Benchmark frameworks and τBench

Stephen W. Thomas[1,*,†], Richard T. Snodgrass[2] and Rui Zhang[2]

[1]*School of Computing, Queen's University, Kingston, ON, Canada*
[2]*Department of Computer Science, University of Arizona, 711 Gould Simpson, AZ 85721, U.S.A.*

SUMMARY

Software engineering *frameworks* tame the complexity of large collections of classes by identifying structural invariants, regularizing interfaces, and increasing sharing across the collection. We wish to appropriate these benefits for families of closely related benchmarks, say for evaluating query engine implementation strategies. We introduce the notion of a *benchmark framework*, an ecosystem of benchmarks that are related in semantically rich ways and enabled by organizing principles. A benchmark framework is realized by iteratively changing one individual benchmark into another, say by modifying the data format, adding schema constraints, or instantiating a different workload. Paramount to our notion of benchmark frameworks are the ease of describing the differences between individual benchmarks and the utility of methods to validate the correctness of each benchmark component by exploiting the overarching ecosystem. As a detailed case study, we introduce τBench, a benchmark framework consisting of ten individual benchmarks, spanning XML, XQuery, XML Schema, and PSM, along with temporal extensions to each. The second case study examines the *Mining Unstructured Data* benchmark framework, and the third examines the potential benefits of rendering the TPC family as a benchmark framework. Copyright © 2013 John Wiley & Sons, Ltd.

1. INTRODUCTION AND MOTIVATION

Benchmarks enable an apples-to-apples comparison of competing software or hardware systems, architectures, algorithms, or other technologies, as well as an evaluation of a given technology under different scenarios. Benchmarks generally consist of *data*, a *schema* to describe the data, and *workload operations* to perform on the data (benchmarks vary on the emphasis each accords these three components). A benchmark provides a *standard* with which to compare and evaluate. For example, the Standard Performance Evaluation Corporation (SPEC) has created a set of commercial benchmarks for testing CPU performance, with the main goal of providing compiler vendors with representative programs [1] (as SPEC elegantly asserts that 'an ounce of honest data is worth a pound of marketing hype' [2]).

Typical uses of a benchmarks include the following [3]:

- Comparing different software and hardware systems.
- Comparing software products on one machine.
- Comparing machines in a compatible family.
- Comparing releases of a product on one machine.
- Comparing implementation strategies of a system.

---

*Correspondence to: Stephen W. Thomas, School of Computing, Queen's University, Kingston, ON, Canada.

†E-mail: sthomas@cs.queensu.ca

It is a challenging task to create datasets, schemas, and workloads that are both simple enough to be understood by a wide audience as well as realistic enough to be useful. Additionally, because major technical and marketing decisions often hinge on the results of a benchmark evaluation, it is critical that the benchmarks be correct. Any method that helps in the creation, maintenance, and description of benchmarks will thus be beneficial to all stake holders involved.

In the software engineering domain, *frameworks* are widely used to manage the complexity of creating and maintaining software functionality [4]. Each framework includes a set of organizing principles that enable it to be assimilated and used effectively [5].

In this paper, we introduce the notion of a *benchmark framework*, which emphasizes the organizational principles of frameworks to achieve many of the advantages that software engineering frameworks provide. Briefly, a benchmark framework is an ecosystem of benchmarks that are related in meaningful ways, along with shared tools and techniques to generate and validate new benchmarks from existing benchmarks. Benchmark framework can help alleviate the difficulties of creating, maintaining, validating, and describing sets of related benchmarks. Benchmark frameworks are built iteratively by using tools to alter existing benchmarks to obtain new benchmarks, followed by a validation phase to ensure that each new benchmark is still correct. This iteration produces a DAG-like structure of benchmarks, where nodes are benchmarks and edges are the relationships between benchmarks. The structure of a benchmark framework encourages the designer to be explicit about the specific relationships between the individual benchmarks, thereby highlighting the commonalities between and distinctions among the benchmarks.

Not every benchmark needs to be in a framework, and benchmark frameworks may not be ideal for every situation. For example, if commercial politics dictate exactly which benchmarks must be used to evaluate a given technology, and how those benchmarks must be structured, then benchmark frameworks may not be appropriate. However, benchmark frameworks have several advantages of which practitioners and researchers can take advantage, such as (i) reusing an extant benchmark without requiring modification of that benchmark, which helps to reduce development cost, (ii) sharing of data, schema, workloads, and tools across benchmarks, (iii) regularizing the execution of the benchmarks, and (iv) clearly communicating the essence of each benchmark. These advantages will encourage practitioners and researchers to create their own benchmarks, on the basis of existing benchmarks, when they may have otherwise avoided doing so because of time and resource constraints. In this paper, we explore these potential advantages of benchmark frameworks.

Consider the TPC benchmarks, a popular family of OLTP and OLAP benchmarks [6]. In their current form, the TPC benchmarks are loosely connected, but understanding these connections is difficult. For example, documentation suggests that the TPC-R benchmark is related to the TPC-H benchmark, but exactly how is not defined (Do the two use the exact same data? Same schema? Same workload? Do the data generators share code?) Further, the TPC-E benchmark seems to be related, in different ways, to both the TPC-C and TPC-H benchmarks. Again, because the relationship is not explicitly defined, potential users of these benchmarks can only guess. As we show in Section 7, additional structuring, documentation, and sharing of tools can realize a TPC benchmark framework that embodies the advantages listed earlier.

To make our proposal concrete, we present τBench, a benchmark framework for temporal databases composed of some ten individual benchmarks. τBench emphasizes the organizational principles of using statement modifiers [7], commutativity validation, and temporal upward compatibility [8] for creating temporal benchmarks from nontemporal ones.

The contributions of this paper are as follows:

- We introduce and describe in detail the notion of a *benchmark framework*, which allows simple, correct, and adaptable sets of related benchmarks (in any domain) to be created and documented.
- We show the utility and generality of benchmark frameworks.
- We present τBench, a benchmark framework that encompasses several XML, relational, and temporal technologies. τBench is distributed according to the GNU Lesser GPL [9] and is available online [10].

- We briefly present the Mining Unstructured Data (MUD) and TPC benchmark frameworks from differing domains, to emphasize the generality and utility of our approach.

The remainder of this paper is structured as follows. We describe software engineering frameworks, our inspiration for benchmark frameworks, in Section 2. We describe benchmarks and benchmark families, as well as introduce benchmark frameworks, in Section 3. We introduce benchmark frameworks in Section 4. We present τBench, a specific instance of a benchmark framework, in Section 5. We briefly present two additional benchmark frameworks from other domains in Sections 6 and 7. With our experience from the case studies, we discuss the advantages and disadvantages of benchmark frameworks in Section 8. Finally, we conclude and outline future work in Section 9.

2. SOFTWARE ENGINEERING FRAMEWORKS

As mentioned previously, developers use software engineering frameworks to manage the complexity of creating and maintaining software functionality [4, 5]. For example, the Java Swing framework provides Java developers with a standardized API for creating graphical user interfaces [11]. This standardization allows developers to easily develop new GUIs and to easily understand existing GUIs written with Swing. Other frameworks include the following: the C++ Standard Template Library (STL) [12], an assemblage of data structures; OpenMS [13], for developing mass spectrometry data analysis tools; the Computational Crystallography Toolbox [14], for helping to advance the automation of macromolecular structure determination; ControlShell [15], for developing real-time software; and log4j [16], for enabling logging with context to Java programs through a hierarchy of loggers.

Software engineering frameworks are typically born from the necessity of taming the complexity that grows naturally over time. Developers begin by writing single classes. Families of related classes emerge as the application grows, such as those for sorting and searching a custom data structure. Over time, the family of related classes grows so large and interconnected that they become unmanageable and unmaintainable without some explicit structure and organization. At this point, developers create frameworks to tame this complexity: structural invariants are identified, interfaces are regularized, constraints are identified and enforced, and the coherency of the family is emphasized.

The creation of frameworks is not mechanical: it involves the manual and creative analysis of the problem domain and desired solution space. The result, however, is often a significant decrease in complexity of the classes and a significant increase in their reusability. In addition, holes in functionality can be easily identified and filled.

Let us examine in some detail the Java Collections framework [17], a large assemblage of data structures. The Collections framework is composed of a sequence of successively elaborated Java interfaces to characterize functionality, a sequence of successively elaborated Java abstract classes that implement these interfaces to provide data structure-independent methods, and a set of concrete classes that extend these abstract classes to provide specific data structures and thus trade-offs in performance between the various operations specified originally in the interfaces. The Collections framework also uses new concepts such as modifiable versus immutable, ordered versus unordered, getters versus mutators, comparables and compareTo, and static classes (e.g., `Arrays`). The framework exploits Java's inheritance, generics, enhanced `for` loops, and dynamic binding language features. Finally, the framework effectively utilizes the Command, Iterator, and Template patterns [18].

Within this framework, the `Iterable` Java interface provides a method to obtain an `Iterator` across the collection (hence, every collection has an iterator). The `Collection` Java interface extends `Iterator` and introduces a dozen methods (functionality) used across the framework. The `List` Java interface extends the `Collection` Java interface and specifies the functionality of an ordered collection through ten additional methods, including the `set(int <E>)` method that explicitly states a position. The `AbstractList` class provides initial implementations of most of the methods of the `List` interface, often utilizing the Template Pattern. Finally, we get

to the `LinkedList` concrete class (which can actually be instantiated); this class provides a specific data structure and refined implementations by overriding some of the implementations in the `AbstractList` class, thereby providing specific performance properties, such as a very fast *next* accessor and minimal storage overhead for highly volatile collections.

The framework thereby presents a unified and carefully structured assemblage of some 75 classes and interfaces supporting several thousand methods, enabling programmers to understand and use these many classes without having to study the details of every individual Java method. As another benefit, this carefully orchestrated dance of Java interfaces, abstract classes, and concrete classes helped to ferret out inconsistencies. For instance, the documentation for the `Vector` class states that 'As of the Java 2 platform v1.2, this class was retrofitted to implement the List interface, making it a member of the Java Collections Framework'. The `Collection` and `List` interfaces make it easier for programmers to use these APIs, as all the data structures in this framework abide by the `Collection` interface and all the ordered data structures abide by the `List` interface.

In this paper, we apply what has been performed repeatedly in software engineering to the domain of benchmarks. Our goal is to achieve the structure and organization that software engineering frameworks provide, except using individual benchmarks and families of related benchmarks instead of individual classes and families of related classes.

3. BENCHMARKS AND BENCHMARK FAMILIES

We now define more thoroughly the concepts and structures of individual benchmarks and benchmark families, on our way to the more internally structured benchmark frameworks in the following section.

3.1. Benchmarks

A *benchmark* is a collection of data, workload, and metrics. Individual benchmarks, such as CloudCMP [19], INEX [20], WSTest [21], X007 [22], XMark [23], and XQBench [24], allow applications and algorithms to be compared and evaluated in a standard, consistent setting, giving researchers and practitioners invaluable guidance as to the relative performance of the systems under consideration [3].

In this paper, we further structure an individual benchmark into ten components, shown in Table I. All but the last component, tools, are required, although some might be only implied by a particular benchmark (we recommend though that a benchmark be explicit about every component.) In the table, a component is ordered after its predecessor(s), so for example, the data model must be specified before the data definition language (DDL), which must in turn be specified before the schema, which must be specified before the actual data of the benchmark.

The *schema* component describes the structure and constraints of the data, such as the relationship between books and authors (from the aforementioned XBench example). The schema is expressed in a DDL, also a component, such as XML Schema, SQL DDL, a new schema language, or informally. The schema component must conform to the benchmark's specified DDL.

The *data* component is the actual data in the benchmark, for example, an XML file containing information for a library catalog of books and authors. The data can be an actual dataset, or a tool that randomly generates data, or a combination of the two. For example, in XBench, XMark, and XOO7, the data are randomly generated using ToXgene, xmlgen, or in-house tools, respectively; in TPC-H (part of the TPC benchmark family, to be discussed in Section 3.2), there is a given single (static) dataset. The data conform to a *data model*, which describes the structure and form of the data. For example, the data in XMark are represented in XML, thus conforming to the hierarchical data model. Within a benchmark, the data must conform to the data model.

The *workload* component contains tasks to perform on the data, such as querying for specific authors, inserting new books, or ensuring that all books are associated with at at least one author. The workload is expressed in a *data manipulation language (DML)*, also a component, such as XQuery, SQL DML, or a new query language.

Table I. The ten components of a benchmark.

Component name	Description	Example
Data model	Describes how the data is stored and organized	XML model
Data definition language	The language in which the schema is written	XML Schema
Schema	Describes the structure of and constraints on the data	A book has at most five authors
Data	The actual data of the benchmark	A catalog with books and authors
Data manipulation language	The language in which the workload is written	XQuery
Workload	A set of tasks (e.g., queries) to perform on the data	A query that selects books by 'Ben'
Metrics	The metrics to be measured and how to measure them	Milliseconds per query
Context	A specification of the execution environment	CPU, memory, and OS version
Documentation	Metadata (authors, date, and version) and execution guide	Author: Bob Smith.
Tools	Executables	Makefiles, data generators and validators, execution scripts, and so on

The *metrics* component is a description of the metrics to be measured, including instructions on how to measure and calculate each one. For example, Document-centric (DC)/Single Document (SD) XBench contains two metrics: bulk load time and query execution time.

The *context* component specifies what characteristics of the software and hardware environments must be reported during the execution of the benchmark. The benchmark creator can be as specific as desired: the context could require that the exact CPU characteristics (e.g., clock speed and cache size) must be reported or simply the number of cores.

The *documentation* component includes metadata and an execution guide. Metadata includes the benchmark name, authors and affiliations, date published and date updated, version, and other descriptive information. The execution guide describes how to configure the software environment, which settings and values are required, order of program execution, and so on.

The metrics, context, and documentation form the *provenance* of both the benchmark and the results of applying that benchmark to a particular hardware/software configuration. We anticipate that benchmarks will evolve through distinct and identified versions, with a benchmark result explicitly associated with a version.

Various kinds of benchmarks will emphasize different components. Database benchmarks emphasize the data and workload and often deemphasize the metrics (e.g., simply transactions per second). Compiler benchmarks such as SPEC have programs as their workload, with little emphasis on the data and even less on the DDL and schema. It appears that all the benchmarks mentioned in this paper contain (or could have contained, if the benchmark was more fully realized) all of the aforementioned components.

Finally, the *tools* are (optional) executables for generating and validating the data, translating and executing the workload, and collecting metrics. Example tools include the following.

Generation. Data generators, such as TOXGENE [25] and τGenerator (Section 5.6.1), turn raw input data into larger and more complex datasets by repeating, combining, and changing the raw data in various ways, perhaps on the basis of user parameters. Data generators are often template based, generating data from a user-specified template.

Validation. Validation scripts ensure that component changes are correct, amongst other things (Section 4.4).

Translation. Data translators are used to easily change data from one format to another, for example, by shredding an XML document into relations [26]. Schema/workload translators such as TXL [27] can be used to translate the schema and workload from one form to another, for example, to translate τXQuery [28] queries into standard XQuery [29] queries [30].

3.2. Benchmark families

A *benchmark family* is a set of loosely related benchmarks. A benchmark family is often developed by a single organization. Benchmark families lack the information necessary to easily identify benchmark relationships or tools necessary to derive new benchmark components from existing benchmark components.

An example of a benchmark family is XBench [31], which contains four benchmarks: DC/SD, DC/Multiple Document (MD), Text-centric (TC)/SD, and TC/MD. The benchmarks are loosely related in that they share a document generation tool. However, the benchmarks do not share schemas or workloads, and the specific relationships between benchmark components are not well defined.

Another benchmark family is TPC, which we described in Section 1. As in XBench, the benchmarks in TPC are loosely related in that they are developed by the same organization, but the relationships between benchmark components are not well defined and no data, schema, or workloads are reused among the currently supported TPC-C, TPC-H, and TPC-E benchmarks.

The SPEC provides a wide rage of evolving benchmarks [1, 32]. For example, the programming language community employs the *SPEC CPU benchmark* to study program performance with intensive computation. This benchmark stresses not only the processor of a machine but also its memory hierarchy and the properties of the compiler. Other SPEC benchmarks include SPECviewperf to evaluate graphics performance, SPECapc to evaluate performance of specific applications, such as LightWave 3D and SolidWorks, SPEC MPI to evaluate MPI-parallel, floating point, compute-intensive across cluster, and SMP hardware performance, SPEC OMP to evaluate OpenMP-parallel shared-memory parallel processing across shared-memory parallel processing, SPECjvm to evaluate performance of the Java Runtime Environment, SPECjbb as well as SPECjEnterprise and SPECjms to evaluate performance of Java business applications, and SPECsfs to evaluate performance of file servers.

The TREC conference series [33] has produced a set of text retrieval benchmarks, each consisting of a set of documents (data), a set of topics (questions), and a corresponding set of relevant judgments (right answers), across a range of domains, such as blog, chemical, genomics, legal, and web, and across question characteristics, such as filtering, interactive, novelty, and query.

4. BENCHMARK FRAMEWORKS

We introduce here a *benchmark framework* as a refined benchmark family. A benchmark framework is an ecosystem of benchmarks that are interconnected in semantically rich ways. For example, the data used in one benchmark might be derived from the data of another (e.g., by increasing the size of the data), whereas the workloads and schemas of the two benchmarks are identical.

Like software frameworks, the creation of a benchmark framework is not mechanical. Instead, they are created iteratively: beginning with a root (or *foundation*) benchmark B, a new benchmark B' is created by changing one or more of the six components of B (components are listed in Table I and described in Section 3.1). A *relationship* between B and B' is made to describe the changes. In this way, a DAG of benchmarks emerges, where nodes are benchmarks and each edge is the relationship between each pair of benchmarks. This graph exactly describes a *benchmark framework*. At this point, benchmark frameworks depart in design from software frameworks such as the Java Collections framework. To build the Collections framework, the designers had to first boil the ocean: almost the entire framework had to be in place before any of it was useful. In contrast, benchmark frameworks are useful immediately, as new benchmarks can easily be derived whether there are two benchmarks or two hundred present in the framework.

Figure 1 shows an example benchmark framework, which we use as a running example in this section. The framework begins with the DC/SD XBench benchmark, a popular XQuery benchmark from the XBench family of benchmarks [31]. The data in the DC/SD XBench benchmark are a library catalog of books and authors contained in a single XML document. In this simple benchmark framework, the τXBench benchmark is derived from the DC/SD XBench benchmark by using a tool to render the data as time varying, whereas the PSM benchmark, which is utilized
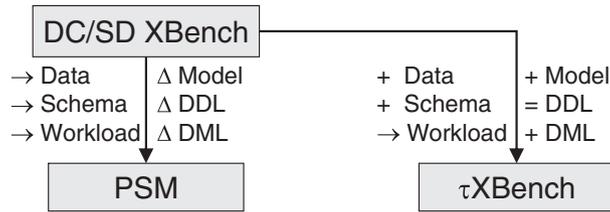
Figure 1. Example benchmark framework.

in a study to evaluate the performance of the SQL/PSM language constructs, is derived from the DC/SD XBench benchmark by using a tool to shred the XML data into relational data. The arcs and annotations in the DAG succinctly denote these relationships; we describe these annotations in detail in Section 4.2.

We now describe the general principles of benchmark frameworks (Section 4.1), define the possible relationships between benchmark components and how to visualize relationships (Section 4.2), define relationship semantics (Section 4.3), and present the methodology for creating a benchmark framework (Section 4.4).

4.1. General principles

Our proposal of benchmark frameworks, detailed in the following sections, is based on several general principles. We define a benchmark framework to be a collection of benchmarks that satisfy all five principles.

First, a benchmark framework must explicitly state one or more *organizing insights* that bring together constituent benchmarks into a coherent whole. For example, the Java Collection Framework discussed in Section 2 utilizes several such insights: successively elaborated and related Java interfaces, abstract classes, and concrete classes, utilizing several powerful design patterns throughout and effectively exploiting particular Java constructs. As we describe three benchmark frameworks as case studies later in this paper, we first state the organizing insight(s) for each.

The second organizing principle is that of *automation*, of producing a component of a child benchmark by executing mapping or conversion programs on a component of a parent benchmark. Whereas some of the components of a benchmark must be defined manually, many of the other components can be generated by a tool that automatically maps a prior component to the new component. We will see many examples of such mapping tools.

Reuse is the third required organizing principle of a benchmark framework. Reusing existing components to create new components is one of the fundamental ideas behind frameworks of any kind, and benchmark frameworks are no different. By reusing benchmark components, developers can save development cost, benchmark relationships become more clear, and cross-validation between components becomes more applicable. Reuse must be applied to several components of each constituent benchmark in the framework. We will see many examples as we examine τBench in detail shortly, as so we will mention just a few applications of reuse here.

The queries comprising a workload can be mapped to a new language with a source-to-source translation, either manually or with a tool. A separate source-to-source translation can map a standard language to a DBMS-specific variant. These translations can be composed, to define further benchmark components.

Similarly, the mapping of the data component between two benchmarks can often be reused elsewhere in the framework to produce yet other components, when two source benchmarks are also related in a different way than to their child benchmarks. An example is the mapping of the data from XBench to PSM being reused in the mapping from τXBench to τPSM. The source-to-source mapping of a schema component can likewise be reused. The third potential for reuse is a tool. For example, τBench employs a single `queryRunner` Java application to time queries in three of its constituent benchmarks (those employing the relational model).

Organizational clarity is the fourth general principle. By clearly defining the relationships between benchmark components, developers and users of the benchmarks can immediately determine the differences and similarities between them. Choosing which benchmark or set of benchmarks to use, or which benchmark to extend, becomes much easier. These relations should be formalized in a representation exemplified in Figure 1.

Providing *consistency and validation checking* is the final general principle required of benchmark frameworks. As mentioned previously, users of benchmarks often have a high stake in the correctness of benchmark components, as they are used to evaluate new technologies and make critical business decisions. Benchmark frameworks allow a new level of consistency and validation checks to be performed across benchmarks, because they share components and tools. Thus, the entire benchmark framework is strengthened with each additional check of each additional component.

These five general principles work in concert to allow benchmark frameworks to be developed faster and to attain higher quality than a more loosely structured family containing the same constituent benchmarks.

4.2. Relationships between benchmark components

A *relationship* between a component of a parent benchmark component and a component of a child benchmark is an explicit link between the two, describing how to change one component into the other. Component relationships characterize the *flow of information* between the benchmarks. Relationships allow new benchmark components to be easily adopted from existing benchmark components and for that adoption to be explicitly and clearly defined. Each component relationship, annotating a line from one benchmark to another, is prefixed by one of the following symbols.

Expand (+). A component is expanded to include additional content. For example, the data in τXBench are expanded from DC/SD XBench by making the data temporal (i.e., adding timestamps to the data and changing random books and authors at various points in time).

Subset ($\supset$). A component is subsetted to remove content. As we will see in Section 5.4, the DDL of τPSM, that of a temporal extension of the SQL DML, is subsetted back to that of SQL DML in the subsequent τPSM-MaxMapped DML.

Retain (=). A component is retained without modification. For example, the DDL remains the same (i.e., XML Schema) between the DC/SD XBench and τXBench benchmarks.

Replace (Δ). A component is replaced (i.e., completely removed) and substituted with a new instance. For example, the data model for PSM, which is the relational data model, replaces DC/SD XBench's hierarchical data model.

Map ($\rightarrow$). A component is mapped from one data model/DDL/DML to another, but the ideas behind the components remain the same. For example, the workload for PSM is logically unchanged from the DC/SD XBench benchmark (i.e., the queries are seeking the same information) but must be mapped from the XQuery syntax to the PSM DML syntax. Here, the *form* of the component changes, but the *content* remains.

We visualize the relationships of a benchmark framework by displaying a DAG, where nodes are benchmarks and edges are relationships between them (e.g., Figure 1). We term this a *Benchmark–Relationship* (or *B-R*) *diagram*.

A B-R diagram enumerates the benchmarks within the framework (or family) as well as between frameworks and families. Thus, the XBench family is included in Figure 1 because one of its benchmarks is used and extended in the τBench framework. An edge between two benchmarks denotes that components are shared between them; the edge is annotated with the exact relationship of each component (we only annotate the first six components, because these best characterize the changes from benchmark to benchmark). For consistency, component annotations are always displayed in the same location. For example, the data component is shown in the upper left of the edge; the DML component is shown in the bottom right. Note the left-right correspondence: the Data accords with

a Data Model, the Schema is expressed in a DDL, and the Workload is expressed in a DML. For brevity, component annotations are omitted from the display if the relationship is retain ($=$), except in Figure 1, in which we aimed for completeness.

4.3. Relationship semantics

A new benchmark can be composed of components from multiple existing benchmarks (e.g., the data and schema from one and the workload from another), which is precisely why a benchmark framework forms a DAG instead of a tree. In these cases, a component annotation existing from both parents implies that the components should be merged (by performing a union) into the child.

Each component can be associated with only one relationship between two benchmarks. If two relationships are needed, such as expanding and mapping a workload, this is better expressed as two steps, yielding the second benchmark derived from the original with one of the relationships, say mapped, and the third benchmark, derived from the second with the other of the relationships, in this case expanding.

When the data are expanded or subsetted, the semantics are clear: data were added or removed, respectively. But when data models are expanded or subsetted, the semantics are domain dependent. In the temporal database domain, for example, a data model can be expanded from nontemporal relational to temporal relational. Although this can also be expressed as a replacement (Δ) of data models, we feel that expanding ($+$) more clearly describes the relationship because the temporal data model offers similar but more information.

In general, the map relationship ($\rightarrow$) is only applicable to the data, schema, or workload components; it is not appropriate to say that a data model, DDL, or DML has been mapped, because these components contain ideas rather than content.

Finally, a change in one component may imply a change in other component(s). For example, when changing the data model from the XML hierarchical data model (where the DDL is likely to be XML Schema or Relax NG [34]) to the relational set-based data model, the DDL also must be changed to an appropriate language (e.g., SQL DDL). Table II lists these and other structural implications that help define the validity of a benchmark framework. This list is not exhaustive but rather illustrative of the semantics of component relationships. As mentioned in Section 9 (future work), we envision an application-specific tool that can automatically validate component relationships, given a list of constrains such as those in Table II. For example, if the relationship between the data component of two benchmarks is subset ($\supset$), then the tool can read the datasets and ensure that one is in fact a subset of the other; similar checks can be performed for schemas and workloads.

Table II. Examples of relationship constraints.

Relationship(s)	Implication	Explanation
$= $ Data $\vee \supset$ Data	$=$ Data model	If the data remain the same, or are subsets, then they must have the same data model, because the data are an instance of the model
$\supset$ Data	$=$ DDL $\wedge$ ($=$ Schema $\vee \supset$ Schema)	Subsetting the data implies constraints on the schema
Δ Data model	Δ DDL $\wedge \Delta$ DML $\wedge$ ($\rightarrow$ Data $\vee \Delta$ Data)	New data models need new DDLs and DMLs, and the data must be mapped or changed to the new model
$\rightarrow$ Data model	($\rightarrow$ DDL $\vee \Delta$DDL) $\wedge$ ($\rightarrow$ DML $\vee \Delta$DML)	Mapping the data model implies fairly whole-scale changes to the DDL and DML
$=$ Schema $\vee \supset$ Schema	$=$ DDL	The schema is an instance of the DDL
$=$ DDL	$=$ Schema $\vee \supset$ Schema	—
Δ DDL	$\rightarrow$ Schema $\vee \Delta$ Schema	—
$=$ Workload $\vee \supset$ Workload	$=$ DML	The workload is an instance of the DML
$\supset$ DML	$\rightarrow$ Workload $\vee \supset$ Workload	Subsetting the DML necessitates changes to the workload that utilizes the removed constructs
Δ DML	$\rightarrow$ Workload $\vee \Delta$ Workload	—

4.4. Methodology

We now describe how to create a benchmark framework, which consists of three steps: (i) begin with a foundation benchmark; (ii) transform the foundation benchmark into a new benchmark to suite your application's needs; and (iii) validate the correctness of the data, schema, and associated constraints. Repeat the process for each additional benchmark that is required. We describe each step in turn.

Foundation. Start with an existing benchmark as a solid foundation. In our example, we use DC/SD XBench as the foundation benchmark, because it provides well-defined schema and datasets, which are based on generally realistic assumptions [31].

Transformation. Build new components from existing components. Either manually or by using automated tools, alter the components of existing benchmarks to create new benchmark components. For example, use the τGENERATOR tool (to be described in Section 5.6.1) to generate temporal data from the nontemporal data of DC/SD XBench. We term the resulting new benchmark the *child* and the existing benchmark the *parent*, respectively.

Validation. Perform cross-validation of benchmark components throughout the process (i.e., after each change). There are no standard set of validation steps applicable to all benchmark frameworks, as these steps are highly domain and task dependent. Validation steps can be simple (e.g., checking referential integrity of two relations using Perl hash tables) or complex (e.g., commutativity validation of multiple benchmarks, described in Section 5.6.3); we provide examples in each of the benchmark frameworks presented in this paper.

By performing continuous validation steps, each component is tested in different ways. Each successful validation adds validity to the entire benchmark framework, because of the rich interconnectedness of the benchmarks: when the data are validated in a child benchmark, the data are also likely to be valid in the parent. These validation steps check more than just the data: they check the schema, the data generation tools, the data and query translation tools, and the query engine. The resulting ecosystem of benchmarks is now stronger than the sum of the individual benchmarks, because of the validation of each component under multiple perspectives.

5. CASE STUDY: τBENCH

Using the methodology just presented, we have created a benchmark framework, called τBench [10, 35]. As we describe in the succeeding texts, τBench makes use of the organizational insight of using statement modifiers, commutativity validation, and temporal upward compatibility when deriving temporal benchmarks from nontemporal benchmarks.

In this section, we first describe the structure of τBench. We then outline the tools we have built to support the creation and validation of the individual benchmarks of τBench. Finally, we provide a detailed example of the creation of one of the benchmarks in the framework.

τBench currently consists of ten individual benchmarks, spanning various technologies including XML, XML Schema [36, 37], XQuery [29], τXSchema [38–43], τXQuery [28, 44, 45], PSM [46, 47], and τPSM [30, 44]; the B-R diagram is given in Figure 2, and the main components are summarized in Table III. The benchmarks are fully described elsewhere [35]; we briefly describe the components of each here.

5.1. The DC/SD XBench *benchmark*

XBench is a family of benchmarks for XML DBMSes, as well as a data generation tool for creating the benchmarks [31]. The XBench family has been widely used as standard datasets to investigate the functionalities of new and existing XML tools and DBMSes (other XML benchmarks include MBench [48], MemBeR [49], Transaction Processing over XML (TPoX) [50], XMach-1 [51], XMark [52], XPathMark [53], and XOO7 [54]. Most of these benchmarks focus on XQuery processing, as does XBench).
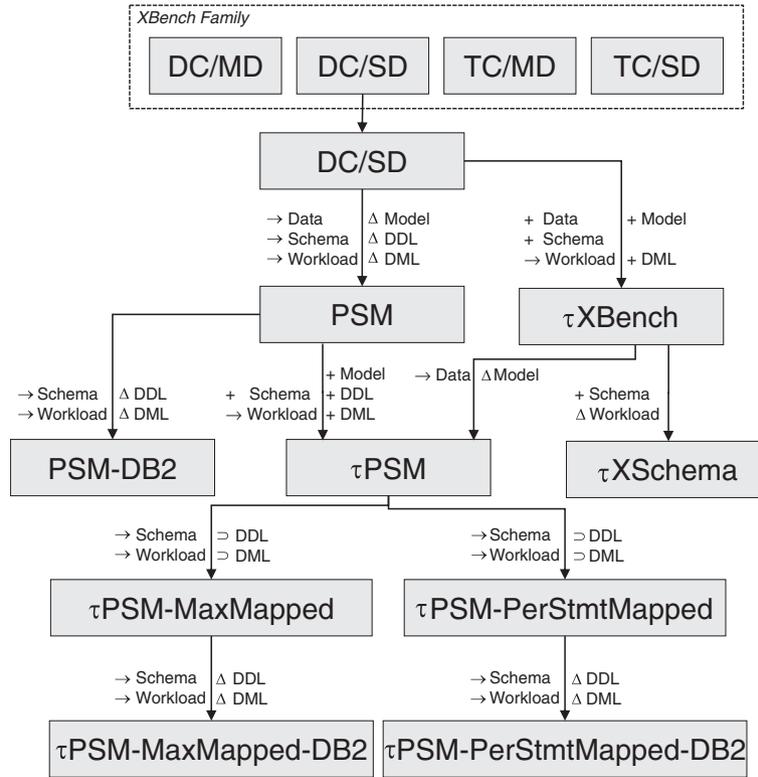
Figure 2. The Benchmark–Relationship diagram of the τBench framework.

Table III. A description of the τBench benchmark framework.

Benchmark name	Derived from	Data model	DDL	DML
DC/SD XBench	—	XML	XML Schema	XQuery
τXBench	DC/SD XBench	XML	τXSchema	τXQuery
τXSchema	τXBench	XML	τXSchema	—
PSM	DC/SD XBench	Relational	PSM DDL	PSM DML
PSM-DB2	PSM	Relational	DB2 DDL	DB2 DML
τPSM	PSM, τXBench	Relational	PSM DDL	τPSM
τPSM-MaxMapped	τPSM	Relational	PSM DDL	PSM DML
τPSM-MaxMapped-DB2	τPSM-MaxMapped	Relational	DB2 DDL	DB2 DML
τPSM-PerStmtMapped	τPSM	Relational	PSM DDL	PSM DML
τPSM-PerStmtMapped-DB2	τPSM-PerStmtMapped	Relational	DB2 DDL	DB2 DML

To create the datasets, the authors of XBench analyzed several real-world XML datasets to quantify their characteristics and relationships. As a result of this analysis, XBench characterizes database applications along two independent dimensions: application characteristics (i.e., whether the database is data centric or text centric) and data characteristics (i.e., single document vs multi-document). Thus, XBench consists of four different categories that cover DC/SD, DC/MD, TC/SD, and TC/MD, respectively.

The DC/SD XBench benchmark is one of the four constituent benchmarks of the XBench family, based strictly on the DC/SD category. We use this benchmark as the foundation of τBench, because of its combination of simplicity and generally realistic assumptions.

Schema. The schema of DC/SD XBench is specified as an XML Schema [36, 37]. The schema defines the library catalog, where <item> (i.e., book) elements contain a list of <author>s, a <publisher>, and a list of <related_items> of related books (Figure 3). Each <item>,

```
<catalog>
  <item id="I2">
    ...
    <authors>
      <author id="A5"> ... </author>
      ...
    </authors>
    <publisher id="P6"> ... </publisher>
    <related_items>
      <related_item id="R1"> ... </related_item>
    </related_items>
  </item>
  ...
</catalog>
```

Figure 3. An outline of the XBench DC/SD schema.

<author>, and <publisher> has a unique id attribute. The schema specifies several constraints, such as the uniqueness of the id attributes and that items can be related to between 0 and 5 other items.

Data. XBench populates each data element and attribute with random data generated from ToXgene [25]. ToXgene is a template-based tool that generates synthetic XML documents according to one or more user-defined templates.

The data in DC/SD XBench consist of a single document, catalog.xml, which contains information about a library catalog, that is, items (books), authors, publishers, and related items. The actual text content of each element or attribute is a randomly generated string or number, depending on the specification in the ToXgene template. For example, an instance of the <last_name> of an <author> could be the string 'BABABABAOGREAT' (however, some content is not random. For example, a <related_ID> element will always point to an existing item ID).

A crucial point to note is that, because of the data generation process, there is a strict one-to-many relationship between items and authors: although an item can have several authors, an author is only the author of a single item (no two items share an author). Similarly, there is a strict one-to-one relationship between items and publishers: an item has exactly one publisher, and each publisher publishes exactly one item. These constraints are not, however, explicit in the schema but must be satisfied by tools that manipulate the data, as we shall see in Section 5.6.1.

Workload. The workload consists of 17 XQuery [29] queries. Each query tests a different feature of the query engine, such as existential quantifiers (query *q6*: 'Return items where some authors are from Canada'), sorting by string types (query *q10*: 'List the item titles ordered alphabetically by publisher name'), and unigram searching (query *q17*: 'Return the ids of items whose descriptions contain the word hockey'). Each such query can thus be considered a *microbenchmark* [48], designed to test the performance of a specific language construct. Other microbenchmarks include IMbench [55] (addressing system latency), MemBeR [49] (also XQuery), and MBench [48] (also XQuery).

5.2. The τXBench benchmark

This benchmark adds a temporal dimension to the XBench DC/SD benchmark, so that temporal databases and query languages can be evaluated.

In particular, τXQuery is a temporal query language that adds temporal support to XQuery by extending its syntax and semantics [28, 44, 45]. τXQuery moves the complexity of handling time from the user/application code to the query processor. τXQuery adds two temporal statement modifiers [7] to the XQuery language: current and sequenced. A *current* query is a query on the current state of the XML data (i.e., elements and attributes that are valid *now*) and has the same semantics as a regular XQuery query applied to the current state of the XML data. *Sequenced* queries, on the other hand, are queries applied to each point in time, resulting in a sequence of temporal elements. Finally, τXQuery also has *representational* (also termed *nonsequenced*) queries, which

query the timeline of the XML data irrespective of time. No statement modifiers are needed for representational queries.

The τXBench benchmark was created to investigate τXQuery and other temporal query languages.

Schema. The schema of τXBench is reused from XBench DC/SD and translated into a *temporal XML schema* [38, 39], which is the same as an XML Schema, except with annotations to specify which elements (i.e., books, authors, and publishers) may vary over time.

Data. The data of τXBench are the result of applying a temporal simulation, τGENERATOR (Section 5.6), to the nontemporal XBench DC/SD dataset. This derived temporal dataset consists of books whose titles, authors, and publishers change over time, authors whose names and addresses change over time, and publishers whose lists of published books changes over time.

Workload. We defined a set of 85 τXQuery queries: a current and nonsequenced query for each of the original 17 queries from XBench DC/SD and three sequenced queries for each of the original 17 XBench queries: a *short period* query that only considers 10% of the time line, a *medium period* query that considers 50%, and a *long period* query that considers 90%.

The queries themselves are the same as those in the XBench DC/SD workload, except with temporal statement modifiers prepended to the query. Semantically, the queries are expanded, but we prefer in the B-R diagram to emphasize that an automated mapping (such as with TXL [27]) could convert the original queries into new queries by simply prepending the temporal statement modifiers. These new queries can then be executed by any tool that understands the statement modifiers, such as the τXQuery engine, or by a source-to-source translator from τXQuery to XQuery.

5.3. The PSM and PSM-DB2 benchmarks

Persistent stored modules (PSM) is the Turing-complete portion of SQL in which stored programs are compiled and stored in a schema and can be executed on the SQL server by queries [46, 47]. PSM consists of *stored procedures* and *stored functions*, which are collectively referred to as *stored routines*. The SQL/PSM standard [46] provides a set of control-flow constructs (*features*) for SQL routines.

We adopted the PSM benchmark from the XBench DC/SD benchmark to assess DBMSes that implement PSM functionality. The PSM-DB2 benchmark is a simple transformation of the schema and workload of the PSM benchmark so that the resulting benchmark is compatible and thus can be executed in IBM DB2.

Data. The data are reused as XBench DC/SD and translated into six relations.

Shredding is the process of transforming an XML tree into a set of two-dimensional DBMS relations [26]. Doing so often requires generating additional tables that contain relational or hierarchical information. For example, an additional table is required to list the relations between items and authors in the XBench DC/SD data. This relationship is implicitly captured in XML by making author a subelement of item but must be explicitly captured in its own relation during the shredding process.

We have created a tool to shred the XBench DC/SD data into six relations: four for the original four temporal elements, one to map author `ids` to item `ids`, and one to map publisher `ids` to item `ids`. In general, each column of each shredded relation corresponds to the text content of a subelement of the corresponding element.

Schema. The schema is the same as XBench DC/SD, except expressed as tables in the relational DDL.

Workload. The workload is derived from the original 17 queries in the XBench DC/SD workload, translated to be in the form of relational DML (i.e., SQL/PSM). Six queries were omitted, because they had no conceptual counterpart in PSM, such as testing the access of relatively ordered XML

elements. The workload was thus designed to be a microbenchmark, in that each PSM query highlights an individual construct of PSM, so that for example the performance of each construct can be studied in isolation, while retaining as much as possible the semantics of the original XBench query from which it was derived. Some queries, such as *q2*, were also changed to highlight a different feature, such as multiple `SET` statements in *q2b* and nested `FETCH` statements in *q7c*. See also *q7b*. Also, we added three queries, *q17*, *q18*, and *q19*, and one variant, *q17b*.

- Query *q2* highlights the construct of `SET` with a `SELECT` row,
- Query *2b* highlights multiple `SET` statements,
- Query *q3* highlights a `RETURN` with a `SELECT` row,
- Query *q5* highlights a function in the `SELECT` statement,
- Query *q6* highlights the `CASE` statement,
- Query *q7* highlights the `WHILE` statement,
- Query *q7b* highlights the `REPEAT` statement,
- Query *q7c* highlights the nested `FETCH` statements,
- Query *q8* highlights a loop name with the `FOR` statement,
- Query *q9* highlights a `CALL` within a procedure,
- Query *q10* highlights an `IF` without a `CURSOR`,
- Query *q11* highlights creation of a temporary table,
- Query *q14* highlights a local cursor declaration with associated `FETCH`, `OPEN`, and `CLOSE` statements,
- Query *q17* highlights the `LEAVE` statement,
- Query *q17b* highlights a nonnested `FETCH` statement,
- Query *q18* highlights a function called in the `FROM` clause, and
- Query *q19* highlights a `SET` statement.

The result is a set of 17 queries.

The PSM-DB2 workload utilizes a source-to-source translation of standard PSM into the syntactic variant supported by IBM DB2.

5.4. The τPSM and related benchmarks

Temporal SQL/PSM is a temporal extension to SQL/PSM that supports temporal relations and queries in conjunction with PSM functionality [30, 44]. The approach is similar to τXQuery in that it requires minor new syntax beyond that already in SQL/Temporal to define temporal PSM routines. Temporal SQL/PSM enables current queries (which are evaluated at the current time), sequenced queries (which are evaluated logically at each point in time independently), and nonsequenced queries (which are the most general of temporal semantics), as well as these three variants of PSM routines, similar to those described earlier for τXQuery.

The τPSM benchmark was created to test and validate Temporal SQL/PSM. There have been two different time-slicing techniques proposed for implementing Temporal SQL/PSM: *maximally fragmented slicing* and *per-statement slicing* [30, 44]. These slicing techniques are each a source-to-source translation of the sequenced queries and PSM routines of the workloads into conventional queries and PSM routines that explicitly manage the timestamps in the data, to optimize the temporal queries in various ways. Thus, we created the τPSM-MaxMapped and τPSM-PerStmtMapped benchmarks for these techniques. We also used the translation discussed earlier to create the schemas and workloads for the two DB2 variants, so that the resulting queries can be executed in DB2 (we note that the SQL:2011 standard [56] and IBM DB2 10 very recently now support temporal tables; our transformations are on conventional, i.e., nontemporal, tables and queries).

Data. The data of τPSM are the result of translating (shredding) the temporal XML data of τXBench into six (temporal) relations (see the previous section). We defined three datasets of different temporal volatility: DS1, DS2, and DS3. DS1 contains weekly changes; thus, it contains 104 slices over 2 years, with each item having the same probability of being changed. Each time

step experiences a total of 240 changes; thus, there are 25K changes in all. DS2 contains the same number of slices but with rows in related tables associated with particular items changed more often (using a Gaussian distribution), to simulate hot-spot items. DS3 returns to the uniform model for the related tuples to be changed, but the changes are carried out on a daily basis, or 693 slices in all, each with 240 changes, or 25K changes in all (the number of slices was chosen to render the same number of total changes). These datasets come in different sizes: SMALL (e.g., DS1.SMALL is 12 MB in six tables), MEDIUM (34 MB), and LARGE (260 MB).

The data for the two τPSM-Mapped and two τPSM-Mapped-DB2 variants are equivalent to those of τPSM.

Schema. The τPSM schemas are derived from the schemas of PSM. We extend each table with a simple temporal statement modifier (i.e., ALTER TABLE ADD VALIDTIME ...) to make the table time varying. In addition, we extend each primary key to include the begin_time column and extend each assertion to be a sequenced assertion, thereby applying independently at each point in time.

The schema for τPSM-Mapped is mapped from the temporal domain to the nontemporal domain (because DBMSes do not yet support the statement modifiers of τPSM), by removing ALTER TABLE ADD VALIDTIME keywords from the table creation commands.

The schema for τPSM-Mapped-DB2 is mapped to meet the syntax requirements of DB2.

Workload. We define 85 queries for τPSM that correspond to the 17 queries in the PSM benchmark: 51 sequenced queries (a short period, medium period, and long period sequenced query for each of the original 17), 17 nonsequenced, and 17 current. The queries are derived directly from the PSM queries by adding simple temporal statement modifiers (i.e., VALIDTIME for sequenced queries, NONSEQUENCED VALIDTIME for nonsequenced queries, and no change for current queries) to each of the queries.

The workload for four τPSM-Mapped variants is the result of an automated transformation, using the previously mentioned maximally fragmented and per-statement slicing techniques [30,44], expressed in TXL, to create 170 queries in total.

5.5. The τXSchema benchmark

τXSchema (Temporal XML Schema) is a language and set of tools that enable the construction and validation of temporal XML documents [38,40–43]. τXSchema extends XML Schema [36,37] with the ability to define temporal element types. A temporal element type denotes that an element can vary over time, describes how to associate temporal elements across slices (or snapshots, which are individual versions of a document), and provides some temporal constraints that broadly characterize how a temporal element can change over time.

In τXSchema, any element type can be rendered as temporal by including in the schema a simple logical annotation stating whether an element or attribute varies over valid time or transaction time, whether its lifetime is described as a continuous state or a single event, whether the item itself may appear at certain times and not at others, and whether its content changes. So, a τXSchema document is just a conventional XML Schema document with a few annotations.

τXSchema provides a suite of tools to construct and validate temporal documents. A temporal document is validated by combining a conventional validating parser with a temporal constraint checker. To validate a temporal document, a temporal schema is first converted to a representational schema, which is a conventional XML Schema document that describes how the temporal information is represented. A conventional validating parser is then used to validate the temporal document against the representational schema. Then, the temporal constraint checker is used to validate any additional temporal constraints specified by the user in the temporal schema.

We have created the τXSchema benchmark to test the functionality and performance of τXSchema.

Schema. The schemas are the same as those in τXBench, augmented with seven additional manually specified constraints. These constraints provide a simple but comprehensive set for testing the

four types of XML constraints (referential integrity, data type, identity, and cardinality) with the various time modifiers (sequenced and nonsequenced). For example, one of the additional constraints specifies that 'over a period of a year, an item may have up to 6 authors', exercising a nonsequenced cardinality constraint.

Data. The data are the same as the data in τXBench.

Workload. Because τXSchema is a schema language and tool system, and because the main goal of testing such a system is to test its schema constraints, there is no workload in this benchmark (in effect, the schema *is* the workload and a microbenchmark at that).

5.6. Supporting tools

To facilitate the creation and validation of the individual benchmarks in τBench, we have developed a suite of tools that generate temporal data, validate the data in various ways, and validate the correctness of the queries. The tool suite, along with the rest of τBench, is available online [10].

The message of this section is that tools are some of the most important glue connecting the benchmarks together into a framework. Some of the tools map a component, for example, data or workload, in one benchmark into another, creating that component from hole cloth. Other tools perform consistency checks *within* a benchmark. Tools that perform cross-validation are especially critical, in that they ensure consistency between say the data components of two related benchmarks in the framework. Such validation leverage the mapping and within-benchmark consistency checking tools, because all the tools much work correctly for the end result to validate. These tools in concert thus tied the benchmarks together into a coherent whole, raising the quality of all of the individual benchmarks within the framework.

5.6.1. τGENERATOR. τGENERATOR is a simulation that creates XML temporal data from the nontemporal XML data (Figure 4). To create the temporal data, τGENERATOR consists of a user-specified number of *time steps* along with a user-specified number of *changes* to the data within each time step. At each time step, a subset of the elements from the original document is changed, creating a new *slice* of the temporal document. The simulation continues until all of the required changes are made.

Simulation description. The details of the simulation are described elsewhere [35]. Here, we give only a brief overview of the issues involved [43].

We designate the `<item>`, `<author>`, `<publisher>`, and `<related_items>` elements of the DC/SD XBench XML data to be *temporal elements*, that is, their values can change over time. The overall goal of the simulation is to randomly change, at specified time intervals, the content of these four temporal elements. The simulation consists of the following four steps.

1. *Setup.* The simulation reads the user-created parameter file to initialize the simulation parameters. The simulation then uses DOM [57] to load the user-specified XBench document (in this case, `catalog.xml`) into memory. The simulation then creates an internal representation of the temporal items (all instances of the four temporal elements), which at this point consist of only a single version.
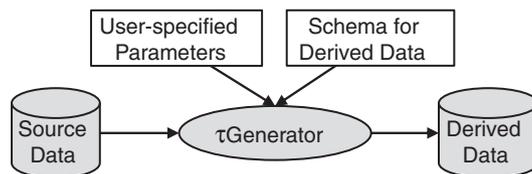

Figure 4. Overview of the τGENERATOR process.

2. *Initial snapshot selection.* The simulation selects a user-specified percentage of `<item>` temporal elements (and all of their subelements) for the initial slice. The simulation selects the elements uniformly at random from `catalog.xml`. Elements that are selected are called active; those that are not selected for the initial slice are called inactive and put into a pool for later use.

   The simulation sets the begin and end time of the first (and only, at this point) version of each selected temporal element to the user-specified begin date and the user-specified forever date, respectively.

3. *Time step loop.* The simulation increases the current time by the user-specified time step. The simulation then

   - Randomly selects a set of active elements to change, on the basis of user-specified parameters.
   - Changes the selected elements, as well as adds new elements and deletes existing elements, on the basis of user-specified parameters. New elements, and the new values for changed elements, come from the pool of inactive elements.
   - Outputs the current slice in XML format. A temporal document, that is, all slices merged together, is also maintained.
   - Checks the stop criteria of the simulation. If the simulation has exceeded the user-specified number of necessary changes, the simulation exits the time step loop and continues to Step 4. Otherwise, it returns to Step 3.

4. *Output.* The simulation outputs the final temporal document.

During the aforementioned simulation process, it is important to maintain the original constraints defined and implied in the DC/SD XBench XML schema. This will ensure that the generated data validate to the sequenced semantics of the original schema. In addition, we add additional checks in the simulation to maintain the set of nonsequenced constraints in the τXSchema benchmark.

We need not consider the original data type constraints, because the simulation does not change the type of data at any point. Similarly, we need not consider identity constraints (which ensure, e.g., that the item's id attribute is globally unique) because the simulation does not generate any ids of its own—all items keep their originally assigned ids. Thus, as long as the original data is correct, this portion of the generated data will also be correct.

However, we do need to consider the referential integrity constraint between item ids and related items. Because only a subset of items is active at any given time, there is no guarantee that related items will point to active item ids. Thus, the simulation must provide this guarantee explicitly. The simulation satisfies this by adding a check at the end of each time step. The check builds a list of all active item ids. Then, for any currently active related items that refer to an item id not in the list, the simulation replaces the value with an id in the list. Note that this ensures both sequenced and nonsequenced referential integrity constraints between item ids and related items.

We also need to consider the cardinality constraints placed on authors. Because the simulation has the ability to add new authors and delete existing authors at any point in time, the simulation must be careful not to violate the maxOccurs and minOccurs constraints in the original schema. This is true for both the sequenced (i.e., at any given time, an item must have between one and four authors) and nonsequenced (i.e., in any given year, an item can have up to six authors) variants of this constraint. The simulation achieves these by maintaining an author counter for each item and each period. For the sequenced constraint, the counter is incremented and decremented as the simulation adds or removes authors. For the nonsequenced, the counter is only incremented when an author is inserted. Finally, when an item is selected by the simulation to insert or remove an author, the simulation first consults this counter; if the counter indicates that such an action will violate a constraint, the simulation chooses another item to change. The simulation will continue trying to choose a viable author for a predefined number of iterations and then will halt to avoid infinite loops.

In this way, the resulting temporal data are made consistent for both sequenced and nonsequenced constraints.

5.6.2. Shredding XML data into relations. As mentioned earlier, we have created a tool to shred the XBench DC/SD data into six relations. If the data to be shredded is temporal, the resulting relations will have two additional columns representing the begin and end time for each row.

5.6.3. Validation tools. We have developed several tools to help us validate our translations and simulations.

Validating the output of τGENERATOR. We have created a tool that validates the XML slices generated by τGENERATOR against their XML schema using XMLLINT [58]. In addition, the tool validates the generated temporal document against its temporal schema using τXMLLINT [38]. These checks ensure that each generated slice is well formed and consistent with the schema constraints and that the generated temporal document is well formed and consistent with the sequenced and nonsequenced semantics of the conventional schema.

Corrupting data. Constraints should be tested on both valid and invalid data. We thus corrupt the data generated by τGENERATOR as a testing phase to determine if the validators and associated schemas can detect the corruptions. We have developed a tool, called τCORRUPTOR, that takes as input a temporal XML document and outputs another temporal XML document that is identical as the input document except for one *victim* element. The victim element is modified to directly contradict one of the constraints specified in the original schema.

Specifically, the input into τCORRUPTOR is a temporal XML document generated by τGENERATOR, as well as a integer parameter specifying which of the seven temporal constraints to invalidate.

1. Sequenced cardinality. Duplicates an entire author four times (but keeps each author `id` unique).
2. Sequenced referential integrity. Changes a related item's `item_id` attribute to a random string.
3. Sequenced identity. Copies the `ISBN` of one item to another.
4. Sequenced datatype. Changes the number of pages to a random string.
5. Nonsequenced identity. Copies the `id` of an item i_1 at time t_1 to the `id` of an item i_2 at time t_2, where $i_1 \neq i_2$ and $t_1 \neq t_2$.
6. Nonsequenced referential integrity. Changes a related item's `item_id` to a random string.
7. Nonsequenced cardinality. Duplicates an entire `<author_RepItem>` six times (but keeps each author `id` unique).

Checking primary keys of shredded relations. We have created a tool that checks each shredded relation to ensure that each primary key is unique across the relation. It does so by building a hash table of primary keys in each relation. If a row is encountered whose key is already defined in the hash table, then the tool reports an error. Relations whose primary key are defined across multiple columns require multidimensional hash tables.

Checking referential integrity of shredded relations. We have created a tool that checks each shredded relation to ensure that each referential integrity constraint is valid across the relation. It does so by first reading the entire relation and creating a hash table of the referenced primary keys. The second pass ensures that each value in a foreign key column is defined in the hash table.

Comparing XML slices to the temporal XML document. We have created a tool that ensures that the generated temporal XML document and the corresponding XML slices are consistent. The last tool makes use of UNSQUASH, which is part of the τXSchema tool suite [38]. UNSQUASH extracts each XML slice from a given temporal XML document. Our tool then compares each generated slice with each extracted slice using X-Diff [59] to ensure that they are equivalent. If all the slices are equivalent, then our tool reports that the temporal XML document and generated slices are equivalent.

Testing queries. Query Executor is a tool providing a simple way to quickly execute multiple queries for a variety of DBMSes. After the DBMS and password are entered, the GUI will display a list of active query directories and allow for additional directories to be added or for old directories to be removed. Each query directory contains a parameter file (`.queryexecutor`) that specifies a connection to the appropriate database. Queries within the selected directory appear in a separate pane, and the tool allows a query to be run individually, all of the queries within the directory to be run, or a selected group of the queries to be executed. The results of each query are stored in files, and Query Executor can run a diff on the last two executions. If one or more of the previous two runs were of more than a single query, the diff will be performed only on the queries from both runs that have the same query number. Query Executor also allows for the results of any run or diff to be exported to a file. This tool has proven very convenient for developing the queries within each benchmark.

Comparing XQuery and PSM queries. We have created a tool that performs a pair of queries and compares their output. First, it runs an XQuery query on the nontemporal XML data and saves the output. Then, it runs the corresponding PSM query on the shredded relational data and saves the output. Because the output formats are different, the tool reformats the PSM output to be of the same format as the XQuery output. If the two results are textually identical, then the queries achieved the same functionality.

This tool is part of our larger commutativity validation process. *Commutativity validation* compares, for example, the output of performing a temporal query on a temporal database with the output of performing a conventional query on a slice of the temporal database (Figure 5). Such a comparison simultaneously validates several tasks and tools: slicing the temporal database; executing temporal queries; slicing temporal relations; and executing conventional queries. Only if all of these steps are working perfectly will the results be the same; thus, a strong level of validation is ensured. This cross-validation relies on snapshot equivalence and temporal upward compatibility [8], as well as the sequenced semantics [7, 60], which relate in very simple ways nontemporal and temporal data, queries, and constraints.

5.6.4. Example: creating the τXBench benchmark. To illustrate the procedure (and power) of benchmark frameworks, we now present a detailed example from τBench. We show how we created the τXBench benchmark from the DC/SD XBench benchmark.

First, we execute τGENERATOR to create the temporal data. We specify an arbitrary starting time of January 2010. We then specify that we want a new slice every week, for a total of 104 weeks (2 years). At each new slice, 240 elements are added, changed, or deleted; thus, there are 25K changes in all. τGENERATOR performs the simulation and outputs the final temporal XML document, `temporal.xml` (34 MB).

τGENERATOR also creates a temporal schema, `temporalSchema.xml`, which simply points to the conventional XML schema and to the temporal document. It also contains a couple of simple
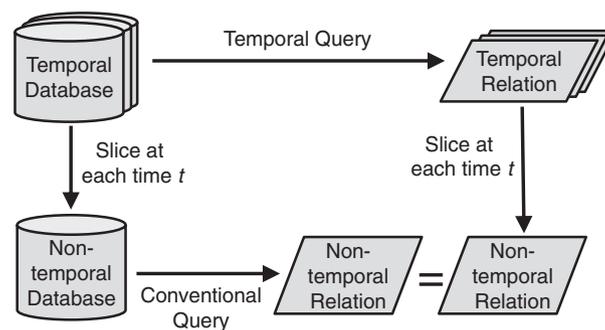

Figure 5. Commutativity validation.

temporal annotations, specifying, for example, that the `<item>`, `<author>`, `<publisher>`, and `<related_items>` elements are temporal elements; that each temporal element has a unique `id`; and that each `<related_item>` points to an existing (and active) `<item>` `id`.

Then, we use our tools to validate the generated output.

Next, we create temporal queries from DC/SD XBench's nontemporal queries simply by adding statement modifiers (i.e., current and sequenced); adding no modifier implies nonsequenced semantics. We validate the results of the current queries by comparing their output with the output of the original nontemporal queries executed on a slice of the `temporalSchema.xml` (using UNSQUASH) at time *now*. We validate the output of the sequenced queries by using the commutativity tool. We do not validate the nonsequenced queries, as they are simply returning the results of the temporal simulation.

After all of these steps, we can be confident that the data are correct, the schema is correct, and the workload is correct. In addition, we know (and can easily describe) the exact origins of the data, schema, and workload.

5.6.5. Running the benchmarks. We have developed a tool, QUERYRUNNER, that presents a graphical interface to run a specified set of queries over a specified data set, collecting the relevant metrics. This tool is applicable to the benchmarks utilizing the relational model.

5.7. Summary and lessons learned

In creating τBench, several tools and concepts were realized. The original XBench benchmark was branched in two different directions: creating a temporal version of XBench (τXBench) and creating a PSM version of XBench (PSM). Each of these was further branched for additional purposes. To do so, we created tools to vary a nontemporal dataset over time, to change hierarchical data into relational data, to validate the data changes, and many other tasks. The resultant benchmarks, their relationships, and the tools used to create them, are all a part of τBench.

As we built τBench, we were surprised by the number of subtle bugs that we found in the temporal data generator (τGENERATOR), which we only discovered by validating and executing the derived benchmarks. With each consistency and validation check that we added to the benchmarks, more bugs were uncovered. In the end, we developed confidence that τGENERATOR, and hence our datasets, was correct and consistent. This process highlights the strength of the component connections, and their ability to be cross checked, of benchmark frameworks.

We found that each additional benchmark required less time to create than the previous. Instead of having to create ten benchmarks from scratch, we created a single benchmark and slightly modified it ten times, using tools to automate as much as possible. In practice, we use each benchmark in τBench for some research project, and we have already reaped the benefits of reuse in terms of reduced development cost. We feel that benchmark frameworks are well worth any initial startup costs.

In addition, we have found that describing the benchmarks has been simplified, because of the added organizational clarity provided by regularized benchmark components and the B-R diagrams.

Finally, we found that because each benchmark component tends to be reused by derived benchmarks, it becomes easier to justify the extra effort of ensuring each component is of the highest quality.

Our experience with developing τBench has been that each aspect of a benchmark framework, including explicating shared and changed and new components in a new benchmark, exploiting the opportunity for cross-validation, and sharing data and tools across constituent benchmarks, brought with it benefits that fully justified the small amount of effort and rigor required. And sometimes, the sharing of components made developing a descendant benchmark substantially easier than developing one from whole cloth.

To further demonstrate the generality and usefulness of benchmark frameworks, we now briefly present two additional case studies from different domains: the MUD framework and the TPC framework, examining how the general principles listed in Section 4.1 apply in these domains.
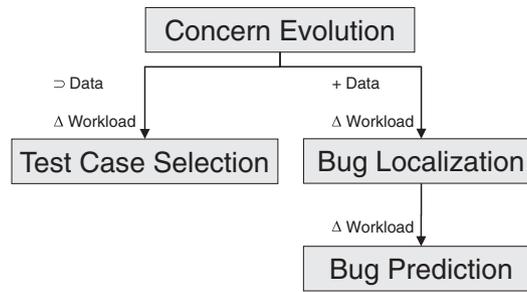
Figure 6. Benchmark–Relationship diagram of the MUD benchmark framework.

6. THE MUD BENCHMARK FRAMEWORK

Mining Unstructured Data is a project from the empirical software engineering domain that uses information retrieval (IR) models to mine the unstructured data found in software repositories such as developer emails, bug report titles, source code identifiers and comments, and requirement documents [61]. An ongoing research project by the first author, the MUD project, considers different software engineering tasks, such as test case selection and bug localization, with benchmarks for each. We have created a benchmark framework to describe the MUD project, which we now present.

The organizational insight of the MUD benchmark framework is sharing (i) the data model, (ii) DDL, (iii) data, (iv) data preprocessing techniques, and (v) validation tools across the constituent benchmarks, to enable a common approach for evaluating different algorithm variations.

Figure 6 shows the B-R diagram for the MUD benchmark framework, which consists of four constituent benchmarks.

6.1. The Concern Evolution *benchmark*

The foundation of the framework is the Concern Evolution benchmark, initially created for measuring how the *concerns* (i.e., technical aspects, such as 'printing', 'database connections', 'logging', and 'XML I/O') in a software system [62] are evolving over time [63]. Here, the data component is the source code of several real-world software projects (Eclipse JDT, Mozilla, Chrome, and AspectJ). The source code is not executed by this benchmark but instead is treated as simple text and input into different text mining algorithms. Specifically, the identifier names, comments, and string literals from the source code are used, as these elements have been shown to represent the developer's thoughts and intentions of the code [64], making them appropriate for automated concern mining algorithms. These code elements can be subjected to different preprocessing steps, such as splitting identifier names (camelCase and under_scores) and removing stop words. Hence, the dataset does not contain a set of executable commands but rather a set of words that happen to be extracted from source code. The workload consists of the concern evolution mining algorithm variations being researched, as algorithms can vary in which preprocessing steps are performed and in which concern mining technique is used, such as latent Dirichlet allocation (LDA) [65] or DiffLDA [66]. The data model and DDL are both the Java programming language, as the data are in the form of Java source code. The schema contains informally expressed constraints, such as 'identifiers have no embedded spaces'. The DML is the R programming environment, because that is how the algorithm variations are implemented.

6.2. The Test Case Collection *benchmark*

The Test Case Selection benchmark is useful for evaluating IR-based test case selection algorithms, which aim to choose a subset of all available test cases while still detecting as many faults in the software as possible [67]. IR-based test case selection algorithms take as input a set of test cases (i.e., the source code of the test cases) and output a subset of test cases (usually much smaller than the total set) that, when executed, will still likely detect a majority of the bugs in the system. The

selection algorithms work by maximizing the distance (in this case, using textual difference metrics, such as the Levenshtein distance, between the source code of the test cases) between the subset of selected test cases, thereby ensuring that all parts of the underlying software system are tested. Here, the workload consists of the selection algorithm variations (i.e., different preprocessing steps, different distance metrics, and different maximization algorithms), and the data are only a subset of the data from the Concern Evolution benchmark, namely the source code files that correspond to test cases. All other components remain the same.

6.3. The Bug Localization benchmark

The IR-based bug localization algorithms [68] aim to identify which source code files likely contain faults (i.e., bugs), given a bug report from a bug repository (e.g., Bugzilla [69]). These algorithms work by building an index on the (preprocessed) source code and then query this index using the content of the bug reports. Again, the workload is the set of bug localization algorithm variations (i.e., different IR models, different similarity measures, and different ranking algorithms). The data are expanded from the Concern Evolution data to include bug report data, such as the bug report title and detailed description entered by user; all other components remain the same.

6.4. The Bug Prediction benchmark

Finally, the Bug Prediction benchmark is used to evaluate bug prediction algorithms, whose aim is to predict the location of bugs in the source code, this time without the help from a bug report [70]. Bug prediction algorithms work by measuring metrics on each source code module (e.g., complexity, coupling/cohesion with other modules, number of changes, and number of previous bugs) and train statistical models to predict the likelihood that each module contains a bug. Here, the workload consists of bug prediction algorithm variations (i.e., which metrics to measure and what statistical models to use), but all other components remain the same.

6.5. Supporting tools

In the creation, validation, and execution of the MUD benchmark framework, the following tools are shared (reused).

Data generation. The data for the original Concern Evolution benchmark is composed of the source code files of several real-world systems, totaling over 70 million lines of code. Specifically, the data contain the source code of Eclipse (the 16 official releases from 2002–2009), IBM Jazz (4 releases from 2007–2008), and Mozilla (10 release from 2004–2006). Typically, when using IR or other text mining algorithms on source code, as MUD does, the source code is first preprocessed to capture the domain-specific information encoded in the identifier names, comments, and string literals in the code. The data generation tool for MUD thus preprocesses each source code file to isolate these elements, removing language-specific syntax, keywords, and so on. Additionally, the data generation tool splits identifier names, removes common English language stop words, and stems each word in an effort to normalize the vocabulary (Figure 7).

Validation tools. Many validation tools are used to check the consistency of the data and data generation tools, as well as the results of the different workloads. For example, the lightweight source
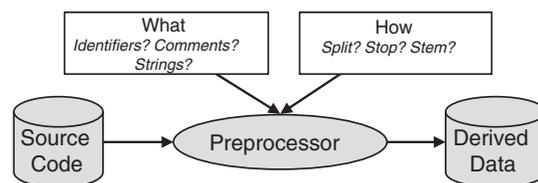

Figure 7. MUD data generation.

code preprocessor (LSCP) tool has internal validation capabilities to ensure that the data preprocessing is performed correctly, by checking the associativity of the tool (e.g., the tool is run three times: once to isolate identifiers only; once to isolate comments only; and once to isolate comments and identifiers. LSCP then compares the union of the first two runs with the third). Similar checks can be performed for other tool parameters.

Additionally, many of the benchmarks train IR models on the data, which have their own class of validity checks. For example, the LDA [65] model assigns topics to each document; we have a validation script in R to check that the sum of each document's assigned topics equals one, by the definition of the LDA model. Dozens of similar checks are performed to ensure the correctness of the data.

6.6. Summary and lessons learned

The MUD benchmark framework is simpler than τBench, but many of the ideas are the same. Starting with a foundation benchmark (i.e., Concern Evolution), the benchmark components are changed to create new benchmarks. In the case of MUD, the data and workload components were changed to create the Test Case Selection and Bug Localization benchmarks; the Bug Localization benchmark was further changed to create the Bug Prediction benchmark. Tools were created to automate the generation (i.e., via preprocessing) of the data and to validate the correctness of several IR models. The resultant benchmarks, their relationships, and the data creation and validation tools are all a part of MUD.

During the creation of MUD, we learned a few important lessons. The first was the power of regularizing the dataset in the root benchmark, so that it can easily be reused in subsequent benchmarks. For example, when starting the bug localization project, instead of having to hunt for appropriate data and create a new benchmark from scratch, we simply reused the data from the Concern Evolution benchmark, added the bug reports, and implemented the algorithm variations.

As was the case with τBench, we also found in MUD that the consistency and validation checks helped to find subtle bugs and corner cases in the data generator. For example, a step in one of the text mining algorithms in the workload of the Bug Localization benchmark is supposed to output a vector that sums to one for every input document in the data. However, because of an inconsistent parameter setting in the algorithm, the algorithm produced vectors containing all zeroes for some input documents. Because we had created a data validation tool to check the vectors for this, we were able to easily uncover and correct this error.

Finally, we found that by describing the research projects as benchmarks, which consist of course of data and the workloads applied to the data, the resulting organizational clarity enable us to more easily implement and test new algorithm variants. For example, with the Bug Localization benchmark in place, a researcher can test a new suite of bug localization algorithms by simply deriving a new benchmark and replacing the workload, leaving the other components the same.

7. TPC: A REPRISE

As mentioned in the introduction, the TPC family of benchmarks, while loosely related, presents complexity for end users in exercising these benchmarks in a unified fashion, because of the lack of explicit relationships. We present a *proposed* benchmark framework for TPC to illustrate the advantages of benchmark frameworks. We propose that the organizational insight is structuring the constituent benchmarks around a single relational database that is both realistic and appropriate for both transactional and analytical applications.

Figure 8 shows our TPC benchmark framework, composed of four individual benchmarks, along with two separate benchmarks derived from them.

The TPC-C benchmark contains a large dataset for a wholesale supplier, consisting of warehouses, sales districts, and customer transactions. The OLTP workload consists of mostly simple queries as well as small modifications such as INSERT and UPDATE operations.
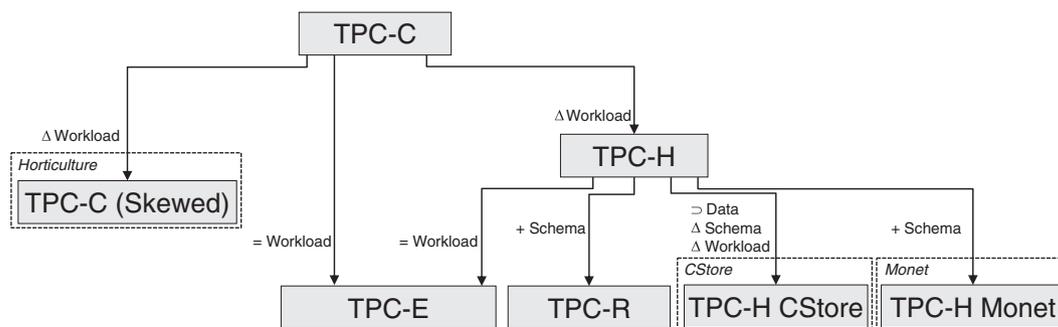
Figure 8. Benchmark–Relationship diagram of a proposed TPC framework.

The TPC-H benchmark utilizes a very similar, but not identical, set of tables. This benchmark could in fact have shared the same data and schema as TPC-C. But because TPC-H emphasizes analytical applications, it uses an entirely new OLAP workload, consisting of sophisticated analytical queries that extensively involve joins, predicates, and aggregations. For that reason, in our proposed framework, we show the only difference between the two benchmarks being a changed workload.

The TPC-E benchmark was designed to evaluate DBMSes with a more generalized workload, that of complicated analytical queries with transaction processing. As a family, this benchmark differs in small but nonessential ways from both of the prior benchmarks. In our proposed framework, TPC-E is simply a merger of the TPC-C and TPC-H benchmarks, containing the same data and schema as both, with a union of their queries.

The TPC-R benchmark is the same as TPC-H, except that it allows for optimizations to be performed before executing the queries, such as creating indices on certain columns to enhance average query performance. Thus, all benchmark components are the same, except the schema (which now contains CREATE INDEX statements).

Even though this benchmark framework is simple, its understandability has been greatly increased because of greater organizational clarity (a B-R diagram of the actual TPC family of benchmarks would be quite a bit more complex, because of superfluous changes of components between benchmarks). For example, an end user may not have been previously known if the TPC-C and TPC-R benchmarks could have used the exact same data; now, it is clear that they do. And tools, such as those for generating sample data at various scale factors and for executing the workload against the database and collecting the metrics, could be reused across the framework. Finally, the workload, data, and DBMS query processing in the TPC-H, TPC-E, and TPC-R benchmarks can be cross-validated by ensuring that the queries return the exact same results in these three benchmarks.

Moreover, we found that in practice, such benchmarks are often modified for evaluating certain features of newly developed DBMSes. For instance, Stonebraker et al. investigated a column-oriented DBMS implementation, namely the *CStore* [71]. They evaluated this system with a modified version of the TPC-H benchmark in which five tables were removed and the schema of the remaining three tables was modified. In addition, they modified one original query from the benchmark and added six new queries. Another example is the evaluation of the *MonetDB/X100* system performed by Boncz et al., in which the schema of the TPC-H benchmark is modified by creating additional indices on specific columns and sorting particular tables (both of which are additional physical schema additions) before queries are run [72]. The third example is a version of TPC-C used to evaluate the automated partitioning facilities in the *Horticulture* system, in which they 'generated a temporally skewed load for TPC-C, where the WAREHOUSE_id used in the transactions' input parameters is chosen...[to present] a significant amount of skew' [73].

If the TPC benchmarks employed the benchmark framework approach, these specific modifications performed by the previously discussed research could be more explicitly depicted in the framework's B-R diagram. Moreover, the change(s) applied on each path (e.g., data or schema) of deriving a new benchmark from an existing one can be more clearly communicated. The TPC framework B-R diagram emphasizes that the TPC-H CStore benchmark is in fact changed in both

schema and workload, making it clear that the experimental results depend in part on these changes and thus may not be applicable to the original TPC-H benchmark. This diagram also makes explicit that the TPC-C (Skewed) benchmark has a skewed *workload*.

8. ADVANTAGES AND DISADVANTAGES OF BENCHMARK FRAMEWORKS

Benchmark frameworks are not meant to suite every purpose and every benchmark. When commercial politics mandate a specific benchmark to be used time and time again, without modifications to the benchmark, benchmark frameworks will offer no advantages.

The initial startup cost of benchmark frameworks (i.e., clearly defining benchmark components and their relationships and creating shared tools and consistency checks) may be high for small families of simple benchmarks or when the benchmarks share little-to-no components (in which case, it is probably best to merge such benchmarks into neither a family nor a framework).

At the same time, the organization and mechanics of benchmark frameworks bring many benefits to both the developers and the users of the individual benchmarks. First, benchmark frameworks encourage the reuse of existing benchmarks and their components, without the need of creating new components from scratch. Reuse is especially helpful with large and complex datasets, schemas, and workloads.

Reuse also brings forth another advantage, that of reducing the development cost of subsequent benchmarks. Components can be reused and iteratively modified to suite the new benchmark's purposes. Researchers and practitioners who may have otherwise not created a new benchmark now enjoy a lower development cost. High quality need not be traded off against greater development effort: the former can be increased at the same time that the latter is decreased through the general principles inherent in benchmark frameworks.

The structure and organization of benchmark frameworks, embodied by the B-R diagrams and explicit connections between components, are easier to understand than traditional benchmark families. When large, distributed groups collaborate to develop related benchmarks, and the organization provided by benchmark frameworks becomes especially important and useful.

As illustrated in Section 7, the B-R diagram emphasizes what is borrowed and what is changed when a benchmark is not used wholesale but rather piecemeal in the evaluation of a software or hardware system.

Finally, and perhaps most importantly, the structuring of the framework and the sharing and cross-validation that this structure encourages result in constituent benchmarks whose correctness is more demonstrable, yielding greater confidence on the part of the developers and users of the correctness of those constituent benchmarks. For benchmarks that are to be actually used for evaluation, the relatively small amount of effort required to create a benchmark framework is generally well worth this added confidence. As evidence, we have already had interest from benchmarking practitioners involved with the TPC family, because of the benefits they see in the framework approach.

9. CONCLUSIONS AND FUTURE WORK

We have introduced the notion of a benchmark framework, which is an ecosystem of related benchmarks having well-defined and well-specified relationships, as well as tools to help transform and validate individual benchmarks, based on organizational principles. These synergistic properties of benchmark frameworks help overcome the difficulties of maintaining, describing, and validating families of related benchmarks. Benchmark frameworks bring *clear communication* for describing the changes made to existing benchmark components, by using B-R diagrams. Benchmark frameworks also encourage *sharing* and *reusing*, allowing an existing benchmark to be quickly adapted to test new systems or strategies, by using automated tools to translate existing components into new ones. Benchmark frameworks encourage *regularization of the execution* of individual benchmarks, by using shared tools and environments. Finally, benchmark frameworks help ensure *validity*, by consistently evaluating the validity of benchmark components from various angles (e.g., commutativity).

We have also introduced τBench, a benchmark framework for evaluating temporal schema and query languages on XML and relational data [35]. Although τBench spans many technologies (from XML Schema to Temporal SQL/PSM) and is used by several distinct projects, the benchmarks are easy to understand, and creating new, related benchmarks is straightforward and reliable; this is one of the powers of benchmark frameworks. An early implementation of τBench is available on our project website [10]. Our experience with this framework is that it was *easier* to develop than a family of only loosely related benchmarks, because of extensive sharing of components and tools between constituent benchmarks.

To emphasize the generality of benchmark frameworks, we also briefly presented two additional benchmark frameworks. Although these were from different domains (those of empirical software engineering and OLTP/OLAP), the advantages of benchmark frameworks were still available: relationships between benchmarks were explicit, and tools helped in the creation and validation of each individual benchmark.

In future work, we would like to augment the relationship constraints in Table II. We have developed an XML schema for B-R diagrams; it would be useful to have a meta-tool to test such constraints. The meta-tool can check constraints by reading the data, schemas, and workloads of two benchmarks and ensuring that they comply with the stated relationships in the B-R diagram. For example, if the diagram shows that the data component of one benchmark is equivalent to the data component of its parent benchmark, then the meta-tool can ensure this equality.

In this paper, we have discussed the generality of benchmark frameworks and the methods to create and validate them. In future work, as we and others develop more benchmark frameworks from other domains, we wish to determine whether the tools could themselves also benefit from general principals. For example, can we find general design choices and techniques for data generation and validation tools, no matter the underlying domain of the benchmark framework? Or are these tools only contributory to such general principals?

The τBench benchmark framework can be expanded in several directions: automating some of the mappings (such as for the τXBench workload), adding new DBMSes for PSM (that should be easier with automated mappings), adding new mapping strategies for τPSM (also aided by automated mappings), adding new schema languages for XML, and adding new implementation strategies for τXSchema. In particular, several DBMSes now offer temporal support. Oracle 10g added support for valid time tables and transaction time tables, and Oracle 11g enhanced support for valid time queries [74]. Teradata recently announced support in Teradata Database 13.10 of most of these facilities as well [75], as did IBM for DB2 10 for z/OS [76].

Designing an individual benchmark as well as a realized benchmark framework is the result of a long sequence of design decisions. For example, in designing the PSM benchmark, we adopted a particular way of shredding the XML document into relations. It would be useful to have a way of capturing such interbenchmark transformations in a more complete manner than just the five relationships listed in Section 4.2, for example, with annotations that name the software that performs that shredding, or better yet, with a formal notation that gives the semantics of that transformation. Doing so would richen the kinds of relationship constraints (e.g., Table II) that could be stated and verified.

Finally, we encourage the XBench, SPEC, and TPC organizations to consider refining their respective benchmark families into more structured and more convenient benchmark frameworks.

REFERENCES

1. Henning J. SPEC CPU2000: measuring CPU performance in the new millennium. *Computer* 2000; **33**(7):28–35.
2. SPEC. SPEC Organizational Information, 2011. Available from: http://www.spec.org/spec, [last accessed 12 March 2011].

3. Gray J (ed.). *The Benchmark Handbook for Database and Transaction Systems*, (2nd edn). Morgan Kaufmann, 1993.
4. Fayad M, Schmidt D. Object-oriented application frameworks. *Communications of the ACM* 1997; **40**(10):32–38.
5. Johnson R. Frameworks = (components + patterns). *Communications of the ACM* 1997; **40**(10):39–42.
6. TPC. TPC Transaction Processing Performance Council. Available from: http://www.tpc.org/tpcc/ [last accessed 14 July 2011].
7. Böhlen MH, Jensen CS, Snodgrass RT. Temporal statement modifiers. *ACM Transactions on Database Systems* 2000; **25**(4):407–456.
8. Bair J, Böhlen MH, Jensen CS, Snodgrass RT. Notions of upward compatibility of temporal query languages. *Business Informatics (Wirtschafts Informatik)* 1997; **39**(1):25–34.
9. Free Software Foundation. GNU General Public License, 2010. Available from: http://www.gnu.org/licenses [last accessed 2 October2010].
10. Thomas SW, 2011. Available from: http://www.cs.arizona.edu/projects/tau/tBench [last accessed 25 February 2013].
11. Walrath K, Campione M, Huml A, Zakhour S. *The JFC Swing Tutorial: A Guide to Constructing GUIs*. Addison Wesley Longman Publishing Co., Inc: Redwood City, CA, USA, 2004.
12. Austern M. *Generic Programming and the STL: Using and Extending the C++ Standard Template Library*. Addison-Wesley Longman Publishing Co., Inc., 1998.
13. Sturm M, Bertsch A, Gröpel C, Hildebrandt A, Hussong R, Lange E, Pfeifer N, Schulz-Trieglaff O, Zerck A, Reinert K, Kohlbacher O. OpenMS: An open-source software framework for mass spectrometry. *BMC Bioinformatics* 2008; **9**.
14. Grosse-Kunstleve RW, Sauter NK, Moriarty NW, Adams PD. The computational crystallography toolbox: crystallographic algorithms in a reusable software framework. *Journal of Applied Crystallography* 2002; **35**:126–136.
15. Schneider SA, Ullman MA, Chen VW. Controlshell: a real-time software framework. In *Proceedings of the International Conference on Systems Engineering*, 1991; 129–134.
16. Apache. Log4J: logging services. Available from: http://logging.apache.org/log4j/1.2/ [last accessed 25 November 2011].
17. Naftalin M, Wadler P. *Java Generics and Collections*. O'Reilly Media, Inc., 2006.
18. Gamma E, Helm R, Johnson R, Vlissides J. *Design Patterns: Elements of Reusable Object-oriented Software*, Vol. 206. Addison-wesley Reading: MA, 1995.
19. Li A, Yang X, Kandula S, Zhang M. CloudCmp: comparing public cloud providers. In *Proceedings of the 10th Conference on Internet Measurement*, Melbourne, Australia, 2010; 1–14.
20. INEX. Initiative for the evaluation of XML retrieval. Available from: https://inex.mmci.uni-saarland.de/ [last accessed 14 July 2011].
21. Microsoft Corporation. WSTest web services benchmark. Available from: http://msdn.microsoft.com/en-us/netframework/cc302396 [last accessed November 25, 2011].
22. Li Y, Bressan S, Dobbie G, Lacroix Z, Lee M, Nambiar U, Wadhwa B. XOO7: applying OO7 benchmark to XML query processing tool. In *Proceedings of the 10th International Conference on Information and Knowledge Management*, Atlanta, Georgia, 2001; 167–174.
23. Schmidt A, Waas F, Kersten M, Carey M, Manolescu I, Busse R. XMark: a benchmark for XML data management. In *Proceedings of the 28th International Conference on Very Large Data Bases*, Hong Kong, China, 2002; 974–985.
24. Fischer P. XQBench, 2011. Available from: http://xqbench.org/ [last accessed 25 November 2011].
25. Barbosa D, Mendelzon A, Keenleyside J, Lyons K. ToXgene: a template-based data generator for XML. In *Proceedings of the International Conference on Management of Data*, Madison, Wisconsin, USA, 2002; 616.
26. Shanmugasundaram J, Shekita E, Kiernan J, Krishnamurthy R, Viglas E, Naughton J, Tatarinov I. A general technique for querying XML documents using a relational database system. *SIGMOD Record* 2001; **30**(3):20–26.
27. Cordy J, Halpern-Hamu C, Promislow E. TXL: a rapid prototyping system for programming language dialects. *Computer Languages* 1991; **16**(1):97–107.
28. Gao D, Snodgrass RT. Temporal slicing in the evaluation of XML queries. In *VLDB '2003: Proceedings of the 29th International Conference on Very Large Data Bases*, Berlin, Germany, 2003; 632–643.
29. W3C. XQuery 1.0: An XML Query Language, W3C Recommendation, 2007. Available from: http://www.w3.org/TR/xquery [last accessed 5 November 2011].
30. Snodgrass RT, Gao D, Zhang R, Thomas SW. Temporal support for persistent stored modules. In *Proceedings of the 28th International Conference on Data Engineering*, Washington DC, USA, April 2012; 114–125.
31. Yao B, Özsu M, Khandelwal N. XBench benchmark and performance testing of XML DBMSs. In *Proceedings of the 20th International Conference on Data Engineering*, Boston, MA, USA, 2004; 621–632.
32. SPEC. Standard Performance Evaluation Corporation. Available from: http://www.spec.org/benchmarks.html [last accessed 14 July 2011].
33. TREC. Text Retrieval Conference. Available from: http://trec.nist.gov/ [last accessed 14 July 2011].
34. Van der Vlist E. *Relax Ng*. O'Reilly Media, 2004.
35. Thomas SW, Snodgrass RT, Zhang R. τBench: extending XBench with time. *Technical Report TR-93*, TimeCenter, 2010.
36. Mendelsohn N. XML Schema part 1: structures second edition, 2004. Available from: http://www.w3.org/TR/2001/REC-xmlschema-1/ [last accessed 25 November 2011].
37. Mendelsohn N. XML Schema part 2: datatypes second edition, 2004. Available from: http://www.w3.org/TR/2001/REC-xmlschema-2 [last accessed 25 November 2011].

38. Currim F, Currim S, Dyreson C, Joshi S, Snodgrass RT, Thomas SW, Roeder E. τXSchema: support for data- and schema-versioned XML documents. *Technical Report TR-91*, TimeCenter, 2009.

39. Currim F, Currim S, Dyreson C, Snodgrass RT, Thomas SW, Zhang R. Adding temporal constraints to XML schema. *IEEE Transactions on Knowledge and Data Engineering* 2012; **24**:1361–1377.

40. Currim F, Currim S, Dyreson CE, Snodgrass RT. A tale of two schemas: creating a temporal XML schema from a snapshot schema with τXSchema. In *Proceedings of the 9th International Conference on Extending Database Technology*, Heraklion, Crete, Greece, 2004; 348–365.

41. Joshi S. τXSchema - support for data- and schema-versioned XML documents. *Master's Thesis*, Computer Science Department, University of Arizona, Tucson, AZ, 2007.

42. Snodgrass RT, Dyreson C, Currim F, Currim S, Joshi S. Validating quicksand: temporal schema versioning in τXSchema. *Data Knowledge and Engineering* 2008; **65**(2):223–242.

43. Thomas SW. The implementation and evaluation of temporal representations in XML. *Master's Thesis*, Computer Science Department, University of Arizona, Tucson, AZ, 2009.

44. Gao D. Supporting the procedural component of query languages over time-varying data. *Ph.D. Thesis*, University of Arizona, Tucson, AZ, 2009.

45. Gao D, Snodgrass RT. Syntax, semantics, and evaluation in the τxquery temporal XML query language. *Technical Report TR-72*, TimeCenter, 2003.

46. Date C, Darwen H. *A Guide to the SQL Standard*. Addison-Wesley New York, 1987.

47. Melton J. *Understanding SQL's Stored Procedures: A Complete Guide to SQL/PSM*. Morgan Kaufmann Publishers, 1998.

48. Runapongsa K, Patel JM, Jagadish HV, Chen Y, Al-Khalifa S. The Michigan benchmark: towards XML query performance diagnostics. *Information Systems* 2006; **31**:73–97.

49. Afanasiev L, Manolescu I, Michiels P. Member: A micro-benchmark repository for xquery. In *Database and XML Technologies*, Vol. 3671, Bressan S, Ceri S, Hunt E, Ives Z, Bellahsne Z, Rys M, Unland R (eds), Lecture Notes in Computer Science. Springer Berlin/Heidelberg, 2005; 578–578. 10.1007/11547273_11.

50. Nicola M, Kogan I, Schiefer B. An xml transaction processing benchmark. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, SIGMOD '07. ACM: New York, NY, USA, 2007; 937–948.

51. Böhme T, Rahm E. Xmach-1: a benchmark for xml data management. In *Proceedings of German Database Conference BTW2001*, March 2001; 264–273.

52. Schmidt A, Waas F, Kersten ML, Carey MJ, Manolescu I, Busse R. Xmark: a benchmark for xml data management. In *International Conference on Very Large Data Bases (VLDB)*, Hong Kong, China, August 2002; 974–985.

53. Franceschet M. Xpathmark – an xpath benchmark for xmark generated data. In *International XML Database Symposium (XSYM)*, Trondheim, Norway, 2005; 129–143.

54. Li YG, Bressan S, Dobbie G, Lacroix Z, Lee ML, Nambiar U, Wadhwa B. XOO7: applying OO7 benchmark to XML query processing tool. *Proceedings of the 10th International Conference on Information and Knowledge Management*, Atlanta, GA, USA, 2001; 167–174.

55. McVoy L, Staelin C. lmbench: portable tools for performance analysis. In *Proceedings of the 1996 Annual Conference on USENIX Annual Technical Conference*, ATEC '96. USENIX Association: Berkeley, CA, USA, 1996; 23–23.

56. SQL Standards. Part 2: foundation. *Technical Report ISO/IEC 9075-2:2011*, ISO, Geneva, Switzerland, 2011.

57. W3C. Document object model, 2007. Available from: http://www.w3.org/DOM [last accessed 26 March 2007].

58. Libxml. The XML C parser and toolkit of Gnome, version 2.7.2, 2008. Available from: http://xmlsoft.org/, [last accessed 5 February 2009].

59. Wang Y, DeWitt D, Cai J. X-Diff: an effective change detection algorithm for XML documents. In *Proceedings of the 19th International Conference on Data Engineering*, Bangalore, India, 2004; 519–530.

60. Snodgrass RT. *Developing Time-oriented Database Applications in SQL*. Morgan Kaufmann Publishers Inc.: San Francisco, CA, USA, 2000.

61. Thomas SW. Mining software repositories with topic models. *Technical Report 2012-586*, School of Computing, Queen's University, 2012.

62. Robillard MP, Murphy GC. Representing concerns in source code. *ACM Transactions on Software Engineering and Methodology* 2007; **16**(1).

63. Thomas SW, Adams B, Hassan AE, Blostein D. Validating the use of topic models for software evolution. In *Proceedings of the 10th International Working Conference on Source Code Analysis and Manipulation*, Timisoara, Romania, 2010; 55–64.

64. Antoniol G, Gueheneuc YG, Merlo E, Tonella P. Mining the lexicon used by programmers during sofware evolution. In *International Conference on Software Maintenance*, Paris, France, 2007; 14–23.

65. Blei DM, Ng AY, Jordan MI. Latent Dirichlet allocation. *Journal of Machine Learning Research* 2003; **3**:993–1022.

66. Thomas SW, Adams B, Hassan AE, Blostein D. Modeling the evolution of topics in source code histories. In *Proceedings of the 8th Working Conference on Mining Software Repositories*, Honolulu, Hawaii, USA, 2011; 173–182.

67. Rothermel G, Untch R, Chu C, Harrold M. Prioritizing test cases for regression testing. *Transactions on Software Engineering* 2001; **27**(10):929–948.

68. Lukins SK, Kraft NA, Etzkorn LH. Bug localization using latent dirichlet allocation. *Information and Software Technology* 2010; **52**(9):972–990.

69. Mozilla Foundation. Bugzilla, 2011. Available from: http://www.bugzilla.org/. (last accessed 25 February 2013).
70. D'Ambros M, Lanza M, Robbes R. Evaluating defect prediction approaches: a benchmark and an extensive comparison. *Empirical Software Engineering* 2012; **17**(4):531–577.
71. Stonebraker M, Abadi DJ, Batkin A, Chen X, Cherniack M, Ferreira M, Lau E, Lin A, Madden S, O'Neil E, O'Neil P, Rasin A, Tran N, Zdonik S. C-store: a column-oriented DBMS. In *Proceedings of the International Conference on Very Large Databases*, Lyon, France, 2005; 553–564.
72. Boncz PA, Zukowski M, Nes N. MonetDB/X100: hyper-pipelining query execution. In *Conference on Innovations in Database Research*, Asilomar, CA, USA, 2005; 225–237.
73. Pavlo A, Jones EPC, Zdonik S. On predictive modeling for optimizing transaction execution in parallel OLTP systems. *Proc. VLDB Endow* October 2011; **5**(2):85–96.
74. Oracle Corporation. Workspace Manager Developer's Guide 11g Release 1 (11.1), Aug. 2008.
75. Teradata Corporation. Teradata Temporal, 2010. Available from: http://www.teradata.com/t/database/Teradata-Temporal/ [last accessed October 2010].
76. IBM Corporation. A Matter of Time: Temporal Data Management in DB2 for z/OS, 2010. accessed December 2010.