner. For the iterative workloads, for the same input data size, the running time can vary depending on the precision required in the results.

4. SYSTEM IMPLICATIONS

Given the varied computational and runtime characteristics of the analytics exemplars, it is clear that a single systems solution for different analytics ap-

Model Exemplar	Computational pattern	Key data types, Data structures, and Functions
Regression Analysis	Matrix inversion, LU decomposition Transpose, Factorization	Double-precision and Complex data Sparse/Dense matrices, Vectors
Clustering	Metric-based iterative convergence	Height-balanced tree, Graph, Distance functions, `log` function
Nearest-Neighbor Search	Non-iterative distance calculations Singular value decomposition, Hashing	Higher-dimensional data structures, Hash tables, Distance functions
Rule Mining	Set intersections, Unions, and Counting	Hash-tree, Prefix trees, Bit vectors
Neural Networks	Iterative Feedback networks Matrix multiplication, Inversion, Factorization	Sparse/dense matrices, Vectors, Double-precision/Complex data Smoothing functions
Support Vector Machines	Factorization, Matrix multiplication	Double-precision Sparse matrices, Vectors Kernel functions (e.g., Linear)
Decision Trees	Dynamic programming Recursive Tree Operations	Integers, Double-precision, Trees, Vectors, `log` function
Time Series Processing	Smoothing via averaging, Correlation Fourier and Wavelet transforms	Integers, Single-/Double-precision, Dense matrices Vectors, Distance and Smoothing functions
Text Analytics	Parsing, Bayesian modeling, String matching Hashing, Singular value decomposition Matrix multiplication, Transpose, Factorization	Integers, Single-/Double-precision, Strings Sparse matrices, Vectors, Inverse indexes, String functions, Distance functions
Monte Carlo Methods	Random number generators Polynomial evaluation, Interpolation	Double-precision, Bit vectors Bit-level operations, `log`, `sqrt` functions
Mathematical Programming	Matrix multiplication, Inversion, Factorization Dynamic programming, Greedy algorithms, Backtracking-based search	Integers, Double-precision, Sparse Matrices, Vectors, Trees, Graphs
On-line Analytical Processing	Grouping and ordering Aggregation over hierarchies	Prefix trees, Relational tables, OLAP Operators Sorting, Ordering, Aggregation operators
Graph Analytics	Graph traversal, Eigensolvers, Matrix-vector, Matrix-matrix multiplication, Factorization	Integer, Single-/Double-precision, Adjacency Lists Trees, Queues, Dense/Sparse matrices

Table 4: Computational characteristics of the analytics exemplars

Model Exemplar	Execution characteristics		Input-Output characteristics	
	Methodology	Memory Issues	**(Read-only)** Input Data	Output Data
Regression Analysis	Iterative	In-memory Disk-based	Large historical Structured	Small Scalar
Clustering	Iterative	In-memory Disk-based	Large historical Unstructured or structured	Small scalar Unstructured or structured
Nearest-Neighbor Search	Non-iterative	In-memory	Large historical Structured	Small Scalar or structured
Association Rule Mining	Iterative Non-iterative	In-memory Disk-based	Large historical Structured	**Larger** Structured
Neural Networks	Iterative Two Stages	In-memory Disk-based	Large Structured	Small Scalar
Support Vector Machines	Iterative Two Stages	In-memory Disk-based	Large Structured	Small Scalar
Decision Tree Learning	Iterative Two Stages	In-memory Disk-based	Large Structured & Unstructured	Small Scalar
Time Series Processing	Non-iterative Real-time	In-memory	High volume streaming Structured or unstructured	Small scalar or streaming Structured or unstructured
Text Analytics	Iterative Non-iterative	In-memory Disk-based	Large historical or streaming Structured or unstructured	Large or small Structured or unstructured
Monte Carlo s Methods	Iterative	In-memory	Small Scalar	Large Scalar
Mathematical Programming	Iterative	In-memory	Small Scalar	Small Scalar
On-line Analytical Processing (OLAP)	Non-iterative	In-memory Disk-based	Large historical Structured	**Larger** Structured
Graph Analytics	Iterative	In-memory Disk-based	Large historical Unstructured	Small Scalar or unstructured

Table 5: Runtime characteristics of the analytics exemplars

plications would be sub-optimal. As Tables 4 and 5 demonstrate, each exemplar has a unique set of computational and runtime features, and ideally, every exemplar would get a system tailor-made to match its requirements. However, we have also observed that different analytic exemplars share many computational and runtime features. Therefore, for a systems designer, the challenge is to customize analytics systems using as many re-usable software and hardware components as possible.

4.1 System Acceleration Opportunities

Table 6 describes system opportunities for accelerating analytics exemplars. Based on the computational and runtime characteristics described in Tables 4 and 5, we first identify key bottlenecks in the execution of analytic exemplars, namely compute-bound, memory-bound, and I/O bound (which covers both disk and network data traffic). As Table 6 illustrates, a majority of the analytics exemplars are compute bound in the in-memory mode and I/O-bound when in the disk-based mode. The compute-bound exemplars can benefit from traditional task-based parallelization approaches on multi-core processors, as well as by hardware-based acceleration via SIMD instructions or using GPUs. When used in the disk-based scenarios, these exemplars can improve their I/O performance by using solid state drives or data compression. Some of the analytics exemplars are memory-bound due to their reliance on algorithms that traverse large in-memory data structures such as trees or sparse matrices. For these exemplars, a better memory sub-system, with faster, larger, and deeper memory hierarchies, would be most beneficial. Once the memory accesses are optimized, these exemplars can also benefit from traditional computational acceleration techniques. Finally, some of the exemplars exhibit unique computational patterns (e.g., bit-level manipulations, pattern matching, or string processing) which could be accelerated using special-purpose processors such as FPGAs or by introducing new instructions in general-purpose processors. In most cases, the exemplars can be accelerated using commodity hardware components (e.g., multi-core processors, GPUs or SSDs). These hardware components can be then used to optimize re-usable software kernel functions (e.g., numerical linear algebra, distance functions, etc.), which themselves can be parallelized by a variety of parallelization techniques such as task parallelism, distributed-memory message-passing parallelism or MapReduce [26, 2]. These functions can be used as a basis of specialized implementations of the exemplars. Such hardware-software co-design enables optimized analytics solutions that can balance customization and commoditization.

4.2 The Netezza Example

An example of hardware-software co-design for database workloads is the Netezza data warehouse and analytics appliance [11]. The Netezza appliance supports both SQL-based OLAP and analytics queries. Netezza uses a combination of FPGA-based acceleration and customized software to optimize data-intensive mixed database and analytics workloads with concurrent queries from thousands of users. The Netezza system uses two key principles to achieve scalable performance: (1) Reduce unnecessary data traffic by moving processing closer to the data, and (2) Use parallelization techniques to improve the processing costs. A Netezza appliance is a distributed-memory system with a host server connected to a cluster of independent servers called the snippet blades (S-Blades). A Netezza host first compiles a query using a cost-based query optimizer that uses the data and query statistics, along with disk, processing, and networking costs to generate plans that minimize disk I/O and data movement. The query compiler generates executable code segments, called snippets which are executed in parallel by S-blades. Each S-blade is a self-contained system with multiple multi-core CPUs, FPGAs, gigabytes of memory, and a local disk subsystem. For a snippet, the S-Blade first reads the data from disks into memory using a technique to reduce disk scans. The data streams are then processed by FPGAs at wire speed. In a majority of cases, the FPGAs filter data from the original stream, and only a tiny fraction is sent to the S-Blade CPUs for further processing. The FPGAs can also execute some additional functions which include decompression, concurrency control, projections, and restrictions. The CPUs then execute either database operations like sort, join, or aggregation or core mathematical kernels of analytics applications on the filtered data streams. Results from the snippet executions are then combined to compute the final result. The Netezza architecture also supports key data mining and machine learning algorithms on numerical data (e.g., matrices) stored in relational tables.

A key lesson learned from the design of Netezza has been the huge value of specializing system design for analytics. Orders of magnitude improvements in efficiency can be achieved by carefully analyzing the system requirements and innovating using a collaborative software-hardware design methodology. As analytics applications become more main-

Model Exemplar	Bottleneck	Acceleration requirements and opportunities
Regression Analysis	Compute-bound I/O-bound	Shared- and Distributed-memory task parallelism Data parallelism via SIMD or GPUs Faster I/O using solid state drives
Clustering		
Nearest-Neighbor Search		
Neural Networks		
Support Vector Machines		
Association Rule Mining	I/O-bound	Shared-memory task parallelism Faster I/O using solid state drives Faster bit operations or tree traversals via FPGAs
Decision Tree Learning	Memory-bound	Larger and deeper memory hierarchies Data parallelism via SIMD
Time Series Processing	Compute-bound Memory-bound	Shared- and Distributed-memory task parallelism Data parallelism via SIMD or GPUs High-bandwidth, low-latency memory subsystem Pattern matching via FPGA
Text Analytics	Memory-bound I/O-bound	Shared- and Distributed-memory task parallelism Data parallelism via SIMD or GPUs Larger and deeper memory hierarchies Faster I/O via solid state drives Pattern matching and string processing via FPGA
Monte Carlo Methods	Compute-bound	Shared- and Distributed-memory task parallelism Data parallelism via SIMD or GPUs Faster bit manipulations using FPGAs or ASICs
Mathematical Programming	Compute-bound	Shared-memory task parallelism Massive data-parallelism via GPUs Larger and deeper memory hierarchies Search-tree traversals via FPGAs
On-line Analytical Processing	Memory-bound I/O-bound	Shared- and Distributed-memory task parallelism Data parallelism via SIMD or GPUs Larger and deeper memory hierarchies Pattern Matching via FPGAs, Faster I/O using solid state drives
Graph Analytics	Memory-bound	Shared-memory task parallelism Larger and deeper memory hierarchies

Table 6: Opportunities for parallelizing and accelerating analytics exemplars

stream, future database systems need to be designed in an integrated manner to support both the classical and analytics workoads.

5. SUMMARY

In this survey paper and the accompanying research report [5], we have reviewed the growing field of analytics that uses mathematical formulations to solve business and consumer problems. We have identified some of the key techniques employed in analytics, called *analytics exemplars*, both to serve as an introduction for the non-specialist, and to explore the opportunity for greater optimization for parallel computer architectures, and systems software. We hope this work spurs follow-on work on analyzing and optimizing analytics workloads.

6. REFERENCES

[1] A. P. Armacost, C. Barnhart, K. A. Ware, and A. M. Wilson. UPS Optimizes Its Air Network. *Interfaces*, 34(1), January-February 2004.

[2] R. Bekkerman, M. Bilenko, and J. Langford, editors. *Scaling Up Machine Learning: Parallel and Distributed Approaches*. Cambridge University Press, 2011. To appear.

[3] R. M. Bell, Y. Koren, and C. Volinsky. All Together Now: A Perspective on the Netflix Prize. *Chance*, 23(1):24–29, 2010.

[4] I. Bhattacharya, S. Godbole, A. Gupta, A. Verma, J. Achtermann, and K. English. Enabling analysts in managed services for CRM analytics. In *In Procs. of the 2009 ACM KDD Intl. Conf. on Knowledge and Data Discovery*, 2009.

[5] R. Bordawekar, B. Blainey, C. Apte, and M. McRoberts. Analyzing analytics, part 1: A survey of business analytics models and algorithms. Technical Report RC25186, IBM T. J. Watson Research Center, July 2011.

[6] R. Cantor, F. Packer, and K. Cole. Split ratings and the pricing of credit risk. Technical Report Research Paper No. 9711, Federal Reserve Bank of New York, March 1997.

[7] P. Crosbie and J. Bohn. Modeling default risk, December 2003. Moody's KMV.

[8] T. Davenport and J. Harris. *Competing on Analytics, The New Science of Winning*. Harvard Business School Press, 2007.

[9] T. Davenport, J. Harris, and R. Morison. *Analytics at Work, Smarter Decisions, Better Results*. Harvard Business School Press, 2010.

[10] J. Dean and S. Ghemawat. Mapreduce: a flexible data processing tool. *Communications of the ACM*, 53(1), 2010.

[11] D. J. Feldman. Netezza Performance Architecture. Keynote Presentation at the DaMoN'13 Workshop.

[12] D. Ferrucci, E. Brown, J. Chu-Carroll, J. Fan, D. Gondek, A. A. Kalyanpur, A. Lally, J. W. Murdock, E. Nyberg, J. Prager, N. Schlaefer, and C. Welty. Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 59(Fall), 2010.

[13] A. Go, R. Bhayani, and L. Huang. Twitter sentiment classification using distant supervision. Technical report, Computer Science Department, Stanford University, December 2009.

[14] E. Goode. Sending the police before there's a crime. The New York Times, August 16, 2011.

[15] P. Grosse, W. Lehner, T. Weichert, F. Faerber, and W.-S. Li. Bridging Two Worlds with RICE: Integrating R into the SAP In-Memory Computing Engine. In *Proc. of the VLDB Endowment*, September 2011.

[16] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, 2006.

[17] IBM Corp. Engineering and Scientific Subroutine Library (ESSL) and Parallel ESSL. `www.ibm.com/systems/software/essl`.

[18] Intel Inc. Intel Math Kernel Library. `software.intel.com`.

[19] J. Joyce. Pandora and the Music Genome Project. *Scientific Computing*, 23(10):40–41, September 2006.

[20] M. Lohatepanont and C. Barnhart. Airline Schedule Planning: Integrated Models and Algorithms for Schedule Design and Fleet Assignment. *Transportation Science*, 38(1), February 2004.

[21] G. Mohler, M. Short, P. Brantingham, F. Schoenberg, and G. Tita. Self-exciting point process modeling of crime. *Journal of American Statistical Association*, 106(493), March 2011.

[22] T. Mutanen. Customer churn analysis- a case study. Technical Report VTT-R-01184-06, Technical Research Centre of Finland, 2006.

[23] A. A. Nanavati, S. Gurumurthy, G. Das, D. Chakraborty, K. Dasgupta, S. Mukherjea, and A. Joshi. On the structural properties of massive telecom call graphs: Findings and implications. In *Proc. of the Conference on*

Information and Knowledge Management (CIKM'06), pages 435–444, November 2006.

[24] E. W. T. Ngai, L. Xiu, and D. C. K. Chau. Application of data mining techniques in customer relationship management: A literature review and classification. *Expert Systems with Applications*, 36:2592–2602, 2009.

[25] Oracle Inc. Oracle Advanced Analytics, 2013. Oracle Data Sheet.

[26] A. Rajaraman and J. Ullman. *Mining Massive Datasets*. Cambridge University Press, 2010. Free version available at `infolab.stanford.edu/~ullman/mmds.html`.

[27] K. Rexer, H. Allen, and P. Gearan. 2010 data miner survey summary. In *Procs. of the Predictive Analytics World*, October 2010.

[28] Y. Richter, E. Yom-Tov, and N. Slonim. Predicting customer churn in mobile networks through analysis of social groups. In *Procs. of the SIAM International Conference on Data Mining, SDM 2010*, pages 732–741, 2010.

[29] Science Special Issue. Dealing with data. *Science*, 331(6018), February 2011.

[30] V. Sindhwani, A. Ghoting, E. Ting, and R. Lawrence. Extracting insights from social media with large-scale matrix approximations. *IBM Journal of Research and Development*, 55(5):9:1–9:13, Sept-Oct 2011.

[31] Sony Pictures Inc. Jeopardy! The IBM Challenge. `www.jeopardy.com/minisites/watson`.

[32] Splunk Inc. Splunk Tutorial, 2011. `www.splunk.com`.

[33] A. Suman. Automated face recognition, applications within law enforcement: Market and technology review, October 2006.

[34] The Economist. Algorithms: Business by numbers. Print Edition, 13th September, 2007.

[35] The Economist. Data, data everywhere. Print Edition, 25th February, 2010.

[36] X. Wu, V. Kumar, J. R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, G. J. McLachlan, A. Ng, B. Liu, P. S. Yu, Z.-H. Zhou, M. Steinbach, H. David J, and D. Steinberg. Top 10 algorithms in data mining. *Knowledge Information Systems*, 14:1–37, 2008.