



HAL
open science

Big Steps Towards Query Eco-Processing - Thinking Smart

Simon Pierre Dembele, Ladjel Bellatreche, Carlos Ordonez, Nabil Gmati, Mathieu Roche, Tri Nguyen-Huu, Laurent Debreu

► **To cite this version:**

Simon Pierre Dembele, Ladjel Bellatreche, Carlos Ordonez, Nabil Gmati, Mathieu Roche, et al.. Big Steps Towards Query Eco-Processing - Thinking Smart. *Revue Africaine de Recherche en Informatique et Mathématiques Appliquées*, 2021, Volume 34 - 2020 - Special Issue CARI 2020, Volume 34 - 2020 - Special Issue CARI 2020 (34), pp.6767. <10.46298/arima.6767>. <hal-02931309v3>

HAL Id: hal-02931309

<https://hal.science/hal-02931309v3>

Submitted on 26 Mar 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Big Steps Towards Query Eco-Processing - Thinking Smart

Simon Pierre Dembele¹, Ladjel Bellatreche¹, Carlos Ordonez²

¹LIAS/ISAE-ENSMA, Poitiers, France

²University of Houston, Houston, USA

*E-mail : simon-pierre.dembele@ensma.fr/bellatreche@ensma.fr

DOI : [10.46298/arima.6767](https://doi.org/10.46298/arima.6767)

Submitted on 08 September 2020 - Accepted on 03 March 2021

Volume : 34 - Year : 2021

Special Issue : **Special Issue CARI 2020**

Editors : Nabil Gmati, Mathieu Roche, Tri Nguyen-Huu, Laurent Debreu

Abstract

Computers and electronic machines in businesses consume a significant amount of electricity, releasing carbon dioxide (CO_2), which contributes to greenhouse gas emissions. Energy efficiency is a pressing concern in IT systems, ranging from mobile devices to large servers in data centers, in order to be more environmentally responsible. In order to meet the growing demands in the awareness of excessive energy consumption, many initiatives have been launched on energy efficiency for big data processing covering electronic components, software and applications. Query optimizers are one of the most power consuming components of a DBMS. They can be modified to take into account the energetical cost of query plans by using energy-based cost models with the aim of reducing the power consumption of computer systems. In this paper, we study, describe and evaluate the design of three energy cost models whose values of energy sensitive parameters are determined using the *Nonlinear Regression* and the *Random Forests* techniques. To this end, we study in depth the operating principle of the selected DBMS and present an analysis comparing the performance time and energy consumption of typical queries in the TPC benchmark. We perform extensive experiments on a physical testbed based on PostgreSQL, MontetDB and Hyrise systems using workloads generated using our chosen benchmark to validate our proposal.

Keywords

Green query processing; DBMS audit; NonLinear Regression technique; Random Forest Technique

I INTRODUCTION

Energy Consumption in data center is becoming one of the most important issues in today's world. The hardware designers are making efforts to produce the hardware that minimizes the power Consumption. But despite of best efforts the providers of data storage and processing solutions still consume more energy. Data are at the heart of the new world order. As mentioned in The

Economist "*the world's most valuable resource is no longer oil, but data*"¹. Data Storage Systems (DSS) have to satisfy at the same time two important, crucial and conflictual Non-Functional Requirements (NFR): (1) a rapid processing of the deluge of data issued by enterprise sources, social networks, Internet of Things, etc. and (2) an optimal energy consumption.

Tackling energy over consumption, focusing on environmental issues and, adopting environmentally sound practices are important elements currently in the agenda of businesses, governments, organizations, associations, scientists, industrialists, ordinary and famous people around the world since climate change flags are flying everywhere. We are obliged to minimize or eliminate where possible the environmental impact of IT structure to help create a more sustainable environment. To reduce environmental problems and to create a sustainable environment, we must rethink our life and work styles by the means of political and economical actions including certainly a review of the process of data deluge. A World Energy Council study found that without any change in our current practice, the world energy demand in 2020 would be 50%-80% higher than the 1990 levels. According to International Energy Outlook-2013, the energy consumption will be grown by 56% between 2010 and 2040. Total world energy use will rise from 630 quadrillion British thermal units (Btu) in 2020 and to 820 quadrillion Btu in 2040[35]. Such ever-increasing demand could place significant strain on the current energy infrastructure and potentially damage world environmental health by CO , CO_2 , SO_2 , and NO_x effluent gas emissions and global warming. Like any oil, data pollutes as mentioned in the latest Blog entry of the Martin Tisné published on July 24, 2019: "*Data isn't the new oil, it's the new CO2*"². This pollution is caused by storing and processing this data.

Data storage and processing are important subsystems in data centers, and they have evolved and increased their capacity due to the advent of new paradigms, such as cloud computing and virtualization technology. As database researchers, we are then obliged to make ourselves, academia, industry, IT companies, funding agencies, and students, aware of actions related to energy savings of the *DSS* by considering *small* and *big* initiatives. The *DSS* infrastructure includes hardware, software, and facility service components that support the delivery of business systems and IT-enabled processes. According to the statistics published by the InfoTech group, IT equipment consumes approximately 50% of the total energy. Figure 1 illustrates the power consumption distribution of major components that consume energy in IT infrastructure. *DSSs* regardless of their types (DBMSs, data centers, and parallel database machines, etc.) have been identified as one of the major energy consumption components in IT infrastructure. This consumption is associated with their servers, storage devices, networks and infrastructure facilities such as cooling and power conditioning systems [28]. The processor consumes a major portion of energy followed by the storage device [68]. The first and major efforts in managing the energy of *DSS* have particularly touched data centers [60] since they have been pointed out by several organization reports as one of the biggest energy consumers. Based on U.S. Environmental Protection Agency, in 2014, U.S. data centers consumed approximately 70 billion kilowatt-hours, totaling about 1.8 percent of domestic electricity consumption³.

Making storage systems green has been addressed in many research efforts in the literature. The 2010s were a decade that marked the beginning of the integration of energy when designing

¹<https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-protect-discretionary{\char\hyphenchar\font}{}{}resource-is-no-longer-oil-but-data>

²<https://luminategroup.com/posts/blog/data-isnt-the-new-oil-its-the-new-co2>

³https://www.epa.gov/sites/production/files/2017-11/documents/2017-06_0.pdf




Figure 1: Energy distribution among different components[63].

query optimizers [34, 43], where hardware and software solutions have been discussed, evaluated, analyzed and implemented. Hardware research efforts got more attention than software ones. This is because several studies consider that the operating systems and firmware (hardware programs) manage energy and consequently save energy power of query processors. This finding is *questionable* since they use techniques covering software (e.g., finding the best query plan satisfying the first *NFR*) and hardware (e.g., executing the appropriate algorithms associated with the selected plan on the target platform hosting the DBMS). DSS can benefit from existing energy-aware hardware (such as removable storage media [27], graphics cards [61], etc.). For software solutions, DBMSs need to accurately estimate hardware power consumption under static and dynamic loads, revisit query execution, scheduling, database systems design, and data management strategies in order to meet the response time and allow better energy management. This poses several challenges for the researchers: it is necessary to develop cost models and metrics that consider the parameters related to the database systems components, such as the database schema (e.g., the size of the tables, the length of tuples, etc. .), the query workload (e.g., type of queries, join selectivity factors and selection predicates, intermediate result sizes, etc.), hardware (e.g., buffer size, disk page size, etc.) and the deployment platform (e.g., RAM, disk size, number of nodes, etc.).

In this paper, we focus on building a green query processor – considered as one of the most important energy consumers of the DBMS [11] by proposing several *costs models estimating the energy consumption* of a query executed by an optimizer in several DBMSs offering parallel mode with different functioning policies. The crucial aspect of our work that differentiates it from previous works [63, 66, 42, 70, 81, 80, 79] is the inclusion of the main memory (RAM) parameter in our cost model for PostgreSQL and MonetDB. Additionally, we propose a novel cost model for an in-memory system named Hyrise. These cost models are proposed via statistical regression and random forest techniques. Our contributions can be summarized as follows:

- We provide extensive experiments to compare the data loading time, queries execution time and power consumption of TPC-H benchmark queries in three DBMS (PostgreSQL, MonetDB, Hyrise).
- We have proposed three energy cost models by taking into account the new parameters (main memory, cache memories) in centralized DBMS to predict the energy cost of query processing. These energy cost models value are estimated via statistical regression and random forest techniques.
- Intensive experiments have been conducted to evaluate the quality of our models (regression and random forest) using TPC-H benchmarks on PostgreSQL, MonetDB and

Hyrise.

We claim that building green query processors passes through a deep audit that allows designers identifying relevant energy-sensitive parameters that are the entries of the mathematical cost models estimating energy when executing a query. The development of such models necessitates a deep understanding of the functioning of the target DBMS hosting the database application (e.g., query execution mode). Note that the value of some energy sensitive parameters cannot be obtained from the statistic module of the DBMS, therefore they have to be computed using machine learning techniques.

Paper outline: The outline of this paper is as follows: Section II summarizes the fundamental notions of energy and describes the design of the different DBMS used in our study. A consequence of energy overusing and the benefit of energy saving are also presented. Section III details a comparative study of DBMS in terms of Data loading time, query processing time and power consumption. Section IV describes the technical details of our mathematical energy cost models and presents the Non-Linear Regression and Random Forest techniques used to set the value of energy-sensitive parameters. Section V presents and interprets our experimental results. Section VI surveys the related works while our conclusions and perspectives are given in Section VII.

II PRELIMINARY

In this section, we introduce the concepts of energy and discuss the consequences of high energy consumption. We will also talk about the benefits of energy efficiency and describe the design of DBMS PostgreSQL, MonetDB and Hyrise.

2.1 Energy concept

Energy is a measure of the ability to do work, to change system state, it comes in many forms (magnetic energy, electrical energy, chemical energy, and nuclear energy...) and can transform from one type to another[76]. In our study, we consider electrical energy. Energy is a physical quantity dependent on time. It's measured in *Joules*. The energy transfer in one second is defined as electrical power. More precisely:

Power is the amount of a system energy per unit of time or the rate of doing a job. *Watts* is the unit of measure. Formally, energy and power can be defined as follows:

$$P(t) = d/dtE(t) ; E(t) = \int_0^t p(\tau)d\tau \quad (1)$$

where P , t , and E represent, respectively, a power, a period, and energy. Since it is hard to guarantee the accuracy of energy measure, in this paper we use the *average power* representing the average power consumed during the execution of the workload.

Energy efficiency (EE) expresses the optimal use of energy to offer the same service. It is expressed by [28]:

$$EE = \frac{\text{Useful Energy output}}{\text{Total Energy input}} = \frac{\text{Performance}}{P} \quad (2)$$

Based on the above equation, we remark that there are two ways to improve EE : (i) by improving performance with the same power and (ii) by reducing power consumption without sacrificing too much performance (i.e. with a reasonable performance degeneration).

2.2 Energy overusing issues

Energy is both a solution and a problem for sustainable development. It contributes to economic development, poverty reduction, education and the general improvement of the quality of life and that of humanity, but it is also one of the main causes of pollution of air and other harmful effects on human health and the environment. Current global energy consumption is estimated at 22×10^9 kWh per year. This consumption is approximately equivalent to an emission of 6.6×10^9 tonnes of carbon dioxide (CO₂) into the atmosphere [19]. The CO₂ is a huge provider of greenhouse gases (GHG). The Organization for Economic Cooperation and Development(OECD) warns that, given current trends, energy-related emissions will increase by 70% by 2050. This can accelerate the negative consequences of climate change, including higher temperatures and an increase in the frequency of extreme events.

2.3 Energy efficiency Benefit

The reduction in energy consumption leads to the reduction of greenhouse gas(GHG) emissions and other pollutants, thus improving the quality of life in the environment. For example at the level of information systems, most computer systems are energy wasters due to the high consumption of certain electronic circuits even in the inactive state. Migration to systems that improve energy efficiency offers great prospects in businesses, as it offers many benefits such as lower energy bills, increase reliability of the electricity distribution system, reduce emissions of components polluting the atmospheric layer. The increasingly efficient and greener use of the energy required to run IT infrastructures, ensures the longevity of systems and promotes the consumption of energy from renewable sources. In fact, energy dissipation has a detrimental effect on the reliability of electronic circuits[50]. The malfunction as well as the deterioration of the components is caused by many factors, including a sharp rise in temperature, over-voltage of the terminals of the electric dipole, etc. Equipment failure rate increases when the temperature exceeds 15° Celsius[4].

2.4 Panorama of our Studied DBMS

In order to demonstrate the effectiveness of our procedure to build a green query processor, we apply our proposed models on three DBMS: PostgreSQL, MonetDB and Hyrise. We chose to work with these systems because they are very widely used, open source and offering a parallel mode for queries processing in centralized DBMS. In addition, they are different in their storage model: Row-Store for PostgreSQL, Colum-Store for MonetDB and Hybrid-store for Hyrise.

2.4.1 PostgreSQL:

It is a row-store DBMS supporting object-relational databases. It uses the server/client model and supports the standard database languages. It offers many advanced functionalities such as user-defined types, table inheritance, sophisticated locking mechanism, foreign key referential integrity, views, rules, sub-query, nested transactions, multi-version concurrency control, and asynchronous replication [82].

The support of a parallel query involves multiple background worker processes. There is a back-end process that handles all queries issued by the connected client. This back-end consists of the following subsystems:

1. **Parser:** it checks the query syntax expressed in a high-level query language like SQL to determine whether it is well formulated according to the grammar rules of the query language.




Figure 2: PostgreSQL query processor

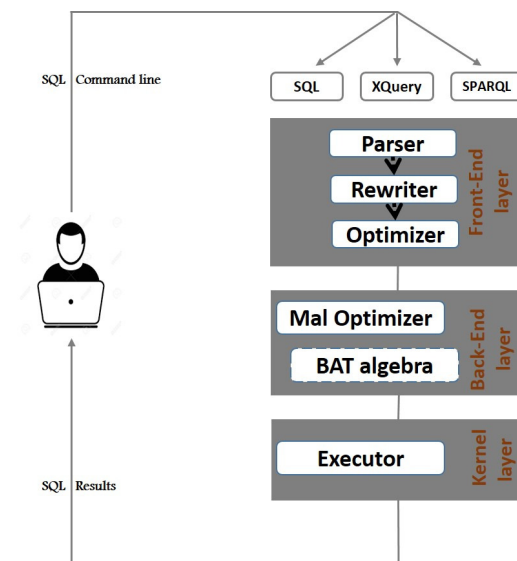


Figure 3: MonetDB query processor.

Figure 4: Internal Architecture of PostgreSQL and MonetBD

2. **Analyzer:** The query must be validated by verifying that all attributes and relationship names are valid and semantically significant in the schema of the database.
3. **Re-writer:** Using transformation rules, an internal representation of the query is then created (query tree).
4. **Planner:** It generates the cheapest plan tree that can be executed from the query tree.
5. **Executor:** It executes the query via accessing the tables that are generally stored in a single heap, with each column stored on a single tuple. A query has many possible execution strategies, and the selection of the best plan is usually conducted by cost model-driven strategies [24].Figure 2 summarizes the different phases of the PostgreSQL query processor.

2.4.2 MonetDB:

It was designed primarily for data warehouse applications. It achieves significant speedup compared to more traditional designs by innovations at all layers of a DBMS, e.g., a storage model based on vertical fragmentation (column-store), a modern CPU-tuned query execution architecture, adaptive indexing, run-time query optimization, and a modular software architecture[38].

Internally, the design, the architecture and the implementation of MonetDB reconsider all aspects and components of classical database architecture and technology by exploiting effectively the potentials of modern hardware.

1. **Storage model:** It is a significant deviation of traditional database systems. It uses the decomposed storage model (DSM) which represents relational tables using vertical fragmentation, by storing each column in a separate *#surrogate, value#* table, called binary association table (BAT). The left column (the surrogate or object-identifier - (oid)) is called the head, whereas, the right column is the tail. MonetDB executes a low-level relational algebra called the BAT algebra. Data in execution is always stored in (intermediate) BATs, and even the result of a query is a collection of BATs[10]. During the query evaluation, all intermediate results are in a column format. Only just before sending

```

mal
-----
function user_s4(():void;
  X_2:void := querylog.define("explain select c_name,c_phone from customer
  where c_mktsegment=\BUILDING\";:str, "default_pipe":str, 42:int);
  X_44:bat[:str] := bat.new(nl:istr);
  X_50:bat[:int] := bat.new(nl:int);
  X_48:bat[:int] := bat.new(nl:int);
  X_47:bat[:str] := bat.new(nl:istr);
  X_46:bat[:str] := bat.new(nl:istr);
  X_5:int := sql.nvc();
  X_26:bat[:str] := sql.bind(X_5:int, "tpcs":str, "customer":str, "c_mktse
gment":str, 0:int);
  C_60:bat[:oid] := sql.tid(X_5:int, "tpcs":str, "customer":str);
  C_38:bat[:oid] := algebra.thetaselect(X_26:bat[:str], C_6:bat[:oid], "BU
ILDING":str, "=":str);
  X_19:bat[:str] := sql.bind(X_5:int, "tpcs":str, "customer":str, "c_phone
":str, 0:int);
  X_41:bat[:str] := algebra.projection(C_38:bat[:oid], X_19:bat[:str]);
  X_9:bat[:str] := sql.bind(X_5:int, "tpcs":str, "customer":str, "c_name":
str, 0:int);
  X_40:bat[:str] := algebra.projection(C_38:bat[:oid], X_9:bat[:str]);
  X_54:bat[:str] := bat.append(X_44:bat[:str], "tpcs.customer":str);
  X_53:bat[:str] := bat.append(X_46:bat[:str], "c_name":str);
  X_55:bat[:str] := bat.append(X_47:bat[:str], "varchar":str);
  X_57:bat[:int] := bat.append(X_48:bat[:int], 25:int);
  X_59:bat[:int] := bat.append(X_50:bat[:int], 0:int);
  X_61:bat[:str] := bat.append(X_51:bat[:str], "tpcs.customer":str);
  X_62:bat[:str] := bat.append(X_53:bat[:str], "c_phone":str);
  X_64:bat[:str] := bat.append(X_55:bat[:str], "char":str);
  X_66:bat[:int] := bat.append(X_57:bat[:int], 35:int);
  X_68:bat[:int] := bat.append(X_59:bat[:int], 0:int);
  sql.resultSet(X_61:bat[:str], X_62:bat[:str], X_64:bat[:str], X_66:bat[:
int], X_68:bat[:int], X_40:bat[:str], X_41:bat[:str]);
end user_s4;

```

Figure 5: MonetDB architecture: MAL plans program produced by the command *explain*.

the final result to the client, $N - ary^4$ tuples are constructed. This approach allows the query engine to exploit CPU- and cache-optimized vector-like operator implementations throughout the whole query evaluation relying on a bulk processing model allowing minimizing function calls, type casting, various metadata handling costs, etc. Intermediate needed results are materialized to increase their reuse[13].

2. **Query execution model:** The MonetDB kernel is an abstract machine, programmed in the MonetDB Assembly Language (MAL). The core of MAL is formed by a closed low-level two-column relational algebra on BATs. N-ary relational algebra plans are translated into two-column BAT algebra and compiled to MAL programs. These MAL programs are then evaluated in an operator-at-a-time manner. Figure 3 shows the internal design of MonetDB.

MonetDB’s query processing scheme is centered around three software layers:

- *The top layer or front-end* provides the user-level data model and query language. The query language is first parsed into an internal representation (e.g., SQL into relational algebra), which is then optimized using domain-specific rules. In general, these domain-specific strategic optimizations aim primarily at reducing the amount of data to be processed (i.e., the size of intermediate results). The optimized logical plan is then translated into MAL[38]. Figure 5 illustrates the MAL plan produced for the following query.
 SELECT c_name, c_phone
 FROM CUSTOMERS WHERE customers.city='London';
- *The middle layer or back-end* consists of the MAL optimizers framework and the MAL interpreter. The MAL optimizers framework consists of a collection of optimizer modules, where each MAL code is transformed in one more efficient by adding resource management directives if necessary. The modules provide facilities ranging from symbolic processing up to just-in-time data distribution and execution. The approach breaks with the hitherto omnipresent cost-based optimizers by recognizing that not all decisions can be cast together in a single cost formula [38].
- *The bottom layer or kernel* provides BATs as MonetDB’s bread-and-butter data structure, as well as the library of highly optimized implementations of the binary relational algebra operators. They maintain properties over the object accessed to gear the selection of subsequent algorithms. For example, the Select operator can benefit both from sorted-ness of the BAT or it may call for a sample to derive the expected sizes[38].

⁴https://www.cs.odu.edu/~toida/nerzic/content/relation/definition/cp_gen/index.html

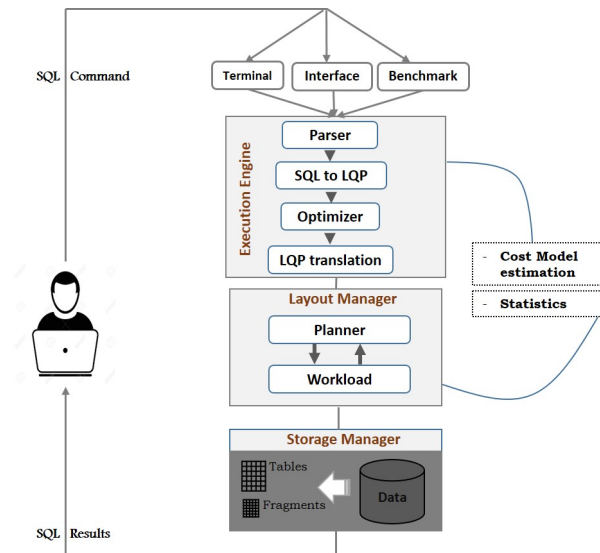


Figure 6: Internal Architecture of Hyrise

For query parallel execution in Monetdb, a sequential execution plan is generated firstly and parallelization is then added in the second optimization phase. The individual MAL operators are marked as either *blocking* or *parallelizable*. The optimizers will alter the plan by splitting up the columns of the largest table into separate chunks, then executing the *parallelizable* operators once on each of the chunks, and finally merging the results of these operators together into a single column before executing the blocking operators[75].

2.4.3 Hyrise:

Hyrise⁵ was first presented in 2010, developed by *Hasso-Plattner-Institute* in Germany under the direction by Professor Hasso Plattner to introduce the concept of the hybrid storage model (row-oriented and column-oriented) for in-memory databases. The Hyrise system is a new implementation based on the relational model, which automatically splits database tables vertically and then horizontally into small partitions, of varying widths depending on how the columns of the table are accessed [25][77]. The system aims at maximizing the performance of the cache memory for analytical(OLAP-type) and transactional (OLTP-type) queries by customizing the dynamic design of storage data. The goal of the Hyrise designers is to provide a flexible in-memory data management platform, allowing experiments to be conducted to validate new concepts around in-memory data management.

The overall architecture of the hyrise system consists of three main components: a storage manager, a query execution engine and a layout manager shown in figure 6. The storage manager is responsible for creating and maintaining the hybrid containers that store data. Storage is based on the vertical partitioning of relational tables into disjoint sets of attributes. Each partition is represented internally as a container. The containers are stored in a single compressed block in main memory (RAM) and can be merged after decompression to build the initial relational table. The layout manager analyzes a given query workload and demands the best possible layout (partitioning) for that workload from the storage manager.

⁵<https://hpi.de/plattner/projects/hyrise.html>

III COMPARATIVE ANALYSIS OF OUR DBMS

This section presents an experimental study comparing our three DBMS in terms of rapid query performance, reduction of power consumption, data loading time and number of caches misses without exploiting our models when queries executing. These experiments are conducted using the first 6 (SPJ) queries of the TPC-H benchmark. The system under which we perform our assessments is described in section 3.1.

3.1 Experimental Setup

Hardware: The experiments are conducted using a DELL PRECISION T1700 with Dell 073MMW motherboard(Socket 1150 LGA), Intel Core i7 4770 CPU@ 3.40GHz Haswell 22nm Technology(1 CPU-4Core-8 Threads), 12Go Dual-Channel DDR3@ 798MHz main memory, NVIDIA NVS 310 511 MB, ATA Disk Western Digital(WD) 500 GB.

OS: We run all the experiments using Ubuntu 18.04 bionic as the operating system with kernel version 5.0.0-27-generic.

Data sets: We use the TPC-H decision support benchmark, which is designed for evaluating data warehouses. It is composed of eight tables (*PART*, *PARTSUPP*, *REGIONAL*, *SUPPLIER*, *CUSTOMER*, *LINEITEM*, *NATION*, and *ORDERS*), we generate data at different scale factors 1 GB, 2GB, 3GB, 5GB, 10GB, 30GB⁶. We chose to use benchmark from TPC because they are known and have many years skilling in the definition of transaction processing and database benchmark.

DBMS Systems: We analyse three systems: an open-source row-store PostgreSQL version release 10.10, an open-source column-store MonetDB release 11.33.11 and the open-source hybrid-store Hyrise release V2.

Profiling: After generated the data from TPC-H benchmark in textual(comma-separated values - CSV) files, we use COPY command from SQL standard to load data in DBMS. We collect statistics about average time and power consumption for each query in the particular DBMS. For power measurement, we use power meters called Watt-UP-PRO⁷ at a 1Hz frequency. For caches misses analyses, we use Cachegrind tool⁸ and Processor Counter Monitor(PCMC)⁹.

3.2 Experimental Evaluation

We conduct several experiments to evaluate the data loading performance, query execution performance and power consumption of the tested systems. We present in this section our finding.

- **Data loading and Power Analysis:** This experiment investigates the behavior of PostgreSQL, MonetDB and Hyrise when loading the two datasets of size 1GB and 3GB. Figure 7 plots the data load time for each DBMS under the two scale factors(SF). As we can see, the loading time increases for each system as the size of the dataset increases. Monet outperforms the rest of the DBMS in the two cases; when considering 3GB datasets, Monet is 3.5 faster than PostgreSQL and 1.5 faster than Hyrise. After Monet, Hyrise exhibits fastest loading time than PostgreSQL. PostgreSQL shows the worst performance, the reason is that the row-stores require more space because they store auxiliary information

⁶Giga Byte

⁷http://www.energyalternatives.ca/pdf/wattsup_TTW.pdf

⁸<https://valgrind.org/docs/manual/cg-manual.html>

⁹<https://github.com/opcm?tab=repositories>




Figure 7: Data loading time.




Figure 8: Energy Consumed when loading data,

in each tuple (e.g., a header) and do not use compression. The increased data loading time into the hyrise system than monet is due to the additional processing required for data compression. Hyrise applies a higher compression ratio than Monet because monet only compresses string values. Like the data loading time, the energy consumed during this phase is shown in figure 8. Energy is a physical quantity influenced by the interval of time that elapses. In the figure, we notice that the energy consumed by Postgresql is more considerable than that of monetdb and hyrise.

- Query Performance and Power Analysis:** In this experiment, we evaluate the runtime and power consumption of the Postgresql and Monetdb systems when performing queries on the two database configurations: 30 GB and 100 GB. Then, we analyze with the dataset of size 5 GB the behavior of the three systems (Postgresql, monetDB,hyrise). Lastly, elapsed time for each SPJ query on the same database has been compared. Figures 9(b) and 9(c) summarize the average execution time results obtained using the SPJ queries for PostgreSQL and MonetDB. Figure 9(a) represents the average execution time for three systems on the 5 GB dataset.




Figure 9: Execution time comparison using $Q_1, Q_2, Q_3, Q_4, Q_5, Q_6$ queries from TPC-H benchmark in parallel mode

These results show that the performance of Hyrise outperforms those obtained by Post-

greSQL and MonetDB on the 2GB dataset for all considered. We can see that for most queries, Hyrise performance is within an order of magnitude of the other databases. The PostgreSQL performance gaps on Figure 9(a), 9(b), 9(c) is mainly due to the fact that it cumulates an important number of page faults when processing queries. The row-store systems have to scan and use the entire n-tuples rather than only the needed columns values. Therefore, the entire rows plus the built-in index tree cannot reside long enough in the main memory or in the cache memories. They must be swapped on the disk this leads to many disks IOs. For column-oriented systems like MonetDB, just the values of the columns required to answer the queries are loaded. hybrid system like hyrise combines the advantages of column and row storage systems and it is well designed to support efficient parallel execution for multi-core processors.

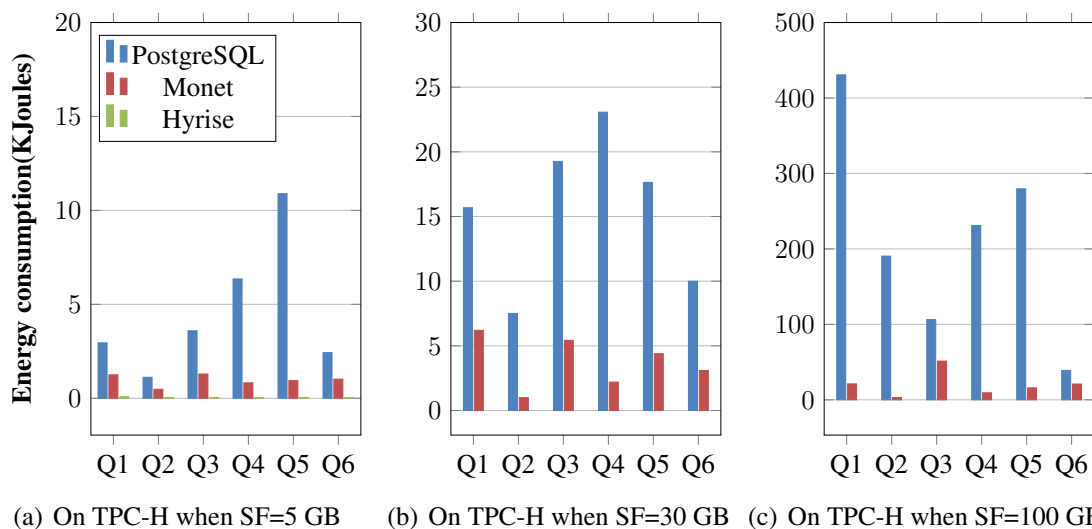


Figure 10: Average energy consumption for PostgreSQL, MonetDB, Hyrise using $Q_1, Q_2, Q_3, Q_4, Q_5, Q_6$ queries from TPC-H benchmark.

To compare the energy consumption on our three systems when executing queries, we measure the power dissipated per second using an amperage apparatus named wattmeter. At the end of these measurements, we calculate the average energy consumed by each query. Figure 10(a) depicts the energy consumption of three DBMS on 5 GB datasets. Hyrise outperforms PostgreSQL and MonetDB under the 5 GB TPC-H dataset because of the rapid processing of queries. Hyrise directly accesses compressed data in main memory and not in secondary memories, this speeds up data processing. Average energy consumed by MonetDB, PostgreSQL during the execution of the queries on SF=30 and SF=100 can be found in Figures 10(b) and 10(c). MonetDB energy consumption during queries execution is reasonable than for PostgreSQL. The high consumption of the PostgreSQL query processor can be explained by the characteristics of this type of row-store DBMSs. MonetDB uses compression techniques to reduce the cost of data scans. This directly impacts query performance due to fewer I/O requests and page faults. To understand *Where does power go* on PostgreSQL DBMS, we captured the last level cache(LLC) misses and Cache level 2(L2) misses. We execute the first three queries of TPC-H benchmark on the 5 GB datasets. The results for these queries are shown in Figure 11. From these results illustrated, the LLC caches misses and L2 caches misses of Hyrise outperforms the row-store DBMS (PostgreSQL) and column-store(MonetDB).




Figure 11: LLC and L2 cache misses for Q_1, Q_2, Q_3 queries from TPC-H benchmark.

IV ENERGY COST MODELS PROCESS

In this section, we describe our models that estimates the power consumption when executing query based on settings of a particular system. An energy model can be represented in the form of equations, graphical models, decision trees, neural networks, etc. The choice of representation affects the accuracy of the model. An accurate model of power consumption is essential to improve energy efficiency optimizations, the best model should be accurate, fast, generic, portable, simple, non-intrusive and low cost [60].

4.1 Modeling process

Inspired from [60], we establish the main steps of building green query processors are: (1) identification of relevant energy-sensitive parameters belonging to hardware and software components, (2) elaboration of mathematical cost models estimating consumed energy when executing a query on a target DBMS and (3) setting of values of the energy-sensitive parameters using machine learning techniques.

4.1.1 Step 1: (Energy-sensitive parameters identification)

Regarding the extraction of energy-sensitive parameters, we have identified a set of parameters that cover the main components of a *DSS*, that we have classified them into four categories as it is sketched in Figure 12.

4.1.2 Step 2: (Energy-cost model construction)

Our work focuses on modeling energy consumption of query operations related to its resources consumption on a single-node database environment, therefore we do not consider the network transmission cost. To construct our model, we investigate the executor tasks to understand how to profile the power consumption on an individual query. To do that, we execute a set of queries (simple and complex) with different scale of factors of the TPC-H benchmark. In our process, we treat energy as the resource consumed by systems operation during queries execution. The input parameters (resources) that we consider are listed in Table 1 for each systems. From Table 1, note that our models consider the main resources responsible for the energy consumption of each system. For in-memory system Hyrise, the cost of access to secondary memories is zero




Figure 12: Energy-sensitive parameters

but it should be noted that the cost of the different caches memories ($L1^{10}$, $L2^{11}$, $L3^{12}$) is not uniform.

On PostgreSQL, we consider that queries are executing in pipelined fashion [57]. On MonetDB the data are processed in an *operator at a time* manner. In this processing model, the operator processes the entire column at once before moving on to the next operator because the intermediate result of every operator has to be materialized. For Monetdb, we adopt an operations-based modeling approach. Hyrise follows a push-based operator model. Each operator generates a valid table, which is the only data entry for the following operator.

For a given query plan of Q_i executing in PostgreSQL, denoted by $PlanPost_i$ consisting by k pipelines noted $\{PL_1^i, PL_2^i, PL_3^i, \dots, PL_k^i\}$, its average power cost is estimated as follows:

¹⁰Cache Memory Level 1

¹¹Cache Memory Level 2

¹²Cache Memory Level 3

| | Item | DBMS Systems | | |
|-------------------------|---------------------|--------------|---------|--------|
| | | PostgreSQL | MonetDB | Hyrise |
| Input Parameters | Processor (cpu) | × | × | × |
| | Main Memory (mem) | × | × | × |
| | Disk (dio) | × | × | - |
| | L1 Cache memory(L1) | - | - | × |
| | L2 Cache memory(L2) | - | - | × |
| | L3 Cache memory(L3) | - | - | × |

Table 1: Parameters used in ours models

$Power(Q_i) = \frac{\sum_{j=1}^k Power(PL_j^i) * Time(PL_j^i)}{Time(Q_i)}$, where $Time(Q_i)$, $Time(PL_j^i)$ represent respectively, the execution time of the query Q_i and the execution time of PL_j^i . The power dissipated when processing the query is the combination of main identified resources (*CPU, Main memory and Disk*) energy consumption.

$$Power(PL_j^i) = W_{cpu} * \sum_{u=1}^n C_{cpu_u} + W_{mem} * \sum_{u=1}^n C_{mem_u} + W_{dio} * \sum_{u=1}^n C_{dio_u} \quad (3)$$

where W_{cpu} , W_{mem} and W_{dio} are the model parameters. W_{cpu} , W_{mio} and W_{dio} are unit-power cost for instructions, read/write operations on memory and read/write operations on disk respectively. The C_{cpu_u} is the number of instructions executed by CPU. C_{mem_u} is the number of read or write operations accessed on memory. C_{dio_u} is the number of read or write operations accessed on disk. The n is the number of operators in the pipeline. u is the summation index.

For MonetDB, the energy cost for a given query plan (Q_i) denoted by $PlanMonetDB_i$ consisting by k operations noted $\{OP_1^i, OP_2^i, OP_3^i, \dots, OP_k^i\}$ is estimated as follows: $Power(Q_i) = \frac{\sum_{j=1}^k Power(OP_j^i) * Time(OP_j^i)}{Time(Q_i)}$, where $Time(Q_i)$, $Time(OP_j^i)$ represent respectively, the execution time of the query Q_i and the execution time of OP_j^i . The power dissipated when processing is the combination of the energy consumption of the main resources identified. The formula is given by the following equation:

$$Power(OP_j^i) = W_{cpu} * C_{cpu_j} + W_{mem} * C_{mem_j} + W_{dio} * C_{dio_j} \quad (4)$$

Each Hyrise plan invokes a GetTable operator at the start. The latter is responsible for loading the table name and data blocks from the storage manager (from main memory to the CPU). The power dissipated when processing is the combination of the energy consumption of the main resources identified(CPU, main memory and caches memories). For a given query plan (Q_i) denoted by $PlanHyrise_i$ consisting by k operations noted $\{OP_1^i, OP_2^i, OP_3^i, \dots, OP_k^i\}$ is estimated as follows: $Power(Q_i) = \frac{\sum_{j=1}^k Power(OP_j^i) * Time(OP_j^i)}{Time(Q_i)}$, where $Time(Q_i)$, $Time(OP_j^i)$ represent respectively, the execution time of the query Q_i and the execution time of OP_j^i . The formula is given by the following equation:

$$Power(OP_j^i) = W_{cpu} * C_{cpu_j} + W_{L2} * C_{L2_j} + W_{L3} * C_{L3_j} + W_{mem} * C_{mem_j} \quad (5)$$

where W_{L2} , W_{L3} , W_{mem} and W_{CPU} are the model parameters. W_{L2} and W_{L3} are unit-power for memory cache level 2 and 3. The C_{L2_j} is the miss count of L1 cache and C_{L3_j} is the miss count of L2 cache.

4.1.3 Step 3: (Machine learning)

To identify the different values of the power unit cost in the equations 3, 4, 5, we use the Regression techniques at first and then the Random Forest(RF) regression.

1. **Regression Technique:** Regression is a predictive technique used in machine learning. It consists in analyzing a relationship between two quantitative variables and using it to estimate the unknown value of one using the known value of the other. It is commonly used in management and marketing techniques, profit forecasts, business planning, financial forecasting, time series prediction, biomedical modeling, and environmental modeling. There are different families of regression functions and different ways of measuring the error rate. They can be of a simple linear model, multiple linear, nonlinear, etc.[1].

In our work, we use a non-linear regression technique involving the identified energy-sensitive variables. To do that, we set to 2 the polynomial degree for PostgreSQL and Hyrise. For MonetDB, we use the degree 3. We find the values of parameters of our equations: W_{cpu} , W_{mem} , W_{L2}, W_{L3} and W_{dio} by applying this setting in the regression package defined in the R language.

2. **Random forest** : Random forests or random decision forests are an ensemble learning method for classification, regression and other tasks that operate by constructing a multitude of decision trees at training time and outputting the class that is the mode of the classes (classification) or mean prediction (regression) of the individual trees. It was introduced by Breiman and Cutler. Recent studies have suggested that Random Forest offers features which make it very attractive algorithm for statistical learning. The method is based upon an ensemble of decision trees, from which the prediction of a continuous variable is provided as the average of the predictions of all trees. In RF regression, an ensemble of regression trees is grown from separate bootstrap samples of the training data using the CART¹³ algorithm. The branches in each tree continue to be subdivided while the minimum number of observations in each leaf is greater than a predetermined value[8]. There are three possible training parameters for Random Forest:
 - **ntree**: the number of trees in the Forest;
 - **mtry**: the number of different descriptors tried at each split;
 - **node size**: the minimum node size below which leaves are not further subdivided.

Random Forest includes a method for assessing the importance of descriptors to the model. Random Forests were trained using the random Forest library in the statistical computing environment R.

4.2 Models training

4.2.1 Data Sets:

The data set was selected from TPC-H Benchmark¹⁴. The TPC-H Benchmark (TPC-H) is a decision support benchmark. It consists of a suite of business oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance. This benchmark illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions. We generate data at different scale factors: 1 GB, 2GB, 3GB, 5GB, 10GB, 30GB. For MonetDB and PostgreSQL models training, we study and collect the characteristics of forty-four (44) Select-Project-Join (SPJ) queries and measure the energy consumed by each of them using the power meter. In the case of Hyrise, we collect the statistics of ninety-nine (99) Select-Project-Join (SPJ) queries. For the test set, we use the QGEN¹⁵ tool from TPC-H benchmark to generate twenty-two (22) Select-Project-Join (SPJ) queries. An analytical study of training and test set prediction is provided at the end of the Model Evaluation in section V.

¹³https://fr.wikipedia.org/wiki/Algorithme_CART

¹⁴<http://www.tpc.org/tpch/>

¹⁵<https://github.com/electrum/tpch-dbgen>




Figure 13: Deployment of our experimental testbed.

4.2.2 System settings:

Our experiments run on the server with an Intel Core i7 4770 processor (L1 cache size is 32KB, L2 cache size is 256KB and L3 cache size is 8MB), 12GB DDR3 main memory and a 500GB SATA hard drive for Hyrise system model training. For MonetDB and PostgreSQL, we use a DELL E6430 (Intel Core i5 3340 processor (1 CPU - 2 Core - 4 Threads), 8GB DDR3 main memory, ATA Disk Toshiba 500 GB). We use power meters called Watt UP PRO at a 1Hz frequency placed between the electrical power source and the database server to quantify the system consumption. The environment topology of our experimentation is shown in Figure 13.

4.2.3 Operations Counting:

Many modern processors contain a performance monitoring unit (PMU). The design and functionality of a PMU is CPU-specific. We use Cachegrind¹⁶ and PCM¹⁷ to get some statistical informations from PMU when querying evaluation. Cachegrind simulates how the program interacts with a machine's cache hierarchy. Processor Counter Monitor (PCM) is an application programming interface (API) and a set of tools based on the API to monitor performance and energy metrics of Intel Core, Xeon, Atom and Xeon Phi processors. We also use tool like explain, trace to collect information about tables accessed. EIST¹⁸ (Enhanced Intel SpeedStep Technology), an Intel DVFS implementation, changes CPU frequency and voltage to improve energy efficiency. This technique can incur errors in energy measurement. We turn off these options and execute our queries under the given frequency and voltage.

V ENERGY MODEL EVALUATION RESULTS

In this section, we will evaluate our energy model and present the results obtained after intensive experimentation on PostgreSQL, MonetDB, Hyrise.

5.1 Parameters values

After having made a series of observations in which the queries chosen for models training are executed one after the others, we use regression and Random Forest in R language to determine parameters values for our models describe in equations 3, 4, 5. The following equations 6, 7,

¹⁶<https://valgrind.org/docs/manual/cg-manual.html>

¹⁷<https://github.com/opcm/pcm>

¹⁸<https://en.wikipedia.org/wiki/SpeedStep>

8 describe our model equations with parameter values for the three systems respectively for MonetDB, PostgreSQL and Hyrise.

$$\begin{aligned}
P(OP_j^i) = & -2,44 * 10^{-7} * C_{cpu} + 9,17 * 10^{-5} * C_{mem} + 1,99 * 10^{-6} * C_{dio} + 1,51 * 10^{-13} * C_{cpu} \\
& * C_{mem} - 7,22 * 10^{-11} * C_{mem} * C_{dio} + 1,67 * 10^{-13} * C_{cpu} * C_{dio} - 2,82 * 10^{-18} \\
& * C_{cpu} * C_{mem} * C_{dio} + 3,27 * 10^{-14} * C_{cpu}^2 - 1,46 * 10^{-9} * C_{mem}^2 - 1,56 * 10^{-12} \\
& * C_{dio}^2 - 3,74 * 10^{-19} * C_{cpu}^2 * C_{mem} + 1,07 * 10^{-16} * C_{cpu} * C_{mem}^2 + 9,67 * 10^{-21} \\
& * C_{cpu}^2 * C_{dio} + 2,93 * 10^{-16} * C_{mem}^2 * C_{dio} + 1,83 * 10^{-20} * C_{cpu} * C_{dio}^2 + 2,58 * 10^{-17} \\
& * C_{mem} * C_{dio}^2 + 3,87 * 10^{-21} * C_{cpu}^3 - 2,30 * 10^{-15} * C_{mem}^3 + 2,07 * 10^{-19} * C_{dio}^3 \\
& + 47,26
\end{aligned} \tag{6}$$

$$\begin{aligned}
P(PL_j) = & 4,53 * 10^{-6} * C_{dio} - 1,46 * 10^{-6} * C_{mem} - 1,12 * 10^{-6} * C_{cpu} + 2,30 * 10^{-14} \\
& * C_{cpu} * C_{mem} - 5,88 * 10^{-12} * C_{mem} * C_{dio} + 5,97 * 10^{-15} * C_{cpu} * C_{dio} \tag{7} \\
& + 2,46 * 10^{-17} * C_{cpu}^2 + 8,01 * 10^{-12} * C_{mem}^2 - 9,25 * 10^{-13} * C_{dio}^2 + 48,8
\end{aligned}$$

$$\begin{aligned}
P(OP_j^i) = & 23,61 - 4,63 * 10^{-10} * C_{cpu} + 3,53 * 10^{20} * C_{cpu}^2 - 7,30 * 10^8 * C_{L2} - 1,69 * 10^{-17} \\
& * C_{cpu} * C_{L2} + 6,59 * 10^{-15} * C_{L2}^2 + 9,02 * 10^{-6} * C_{L3} + 1,96 * 10^{-16} * C_{cpu} * C_{L3} \\
& - 1,43 * 10^{-13} * C_{L2} * C_{L3} + 8,67 * 10^{-14} * C_{L3}^2 + 1,85 * 10^{-7} * C_{mem} + 1,99 * 10^{-17} \\
& * C_{cpu} * C_{mem} - 2,20 * 10^{-14} * C_{L2} * C_{mem} + 2,39 * 10^{-13} * C_{L3} * C_{mem} \\
& + 1,92 * 10^{-14} * C_{mem}^2
\end{aligned} \tag{8}$$

The Random Forest was trained upon all the statistics collected from the queries. Training parameters were set by changing each parameter one-by-one and by regenerating the regression model. The optimum value of each parameter was selected from the following ranges: mtry-from 1 to 300; ntree-from 1 to 3. For nodesize value, it is fixed to the default. After analysis of output, the parameters were selected with values : for Hyrise model training(mtry=3 and ntree=260), for PostgreSQL(mtry=2 and ntree=30) and for monetdb(mtry=2 and ntree=40). To assess the quality of our models with the different machine learning techniques, the fit to the training data was taken into account. Table 2 gives more information about the fit to the training data. We give for each models the correlation's coefficient (squared correlation) denoted by R^2 , Adjusted R-squared and Root Mean of Squared Error(RMSE).

5.2 Models Validation

In this subsection, we present the results of a series of experiments running the standard workload provided by the two benchmarks TPC-H and TPC-DS¹⁹. We used our energy models (Non-Linear Regression or Random Forest) defined in equations 6, 7 and 8 to execute the 22 queries on a 5GB data size for PostgreSQL and Monetdb. for Hyrise, we run the queries on 2GB and 3GB. We compare the energy estimated (EE) by our cost models with the real energy (ER) consumed by the database server during the processing of the query in parallel mode. To validate the

¹⁹<http://www.tpc.org/tpcds/>

| Variables | PostgreSQL | | Monet | | Hyrise | |
|--------------------|------------|------|-----------|------|-------------|------|
| | NLR | RF | NLR | RF | NLR | RF |
| Multiple R-squared | 0.83 | - | 0.69 | - | 0.61 | - |
| Adjusted R-squared | 0.78 | - | 0.45 | - | 0.54 | - |
| Residual interval |] - 2, 2[| - |] - 3, 3[| - |] - 10, 15[| - |
| Root Mean (RMSE) | - | 1.62 | - | 0.79 | - | 4.35 |

Table 2: Regression statistics and residual outputs

energy model as an accurate prediction of cost for database system, we use the metric called Estimation Error (ERR) to quantify the model accuracy. The metric is defined by the following formula: $ERR = \frac{|ER-EE|}{ER}$. ER denotes the real values of power measured by the power meter and EE denotes the prediction values of our models using Non Linear Regression (NLR) or Random Forest Regression(RFR).

5.2.1 Results Analysis on TPC-H benchmark:

The tables 3 and 4 present the estimation error rates obtained using our energy cost model described in 7 on the PostgreSQL system using the Non Linear Regression(NLR) technique and the Random Forest Regression(RFR) technique. Note that in Table 3, the average of the estimation errors is 4.1% and the maximum error is 10.4% using the Non Linear Regression technique and in Table 4, the average estimate errors is 4.3% and the maximum error is 13.1% in the case of Random Forest Regression(RFR).

Table 3: EER for the TPC-H benchmark queries with Postgresql DBMS SF5 - NLR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|-------|
| 1 | 56.37 | 56.66 | 0.5% |
| 2 | 50.09 | 48.86 | 2.5% |
| 3 | 51.22 | 51.98 | 1.5% |
| 4 | 51.21 | 48.51 | 5.6% |
| 5 | 51.17 | 47.36 | 8.0% |
| 6 | 51.74 | 50.67 | 2.1% |
| 7 | 52.51 | 51.62 | 1.7% |
| 8 | 51.13 | 47.96 | 6.6% |
| 9 | 46.57 | 51.48 | 9.5% |
| 10 | 50.49 | 54.18 | 6.8% |
| 11 | 50.69 | 52.07 | 2.6% |
| 12 | 51.74 | 50.76 | 1.9% |
| 13 | 52.81 | 58.97 | 10.4% |
| 14 | 52.13 | 51.71 | 0.8% |
| 15 | 49.84 | 53.19 | 6.3% |
| 16 | 50.32 | 53.19 | 5.4% |
| 17 | 50.86 | 53.67 | 5.2% |
| 19 | 52.25 | 52.24 | 0.0% |
| 20 | 51.82 | 50.93 | 1.7% |
| 22 | 50.64 | 49.40 | 2.5% |

Table 4: EER for the TPC-H benchmark queries with Postgresql DBMS SF5 - RFR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|-------|
| 1 | 56.39 | 56.66 | 0.5% |
| 2 | 54.40 | 48.86 | 11.4% |
| 3 | 51.22 | 51.98 | 1.5% |
| 4 | 51.45 | 48.51 | 6.1% |
| 5 | 52.29 | 47.36 | 10.4% |
| 6 | 51.72 | 50.67 | 2.1% |
| 7 | 50.96 | 51.62 | 1.3% |
| 8 | 51.52 | 47.96 | 7.4% |
| 9 | 52.71 | 51.48 | 2.4% |
| 10 | 51.01 | 54.18 | 5.8% |
| 11 | 54.42 | 52.07 | 4.5% |
| 12 | 51.33 | 50.76 | 1.1% |
| 13 | 51.26 | 58.97 | 13.1% |
| 14 | 51.04 | 51.71 | 1.3% |
| 15 | 51.02 | 53.19 | 4.1% |
| 16 | 53.96 | 53.19 | 1.5% |
| 17 | 51.37 | 53.67 | 4.3% |
| 19 | 50.95 | 52.24 | 2.5% |
| 20 | 51.39 | 50.93 | 0.9% |
| 22 | 51.50 | 49.40 | 4.3% |

When collecting system analysis data, the *Explain Analysis* command on queries 18 and 21 took a long time without giving any answers. We did not integrate them in the case of *PostgreSQL*.

Table 5: EER for the TPC-H benchmark queries with Monet DBMS SF5 - NLR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|-------|
| 1 | 48.85 | 49.82 | 2.0% |
| 2 | 48.10 | 48.30 | 0.4% |
| 3 | 48.72 | 47.83 | 1.9% |
| 4 | 48.40 | 48.15 | 0.5% |
| 5 | 47.83 | 47.34 | 1.0% |
| 6 | 57.71 | 51.50 | 13% |
| 7 | 48.24 | 47.90 | 0.5% |
| 8 | 47.98 | 47.96 | 0.01% |
| 9 | 47.41 | 48.27 | 1.8% |
| 10 | 50.13 | 50.28 | 0.3% |
| 11 | 47.27 | 48.99 | 3.5% |
| 12 | 48.62 | 47.16 | 3.1% |
| 13 | 48.07 | 48.01 | 0.1% |
| 14 | 48.03 | 46.86 | 2.5% |
| 15 | 47.51 | 45.55 | 4.3% |
| 16 | 48.80 | 47.71 | 2.3% |
| 17 | 48.92 | 47.57 | 2.8% |
| 18 | 48.56 | 47.82 | 1.6% |
| 19 | 44.08 | 48.19 | 8.5% |
| 20 | 48.40 | 47.3 | 2.3% |
| 21 | 48.70 | 48.27 | 0.9% |
| 22 | 48.21 | 47.36 | 1.8% |

Table 6: EER for the TPC-H benchmark queries with Monet DBMS SF5 - RFR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|------|
| 1 | 48.60 | 49.82 | 2.4% |
| 2 | 48.28 | 48.30 | 0.0% |
| 3 | 49.15 | 47.83 | 2.8% |
| 4 | 48.09 | 48.15 | 0.1% |
| 5 | 47.90 | 47.34 | 1.2% |
| 6 | 49.27 | 51.50 | 3.5% |
| 7 | 48.00 | 47.90 | 0.0% |
| 8 | 47.94 | 47.96 | 0.0% |
| 9 | 47.63 | 48.27 | 1.3% |
| 10 | 49.38 | 50.28 | 1.8% |
| 11 | 48.04 | 48.99 | 1.9% |
| 12 | 49.18 | 47.16 | 4.3% |
| 13 | 48.01 | 48.01 | 0.0% |
| 14 | 47.96 | 46.86 | 2.3% |
| 15 | 47.94 | 45.55 | 5.3% |
| 16 | 48.28 | 47.71 | 1.2% |
| 17 | 49.19 | 47.57 | 3.4% |
| 18 | 48.08 | 47.82 | 0.6% |
| 19 | 48.94 | 48.19 | 1.6% |
| 20 | 48.32 | 47.3 | 2.2% |
| 21 | 48.06 | 48.27 | 0.4% |
| 22 | 48.09 | 47.36 | 1.6% |

In tables 5 and 6, we give the estimation error rates obtained on the MonetDB system by using the two techniques. The average of the estimation errors is 2.5% and the maximum error is 13% by using the Non Linear Regression technique. By using the Random Forest technique, the average of errors is 1.7% and the maximum error is 5.3%.

The estimation error rates for Hyrise system are presented in tables 7 and 8. The average of the estimation errors is 9.0% and the maximum error is 22.14% using the Non Linear Regression technique and for Random Forest Regression, it is 7.4% and the maximum error is 17,9%.

To summarize, we present in Figure 14 the average errors obtained on the three database

Table 7: EER for the TPC-H benchmark queries with Hyrise DBMS SF2 - NLR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|-------|
| 1 | 70.22 | 61.1 | 14.9% |
| 2 | 46.42 | 43.93 | 5.7% |
| 3 | 47.08 | 42.10 | 11.8% |
| 4 | 46.95 | 48.40 | 3.0% |
| 5 | 46.92 | 48.83 | 3.9% |
| 6 | 48.44 | 39.66 | 22.1% |
| 7 | 46.95 | 42.15 | 11.4% |
| 8 | 45.42 | 42.90 | 5.9% |
| 9 | 45.57 | 50.03 | 8.9% |
| 10 | 48.80 | 46.63 | 4.6% |
| 11 | 45.93 | 41.66 | 10.2% |
| 12 | 46.92 | 43.90 | 6.9% |
| 13 | 48.25 | 46.50 | 3.8% |
| 14 | 47.02 | 41.90 | 12.2% |
| 15 | 46.74 | 42.62 | 9.7% |
| 16 | 46.53 | 48.85 | 4.8% |
| 17 | 47.63 | 43.92 | 8.4% |
| 18 | 45.41 | 52.30 | 13.1% |
| 19 | 46.70 | 50.06 | 6.7% |
| 20 | 46.93 | 42.30 | 10.9% |
| 21 | 48.83 | 55.02 | 11.2% |
| 22 | 46.99 | 43.30 | 8.5% |

Table 8: EER for the TPC-H benchmark queries with Hyrise DBMS SF2 - RFR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|-------|
| 1 | 63.84 | 61.1 | 4.5% |
| 2 | 45.90 | 43.93 | 4.5% |
| 3 | 47.51 | 42.10 | 12.9% |
| 4 | 48.02 | 48.40 | 0.8% |
| 5 | 48.98 | 48.83 | 0.3% |
| 6 | 44.75 | 39.66 | 12.9% |
| 7 | 47.34 | 42.15 | 12.3% |
| 8 | 50.56 | 42.90 | 17.9% |
| 9 | 52.37 | 50.03 | 4.7% |
| 10 | 49.82 | 46.63 | 6.9% |
| 11 | 45.74 | 41.66 | 9.8% |
| 12 | 47.30 | 43.90 | 7.8% |
| 13 | 50.63 | 46.50 | 8.9% |
| 14 | 44.67 | 41.90 | 6.6% |
| 15 | 44.95 | 42.62 | 5.5% |
| 16 | 46.84 | 48.85 | 4.1% |
| 17 | 46.96 | 43.92 | 6.9% |
| 18 | 51.80 | 52.30 | 1.0% |
| 19 | 46.91 | 50.06 | 6.3% |
| 20 | 46.89 | 42.30 | 10.9% |
| 21 | 59.44 | 55.02 | 8.0% |
| 22 | 47.26 | 43.30 | 9.1% |




Figure 14: Queries energy consumption average errors

systems by exploiting the two learning techniques: Non Linear Regression (NLR) and Random Forest(RF).

Our finding demonstrates the ability of our models to be exploitable. Almost, in all cases, the estimated energy is very close to the real values obtained by the measuring equipment. These finding allow us to pronounce on the quality and accuracy of our models on the TPC-H benchmark, however this is not an indication of predictive ability on other benchmark or on the other hardware configuration. As shown in the figure 14, by comparing the estimation errors obtained using the two training techniques on our models, we find that the Random Forest technique fits better than the nonlinear regression in our study. However, in the case of the PostgreSQL system, the average error of the random forest is similar to that of the Non Linear Regression. The reasons for the performance of Random Forest can be explained by that it is easier to train as it includes a descriptor selection procedure and is not strongly dependent up on training parameters. Random Forests are immune to the problems of overfitting[8].

5.2.2 Results Analysis on TPC-DS benchmark

To evaluate the quality of the model proposed for the Hyrise system, we used it on another benchmark (new DB diagram) in order to evaluate its portability. For this, we use the benchmark TPC-DS which is more complex than the TPC-H benchmark because of its diverse scheme, its distribution data and its decision support request load. We used a dataset with a size of 3 GB. We selected 11 queries out of the 99 random queries so that the aggregation functions defined within the queries are implemented in the Hyrise system.

The results of queries execution are shown in Table 9. Note that the results of this new configuration are also of good quality with the Random Forest learning technique, as we can see in Table 9. The average error is 10% and the maximum error is below 19%. The model on the other hand with the Non Linear Regression technique was not able to predict the energy values accurately. We found an average error greater than 50%.

Energy modelization and optimization in database systems is the basic work to design energy-efficient database systems. The methods used in the literature to predict the energy consumption of queries in databases are roughly divided between those, which use linear regression methods, and those, which employ Non Linear Regression. Here we have used Random Forest Regression technique and we found that it has performed better than Non-linear Regression for the data set used.

Table 9: EER for the TPC-DS benchmark queries with Hyrise DBMS SF3 - RFR Technics

| Queries | Prediction | Measured | EER |
|---------|------------|----------|-------|
| 1 | 46.67 | 43,06 | 8.4% |
| 2 | 46.63 | 42,74 | 9.1% |
| 3 | 46.65 | 41,9 | 11.3% |
| 4 | 46.63 | 43,28 | 7.7% |
| 5 | 46.63 | 47,66 | 2.2% |
| 6 | 46.63 | 39,34 | 18,5% |
| 7 | 46.63 | 40,82 | 14.2% |
| 8 | 46.63 | 41,98 | 10.8% |
| 9 | 46.63 | 42,08 | 11.1% |
| 10 | 46.63 | 42,9 | 8.7% |
| 11 | 46.63 | 43,02 | 8.4% |

VI RELATED WORK

In this section, we review then main studies and research efforts dedicated to reduce the energy consumption when processing data. This review covers major elements of the DBMS environment. Figure 15 lists the major approaches explored to increase the EE.

6.1 Hardware level

6.1.1 Dynamique Power Management

DPM techniques applied at the hardware and firmware level can be broadly divided into two categories: Dynamic Component Deactivation (DCD) and Dynamic Performance Scaling (DPS). Computer components that do not have the ability to automatically adjust their performance must be completely deactivated to save power consumption, this approach is called Dynamic Component Deactivation. Dynamic Performance Scaling applies different techniques like Dynamic Voltage and Frequency Scaling (DVFS) technique to computer components in the aims to adjust dynamically their performance to the power consumption. Adopting dynamic power-management could help improve energy-efficiency. For example the studies in [16],[52],[43] use the frequency and/or voltage mechanism that allow DBMS to trade energy consumption for performance. In order to save energy, they manage the DVFS level based on throughput target and workload characteristics (e.g., I/O or CPU intensive). In the aims to save power consumption of online transaction processing (OLTP) in a multicore environment, the authors in [33] have




Figure 15: Overview of techniques to improve the EE

proposed an application-aware that dynamically scale the operating frequency of processors based on response time. Application-aware use a trade-off between power consumption and response time accepted as system level agreement. The authors report a reduction of 7.6% of total power consumption on TPC-C benchmark queries. The authors of [67], experimentally studied the impact of reducing the clock frequency of the processor on the energy efficiency of common database algorithms such as scans, aggregations, hash joins. They showed that reducing unused computing power significantly improves the energy efficiency of memory-bound database algorithms. [65] proposed the Predictive Energy Saving Online Scheduling (PESOS) algorithm in the context of Web search. Using the DVFS technique, PESOS selects the most appropriate CPU frequency to process a query. The goal of PESOS is to reduce the CPU energy consumption of a query-processing node while imposing a required tail latency on the query response times. The paper of [73] presents a technique called *POLARIS* for reducing the power consumption of transactional database systems. *POLARIS* directly manages processor DVFS within the DBMS and controls database transaction scheduling. Its goal is to minimize power consumption while ensuring the transactions are completed within a specified latency target. The work of [5] proposes a control-theoretic approach for DVS in multiprocessor system-on-chip pipelined architectures. It uses linear and non-linear feedback control policies to adapt frequencies using queues occupancy for incoming streams. It aims at controlling the inter-processor queue occupancy. In the linear control technique, the same feedback strategy can be iterated backward and applied in sequence to each previous stage using linear functions. Non-linear techniques use non-linear feedback functions to control the operation of systems that are either non-linear, time-varying, or a combination of both.

6.1.2 Microarchitectural Power Saving

Microarchitectural techniques allowing energy saving in specific components (eg main memory, cache memories, etc.) are widely known and studied in the literature. the authors in [53] used DRAM frequency scaling and power-down modes to improve power consumption without sacrificing performance under both transactional and analytical workloads. They start by identifying the hardware features that help reduce DRAM power consumption, and then they perform an experimental analysis of the impact of those features on power and performance. The work of [55] proposes a hybrid memory architecture consisting of both Dynamic Random Access Memory (DRAM) and Non-Volatile Memory (NVM) to reduce database energy consumption, through an application-level data management policy that decides to place data on DRAM or NVM. In a recent work, the authors of [72] proposed the Energy-Control Loop (*ECL*) as a DBMS-integrated approach for adaptive energy-control on in-memory database systems that obeys a query latency limit as a soft constraint and actively optimizes energy efficiency and performance of the DBMS. The *ECL* relies on adaptive workload-dependent energy profiles that are continuously maintained at runtime. In[44] authors present survey on cache power/energy-consumption reduction techniques. They have surveyed various cache tuning techniques based on two categories, offline static cache tuning and online dynamic cache tuning, ranging from hardware support to cache tuning techniques. In[46] the effect of power and performance are analyzed by reducing cache capacity. Reducing cache size in some applications give significant energy saving in order of 3.17% and 24.16% in dynamic and static energy respectively while degrade the performance by increasing cache misses.

6.1.3 Telecommunication Network Power Management

Telecommunication networks constitute a major sector of ICT and they undergo a tremendous growth. Energy is mostly consumed in network transmission and switching equipment such as routers. In the literature, we observe that the focus of many research is on control and optimization strategies for computer network equipment enabling energy saving, by adapting network capacities and computing resources to the current traffic load and demands, while ensuring end-to-end quality of service (QoS)[45]. The approaches to reduce energy consumption in networks is divided into four categories: (i) selectively turning off network elements,(ii) energy-efficient network design, (iii) energy-efficient IPpacket forwarding, and (iv) green routing. by Zhang and al. [30]. In [7], the authors propose that WLANs be re-designed so that the power consumed by the WLAN scales with its offered load. Such strategies allow powering on or off WLAN access points dynamically, based on the volume and location of user demand. To achieve a power saving router, Authors in [21] proposed to follow two approaches; power efficient designing and power saving designing. The former is an approach to create a high performance router at low power consumption and the latter is an approach to save wasted power. For power efficient designing, they have developed technologies for integrating the ASICs/FPGAs and memories of routers. for power saving design, the authors worked on static performance control, which allows turning off unused ports and modules. Furthermore, they proposed a technology that lowers the frequency of lightly utilized modules to save the wasted power. Running under the low frequency mode, the power efficiency improved by 10-20%.

6.2 Software Level

6.2.1 OS-Based Power Management

The power-efficient resource management at the OS level plays an important role in energy efficiency and it completes the effort provided at the hardware level. The authors in [6] have developed a real-time power manager in the kernel for the Linux OS called the On-Demand Governor (LGD). The goal of the on-demand governor is to keep the performance loss due to reduced frequency to the minimum. [20] discusses the behavior of the current task management subsystems (scheduler and load balancer) in the Linux kernel on a multi-core SMP system and also it covers its effectiveness in saving energy consumption under several situations (e.g. idle, moderate load). It then describes several techniques such as timer migration, task wakeup biasing and related heuristics for reducing energy consumption. Rajkumar et al. [2] have proposed several algorithms for application of DVFS in real-time systems and have implemented a prototype as a modified Linux kernel Linux/Resource Kernel (Linux/RK). The objective is to minimize the energy consumption, while maintaining the performance isolation between applications.

6.2.2 Energy-Efficiency in DBMS query processing

In those Approaches, most important studies concern (i) the definition of cost models to predict the energy and (ii) the proposition of cost-driven techniques for reducing energy. In [43, 29] the authors build a static power profile for each basic database operation using a simple linear regression technique. In [40], the authors use pipeline-based modeling to the sources of peak power consumption for a query and to recommend plans with low peak power. the work of [34] proposes a framework for energy-aware database query processing. It extends query plans produced by traditional query optimizer with an energy consumption prediction for some specific database operators like select, project and join using linear regression technique. In [58], the authors proposed cost models to predict the power consumption of single running queries. the proposed model relies on pipeline segmenting of the query. The model is built based on the CPU, I/O cost of each pipeline and measured energy in an offline task, using non-linear regression technique (polynomial regression). The authors in [66] analyze the effect of the three main cache structures (Database Buffer Cache, Dictionary Cache, and Library Cache). Based on this, they have taken the cost of memory into account in their linear model cost dedicated to sequential query processing mode. In [70][81], the authors confront the sequential and parallel execution modes and study their impact on *EE*. They extended the sequential cost model proposed in [58] in order to predict the energy consumption of queries when they are executed in parallel.

6.3 Environmental management and conventional rules

In response to the over energy consumption in datastores and the need to reduce the related environmental, economic and energy supply security impacts, the governments, and corporate sector impose regulations and acts. The European Commission (EC) created a Code of Conduct for Energy Efficiency²⁰ in Data Centers in 2008 with the aim of improving the energy efficiency. Like EC, the U.S. Department of Energy has the ENERGY STAR²¹ program, which offers energy efficiency guidelines for all types of buildings including data centers. In addition to the Code of Conduct for Energy Efficiency, there are several energy-efficiency standards imposed to control the manufacture and use of equipment.

VII CONCLUSION

In this paper, we introduced a methodology to accurately estimate the power and energy consumption of queries in a database server. Our proposals center around four main steps for the construction of an energy model: (1) establishment of a deep audit that allows understanding the query processor functioning, (2) identification of relevant energy-sensitive parameters belonging to hardware and software components, (3) elaboration of mathematical cost models estimating the consumed energy when executing a query on a target DBMS and, (4) setting the values of the energy-sensitive parameters using Non-Linear Regression and Random Forest techniques. One important feature of our methodology is the fact that we use novel energy-sensitive parameters and different machine learning techniques. This procedure highlighted also the importance of audit of DBMSs by understanding their functioning, before proposing any energy modeling. For that purpose, we perform thorough experiments on the power consumption patterns of various workloads generated from the TPC-H and TPC-DS benchmarks on three DBMSs (PostgreSQL, MonetDB and Hyrise) to validate our proposal. In these experiments, we first evaluated and compared the data loading and query execution performances versus the power consumption of

²⁰<https://ec.europa.eu/jrc/en/energy-efficiency/code-conduct/datacentres>

²¹<https://www.energystar.gov/about>

our tested systems. Then we proposed three comprehensive energy cost models, one for each of the tested DBMSs. Later, we validated the models by running a second set of queries, and compared their power cost versus those predicted by the models in each system. Our results show that the predicted power is very close to the real values obtained by the measuring equipment. By comparing the estimation errors obtained using the two training techniques on our models, we find that the Random Forest technique gives better predictions than the Non-Linear regression.

We plan in our future work to use other machine learning techniques and compare them with those used in this work on much larger datasets to determine their impact on our energy models's accuracy. Also, we plan to apply the same approach to other systems and compare them with monetDB, PostgreSQL, and Hyrise systems.

REFERENCES

References

- [1] J. Rawlings. *Applied regression analysis: a research tool*. Wadsworth & Brooks/Cole statistics/probability series. Wadsworth & Brooks/Cole Advanced Books & Software, 1988. ISBN: 9780534092467.
- [2] R. Rajkumar, K. Juvva, A. Molano, and S. Oikawa. “Resource kernels: a resource-centric approach to real-time and multimedia systems”. In: *Multimedia Computing and Networking 1998*. Vol. 3310. SPIE, 1997, pp. 150–164.
- [3] S. Manegold, P. A. Boncz, and M. L. Kersten. “Generic Database Cost Models for Hierarchical Memory Systems”. In: 2002, pp. 191–202.
- [4] D. Anderson, J. Dykes, and E. Riedel. “More Than an Interface SCSI vs ATA”. In: ed. by J. Chase. USENIX, 2003.
- [5] A. Alimonda, A. Acquaviva, S. Carta, and A. Pisano. “A control theoretic approach to run-time energy optimization of pipelined processing in MPSoCs”. In: *Proceedings of the Conference on Design, Automation and Test in Europe*. 2006, pp. 876–877.
- [6] V. Pallipadi and A. Starikovskiy. “The ondemand governor”. In: *Proceedings of the Linux Symposium*. Vol. 2. sn. 2006, pp. 215–230.
- [7] A. P. Jardosh, G. Iannaccone, K. Papagiannaki, and B. Vinnakota. “Towards an Energy-Star WLAN Infrastructure”. In: ed. by E. de Lara and N. Bhatti. IEEE Computer Society, 2007, pp. 85–90.
- [8] D. S. Palmer, N. M. O’Boyle, R. C. Glen, and J. B. O. Mitchell. “Random Forest Models To Predict Aqueous Solubility”. In: *J. Chem. Inf. Model.* 47.1 (2007), pp. 150–158.
- [9] R. Agrawal, A. Ailamaki, P. A. Bernstein, E. A. Brewer, M. J. Carey, S. Chaudhuri, A. Doan, D. Florescu, M. J. Franklin, H. Garcia-Molina, et al. “The Claremont report on database research”. In: *ACM Sigmod Record* 37.3 (2008), pp. 9–19.
- [10] P. A. Boncz, M. L. Kersten, and S. Manegold. “Breaking the Memory Wall in MonetDB”. In: *Commun. ACM* 51.12 (Dec. 2008), pp. 77–85.
- [11] M. Poess and R. O. Nambiar. “Energy cost, the key challenge of today’s data centers: a power consumption analysis of TPC-C results”. In: *PVLDB* 1.2 (2008), pp. 1229–1240.
- [12] S. Harizopoulos, M. A. Shah, J. Meza, and P. Ranganathan. “Energy Efficiency: The New Holy Grail of Data Management Systems Research”. In: *ArXiv abs/0909.1784* (2009).
- [13] M. G. Ivanova, M. L. Kersten, N. J. Nes, and R. A. Gonçalves. “An Architecture for Recycling Intermediates in a Column-store”. In: *ACM SIGMOD*. 2009, pp. 309–320.

- [14] K. Kant. “Data center evolution: A tutorial on state of the art, issues, and challenges”. In: *Computer Networks* 53.17 (2009), pp. 2939–2965.
- [15] W. Lang and J. Patel. “Towards eco-friendly database management systems”. In: *arXiv preprint arXiv:0909.1767* (2009).
- [16] W. Lang and J. Patel. “Towards eco-friendly database management systems”. In: *arXiv preprint arXiv:0909.1767* (2009).
- [17] A. B. Lekov, L. Thompson, A. T. McKane, A. Rockoff, and M. A. Piette. *Opportunities for Energy Efficiency and Automated Demand Response in Industrial Refrigerated Warehouses in California*. Report. Lawrence Berkeley National Laboratory, 2009.
- [18] L. Minas and B. Ellison. *Energy efficiency for information technology: How to reduce power consumption in servers and data centers*. Intel Press, 2009.
- [19] A. Omer. “Energy use and environmental impacts: A general review”. In: *Journal of Renewable and Sustainable Energy* 1 (Sept. 2009).
- [20] V. Srinivasan, G. R. Shenoy, S. Vaddagiri, and D. Sarma. “Energy-aware task and interrupt management in Linux”. In: *Ottawa Linux Symposium*. 2009.
- [21] M. Yamada, T. Yazaki, N. Matsuyama, and T. Hayashi. “Power Efficient Approach and Performance Control for Routers”. In: *2009 IEEE International Conference on Communications Workshops*. 2009, pp. 1–5.
- [22] F. Ahmad and T. N. Vijaykumar. “Joint Optimization of Idle and Cooling Power in Data Centers While Maintaining Response Time”. In: *Proceedings of the Fifteenth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS XV. Pittsburgh, Pennsylvania, USA: ACM, 2010, pp. 243–256. ISBN: 978-1-60558-839-1.
- [23] A. M. Amin and Z. A. Chishti. “Rank-aware Cache Replacement and Write Buffering to Improve DRAM Energy Efficiency”. In: *Proceedings of the 16th ACM/IEEE International Symposium on Low Power Electronics and Design*. ISLPED ’10. New York, NY, USA: ACM, 2010.
- [24] R. Elmasri and S. Navathe. *Fundamentals of Database Systems*. 6th. USA: Addison-Wesley Publishing Company, 2010.
- [25] M. Grund, J. Krüger, H. Plattner, A. Zeier, P. Cudre-Mauroux, and S. Madden. “HYRISE: A Main Memory Hybrid Storage Engine”. In: *Proc. VLDB Endow.* 4.2 (2010).
- [26] M. Iyengar, R. Schmidt, and J. Caricari. “Reducing energy usage in data centers through control of Room Air Conditioning units”. In: 2010, pp. 1–11.
- [27] D. Schall, V. Hudlet, and T. Härder. “Enhancing energy efficiency of database applications using SSDs”. In: *Proceedings of the Third C* Conference on Computer Science and Software Engineering*. ACM. 2010, pp. 1–9.
- [28] D. Tsirogiannis, S. Harizopoulos, and M. A. Shah. “Analyzing the energy efficiency of a database server”. In: *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2010, Indianapolis, Indiana, USA, June 6-10, 2010*. ACM, 2010, pp. 231–242.
- [29] Z. Xu, Y.-C. Tu, and X. Wang. “Exploring power-performance tradeoffs in database systems”. In: *ICDE*. 2010, pp. 485–496.
- [30] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee. “Energy Efficiency in Telecom Optical Networks”. In: *IEEE Commun. Surv. Tutorials* 12.4 (2010), pp. 441–458.
- [31] A. Beloglazov, R. Buyya, Y. C. Lee, A. Zomaya, et al. “A taxonomy and survey of energy-efficient data centers and cloud computing systems”. In: *Advances in Computers* 82.2 (2011), pp. 47–111.

- [32] R. Goncalves and M. Kersten. “The Data Cyclotron Query Processing Scheme”. In: *ACM Trans. Database Syst.* 36.4 (2011), pp. 1–35.
- [33] Y. Hayamizu, K. Goda, M. Nakano, and M. Kitsuregawa. “Application-aware power saving for online transaction processing using dynamic voltage and frequency scaling in a multicore environment”. In: *International Conference on Architecture of Computing Systems*. Springer. 2011, pp. 50–61.
- [34] W. Lang, R. Kandhan, and J. M. Patel. “Rethinking Query Processing for Energy Efficiency: Slowing Down to Win the Race.” In: *IEEE Data Eng. Bull.* 34.1 (2011), pp. 12–23.
- [35] H. Ong, T. Mahlia, and H. Masjuki. “A review on energy scenario and sustainable energy in Malaysia”. In: *Renewable and Sustainable Energy Reviews* 15.1 (2011), pp. 639–647.
- [36] N. Rasmussen. “Determining total cost of ownership for data center and network room infrastructure”. In: *Relatório técnico, Schneider Electric, Paris* 8 (2011).
- [37] M. Etinski, J. Corbalán, J. Labarta, and M. Valero. “Understanding the future of energy-performance trade-off via DVFS in HPC environments”. In: *J. Parallel Distrib. Comput.* 72.4 (2012), pp. 579–590.
- [38] S. Idreos, F. Groffen, N. Nes, S. Manegold, K. S. Mullender, and M. L. Kersten. “MonetDB: Two Decades of Research in Column-oriented Database Architectures”. In: *IEEE Data Eng. Bull.* 35.1 (2012), pp. 40–45.
- [39] M. Kunjir, P. K. Birwa, and J. R. Haritsa. “Peak power plays in database engines”. In: ACM. 2012, pp. 444–455.
- [40] M. Kunjir, P. K. Birwa, and J. R. Haritsa. “Peak power plays in database engines”. In: *EDBT*. ACM. 2012, pp. 444–455.
- [41] X. Liu, J. Wang, H. Wang, and H. Gao. “Generating power-efficient query execution plan”. In: *2nd International Conference on Advances in Computer Science and Engineering (CSE 2013)*. Atlantis Press. 2013.
- [42] Z. Xu, Y.-C. Tu, and X. Wang. “Dynamic energy estimation of query plans in database systems”. In: *2013 IEEE 33rd International Conference on Distributed Computing Systems*. IEEE. 2013, pp. 83–92.
- [43] Z. Xu, X. Wang, and Y.-c. Tu. “Power-aware throughput control for database management systems”. In: *Proceedings of the 10th International Conference on Autonomic Computing ({ICAC} 13)*. 2013, pp. 315–324.
- [44] W. Zang and A. Gordon-Ross. “A survey on cache tuning from a power/energy perspective”. In: *ACM Comput. Surv.* 45.3 (2013), 32:1–32:49.
- [45] K. Bilal, S. U. R. Malik, O. Khalid, A. Hameed, E. Alvarez, V. Wijaysekara, R. Irfan, S. Shrestha, D. Dwivedy, M. Ali, et al. “A taxonomy and survey on green data center networks”. In: *Future Generation Computer Systems* 36 (2014), pp. 189–208.
- [46] S. Chakraborty, D. Deb, D. Buragohain, and H. K. Kapoor. “Cache capacity and its effects on power consumption for tiled chip multi-processors”. In: *2014 International Conference on Electronics and Communication Systems (ICECS)*. IEEE. 2014, pp. 1–6.
- [47] B. He. “When data management systems meet approximate hardware: challenges and opportunities”. In: *Proceedings of the VLDB Endowment* 7.10 (2014), pp. 877–880.
- [48] P. Llopis, J. G. Blas, F. Isaila, and J. Carretero. “Survey of Energy-Efficient and Power-Proportional Storage Systems”. In: *The Computer Journal* 57.7 (2014), pp. 1017–1032.
- [49] I. Manotas, L. Pollock, and J. Clause. “SEEDS: a software engineer’s energy-optimization decision support framework”. In: *ICSE*. 2014, pp. 503–514.
- [50] S. Mittal. “A Survey of Techniques For Improving Energy Efficiency in Embedded Computing Systems”. In: *CoRR* abs/1401.0765 (2014).

- [51] Y.-C. Tu, X. Wang, B. Zeng, and Z. Xu. “A system for energy-efficient data management”. In: *ACM SIGMOD Record* 43.1 (2014), pp. 21–26.
- [52] Y.-C. Tu, X. Wang, B. Zeng, and Z. Xu. “A system for energy-efficient data management”. In: *ACM SIGMOD Record* 43.1 (2014), pp. 21–26.
- [53] R. Appuswamy, M. Olma, and A. Ailamaki. “Scaling the Memory Power Wall With DRAM-Aware Data Management”. In: *Proceedings of the 11th International Workshop on Data Management on New Hardware*. ACM. 2015, p. 3.
- [54] X. Cheng, B. He, and C. T. Lau. “Energy-Efficient Query Processing on Embedded CPU-GPU Architectures”. In: *DaMoN Workshop*. 2015, 10:1–10:7.
- [55] A. Hassan, H. Vandierendonck, and D. S. Nikolopoulos. “Energy-Efficient In-Memory Data Stores on Hybrid Memory Hierarchies”. In: *Proceedings of the 11th International Workshop on Data Management on New Hardware*. ACM. 2015, p. 1.
- [56] A. Nouredine and A. Rajan. “Optimising Energy Consumption of Design Patterns”. In: *Proceedings of the 37th International Conference on Software Engineering - Volume 2*. Florence, Italy: IEEE Press, 2015, pp. 623–626.
- [57] A. Roukh and L. Bellatreche. “Eco-processing of OLAP complex queries”. In: *International Conference on Big Data Analytics and Knowledge Discovery*. Springer. 2015.
- [58] A. Roukh and L. Bellatreche. “Eco-processing of olap complex queries”. In: *International Conference on Big Data Analytics and Knowledge Discovery*. Springer, 2015, pp. 229–242.
- [59] D. Abadi, R. Agrawal, A. Ailamaki, M. Balazinska, P. A. Bernstein, M. J. Carey, S. Chaudhuri, J. Dean, A. Doan, M. J. Franklin, et al. “The beckman report on database research”. In: *Communications of the ACM* 59.2 (2016), pp. 92–99.
- [60] M. Dayarathna, Y. Wen, and R. Fan. “Data Center Energy Consumption Modeling: A Survey”. In: *IEEE Communications Surveys Tutorials* 18.1 (2016), pp. 732–794.
- [61] A. Hurson and H. Azad. *Energy Efficiency in Data Centers and Clouds*. Academic Press, 2016.
- [62] F. Moore. *Data center energy consumption: Enormous data centers creating a hyperscale heat wave*. EN. White Paper, https://www.fujifilmusa.com/products/tape_data_storage/case_studies/pdf/Hyperscale_Heat_Wave.pdf. 2016.
- [63] A. Roukh, L. Bellatreche, and C. Ordonez. “EnerQuery: Energy-Aware Query Processing”. In: *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. Oct. 2016.
- [64] A. Shehabi, S. J. Smith, D. A. Sartor, R. E. Brown, M. Herrlin, J. G. Koomey, E. R. Masanet, N. Horner, I. L. Azevedo, and W. Lintner. *United States Data Center Energy Usage Report*. Report. Energy Technology Area, June 2016.
- [65] M. Catena and N. Tonellotto. “Energy-Efficient Query Processing in Web Search Engines”. In: *IEEE Transactions on Knowledge and Data Engineering* 29.7 (July 2017), pp. 1412–1425. ISSN: 1041-4347.
- [66] B. Guo, J. Yu, B. Liao, D. Yang, and L. Lu. “A green framework for DBMS based on energy-aware query optimization and energy-efficient query processing”. In: *Journal of Network and Computer Applications* 84 (2017), pp. 118–130.
- [67] S. Noll, H. Funke, and J. Teubner. “Energy Efficiency in Main-Memory Databases”. In: *Datenbank-Spektrum* 17.3 (2017), pp. 223–232.
- [68] A. Roukh, L. Bellatreche, A. Boukorca, and S. Bouarar. “Eco-Physic: Eco-Physical Design Initiative for Very Large Databases”. In: *Information Systems* 68 (2017), pp. 44–62.

- [69] R. Sadler. *Video demand drives up global CO2 emissions*. EN. Climate news network. 2017.
- [70] S. P. Dembele, A. Roukh, and L. Bellatreche. “Vers des Optimiseurs Verts de Requêtes en Mode Parallèle”. In: *Business Intelligence & Big Data, 14ème Edition de la conférence EDA, Tanger, Maroc*. Vol. B-14. RNTI. Editions RNTI, 2018, pp. 179–194.
- [71] L. Garris. *THE CURRENT STATE OF DATA CENTER ENERGY EFFICIENCY IN EUROPE*. Aug. 2018.
- [72] T. Kissinger, D. Habich, and W. Lehner. “Adaptive energy-control for in-memory database systems”. In: *Proceedings of the 2018 International Conference on Management of Data*. ACM. 2018, pp. 351–364.
- [73] M. Korkmaz, M. Karsten, K. Salem, and S. Salihoglu. “Workload-Aware CPU Performance Scaling for Transactional Database Systems”. In: *Proceedings of the 2018 International Conference on Management of Data*. ACM. 2018, pp. 291–306.
- [74] A. Ouared, Y. Ouhammou, and L. Bellatreche. “QoS MOS: QoS metrics management tool suite”. In: *Computer Languages, Systems & Structures* 54 (2018), pp. 236–251.
- [75] M. Raasveldt and H. Mühleisen. “MonetDBLite: An Embedded Analytical Database”. In: *CoRR* abs/1805.08520 (2018), pp. 1837–1838.
- [76] N. SCIENCE. *Anatomy of an Electromagnetic Wave*. NASA. Dec. 2018.
- [77] M. Dreseler, J. Kossmann, M. Boissier, S. Klauck, M. Uflacker, and H. Plattner. “Hyrise Re-engineered: An Extensible Database System for Research in Relational In-Memory Data Management.” In: *EDBT*. 2019, pp. 313–324.
- [78] S. N. Livieratos, S. Panetsos, A. Fotopoulos, and M. Karagiorgas. “A New Proposed Energy Baseline Model for a Data Center as a Tool for Energy Efficiency Evaluation”. In: 2019.
- [79] S. P. Dembele, L. Bellatreche, and C. Ordonez. “Towards Green Query Processing - Auditing Power Before Deploying”. In: *2020 IEEE International Conference on Big Data (Big Data)*. Atlanta, GA, USA: IEEE, 2020, pp. 2492–2501.
- [80] S. P. Dembele, L. Bellatreche, and C. Ordonez. “Vers le traitement et optimisation écologique des requêtes”. In: *CARI 2020 - Colloque Africain sur la Recherche en Informatique et en Mathématiques Appliquées*. Thiès, Senegal, Oct. 2020.
- [81] S. P. Dembele, L. Bellatreche, C. Ordonez, and A. Roukh. “Think big, start small: a good initiative to design green query optimizers”. In: *Cluster Computing* 23.3 (2020), pp. 2323–2345.
- [82] PostgreSQLfr. *PostgreSQL Documentation The PostgreSQL Global Development Group*. <https://www.postgresql.org/docs/11/intro-what-is.html>. Accessed: 2019-09-26.

A BIOGRAPHY



Simon Pierre Dembele is a third-year PhD student at ISAE-ENSMA, Poitiers, France, and a part-time faculty member in the Mathematics and Computer Science Department at University of Sciences, Techniques and Technologies of Bamako (USTTB), Mali. He got a Master Degree in computer science from the University of Boumerdes, Algeria. His research interest focuses on green database design.



Ladjel Bella-treche is a Full Professor at National Engineering School for Mechanics and Aerotechnics (ISAE-ENSMA), Poitiers, France. He leads the Data and Model Engineering Team of the Laboratory of Computer Science and Automatic Control for Systems (LIAS). He was a Visiting Professor of Harbin Institute of Technology (HIT) - China (2018, 2019), the Québec en Outaouais - Canada (2009), a Visiting Researcher at Purdue University - USA (2001) and Hong Kong University of Science and Technology, China (1997-1999). His research interest focuses on Data Management Systems and Semantic Web. He serves as an Associate Editor of the Data & Knowledge (DKE) Journal, Elsevier. His research projects have been funded by 2 EU projects.



Carlos Ordonez studied at UNAM University in Mexico, getting a B.Sc. in applied mathematics and an M.S. in computer science. He continued PhD studies at the Georgia Institute of Technology advised by Edward Omiecinski, focusing on accelerating machine learning algorithms with database systems techniques, getting the PhD in 2000. Carlos worked at NCR from 1998 to 2006, collaborating in the optimization of machine learning and cube query processing algorithms on the Teradata parallel DBMS. In 2006 Carlos joined the Department of Computer Science at the University of Houston, where he currently leads the Parallel Data Systems lab. From 2013 to 2015 Carlos collaborated with Michael Stonebraker, regularly visiting MIT. From July 2014 to July 2015 Carlos worked as a visiting researcher with ATT Labs-Research (formerly ATT Bell Labs), where he conducted research on stream analytics, the R language and data quality with Divesh Srivastava. His research projects have been funded by 3 NSF grants.