

Benchmarking Databases with History Support

Martin Kaufmann^{†§}, Peter M. Fischer[#], Norman May[§], Donald Kossmann[†]

[†]Systems Group, ETH Zürich, Switzerland, {martinka,donaldk}@inf.ethz.ch

[#]Albert-Ludwigs-Universität, Freiburg, Germany, peter.fischer@cs.uni-freiburg.de

[§]SAP AG, Walldorf, Germany, norman.may@sap.com

Abstract

An increasing number of applications such as risk evaluation in banking or inventory management require support for temporal data. After more than a decade of standstill, the recent adoption of some bitemporal features in SQL:2011 has reinvigorated the support among commercial database vendors, who incorporate an increasing number of relevant bitemporal features. Naturally, assessing the performance and scalability of temporal data storage and operations is of great concern for potential users. The cost of keeping and querying history with novel operations (such as time travel, temporal joins or temporal aggregations) is not adequately reflected in any existing benchmark. In this paper, we present a benchmark proposal which provides comprehensive coverage of the bitemporal data management. It builds on the solid foundations of TPC-H but extends it with a rich set of queries and update scenarios. This workload stems both from real-life temporal applications from SAP’s customer base and a systematic coverage of temporal operators proposed in the academic literature. In the accompanying paper [6] we present preliminary results of our benchmark on a number of temporal database systems, also highlighting the need for certain language extensions. In the appendix of this technical report we provide all details required to implement the benchmark.

1 Introduction

Temporal information is widely used in real-world database applications, e.g., to plan for the delivery of a product or to record the time a state of an order changed. Particularly the need for tracing and auditing the changes made to a data set and the ability to make decisions based on past or future assumptions are important use cases for temporal data. As a consequence, temporal features were included into the SQL:2011 standard [8], and an increasing number of database systems offer temporal features, e.g., Oracle, DB2, SAP HANA, or Teradata. As temporal data is often stored in an append-only mode, temporal tables quickly grow very large. This makes temporal processing a performance-critical aspect of many analysis tasks. Clearly, an understanding of the performance characteristics of different implementations of temporal queries is required to select the most appropriate database system for the desired workload. Unfortunately, at this time there is no generally accepted benchmark for temporal workloads.

For non-temporal data the TPC has defined TPC-H and TPC-DS for analytical tasks and TPC-C and TPC-E for transactional workloads. Especially TPC-H and TPC-C are popular for comparing database systems. These benchmarks query only the most recent version of

the data. We propose to leverage the insights gained with TPC-H and to TPC-C while widening the scope for temporal data. In particular, it should be possible to evaluate all TPC-H queries at different system times. This allows us to compare results on temporal data with those on non-temporal data. We carefully introduce additional parameters to examine the temporal dimension. Furthermore, we propose additional queries that resemble typical use cases we encountered in real-world use cases at SAP but also during literature review. In some cases, the expressiveness of SQL:2011 is not sufficient to express these queries in a succinct way. For example, the simulation of temporal aggregation in SQL:2011 results in rather complex queries.

More precisely, we propose a novel benchmark for temporal queries which are based on real-world use cases. As such, these queries retrieve both previous states of the system (i.e., a certain system time) but they also examine time intervals defined in the business domain (i.e., application time); this concept was introduced as the bitemporal data model by Snodgrass [11]. The benchmark we propose contains a data generator which first generates a TPC-H data set extended with some temporal data. In contrast to previous related work (such as [1] and [2]) it also generates a history of values using various business transactions on this data to generate system times. These transactions are inspired by the TPC-C benchmark, and they are designed to keep the characteristics generated by TPC-H `dbgen` at every point in time. Consequently, all TPC-H queries can be executed on the generated data, and their result properties for certain system times are comparable to those in the standard TPC-H benchmark. However, over time the overall data set grows as the previous versions are preserved in order to support time travel to a previous state of the system. In order to evaluate the time dimension, for this benchmark, we define additional queries which retrieve data at different points in time.

The remainder of this paper is structured as follows: In Section 2 we summarize the design goals for our proposed benchmark TPC-BiH. We survey related work on benchmarking temporal databases in Section 3. In the core part of the paper (Section 4), we define the schema, the data generator for temporal data, and the queries comprising the benchmark. In Section 5, we summarize our findings and point out future work. In Appendix A we show implementation details of the data generator and Appendix B includes the queries of the benchmark.

2 Goals and Methodology

The goal of this paper is to present a comprehensive benchmark for bitemporal query processing. This benchmark includes all necessary definitions as well as the relevant tools such as data generators. The benchmark setting reflects real-life customer workloads (which have typically not been formalized to match the current expression of the bitemporal model) and is complemented by synthetic queries to test certain operations. The benchmark is targeted towards SQL:2011, which has recently adopted core parts of the temporal data model. Since the expressiveness of SQL:2011 is limited (no complex temporal join, no temporal aggregation), we provide alternative versions of the queries using language extensions. Similarly, in order to support DBMS's which provide temporal support, but have not (yet) adopted SQL:2011 (like Oracle or Teradata), we provide alternative queries.

The schema builds on a well-understood existing non-temporal analytics benchmark: TPC-H. Its tables are extended with different types of history classes, such as degenerated, fully bitemporal or multiple user times. The benchmark data is designed to provide a range of different temporal update patterns, varying the ratio of history vs. initial data, the types of operations (UPDATE, INSERT and DELETE) as well as the temporal distributions within and between the temporal dimensions. The data distributions and correlations stay stable with regard to system time updates and evolve according to well-defined update scenarios in the application time domain. The data generator we developed can be scaled in the dimensions of initial data size and history length independently, providing support for many different scenarios.

Our query workload provides a coverage of common temporal DB requirements. It covers operations such as *time travel*, *key in time*, *temporal joins*, and *temporal aggregations* – the latter is not directly expressible in SQL:2011. Similarly, we investigate many patterns of storage access and time- vs. key-oriented access with varying ranges and selectivity. The query workload also covers the different temporal dimensions (system and application time): The focus of the queries is on stressing the system for individual time dimensions while considering correlations among the dimensions whenever relevant.

In summary, our benchmark fulfills the requirements mentioned in the benchmark handbook by Jim Gray [3], i.e., it is

- *relevant*, since it covers all typical temporal operations.
- *portable*, since it targets SQL:2011 and provides extensions for systems not completely supporting SQL:2011.
- *scalable*, since it provides well-defined data which can be generated in different sizes for base data and history.
- *understandable*, since all queries have a meaning in application scenarios and in terms of operator/system “stress”.

3 Related Work on Temporal Benchmarks

The foundation of the bitemporal data model was established in the proposal for TSQL2 [11]. For a single row the system time – in the original paper called *transaction time* – defines different versions as they were created by DML statements on a row. The system time is immutable, and the values are implicitly generated during transaction commit. Orthogonal to that, validity intervals can be defined on the application level – called *valid time* in the original paper. An example is the specification of the visibility of some marketing campaigns to users. Unlike the system time, the application time can be updated, and both interval boundaries may refer to times in the past, present, or future. The concept of the bitemporal model is now also applied in the SQL:2011 standard [8]. This standard focuses on basic operations like time travel on a single table. Complex temporal joins or aggregations are out of scope, but they are acknowledged as relevant scenarios for future versions of the SQL standard.

The benchmarks published by the TPC are the most commonly used benchmarks for databases. While these benchmarks used to focus either on analytical or transactional workloads, recently a combination has been proposed: The CH-benCHmark [2] extends the TPC-C schema by adding three tables from the TPC-H schema. Yet, no time dimension is included in these benchmarks.

Benchmarking the temporal dimension has been the focus of several studies: In 1995, a research proposal by Dunham et al. outlined possible directions and requirements for such a benchmark. The approach for building a temporal benchmark and the query classes come close to our methods.

A later work by Kuala and Robertson [5] provides logical models of several temporal database application areas alongside with queries expressed in an informal manner. The test suite of temporal database queries [4] from the TSQL2 editors provides a large number of temporal queries focused on functional testing rather than performance evaluation.

A study on spatio-temporal databases by Werstein [12] evaluates existing benchmarks and concludes that the temporal aspects of these benchmarks are insufficient. In turn, a number of queries are informally defined to overcome this limitation.

The work that is most closely related to ours was presented at TPCTC 2012 and includes a proposal to add a temporal dimension to the TPC-H Benchmark [1]. The authors also use TPC-H as a starting point, extend some tables with temporal columns to express bitemporal data, and rely on the data generator and the original queries of TPC-H as part of their workload. Yet, this work seems to be more focused on sketching the possibilities for the bitemporal data model rather than providing explicit definitions of data and queries. Specific differences exist in the language used (we focus on SQL:2011, in [1] a variant of TSQL2 is applied) as well as the derivation of application timestamps (we use existing temporal information in TPC-H for the initial version). Our update scenarios and queries cover a broader range of cases and aim to provide more properties on data and queries.

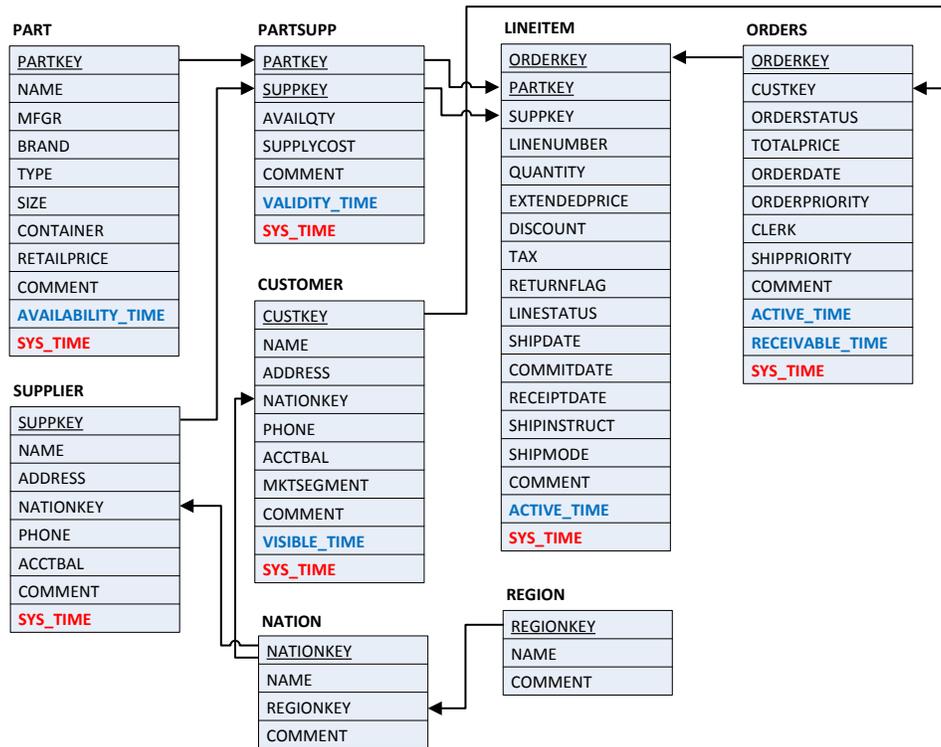


Figure 4.1: Schema

4 Definition of the TPC-BiH Benchmark

The definition of our benchmark consists of a schema, properties of the benchmark data and a range of queries. Our benchmark mainly targets the current SQL:2011 standard, but we also show examples how it can be translated to a system with other temporal expressions.

All queries and the implementation details of the data generator are given the Appendix of this technical report.

4.1 Schema

The schema we use in our benchmark is shown in Figure 4.1. As stated before, it is based on the TPC-H schema and adds temporal columns in order to express system and application times. Each of these time dimensions is stored as an interval and represented physically as two columns, e.g., `sys_time_begin` and `sys_time_end`. This means that any query defined on the TPC-H schema can run on our data set, and will give meaningful results, reflecting the current system time and the full range of application time versions. Specific other temporal dimensions can be added in a fairly straightforward manner. The additional temporal columns are chosen in such a way that the schema contains tables with different temporal properties: Some tables are kept unversioned, some express a correlated/degenerated behavior. Most tables are fully bitemporal, and we also consider the case in which a table has

multiple “user” times. Even if the latter is not well specified in the standard, we observe it a lot in customer use cases.

More specifically, we do not add any temporal columns to REGION and NATION. This is also plausible from application semantics, since this kind of information rarely ever changes. All other relations at least include a system time dimension. For SUPPLIER we simulate a degenerated table by only giving a system time. Since this single time dimension is determined by the loading/updating timing, we do not use any temporal correlation queries between this table and truly bi-dimensional tables. For all the remaining relations, we determine the application time from the existing information present in the data: Tuples in LINEITEM are valid as long as any operation like shipping them is pending. Likewise, tuples from PART are valid when they can be ordered, tuples from CUSTOMER when the customer is visible to the system, tuples from PARTSUPP when the price and the amount are valid. Finally, ORDERS has two time dimensions: *ACTIVE_TIME*: when was the order “active” (i.e., placed, but not delivered yet) and *RECEIVABLE_TIME*: when the bill for the order can be paid (i.e., invoice sent to customer, but not paid yet). Both application times become part of the schema. Since current DBMSs only support a single application time, we designate *ACTIVE_TIME* as such, and keep *RECEIVABLE_TIME* as a “regular” timestamp column. Likewise, if a DBMS does not provide any support for application time, application times are mapped to normal timestamp columns.

4.2 Benchmark Data

Complementing TPC-H with an extensive update workload has been proposed before. Given the structural similarity and the wide recognition, TPC-C has been used for this purpose, e.g., in [2]. We also used a similar approach (with additional timestamp assignment) in a previous version of the benchmark [7], but this proved to not be fully adequate: The set of update scenarios is quite small, and does not provide much emphasis on temporal aspects such as timestamp correlations. The query mix also constrains the flexibility in terms of temporal properties, e.g., since a fixed ratio of updates needs to go to specific tables.

The standard TPC-H has only a very limited number of “refresh” queries, which furthermore do not contain any updates to values. Nonetheless, the data produced by the data generator serves a good “initial” data set. The application time columns defined in the schema are initialized with the temporal data already present in this data: Extreme values of shipdate, commitdate and receiptdate define the validity interval of LINEITEM. Given the dependencies among the data items (e.g., LINEITEMS in an ORDER), we can now derive plausible application times for all bitemporal tables. Where needed, we complement this information with random distributions. The resulting data will contain data tuples with “open” time intervals, since customers or parts may have a validity far into the future.

To express the evolution of data, we define nine update scenarios, stressing different aspects among tables, values, and times:

1. *New Order*: Choose or create a customer, choose items and create an order on them.

2. *Cancel Order*: Remove an order, its dependent lineitems and adapt the number of available parts
3. *Deliver Order*: Update the order status and the lineitem status, adapt the available parts and the customer's balance.
4. *Receive Payment*: Update currently pending orders and the related customers' balances.
5. *Update Stock*: Increase available parts of a supplier.
6. *Delay Availability*: Postpone the date after which items are available from a supplier to a later date, e.g., due to a shipping backlog.
7. *Price Change*: Adapt the price of parts, choosing times from a range spanning from past to future application time.
8. *Update Supplier*: Update the supplier balance. This update stresses a degenerated table.
9. *Manipulate Order Data*: Choose an "old" order (with the application time far before system time) and update its price. This update changes values while keeping the application times (i.e., trying to hide this change).

Since the initial data generation and the data evolution mix are modeled independently, we can control the size of the initial data (called h like in TPC-H) and the length of the history (called m) separately, thus permitting cases like large initial data with a short history ($h \gg m$), small initial data with a long history ($h \ll m$) or any other combination. Similarly to the scaling settings in TPC-H, where $h = 1.0$ corresponds to 1 GB of data, we normalize $m = 1.0$ to the same size, and use the same (linear) scaling rules.

Table 4.1 describes the outcome of applying a mix of these queries on the various tables. The history growth ratio describes how many update operations per initial tuple happen when $h = m$. As we can see, CUSTOMER and SUPPLIER get a fairly high number of history entries per tuple, while ORDERS and LINEITEM see proportionally fewer history operations. When taking the sizes of the initial relations into account, the bulk of history operations is still performed on LINEITEM and ORDER. A second aspect on which the tables differ is the kind of history operations: SUPPLIER, CUSTOMER and PARTSUPP only receive UPDATE statements, whereas the remaining bitemporal relations will see a mix of operations. LINEITEM is strongly dominated by INSERT operations (> 60 percent), ORDERS less so (50 percent inserts and 42 percent updates). CUSTOMERS in turn see mostly UPDATE operations (> 70 percent). The temporal specialization follows the specification in the schema, providing SUPPLIER as a degenerate table. Finally, existing application time periods can be overwritten with new values for CUSTOMER, PART, PARTSUPP and ORDERS which refers to the use case of updating application time, which is an important feature of the bitemporal data model.

Table	History growth ratio	Dominant Operations	Temporal Specialization	Overwrite App.Time
NATION	None	None	non-versioned	no
REGION	None	None	non-versioned	no
SUPPLIER	5	Update	degenerate	no
CUSTOMER	3.7	Update	fully bitemporal	yes
PART	0.25	Update	fully bitemporal	yes
PARTSUPP	0.72	Update	fully bitemporal	yes
LINEITEM	0.32	Insert	fully bitemporal	no
ORDER	0.4	Insert	fully bitemporal	yes

Table 4.1: Properties of the History for each Table

We implemented a generator to derive the application times from the TPC-H `dbgen` output for the initial version and generate the data evolution mix. Details of this generator are given in Appendix A. The generator accounts for the different ways temporal data is supported by current temporal DBMS. Initial evaluations show that this generator can generate 0.6 Million tuples/sec, compared to 1.7 Million tuples/sec of `dbgen` on the same machine. The data generator can also be configured to compute a data set consisting purely of tuples that are valid at the end of the generation interval. This is useful when comparing the cost of temporal data management on the latest version against a non-temporal database.

4.3 Queries

Given the multi-dimensional space of possible temporal query classification, we cluster the queries among common dimensions: Data access [9], temporal operators and specific temporal correlations.

4.3.1 Synthetic Time Travel.

The first group of queries is concerned with testing “slices” of time, i.e., establishing the state concerning a specific time for a table or a set of tables. The details for all queries of this group can be found in Appendix B.1. Also known as *Time Travel*, this is the most commonly supported and used class of temporal queries. Given that time in a bitemporal database has more than one dimension, one can specify different slicing options for each of these dimensions: Each dimension could be treated as a point or as complete slice, e.g., fixing the application time to June 1st, 2013, while considering the full evolution through system time. Further aspects to study are the combination of time travel operations (e.g., to compare values at different points in time), implicit vs. explicit expressions for time and the impact of underlying data/temporal update patterns. The first set of queries is targeted for testing various aspect of time travel in isolation, consisting of nine queries with variants.

T1 and T2 are our baseline queries, performing a point-point access for both temporal dimensions. By varying both timestamps accordingly, particular combinations can easily be specified, e.g., tomorrow’s state in application time, as recorded yesterday. The difference between T1 and T2 is according to the underlying data: T1 uses CUSTOMER, a table with many update operations and large history, but stable cardinalities. T2, in turn, uses ORDERS, a table with a generally smaller history and a focus on insertions. This way, we can study the cost of time travel operations on significantly different histories. T3 and T4

correlate data from two time travel operations within the same table. Comparing their results with T2 (very selective) and T5 (entire history) gives an insight into whether any sharing of history operations is possible. T4 adds a TOP N condition, providing possible room for optimization in the database system. T5 retrieves the complete history of the ORDERS table. Given that all data is requested, it should serve as a yardstick for the maximal cost of simple time travel operations. T6 performs temporal slicing, i.e., retrieving all data of one temporal dimension, while keeping the other to a point. This provides insights if the DBMS prefers any dimension, and a comparison of T2 and T5 yields insights if any optimization for points vs. slices are available. T7 complements T6 by implicitly specifying current system time, providing an understanding as to if different approaches of specifying current time work equally well. T8 and T9 investigate the behavior of additional application times, as outlined in Section 4.1. Since the standard currently only allows a single, designated application time, we can study the benefits of explicit vs. implicit application times. In that context, T8 uses point data (like T2), while T9 uses slicing (like T6).

4.3.2 Time Travel for Application Queries.

Whereas the first group of queries focuses on systematically testing the query performance for combinations of time travel in different time dimensions, the second set of timeslice queries (see Appendix B.2) focuses on application-oriented workloads, providing insights on how well synthetic time travel performance translates into complex analytical queries, e.g., accounting for additional data access cost and possibly disabled optimizations. For this purpose, we use the 22 standard TPC-H queries (similar to what [1] proposes) and extend them to allow the specification of both a system and an application time point. Possible evaluations might contain determining the cost of accessing the current version (in both system as well as current application time) compared against the logically same data stored in a non-temporal table (see Section 4.2).

4.3.3 Pure-Key Queries (Audit).

The next class of queries we study (see Appendix B.3) poses an orthogonal problem: Instead of retrieving all tuples for a particular point in time, we process the history of a specific tuple or a small set of tuples. This way, we can investigate how tuples evolve over time, e.g., for auditing or trend detection. This evolution can be considered along the system time, the application time(s) or both. Additional aspects to study are the effects of constraints on the version range (complete time range, some time period, some versions) and type of tuple selection, e.g., keys or predicates. In total, we specify 6 queries, each with small variants to account for the different time dimensions: K1 selects the tuple using a primary key, returns many columns and does not place any constraints on the temporal range. For key-based histories, this should provide the yardstick, and also offers clear insights into the organization of the storage of temporal data. The cost of this operation can also be compared against T5 and T6, which retrieve all versions of all tuples (for both dimensions or each time dimension, each). To allow easy comparison with the T queries, all queries are executed on the ORDERS relation. K2 alters K1 by placing a constraint on the temporal range.

Compared to K1, this additional information should provide an optimization possibility. K3 alters K2 even further by only retrieving a single column, providing optimization potential for decomposition or column stores. K4 complements K2 by constraining not the temporal range (by a time interval), but the number of versions (by using TOP N). While the intent is quite similar to K2, the semantics and possible execution strategies are quite different. K5 constitutes a special case of K4 in which only the immediately preceding version is retrieved, employing no TOP N expression, but a timestamp correlation. From a technical point of view, this provides additional potential for optimization. From a language point of view, such an access is required for queries that perform change detection. K6 chooses the tuples not via a key of the underlying table, but using a range predicate on a value (`o_totalprice`). Besides a general comparison to key-based access, choosing the value of this parameters allows us to study the impact of the selectivity on the computation cost.

4.3.4 Range-Timeslice Queries.

As the most general access pattern, range-timeslice queries (see Appendix B.4) permit any combination of constraints on both value and temporal aspects. As a result, a broad range of queries falls into that range. We will provide a set of application-derived workloads here, highlighting the variety and the different challenges it brings. As before, these queries contain variants which restrict one time dimension to a point, while varying the other.

R1 considers state change modeling by querying those customers who moved to the US at a particular point in time and still live there. The SQL expression involves two temporal evaluations on the same relation and a join of the results. R2 also handles state modeling, but instead of detecting changes, it computes state durations for `LINEITEMS` (the shipping time). Compared to R1, the intermediate results are much bigger, but no temporal filters are applied when combining them. R3 expresses temporal aggregation, i.e., computing aggregates for each version or time range of the database. At SAP, this turned out to be one of the most sought-after analyses of temporal data. However, SQL:2011 does not provide much support for this use case. The first query (R3.a) computes the greatest number of unshipped items in a time range. In SQL:2011, this requires a rather complex and costly join over the time interval boundaries to determine change points, followed by a grouping on these boundaries for the aggregates. The second query (R3.b) computes the maximum value of unshipped orders within one year. As before, interval joins and grouping are required. R4 computes the products with the smallest difference in stock levels over the history. While the temporal semantics are rather easy to express, the same tables need to be accessed multiple times, and significant amount of post-processing is required. R5 covers temporal joins by computing how often a customer had a balance of less than 5000 while also having orders with a price greater than 10. The join therefore not only includes value join criteria (on the keys), but also time correlation. R6 combines two temporal operators in one query by computing a temporal aggregation for an intermediate table that is retrieved by joining two temporal tables. R7 computes changes between versions over a full set, retrieving those suppliers who increased their prices by more than 7.5 percent in a single update. R7 thus generalizes K4/K5 by determining previous versions for all keys, not just specific ones.

Name	App Time	System Time	System Time value
B3.1	Point	Point	Current
B3.2	Point	Point	Past
B3.3	Correlation	Point	Current
B3.4	Point	Correlation	-
B3.5	Correlation	Correlation	-
B3.6	Agnostic	Point	Current
B3.7	Agnostic	Point	Past
B3.8	Agnostic	Correlation	-
B3.9	Point	Agnostic	-
B3.10	Correlation	Agnostic	-
B3.11	Agnostic	Agnostic	-

Table 4.2: Bitemporal Dimension Queries

4.3.5 Bitemporal Queries.

Nearly all queries so far have treated the two temporal dimensions in the same way: Keeping one dimension fixed, while performing different operations types of operations on the other. While this is a fairly common pattern in real-life queries, we also want to gain a more thorough understanding of queries stressing both time dimensions (see Appendix B.5). Snodgrass [10] provides a classification of bitemporal queries. Our first set of bitemporal queries follows this approach and creates complementary query variants to cover all relevant combinations. These variants span both time dimensions and vary the usage of each time dimension: a) current/(extended to) time point, b) sequenced/time range, c) non-sequenced/agnostic of time. The non-temporal baseline query B3 is a standard self-join: What (other) parts are supplied by the suppliers who supplies part 55? Table 4.2 describes the semantics of each query.

5 Conclusion

In this paper, we presented a benchmark for bitemporal databases which builds on existing benchmarks and presents a comprehensive coverage of temporal data and queries. Preliminary results on existing temporal database systems highlight significant optimization potential and insufficient support for common application use cases in the current SQL:2011 standard. We currently consider the following directions for future work: First, we want to broaden our evaluation, including larger data sets, DBMSs which we have not covered so far and possible tuning guidelines. This will give us further insights into which queries to add and possibly remove for a complete, yet concise coverage of the temporal DBMS workloads. Furthermore, we would like to incorporate explicit update queries, which we can evaluate in their performance characteristics.

References

- [1] Al-Kateb, M., Crolotte, A., Ghazal, A., Rose, L.: Adding a Temporal Dimension to the TPC-H Benchmark. In: Nambiar, R.O., Poess, M. (eds.) *Selected Topics in Performance Evaluation and Benchmarking - 4th TPC Technology Conference*. LNCS, vol. 7755, pp. 51–59. Springer (2012), ISBN: 978-3-642-36726-7
- [2] Cole, R., et al.: The Mixed Workload CH-benCHmark. In: *DBTest*. p. 8 (2011)
- [3] Gray, J.: *Benchmark Handbook: For Database and Transaction Processing Systems*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1992)
- [4] Jensen, C.S., et al.: A consensus test suite of temporal database queries. Tech. rep., Department of Computer Science, Aarhus University (1993)
- [5] Kalua, P.P., Robertson, E.L.: *Benchmarking Temporal Databases - A Research Agenda*. Tech. rep., Indiana University, Computer Science Department (1995)
- [6] Kaufmann, M., Fischer, P.M., May, N., Tonder, A., Kossmann, D.: TPC-BiH: A Benchmark for Bi-Temporal Databases. In: *TPCTC (2013)*, accepted for publication
- [7] Kaufmann, M., Manjili, A., Vagenas, P., Fischer, P., Kossmann, D., Faerber, F., May, N.: Timeline index: A unified data structure for processing queries on temporal data in SAP HANA. In: *SIGMOD (2013)*
- [8] Kulkarni, K.G., Michels, J.E.: Temporal Features in SQL: 2011. *SIGMOD Record* 41(3) (2012)
- [9] Salzberg, B., Tsotras, V.J.: Comparison of access methods for time-evolving data. *ACM Comput. Surv.* 31(2), 158–221 (1999)
- [10] Snodgrass, R.T.: *Developing Time-Oriented Database Applications in SQL*. Morgan Kaufmann (1999)

- [11] Snodgrass, R.T., et al.: TSQL2 language specification. SIGMOD Record 23(1) (1994)
- [12] Werstein, P.: A Performance Benchmark for Spatiotemporal Databases. In: In: Proc. of the 10th Annual Colloquium of the Spatial Information Research Centre. pp. 365–373 (1998)

A Data Generator

The data generator of the TPC-BiH benchmark has 2 independent scaling factors:

- TPC-H scaling factor (h): for version 1 (h=1 for 1 GB data)
- Modification scaling factor (m): update operations (m=1 for 1 GB written data)

A.1 Initialization for Version 1

We use the `dbgen` generator of the TPC-H benchmark for producing version 1 of the data set. As an extension to the TPC-H benchmark we need to initialize also the application time. We initialize the corresponding application time periods of the temporal tables according to the following rules:

- For each Lineitem
 - `ACTIVE_TIME_BEGIN := min(SHIPDATE, COMMITDATE, RECEIPTDATE)`
 - `ACTIVE_TIME_END := max(SHIPDATE, COMMITDATE, RECEIPTDATE)`
- For each Order
 - `$active_time_begin := ∞`
 - `$active_time_end := 0`
 - For each LINEITEM of this ORDER
 - * `$active_time_begin := min($active_time_begin, LINEITEM.ACTIVE_TIME_BEGIN)`
 - * `$active_time_end := max($active_time_end, LINEITEM.ACTIVE_TIME_END)`
 - Set `ORDER.ACTIVE_TIME_BEGIN := min(ORDERDATE, $active_time_begin)`
 - Set `ORDER.ACTIVE_TIME_END := $active_time_end`
 - `$receivable_time_begin := uniform($active_time_begin, $active_time_end)`
 - Set `ORDER.RECEIVABLE_TIME_BEGIN := $receivable_time_begin`
 - Set `ORDER.RECEIVABLE_TIME_END := uniform($receivable_time_begin, $active_time_end)`
- For each Customer
 - `$active_time_begin := ∞`
 - For each ORDER of this CUSTOMER
 - * `$active_time_begin := min($active_time_begin, ORDER.ACTIVE_TIME_BEGIN)`
 - Set `CUSTOMER.ACTIVE_TIME_BEGIN := $active_time_begin`
 - Set `CUSTOMER.ACTIVE_TIME_END := ∞`

- For each Part
 - \$active_time_begin := ∞
 - \$active_time_end := ∞
 - For each LINEITEM of this PART
 - * \$active_time_begin := min(\$active_time_begin, LINEITEM.ACTIVE_TIME_BEGIN)
 - Set AVAILABILITY_TIME_BEGIN := \$active_time_begin
 - Set AVAILABILITY_TIME_END := ∞
- For each PARTSUPP
 - Set VALIDITY_TIME_BEGIN := PART.AVAILABILITY_TIME_BEGIN
 - Set VALIDITY_TIME_END := ∞

A.2 History Generator

The history generator simulates the data set being updated during a time interval of 10 years, i.e., from 1.1.2000 until 31.12.2009. The motivation of introducing a new data generator is to achieve a defined number of updates and deletes for the history. This history generator can be executed several times on an initialized data set, e.g., TPC-H with h=1.0 and 10 times history generator with m=0.5 is similar to a data set TPC-H with h=6.0, because the values for the business times are generated based on the generated values.

The history generator supports different granularities for the application time (e.g., one month or one day). For the experiments, one day is chosen by default. The current timestamp for system time moves linearly from 2000-01-01 to 2009-12-31.

The data is generated by different update scenarios. At each point in time, the next scenario is chosen randomly. Each scenario has a defined probability: These probabilities are chosen such that the state of all tables remains stable, e.g., there are not more placed new orders than can be delivered or too few stock level updates resulting in empty stocks.

A.2.1 New Order [probability1 = 0.30]

- Customer
 - Uniform distribution of one of the following cases:
 1. Uniformly select an existing customer that is visible from CUSTOMER [probability1.1 = 0.5]
 - a) Update customer data (e.g., customer moved to other location) [probability1.1.a = 0.5]
 - * \$visibility_begin := CURRENT_TIME + uniform(1 day, 30 days)
 - * Set VISIBILITY_TIME := [\$visibility_begin, ∞]
 - * Select nation randomly

- * Change address, phone
 - b) Do not change customer data [probability1.1.b = 0.5]
- 2. Insert new Customer [probability1.2 = 0.5]
 - * Generate customer data randomly
 - * Set `VISIBLE_TIME` := [`CURRENT_TIME`,]
 - * Set `ACCTBAL` := 0
- Order
 - Insert a new Order into `ORDERS`
 - * Set `ORDERSTATUS` := ‘O’
 - * Set `ACTIVE_TIME` := [`CURRENT_TIME`, ∞]
 - * `$receivable_begin` := `CURRENT_TIME` + uniform(1 day, 14 days)
 - * Set `RECEIVABLE_TIME` := [`$receivable_begin`, ∞]
- Lineitems
 - Uniformly select number of ordered items [1..7]
 - For each ordered item:
 1. Select one part (which is either visible or becomes visible in the future) according to Uniform distribution
 2. Insert a tuple into `LINEITEM`
 - * Set `LINESTATUS` := ‘O’
 - * Set `ACTIVE_TIME` := [`CURRENT_TIME`, ∞]

A.2.2 Cancel Order [probability2 = 0.01]

- Uniformly select one Order with `ORDERSTATUS` != ‘F’
 1. If(`ORDERSTATUS` = ‘P’) then
 - Update `CUSTOMER`
 - * Set `ACCTBAL` := `ACCTBAL` + `TOTALPRICE`
 2. For each Lineitem of this Order:
 - If `LINESTATUS` = ‘F’ then
 - * update `PARTSUPP`
 - Set `AVAILQTY` := `AVAILQTY` + `QUANTITY`
 - * Delete Lineitem
 3. Delete Order

A.2.3 Deliver Order [probability3 = 0.20]

- Uniformly select Order with ((ORDERSTATUS = 'O' and ACCTBAL TOTALPRICE \geq 0) or (ORDERSTATUS = 'P'))
 - If(ORDERSTATUS = 'O') then
 1. Update Customer
 - * Set ACCTBAL := ACCTBAL TOTALPRICE
 2. Update Order
 - * Set ORDERSTATUS := 'P'
 - For each Lineitem of this Order
 - * If(LINESTATUS = 'O' and AVAILQTY QUANTITY > 0 and CURRENT_DATE between PART.AVAILABILITY_TIME)
 1. Update Partsupp
 - Set AVAILQTY := AVAILQTY QUANTITY
 2. Update Lineitem
 - Set LINESTATUS := 'F'
 - Set ACTIVE_TIME_END := CURRENT_TIME
 - If (CURRENT_DATE between RECEIVABLE_TIME and for each lineitem of this order: LINESTATUS = 'F')
 - * Update Order
 1. Set ACTIVE_TIME_END := CURRENT_TIME
 2. If RECEIVABLE_TIME_END is not set
 - Set RECEIVABLE_TIME_END := CURRENT_TIME
 - Set ORDERSTATUS := 'F'

A.2.4 Receive Payment [probability4 = 0.20]

- Uniformly select order with (CURRENT_DATE between RECEIVABLE_TIME and ORDERSTATUS = 'O' and ACCTBAL TOTALPRICE \geq 0)
 1. Update Order
 - Set RECEIVABLE_TIME_END := CURRENT_TIME
 2. Update Customer
 - Set ACCTBAL := ACCTBAL + TOTALPRICE

A.2.5 Update Stock [probability5 = 0.05]

- Uniformly select lineitem with LINESTATUS = 'O' and AVAILTY \geq QUANTITY and CURRENT_DATE within PART.VISIBILITY_TIME

1. Update Partsupp
 - Set AVAILQTY := AVAILQTY + 2*QUANTITY

A.2.6 Delay Availability [probability6 = 0.05]

- Uniformly select part
 1. Update Part
 - \$availability_begin := CURRENT_TIME + uniform(1 day, 14 days)
 - Set PART.AVAILABILITY_TIME := [\$availability_begin, ∞]

A.2.7 Change Price by Supplier [probability7 = 0.05]

- Uniformly select part and supplier
 1. Update Partsupp
 - \$validity_begin := CURRENT_TIME + gauss(-15 days, 30 days)
 - \$price_change := uniform(-100, +100)
 - Set PARTSUPP.SUPPLYCOST := abs(PARTSUPP.SUPPLYCOST + \$price_change)
 - Set PARTSUPP.VALIDITY_TIME := [\$validity_begin, ∞]

A.2.8 Update Supplier [probability8 = 0.049]

- Uniformly select supplier
 1. Update Supplier
 - balance_change := uniform(-100, +100)
 - Set SUPPLIER.ACCTBAL := abs(SUPPLIER.ACCTBAL + balance_change)

A.2.9 Manipulate Order Data [probability9 = 0.001]

- Uniformly select an order with (CURRENT_DATE > RECEIVABLE_TIME_END + 1 Month and ORDERSTATUS = 'F')
 - \$manipulated_price := 0
 - For each lineitem of this order
 1. Update Lineitem
 - * Set EXTENDED_PRICE := EXTENDED_PRICE uniform(1, 10)
 - * \$manipulated_price = \$manipulated_price + EXTENDED_PRICE
 2. Update Order
 - * Set TOTAL_PRICE := \$manipulated_price

B Queries

In this part of the appendix, we give the details of all queries of the TPC-BiH benchmark. These queries are available in various versions as the vendors of commercial systems use different SQL dialects to express temporal queries. In this paper, we show the queries using the syntax by DB2 as this representation currently is the version closest to the SQL:2011 standard.

For each query we give both the description in natural language, the SQL representation and how the query parameters should be populated.

B.1 Synthetic Time Travel

B.1.1 Simple Time-Travel Query on stable relation (T.1)

“What was the average price of all supplies at a given business time (APP_TIME) known as recorded at a certain previous time in history (SYS_TIME)?”

```
SELECT AVG(ps_supplycost), count(*)
FROM partsupp
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
FOR BUSINESS_TIME AS OF '[APP_TIME]'
```

Parameters:

- SYS_TIME := uniform(PARTSUPP.sys_time_start, PARTSUPP.sys_time_end)
- APP_TIME := uniform(PARTSUPP.validity_time_start, PARTSUPP.validity_time_end)

B.1.2 Simple Time-Travel Query on growing relation (T.2)

“What was the total revenue of all orders at a given business time (APP_TIME) known as recorded at a certain previous time in history (SYS_TIME)?”

```
SELECT SUM(o_totalprice) AS revenue, count(*)
FROM orders
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
FOR BUSINESS_TIME AS OF '[APP_TIME]'
WHERE o_orderstatus = '0'
```

Parameters:

- SYS_TIME := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- APP_TIME := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

B.1.3 Aggregation at Two Defined Points in Time (T.3)

T.3sys: System Time.

“Compare the estimated revenue of the open orders at a given business time (APP_TIME) recorded at two different times in history (SYS_TIME1 and SYS_TIME2).”

```
SELECT agg_t1.revenue old_revenue, agg_t2.revenue new_revenue
FROM
(
  SELECT SUM(o_totalprice) as revenue
  FROM orders
  FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME1]'
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME]'
  WHERE o_orderstatus = '0'
) agg_t1
,
(
  SELECT SUM(o_totalprice) as revenue
  FROM orders
  FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME2]'
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME]'
  WHERE o_orderstatus = '0'
) agg_t2
```

Parameters:

- SYS_TIME1 := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- SYS_TIME2 := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- APP_TIME := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

T.3app: Application Time.

“Compare the estimated revenue of the open orders at two given business times (APP_TIME1 and APP_TIME2).”

```
SELECT agg_t1.revenue old_revenue, agg_t2.revenue new_revenue
FROM
(
  SELECT SUM(o_totalprice) as revenue
  FROM orders
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME1]'
  WHERE o_orderstatus = '0'
) agg_t1
,
(
  SELECT SUM(o_totalprice) as revenue
```

```

FROM orders
FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME2]'
WHERE o_orderstatus = '0'
) agg_t2

```

Parameters:

- APP_TIME1 := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)
- APP_TIME2 := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

B.1.4 Aggregation of a Limited Number of Items at Two Defined Points in Time (T.4)

T.4sys: System Time.

“Get the 10 products for which the revenue at a given business time (APP_TIME) increased most within two times in history (SYS_TIME1 and SYS_TIME2).”

```

SELECT revenue2 - revenue1 AS inc_revenue, agg_t1.o_custkey
FROM
(
  SELECT SUM(o_totalprice) as revenue1, o_custkey
  FROM orders
  FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME1]'
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME]'
  WHERE o_orderstatus = '0'
  GROUP BY o_custkey
) agg_t1
,
(
  SELECT SUM(o_totalprice) as revenue2, o_custkey
  FROM orders
  FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME2]'
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME]'
  WHERE o_orderstatus = 'F'
  GROUP BY o_custkey
) agg_t2
WHERE agg_t1.o_custkey = agg_t2.o_custkey
ORDER BY inc_revenue DESC
FETCH FIRST 10 ROWS ONLY

```

Parameters:

- SYS_TIME1 := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- SYS_TIME2 := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- APP_TIME := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

T.4app: Application Time.

“Get the 10 products for which the revenue increased most within two points in business time (APP_TIME1 and APP_TIME2).”

```
SELECT revenue2 - revenue1 AS inc_revenue, agg_t1.o_custkey
FROM
(
  SELECT SUM(o_totalprice) as revenue1, o_custkey
  FROM orders
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME1]'
  WHERE o_orderstatus = '0'
  GROUP BY o_custkey
) agg_t1
,
(
  SELECT SUM(o_totalprice) as revenue2, o_custkey
  FROM orders
  FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME2]'
  WHERE o_orderstatus = '0'
  GROUP BY o_custkey
) agg_t2
WHERE agg_t1.o_custkey = agg_t2.o_custkey
ORDER BY inc_revenue DESC
FETCH FIRST 10 ROWS ONLY
```

Parameters:

- APP_TIME1 := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)
- APP_TIME2 := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

B.1.5 Retrieving all versions (T.5)

“What is the maximum order price per customer ever recorded in system (regardless of the application time)?”

```
SELECT o_custkey, MAX(o_totalprice), count(*) AS supp_revenue
FROM orders
FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30'
WHERE o_orderstatus = '0'
GROUP BY o_custkey
```

B.1.6 Slicing one temporal dimension (T.6)

T.6sys: System Time.

“What is the average revenue per customer over business time, at a given time in history (SYS_TIME)?”

```

SELECT o_custkey, AVG(o_totalprice) AS supp_revenue
FROM orders
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
WHERE o_orderstatus = '0'
GROUP BY o_custkey

```

Parameters:

- SYS_TIME := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)

T.6app: Application Time.

“What is the average order revenue per customer over history at a given point in business time (APP_TIME)?”

```

SELECT o_custkey, AVG(o_totalprice) AS supp_revenue
FROM orders
FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME]'
FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30'
WHERE o_orderstatus = '0'
GROUP BY o_custkey

```

Parameters:

- APP_TIME := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

B.1.7 Implicit vs. explicit current time (T.7)

“What is the average revenue per customer over time, as recorded today?”

```

SELECT PS_PARTKEY, AVG(ps_availqty) AS avail
FROM partsupp
GROUP BY PS_PARTKEY

```

B.1.8 Time Travel among multiple application times (T.8)

“What is the quantity of the orders which were payable and in transit at APP_TIME, as recorded at a previous time in history (SYS_TIME)?”

```

SELECT count(*)
FROM orders
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
FOR BUSINESS_TIME AS OF '[APP_TIME]'
WHERE '[APP_TIME]' BETWEEN receivable_time_start AND receivable_time_end

```

Parameters:

- SYS_TIME := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- APP_TIME := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

B.1.9 Time Travel to simulated/secondary application time (T.9)

“What is the quantity of the orders which have been in transit at APP_TIME, as recorded at SYS_TIME?”

```
SELECT count(*)
FROM orders
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
WHERE '[APP_TIME]' BETWEEN receivable_time_start AND receivable_time_end
```

Parameters:

- SYS_TIME := uniform(ORDERS.sys_time_start, ORDERS.sys_time_end)
- APP_TIME := uniform(ORDERS.validity_time_start, ORDERS.validity_time_end)

B.2 Time Travel for Application Queries

For each time dimension of a table that occurs in a TPC-H query, a time travel to the current version is performed. Thus, we can measure the overhead of a temporal table compared to storing the current version only in a non-temporal table. As the corresponding extension of the queries is analog for all TPC-H queries, we only show the first TPC-H query as an example.

B.2.1 Pricing Summary Report Query (H.1)

“Compute the Pricing Summary Report for a given point in time APP_TIME when an ordered item was valid as it was reported at a previous time SYS_TIME in history.”

```
SELECT
    l_returnflag,
    l_linestatus,
    SUM(l_quantity) AS sum_qty,
    SUM(l_extendedprice) AS sum_base_price,
    SUM(l_extendedprice*(1-l_discount)) AS sum_disc_price,
    SUM(l_extendedprice*(1-l_discount)*(1+l_tax)) AS sum_charge,
    AVG(l_quantity) AS avg_qty,
    AVG(l_extendedprice) AS avg_price,
    AVG(l_discount) AS avg_disc,
    COUNT(*) AS count_order
FROM
    lineitem FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
    lineitem FOR BUSINESS_TIME AS OF TIMESTAMP '[APP_TIME]'
WHERE
    l_shipdate <= DATE '1998-11-01'
GROUP BY
    l_returnflag,
```

```

    l_linestatus
ORDER BY
    l_returnflag,
    l_linestatus;

```

Parameters:

- SYS_TIME := currentVersion(LINEITEM.sys_time)
- APP_TIME := currentVersion(LINEITEM.active_time)

B.3 Pure-Key Queries (Audit)

B.3.1 Retrieval of all Time Versions for a Given Tuple (K.1)

K.1sys: System Time.

“How did a given order as it is valid at a certain business time APP_TIME evolve in history?”

```

SELECT
    o_orderkey, o_custkey, o_orderstatus, o_totalprice,
    o_orderdate, o_orderpriority, o_shippriority,
    sys_time_start
FROM
    orders
FOR
    SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
FOR
    BUSINESS_TIME AS OF '[APP_TIME]'
WHERE
    o_orderkey = [ORDER_KEY]
ORDER BY
    sys_time_start

```

Parameters:

- ORDER_KEY := uniform({o_orderkey})
- APP_TIME := uniform(orders.active_time_start, orders.active_time_end)

K.1app: Application Time.

“What are all changes to an order as it is recorded today?”

```

SELECT
    o_orderkey, o_custkey, o_orderstatus, o_totalprice,
    o_orderdate, o_orderpriority, o_shippriority,
    active_time_start
FROM

```

```

orders
FOR
  BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
WHERE
  o_orderkey = [ORDER_KEY]
ORDER BY
  active_time_start

```

Parameters:

- ORDER_KEY := uniform({o_orderkey})

B.3.2 Retrieval of a time version range for a given Tuple (K.2)

K.2sys: System Time.

“How did the data of a given customer as it is valid at a certain business time APP_TIME evolve in history?”

```

SELECT
  c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
  sys_time_start
FROM
  customer
FOR
  SYSTEM_TIME FROM TIMESTAMP '[SYS_TIME_BEGIN]' TO TIMESTAMP '[SYS_TIME_END]'
FOR
  BUSINESS_TIME AS OF '[APP_TIME]'
WHERE
  c_custkey = [CUSTOMER_KEY]
ORDER BY
  sys_time_start

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- SYS_TIME_BEGIN := uniform(customer.sys_time_start, customer.sys_time_end)
- SYS_TIME_END := uniform(customer.sys_time_start, customer.sys_time_end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

K.2app: Application Time.

“What are all changes to a certain customer in a time interval as it is recorded today?”

```

SELECT
    c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
    visible_time_start
FROM
    customer
FOR
    BUSINESS_TIME FROM '[APP_TIME_BEGIN]' TO '[APP_TIME_END]'
WHERE
    c_custkey = [CUSTOMER_KEY]
ORDER BY
    visible_time_start

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- APP_TIME_BEGIN := uniform(customer.visible_time_start, customer.visible_time_end)
- APP_TIME_END := uniform(customer.visible_time_start, customer.visible_time_end)

Constraints:

- APP_TIME_BEGIN < APP_TIME_END

B.3.3 Aggregation over System Time (K.3)

K.3sys: System Time.

“What is the minimum account balance of a given customer as it is valid at a certain business time APP_TIME over a time interval in history?”

```

SELECT
    MIN(c_acctbal)
FROM
    customer
FOR
    SYSTEM_TIME FROM TIMESTAMP '[SYS_TIME_BEGIN]' TO TIMESTAMP '[SYS_TIME_END]'
FOR
    BUSINESS_TIME AS OF '[APP_TIME]'
WHERE
    c_custkey = [CUSTOMER_KEY]

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- SYS_TIME_BEGIN := uniform(customer.sys_time_start, customer.sys_time_end)
- SYS_TIME_END := uniform(customer.sys_time_start, customer.sys_time_end)

- APP_TIME := uniform(customer.visible_time_start, customer.visible_time_end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

K.3app: Application Time.

“What is the minimum account balance of a given customer in a business time interval as it is recorded today?”

```
SELECT
    MIN(c_acctbal)
FROM
    customer
FOR
    BUSINESS_TIME FROM '[APP_TIME_BEGIN]' TO '[APP_TIME_END]'
WHERE
    c_custkey = [CUSTOMER_KEY]
```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- APP_TIME_BEGIN := uniform(customer.visible_time_start, customer.visible_time_end)
- APP_TIME_END := uniform(customer.visible_time_start, customer.visible_time_end)

Constraints:

- APP_TIME_BEGIN < APP_TIME_END

B.3.4 Retrieval of a version count range for a given Tuple (K.4)

K.4sys: System Time.

“What are the 10 latest versions of the customer data which are valid at a certain business time APP_TIME?”

```
SELECT
    c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
    sys_time_start, sys_time_end
FROM
    customer
FOR
    SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
FOR
    BUSINESS_TIME AS OF '[APP_TIME]'
WHERE
```

```

    c_custkey = [CUSTOMER_KEY]
ORDER BY
    sys_time_start
FETCH FIRST [HSIZE] ROWS ONLY

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- HSIZE := 10
- APP_TIME := uniform(customer.visible_time_start, customer.visible_time_end)

K.4app: Application Time.

“What are the 10 latest versions of the customer in business time as it is recorded today?”

```

SELECT
    c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
    visible_time_start, visible_time_end
FROM
    customer
FOR
    BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
WHERE
    c_custkey = [CUST_KEY]
ORDER BY
    visible_time_start
FETCH FIRST [HSIZE] ROWS ONLY

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- HSIZE := 10
- APP_TIME := uniform(customer.visible_time_start, customer.visible_time_end)

B.3.5 Retrieval of the preceding version for a given Tuple (K.5)

K.5sys: System Time.

“What is preceding version in history of the customer which is valid at a certain point in business time APP_TIME as it has been recorded at a certain point SYS_TIME in history?”

```

SELECT
    c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
    sys_time_start
FROM
    customer

```

```

FOR
  SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
FOR
  BUSINESS_TIME AS OF '[APP_TIME]'
WHERE
  c_custkey = [CUSTOMER_KEY] AND
  sys_time_end =
  (
    SELECT max(sys_time_end)
    FROM customer
    FOR
      SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
    FOR
      BUSINESS_TIME AS OF '[APP_TIME]'
    WHERE
      c_custkey = [CUSTOMER_KEY] and sys_time_END <= TIMESTAMP '[SYS_TIME]'
  )

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- SYS_TIME := uniform(customer.sys_time_start, customer.sys_time_end)
- APP_TIME := uniform(customer.visible_time_start, customer.visible_time_end)

K.5app: Application Time.

“What is preceding version of the customer visible at business time APP_TIME as it is recorded today?”

```

SELECT
  c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
  visible_time_start
FROM
  customer
WHERE
  c_custkey = [CUSTOMER_KEY] AND
  visible_time_end =
  (
    SELECT max(visible_time_end)
    FROM customer
    WHERE
      c_custkey = [CUSTOMER_KEY] and VISIBLE_TIME_END <= '[APP_TIME]'
  )

```

Parameters:

- CUSTOMER_KEY := uniform({c_custkey})
- APP_TIME := uniform(customer.visible_time_start, customer.visible_time_end)

B.3.6 Retrieval of all Versions for Tuples selected by Value (K.6)

K.6sys: System Time.

“Get all versions in history of the customer data which has been visible at a certain point in business time APP_TIME.”

```
SELECT
    c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
    sys_time_start
FROM
    customer
FOR
    SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
FOR
    BUSINESS_TIME AS OF '[APP_TIME]'
WHERE
    c_acctbal > [CUST_BAL]
ORDER BY
    c_custkey, c_acctbal, sys_time_start
```

Parameters:

- APP_TIME := uniform(customer.validity_time_start, customer.validity_time_end)
- CUST_BAL := uniform(0, 5000)

K.6app: Application Time.

“Get all versions in business time of the customer data as it is recorded today.”

```
SELECT
    c_custkey, c_name, c_address, c_nationkey, c_phone, c_acctbal,
    visible_time_start
FROM
    customer
FOR
    BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
WHERE
    c_acctbal > [CUST_BAL]
ORDER BY
    c_custkey, c_acctbal, visible_time_start
```

Parameters:

- CUST_BAL := uniform(0, 5000)

B.4 Range-Timeslice Queries

B.4.1 Restrict Interval for State Transition (R.1)

R.1sys: System Time.

“Which customers moved to US and still live there as recorded within a time intervall (SYS_TIME_BEGIN, SYS_TIME_END)?”

```
SELECT now.C_NAME
FROM (
  SELECT MAX(sys_time_end) time_until, c_custkey, c_name, c_nationkey
  FROM customer
  FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
  WHERE c_nationkey != (SELECT n_nationkey FROM ${NATION} WHERE n_name = 'US')
  group by c_custkey, c_name, c_nationkey
) prv, (
  SELECT MIN(sys_time_start) time_from, c_custkey, c_name, c_nationkey
  FROM customer
  FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
  WHERE c_nationkey = (SELECT n_nationkey FROM ${NATION} WHERE n_name = 'US')
  GROUP BY c_custkey, c_name, c_nationkey
) now
WHERE
  prv.c_custkey = now.c_custkey AND
  prv.time_until < '[SYS_TIME_BEGIN]' AND
  now.time_from BETWEEN '[SYS_TIME_BEGIN]' AND '[SYS_TIME_END]'
```

Parameters:

- SYS_TIME_BEGIN := uniform(customer.sys_time.start, customer.sys_time.end)
- SYS_TIME_END := uniform(customer.sys_time.start, customer.sys_time.end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

R.1app: Application Time.

“Which customers moved to US and still live there as registered within a time intervall (SYS_TIME_BEGIN, SYS_TIME_END)?”

```
SELECT now.C_NAME
FROM (
  SELECT MAX(visible_time_end) time_until, c_custkey, c_name, c_nationkey
  FROM customer
  FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
```

```

WHERE c_nationkey != (SELECT n_nationkey FROM nation WHERE n_name = 'US')
group by c_custkey, c_name, c_nationkey
) prv, (
SELECT MIN(visible_time_start) time_from, c_custkey, c_name, c_nationkey
FROM customer
FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
WHERE c_nationkey = (SELECT n_nationkey FROM nation WHERE n_name = 'US')
GROUP BY c_custkey, c_name, c_nationkey
) now
WHERE
prv.c_custkey = now.c_custkey AND
prv.time_until < '[SYS_TIME_BEGIN]' AND
now.time_from BETWEEN '[SYS_TIME_BEGIN]' AND '[SYS_TIME_END]'

```

Parameters:

- SYS_TIME_BEGIN := uniform(customer.sys_time.start, customer.sys_time.end)
- SYS_TIME_END := uniform(customer.sys_time.start, customer.sys_time.end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

B.4.2 Average Duration of a Certain State (R.2)

R.2sys: System Time.

“What is the mean time to ship an item as recorded over history?”

```

SELECT AVG(shp.ship_time - ord.order_time)
FROM
(
SELECT MIN(sys_time_start) AS order_time, o_orderkey
FROM orders
FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
WHERE o_orderstatus = 'O'
GROUP BY o_orderkey
) ord
,
(
SELECT MIN(sys_time_start) AS ship_time, o_orderkey
FROM orders
FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
WHERE o_orderstatus = 'F'
GROUP BY o_orderkey
) shp
WHERE ord.o_orderkey = shp.o_orderkey

```

R.2app: Application Time.

“What is the mean time to ship an item as recorded today?”

```
SELECT AVG(shp.ship_time - ord.order_time)
FROM
(
  SELECT MIN(active_time_start) AS order_time, o_orderkey
  FROM orders
  FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
  WHERE o_orderstatus = 'O'
  GROUP BY o_orderkey
) ord
,
(
  SELECT MIN(active_time_start) AS ship_time, o_orderkey
  FROM orders
  FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30'
  WHERE o_orderstatus = 'F'
  GROUP BY o_orderkey
) shp
WHERE ord.o_orderkey = shp.o_orderkey
```

B.4.3 Time-Slider Query (R.3)

R.3a: Time of Maximum Number of Unshipped Items.

“At which points in time in an interval did we have the greatest number of unshipped items?”

```
SELECT MAX(total)
FROM
(
  SELECT COUNT(*) AS total, times.event_time as timestmp
  FROM
  (
    SELECT * FROM orders
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
  ) ord,
  (
    SELECT sys_time_start AS event_time FROM orders
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
  UNION
    SELECT sys_time_end AS event_time FROM orders
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
  ) times
  WHERE ord.sys_time_start <= times.event_time AND
  times.event_time < ord.sys_time_end
```

```

    GROUP BY times.event_time
) agg

```

Parameters:

- SYS_TIME_BEGIN := uniform(customer.sys_time_start, customer.sys_time_end)
- SYS_TIME_END := uniform(customer.sys_time_start, customer.sys_time_end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

R.3b: Value for Maximum Number of Unshipped Items.

“What was the maximum value of unshipped orders within a time interval?”

```

SELECT MAX(total)
FROM
(
    SELECT SUM(ord.o_totalprice) AS total, times.event_time as timestmp
    FROM
    (
        SELECT * FROM orders
        FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
        WHERE o_orderstatus = '0'
    ) ord,
    (
        SELECT sys_time_start AS event_time FROM orders
        FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
    UNION
        SELECT sys_time_end AS event_time FROM orders
        FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
    ) times
    WHERE ord.sys_time_start <= times.event_time AND
           times.event_time < ord.sys_time_end
    GROUP BY times.event_time
) agg

```

Parameters:

- SYS_TIME_BEGIN := uniform(customer.sys_time_start, customer.sys_time_end)
- SYS_TIME_END := uniform(customer.sys_time_start, customer.sys_time_end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

B.4.4 Comparison of two Extreme Points of Aggregation (R.4)

“What is the minimum difference of min. and max. stock level of a product?”

```
SELECT MIN(diff), ps_partkey, ps_suppkey
FROM
(
  -- Select positive value
  SELECT (CASE WHEN (diff1 > diff2) THEN diff1 ELSE diff2 END) AS diff,
         ps_partkey, ps_suppkey
  FROM
  (
    SELECT (max_time_from - min_time_to) AS diff1,
           (min_time_from - max_time_to) AS diff2, ps_partkey, ps_suppkey
    FROM
    (
      -- Do not know if minimum occurs before maximum or vice versa
      SELECT min_time_from, min_time_to, max_time_from, max_time_to,
             smin.ps_availqty, smax.ps_availqty, glob.ps_partkey
      FROM
      (
        -- All cases where stock is minimal
        SELECT sys_time_start AS min_time_from,
               sys_time_end AS min_time_to,
               ps_availqty, ps_partkey, ps_suppkey
        FROM partsupp
        FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
      ) smin, (
        -- All cases where stock is maximal
        SELECT sys_time_start as max_time_from,
               sys_time_end AS max_time_to,
               ps_availqty, ps_partkey, ps_suppkey
        FROM partsupp
        FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
      ) smax, (
        -- Get global minimum and maximum of this product in history
        SELECT MIN(ps_availqty) AS gmin, MAX(ps_availqty) AS gmax,
               ps_partkey, ps_suppkey
        FROM partsupp
        FOR SYSTEM_TIME FROM '0001-01-01' TO '9999-12-30'
        GROUP BY ps_partkey, ps_suppkey
      ) glob
    )
  )
  WHERE
    glob.ps_partkey = smin.ps_partkey AND
    glob.ps_partkey = smax.ps_partkey AND
    glob.ps_suppkey = smin.ps_suppkey AND
```

```

        glob.ps_suppkey = smax.ps_suppkey AND
        smin.ps_availqty = glob.gmin AND
        smax.ps_availqty = glob.gmax AND
        glob.gmax != glob.gmin
    ) minmax
) aggr
)
GROUP BY ps_partkey, ps_suppkey
FETCH FIRST [HSIZE] ROWS ONLY

```

Parameters:

- HSIZE := 10

B.4.5 Temporal Join Query (R.5)

“How many times did a customer with a balance smaller than 5000 have an open order with total price more than 10?”

```

SELECT COUNT(DISTINCT(o_orderkey))
FROM
(
    SELECT * FROM customer
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]',
    WHERE c_acctbal < 5000
) cust,
(
    SELECT * FROM orders
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]',
    WHERE o_orderstatus = '0' AND o_totalprice > 10
) ord
WHERE
    c_custkey = o_custkey
    AND NOT (cust.sys_time_end <= ord.sys_time_start
            OR ord.sys_time_end <= cust.sys_time_start)

```

Parameters:

- SYS_TIME_BEGIN := uniform(customer.sys_time.start, customer.sys_time.end)
- SYS_TIME_END := uniform(customer.sys_time.start, customer.sys_time.end)

Constraints:

- SYS_TIME_BEGIN < SYS_TIME_END

B.4.6 Aggregation over Temporal Join Query (R.6)

“What was the maximum value of unpaid orders of customers with an account balance of less than 100?”

```
-- compute temporal join of customers and orders
CREATE VIEW high_risk_orders AS
SELECT
    o_totalprice AS orderprice,
CASE WHEN cust.sys_time_start > ord.sys_time_start
    THEN cust.sys_time_start ELSE ord.sys_time_start END
AS time_start,
CASE WHEN cust.sys_time_end < ord.sys_time_end
    THEN cust.sys_time_end ELSE ord.sys_time_end END
AS time_end
FROM
(
    SELECT * FROM customer
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
    WHERE c_acctbal < 100
) cust,
(
    SELECT * FROM orders
    FOR SYSTEM_TIME FROM '[SYS_TIME_BEGIN]' TO '[SYS_TIME_END]'
    WHERE o_orderstatus = '0'
) ord
WHERE
    c_custkey = o_custkey
    AND NOT (cust.sys_time_end <= ord.sys_time_start
        OR ord.sys_time_end <= cust.sys_time_start)

-- compute temporal aggregation
SELECT MAX(total) AS aggr_result
FROM
(
    SELECT SUM(ord.orderprice) AS total, times.event_time as timestmp
    FROM
    (
        SELECT * FROM high_risk_orders
    ) ord,
    (
        SELECT time_start AS event_time
        FROM high_risk_orders
    UNION
        SELECT time_end AS event_time
        FROM high_risk_orders
```

```

    ) times
  WHERE ord.time_start <= times.event_time AND
         times.event_time < ord.time_end
  GROUP BY times.event_time
) agg

```

B.4.7 Changes between versions (R.7)

“Which changes among the suppliers resulted in a price increase of more than 7.5 percent?”

```

SELECT ps1.ps_partkey, ps1.ps_suppkey, ps1.ps_supplycost,
       ps1.validity_time_start, ps1.validity_time_end,
       ps2.validity_time_start, ps2.validity_time_end,
       ps2.ps_supplycost
FROM
  partsupp
  FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30' AS PS1,
  partsupp
  FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30' AS PS2
WHERE
  ps1.ps_suppkey = ps2.ps_suppkey AND
  ps1.ps_partkey=ps2.ps_partkey AND
  ps1.validity_time_end > ps2.validity_time_end AND
  ps2.validity_time_end = (SELECT max(validity_time_end)
FROM partsupp
  FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30' as ps3
WHERE
  ps3.ps_suppkey = ps1.ps_suppkey AND
  ps3.ps_partkey=ps1.ps_partkey AND
  ps1.validity_time_end > ps3.validity_time_end) AND
  ps1.ps_supplycost > ps2.ps_supplycost*1.075
ORDER BY
  ps1.validity_time_end, ps2.validity_time_end DESC

```

B.5 Bitemporal Queries

B.5.1 App Time Current/ System Time Current (B.1)

“What (other) parts are currently supplied by the suppliers who supply a given part at the same business time APP_TIME, as currently recorded?”

```

SELECT ps2.*
FROM partsupp FOR BUSINESS_TIME AS OF '[APP_TIME]' as ps1,
  partsupp FOR BUSINESS_TIME AS OF '[APP_TIME]' as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
       ps1.ps_suppkey = ps2.ps_suppkey AND
       ps1.ps_partkey <> ps2.ps_partkey

```

Parameters:

- PART_KEY := uniform({ps_partkey})
- APP_TIME := uniform(partsupp.visible_time_start, partsupp.visible_time_end)

B.5.2 App Time Past/ System Time Past (B.2)

“What (other) parts were at time APP_TIME supplied by the suppliers who supplied part of a given part two week ago, as recorded at a previous time SYS_TIME in history?”

```
SELECT ps2.*, ps1.ps_suppkey, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
FOR BUSINESS_TIME AS OF '[APP_TIME1]' as ps1,
partsupp
FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]'
FOR BUSINESS_TIME AS OF '[APP_TIME2]' as ps2
WHERE ps1.ps_partkey = [PARTKEY] AND
       ps1.ps_suppkey = ps2.ps_suppkey AND
       ps1.ps_partkey <> ps2.ps_partkey
```

Parameters:

- PART_KEY := uniform({ps_partkey})
- SYS_TIME := uniform(partsupp.sys_time_start, partsupp.sys_time_end)
- APP_TIME1 := uniform(partsupp.validity_time_start, partsupp.validity_time_end)
- APP_TIME2 := uniform(partsupp.validity_time_start, partsupp.validity_time_end)

B.5.3 App Time Correlation / System Time Current (B.3)

“What other parts are or were supplied by the suppliers who supply/ied a given part at the same time, as currently recorded?”

```
SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30' as ps1,
partsupp FOR BUSINESS_TIME FROM '0001-01-01' TO '9999-12-30' as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
       ps1.ps_suppkey = ps2.ps_suppkey AND
       ps1.ps_partkey <> ps2.ps_partkey AND
       ps1.VALIDITY_TIME_START <= ps2.VALIDITY_TIME_END AND
       ps1.VALIDITY_TIME_END >= ps2.VALIDITY_TIME_START
```

Parameters:

- PART_KEY := uniform({ps_partkey})

B.5.4 App Time Current / System Time Correlation (B.4)

“What (other) parts were recorded at some time to be supplied by the suppliers who currently supplies a given part”

```
SELECT ps2.*, ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR BUSINESS_TIME AS OF '[APP_TIME]'
FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as PS1,
partsupp FOR BUSINESS_TIME AS OF '[APP_TIME]'
FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as PS2
WHERE ps1.ps_partkey = [PART_KEY] AND
      ps1.ps_suppkey = ps2.ps_suppkey AND
      ps1.ps_partkey <> ps2.ps_partkey AND
      ps1.sys_time_start < ps2.sys_time_END AND
      ps2.sys_time_start < ps1.sys_time_END
```

Parameters:

- PART_KEY := uniform({ps_partkey})
- APP_TIME1 := uniform(partsupp.validity_time_start, partsupp.validity_time_end)

B.5.5 App Time Correlation / System Time Correlation (B.5)

“When did we record that some (other) parts are or were supplied by the suppliers who supplies/ied a given part at the same time?”

```
SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
      ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps1,
      partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps2
WHERE ps1.ps_partkey = [PART_KEY]
      ps1.ps_suppkey = ps2.ps_suppkey AND
      ps1.ps_partkey <> ps2.ps_partkey AND
      ps1.sys_time_start < ps2.sys_time_END AND
      ps2.sys_time_start < ps1.sys_time_END AND
      ps1.VALIDITY_TIME_START <= ps2.VALIDITY_TIME_END AND
      ps1.VALIDITY_TIME_END >= ps2.VALIDITY_TIME_START
```

Parameters:

- PART_KEY := uniform({ps_partkey})

B.5.6 App Time Correlation / System Time Correlation (B.6)

“What (other) parts were supplied at any time by the supplier who supplied at any time a given part, as far as we currently know?”

```

SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp as ps1, partsupp as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
       ps1.ps_suppkey = ps2.ps_suppkey AND
       ps1.ps_partkey <> ps2.ps_partkey

```

Parameters:

- PART_KEY := uniform({ps_partkey})

B.5.7 App Time Agnostic/ System Time Past (B.7)

“What (other) parts were supplied at any time by the supplier who supplied at any time a given part, as far as we knew it at a previous point in time SYS_TIME?”

```

SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]' as ps1,
     partsupp FOR SYSTEM_TIME AS OF TIMESTAMP '[SYS_TIME]' as ps2
WHERE ps1.ps_partkey = [PART_KEY]
     AND ps1.ps_suppkey = ps2.ps_suppkey
     AND ps1.ps_partkey <> ps2.ps_partkey

```

Parameters:

- PART_KEY := uniform({ps_partkey})
- SYS_TIME := uniform(partsupp.sys_time_start, partsupp.sys_time_end)

B.5.8 App Time Agnostic/ System Time Correlation (B.8)

“When was it recorded that some (other) part was supplied by the supplier who supplied at any time a given part?”

```

SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps1,
     partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
       ps1.ps_suppkey = ps2.ps_suppkey AND
       ps1.ps_partkey <> ps2.ps_partkey AND
       ps1.sys_time_start < ps2.sys_time_END AND
       ps2.sys_time_start < ps1.sys_time_END

```

Parameters:

- PART_KEY := uniform({ps_partkey})

B.5.9 App Time Current/ System Time Agnostic (B.9)

“When was it recorded that some (other) part is supplied by the supplier who supplies a given part?”

```
SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp
   FOR BUSINESS_TIME AS OF '[APP_TIME]',
   FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps1,
partsupp
   FOR BUSINESS_TIME AS OF '[APP_TIME]',
   FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
      ps1.ps_suppkey = ps2.ps_suppkey AND
      ps1.ps_partkey <> ps2.ps_partkey AND
      ps1.sys_time_start <= ps2.sys_time_START AND
      ps2.sys_time_start < ps1.sys_time_END
```

Parameters:

- PART_KEY := uniform({ps_partkey})
- APP_TIME := uniform(partsupp.validity_time_start, partsupp.validity_time_end)

B.5.10 App Time Correlation/ System Time Agnostic (B.10)

“When was it recorded that some (other) part is or was supplied by the supplier who at the same the supplied a given part?”

```
SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps1,
     partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
      ps1.ps_suppkey = ps2.ps_suppkey AND
      ps1.ps_partkey <> ps2.ps_partkey AND
      ps1.sys_time_start <= ps2.sys_time_START AND
      ps2.sys_time_start < ps1.sys_time_END AND
      ps1.VALIDITY_TIME_START <= ps2.VALIDITY_TIME_END AND
      ps1.VALIDITY_TIME_END >= ps2.VALIDITY_TIME_START
```

Parameters:

- PART_KEY := uniform({ps_partkey})

B.5.11 App Time Agnostic/ System Time Agnostic (B.11)

“When was it recorded that a part was supplied by the supplier who supplied at some time a given part?”

```
SELECT ps2.*, ps1.sys_time_start, ps1.sys_time_end,
       ps1.validity_time_start, ps1.validity_time_end
FROM partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps1,
     partsupp FOR SYSTEM_TIME BETWEEN '0001-01-01' AND '9999-12-30' as ps2
WHERE ps1.ps_partkey = [PART_KEY] AND
     ps1.ps_suppkey = ps2.ps_suppkey AND
     ps1.ps_partkey <> ps2.ps_partkey AND
     ps1.sys_time_start <= ps2.sys_time_START AND
     ps2.sys_time_start < ps1.sys_time_END
```

Parameters:

- PART_KEY := uniform({ps_partkey})