Relational Algebra by Way of Adjunctions

JEREMY GIBBONS, University of Oxford, United Kingdom

FRITZ HENGLEIN, University of Copenhagen, Denmark

RALF HINZE, Universität Kaiserslautern, Germany

NICOLAS WU, University of Bristol, United Kingdom

Bulk types such as sets, bags, and lists are monads, and therefore support a notation for database queries based on comprehensions. This fact is the basis of much work on database query languages. The monadic structure easily explains most of standard relational algebra—specifically, selections and projections—allowing for an elegant mathematical foundation for those aspects of database query language design. Most, but not all: monads do not immediately offer an explanation of relational join or grouping, and hence important foundations for those crucial aspects of relational algebra are missing. The best they can offer is cartesian product followed by selection. Adjunctions come to the rescue: like any monad, bulk types also arise from certain adjunctions; we show that by paying due attention to other important adjunctions, we can elegantly explain the rest of standard relational algebra. In particular, graded monads provide a mathematical foundation for indexing and grouping, which leads directly to an efficient implementation, even of joins.

CCS Concepts: • **Information systems** → **Relational database query languages**; • **Theory of computation** → **Categorical semantics**; • **Software and its engineering** → **Semantics**;

Additional Key Words and Phrases: SQL, comprehension, adjunction, monad, graded monad.

ACM Reference Format:

Jeremy Gibbons, Fritz Henglein, Ralf Hinze, and Nicolas Wu. 2018. Relational Algebra by Way of Adjunctions. *Proc. ACM Program. Lang.* 2, ICFP, Article 86 (September 2018), 28 pages. <https://doi.org/10.1145/3236781>

1 INTRODUCTION

Consider the following SQL query, which computes the names and amounts outstanding of customers with overdue invoices, by joining separate customer and invoice tables:

```
SELECT name, amount
FROM customers, invoices
WHERE cid = cust AND due < today
```

Standard database techniques [Date 2004] translate this query into applications of relational algebra operations such as projection π , selection σ , and equijoin $\bowtie$:

$$\pi_{name, amount} (\sigma_{due < today} (customers \text{ cid} \bowtie_{cust} invoices))$$

About 25 years ago, several people observed [Buneman et al. 1994; Trinder 1991; Wadler 1992] that bulk types such as sets and bags (sometimes called ‘multisets’) form monads, and so support a comprehension syntax that provides a convenient notation and clear and well-understood mathematical

Authors’ addresses: Jeremy Gibbons, Department of Computer Science, University of Oxford, Oxford, United Kingdom, jeremy.gibbons@ox.ac.uk; Fritz Henglein, DIKU, University of Copenhagen, Copenhagen, Denmark, henglein@diku.dk; Ralf Hinze, Fachbereich Informatik, Universität Kaiserslautern, Kaiserslautern, Germany, ralf@cs.uni-kl.de; Nicolas Wu, Department of Computer Science, University of Bristol, Bristol, United Kingdom, nicolas.wu@bristol.ac.uk.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2018 Copyright held by the owner/author(s).

2475-1421/2018/9-ART86

<https://doi.org/10.1145/3236781>

foundations for database queries. Many modern programming languages provide built-in syntax for such comprehensions, yielding for free an embedded domain-specific language for queries—for example, in Haskell (writing a lowered dot for record field selection):

$$[(c.name, i.amount) \mid c \leftarrow customers, i \leftarrow invoices, c.cid == i.cust, i.due < today]$$

Comprehensions have been used for many years as a programming construct for collections, dating back at least to Schwartz’s SETL from the late 1960s [Schwartz et al. 1986] and Darlington’s NPL from the 1970s [Darlington 1975]. They are the basis of queries in the Kleisli [Wong 2000], LINQ [Meijer 2007; Syme 2006], and Links [Cheney et al. 2013; Cooper et al. 2006] database languages, and also feature in proposals for query notations in OQL [Grust and Scholl 1999], XQuery [Fernandez et al. 2001], and MapReduce-style data analysis [Alexandrov et al. 2015].

However, there is a catch. Monads and monad comprehensions provide no direct way to express an equijoin; this can only be expressed indirectly, as a cartesian product followed by a selection. This becomes clear if we desugar the Haskell comprehension notation, yielding something equivalent to the following:

$$\begin{aligned} &fmap (\lambda(c, i) \rightarrow (c.name, i.amount)) (\\ &\quad filter (\lambda(c, i) \rightarrow i.due < today) (\\ &\quad \quad filter (\lambda(c, i) \rightarrow c.cid == i.cust) (\\ &\quad \quad \quad cp (customers, invoices)))) \end{aligned}$$

This amounts to a very inefficient query plan, computing the full cartesian product (via *cp*) of the *customers* and *invoices* tables before immediately discarding the unmatching pairs—typically the great majority—via the innermost *filter*. Of course, there is a much better plan, computing the same subset of customer/invoice pairs in a more circumspect manner:

$$\begin{aligned} &fmap (fmap name \times fmap amount) (cod (\\ &\quad fmap (id \times filter (\lambda i \rightarrow i.due < today)) (\\ &\quad \quad merge (customers \text{ ‘indexBy’ } cid, invoices \text{ ‘indexBy’ } cust)))) \end{aligned}$$

Here, *indexBy* groups a collection into an indexed table, and is written as an infix operator; *merge* combines two indexed tables with matching indices; and *cod* retrieves the elements—the ‘codomain’—from an indexed table. The types are something like the following:

$$\begin{aligned} indexBy &:: Key\ k \Rightarrow Bag\ v \rightarrow (v \rightarrow k) \rightarrow (k \rightarrow Bag\ v) \\ merge &:: Key\ k \Rightarrow (k \rightarrow v, k \rightarrow v') \rightarrow (k \rightarrow (v, v')) \\ cod &:: Key\ k \Rightarrow (k \rightarrow Bag\ v) \rightarrow Bag\ v \end{aligned}$$

parametric in the type *v* of values, but with a qualification *Key* on the type *k* of keys—only certain types are suitable as keys—so that *k*-indexed tables of *v*-values are represented by the type *k* $\rightarrow$ Bag *v*. (We will explain the details later in the paper.)

This alternative program yields a result of type Bag (Bag *Name* $\times$ Bag *Amount*), and admits a straightforward linear-time implementation. (Going one step further and flattening the result to a Bag (*Name* $\times$ *Amount*) may blow up the size of the result from linear in the input size to quadratic, so does not admit a linear-time implementation.) However, this program does not directly correspond to anything that can be written solely in the more accessible comprehension syntax.

Do we have to sacrifice the mathematical elegance of comprehension notation on the altar of efficient execution? In this paper, we show how to sidestep such an unfortunate dilemma. Specifically, we reconsider the mathematics that gives rise to the monads of bulk types like sets and bags—*category theory*, the “mathematics of mathematics” [Cheng 2015]. The central notion is that of *adjunction*, the categorical generalisation of Galois connections and “a concept of fundamental logical and mathematical importance that is not captured elsewhere in mathematics” [Awodey

2006, p179]. Adjunctions capture *universal properties*, which embody essentially all the important equivalences that justify query transformations—in particular, query optimizations.

For example, the fact that bags are the *free commutative monoid* on a type of elements is embodied in an adjunction between two categories: the category of commutative monoids and their homomorphisms, and the category of sets and total functions. The equational properties of bag homomorphisms such as maps and filters that underlie many query optimizations all arise from this adjunction.

Of course, we are not the first to argue that adjunctions are a crucial ingredient of programming language design, even in the more specific domain of query languages; for example, Libkin and Wong [1997] present an approach “based on turning universal properties of collections into programming syntax”. The novelty in this paper is to show that adjunctions induce not only monads and comprehensions, which explain relational database selections and projections, but also all the other abstract gadgets needed to complete the explanation of the efficient query plan—in particular, of the *grouping and merging* required for linear-time equijoins. One might say that monads give rise to the accessible point-wise comprehension notation that exemplifies the relational calculus, whereas the underlying adjunctions give rise to the more flexible point-free combinators of the relational algebra.

In summary, we revisit the story of how a comprehension-style notation for collection processing falls out of certain adjunctions, and extend that story to explain how equijoins can also fit into the picture without sacrificing efficiency. Specifically, our contributions are:

- elaborating how *adjunctions* underlie collection processing, including the convenient notation based on comprehensions for expressing queries over collections;
- deriving the *equational properties* of collection-processing operations from the adjunction identities;
- extending the operations to support *grouping* into and *merging* of indexed collections, and identifying the relevant adjunctions;
- adapting the *comprehension notation* to accommodate also equijoins, by translation into grouping and merging;
- exploiting *graded monads* as an appropriate abstraction for capturing the *finiteness* required for aggregations;
- justifying a body of *query transformations* that underlie query optimization, by calculation from the equational properties of collection-processing operations.

We do not claim to present any new query transformations; rather, we see it as a strength that we explain existing well-known transformations from a new perspective.

The rest of this paper is structured as follows. Section 2 recalls some fundamental concepts from category theory on which we build, mainly for the purposes of fixing notation. Section 3 builds a simplified version of the relational algebra in which the containers are all finite bags. We then develop a series of richer models of containers, starting with finite maps in Section 4, then moving to indexed tables in Section 5; in Section 6 we prove the correctness of a linear-time implementation of equijoin on indexed tables. We then introduce graded monads in Section 7 as a means of modelling indexed tables with finite domains; this serves as our final model of databases. With the machinery in place, we show in Section 8 how the transformations used for query optimization arise naturally from the theory. We conclude in Section 9. Appendix A (in the online supplementary materials) gives more details of the adjunctions underlying graded monads; Appendix B provides a prototype implementation in Haskell—intended more as an alternative specification of behaviour than something to be executed.

2 CATEGORY THEORY

The mathematical structures we present as a foundation for relational algebra are rooted in category theory, so we present a very quick overview of the relevant definitions. A full tutorial is beyond the scope of this paper, but this section serves to fix notation for those who have some familiarity. Nevertheless, we hope this provides at least some intuition for readers otherwise new to the material. For more details, see any standard textbook [Awodey 2006; Pierce 1991]. Conversely, the categorically adept reader might simply skip this section.

2.1 Categories

A *category* $\mathcal{C}$ consists of a collection of *objects* and a collection of *arrows*; each arrow $f : A \rightarrow B$ goes from a *source* object A to a *target* object B . For each object A there is an identity arrow $id : A \rightarrow A$. Adjacent arrows $f : A \rightarrow B$ and $g : B \rightarrow C$ compose to form $g \cdot f : A \rightarrow C$. Composition is associative, with appropriate identity arrows as neutral elements. The *hom-set* $\mathcal{C}(A, B)$ consists of the arrows in $\mathcal{C}$ from A to B . An *initial* object 0 has a unique arrow from it to any other object. Dually, a *final* object 1 has a unique arrow to it from any other.

Think of the objects as ‘types’ and the arrows as ‘functions’ or ‘programs’ from one type to another. A fundamental instance is the category **Set**, in which the objects are sets and the arrows total functions; the initial object $0 = \{\}$ is the empty set, and any singleton set $1 = \{\{\}\}$ is final. But we will also consider instances with more structured objects (such as ordered sets, or monoids); then the arrows are typically structure-preserving mappings (in these cases, monotonic functions and monoid homomorphisms, respectively). For instance, the category **CMon** in Section 3 has commutative monoids M as objects, represented by triples $(M, \epsilon, \otimes)$ where M is a set and the operation $\otimes : M \times M \rightarrow M$ is associative and commutative with neutral element $\epsilon : M$. Arrows $h : (M_1, \epsilon_1, \otimes_1) \rightarrow (M_2, \epsilon_2, \otimes_2)$ in this category are homomorphisms: functions $h : M_1 \rightarrow M_2$ between the underlying sets satisfying $h \epsilon_1 = \epsilon_2$ and $h (x \otimes_1 y) = h x \otimes_2 h y$. The underlying set of a commutative monoid is called its *carrier*, and when the context is clear is often used to refer to the whole structure, leaving the operation and neutral element implicit.

2.2 Functors

Categories are themselves structured objects; *functors* are the structure-preserving mappings between categories. Functor $F : \mathcal{C} \rightarrow \mathcal{D}$ from category $\mathcal{C}$ to category $\mathcal{D}$ maps the objects and arrows of $\mathcal{C}$ to objects and arrows of $\mathcal{D}$, respecting the structure by preserving sources and targets ($F f : F A \rightarrow F B$ when $f : A \rightarrow B$) and identities and composition ($F (id A) = id (F A)$ and $F (g \cdot f) = F g \cdot F f$). There is an identity functor Id , and functors compose $G \circ F$; so categories and functors themselves form a category **Cat**. For example, functor $U : \mathbf{CMon} \rightarrow \mathbf{Set}$ maps a monoid to its carrier, and a monoid homomorphism to the underlying function between carriers. When $\mathcal{C}$ and $\mathcal{D}$ coincide, we call F an *endofunctor*. In **Set**, endofunctors can be seen as parametric type constructors such as **Bag**, which maps an object A (a set of elements) to the object **Bag** A (the set of finite bags of those elements), and an arrow $f : A \rightarrow B$ to the arrow **Bag** $f : \mathbf{Bag} A \rightarrow \mathbf{Bag} B$ that applies f to each element of a bag. It is an instructive exercise to verify that **Bag** indeed respects structure.

2.3 Natural Transformations

Given two functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$ between the same two categories, a *natural transformation* $\phi : F \rightarrow G$ is a family of arrows $\phi A : F A \rightarrow G A$ in $\mathcal{D}$, one for each object A in $\mathcal{C}$, respecting the arrow structure of $\mathcal{C}$, in the sense that $G f \cdot \phi A = \phi B \cdot F f$ for each $f : A \rightarrow B$ (see Figure 1(a)). A natural transformation between endofunctors in **Set** is a polymorphic function; for example,

Fig. 1. (a) natural transformation, (b) algebra homomorphism

$\{-\} : \text{Bag} \rightarrow \text{Set}$ is the polymorphic family of functions, one per element type, flattening a bag to its underlying set by forgetting multiplicity of elements.

Natural transformations can be composed. Given the functors $E, F, G : \mathcal{C} \rightarrow \mathcal{D}$, and natural transformations $\psi : E \rightarrow F$ and $\phi : F \rightarrow G$, the *vertical composition* $\phi \cdot \psi : E \rightarrow G$ has components $(\phi \cdot \psi) A = \phi A \cdot \psi A$. Given the functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$ and $H, K : \mathcal{D} \rightarrow \mathcal{E}$, and the natural transformations $\phi : F \rightarrow G$ and $\gamma : H \rightarrow K$, the *horizontal composition* $\gamma \phi : H \circ F \rightarrow K \circ G$ has components $(\gamma \phi) A = \gamma (G A) \cdot H (\phi A) = K (\phi A) \cdot \gamma (F A)$. It is convenient to mix these natural transformations and functors: the natural transformation $H \phi$ is given by $(H \phi) A = H (\phi A)$, and γF is given by $(\gamma F) A = \gamma (F A)$.

2.4 Adjunctions

Functors $L : \mathcal{D} \rightarrow \mathcal{C}$ and $R : \mathcal{C} \rightarrow \mathcal{D}$ form an *adjunction* $L \dashv R : \mathcal{D} \rightarrow \mathcal{C}$ if the arrows $L A \rightarrow B$ in $\mathcal{C}$ are in natural 1-to-1 correspondence with the arrows $A \rightarrow R B$:

$$[-] : \mathcal{C}(L A, B) \cong \mathcal{D}(A, R B) : [-]$$

This is ‘natural’ in the sense that there are natural transformations $\eta : \text{Id} \rightarrow R \circ L$, defined so that $\eta A = [-] id (L A)$ is the correspondent in $\mathcal{D}$ of $id (L A)$ in $\mathcal{C}$, and symmetrically $\epsilon : L \circ R \rightarrow \text{Id}$ with $\epsilon B = [-] id (R B)$ the correspondent in $\mathcal{C}$ of $id (R B)$ in $\mathcal{D}$, subject to the so-called triangle identities $R \epsilon \cdot \eta R = id$ and $\epsilon L \cdot L \eta = id$. We call L the ‘left adjoint’, R the ‘right adjoint’, η the ‘unit’, and ϵ the ‘counit’. In diagrams (such as in Figure 2), it is often convenient to draw L and R as parallel horizontal arrows in opposite directions; then it is conventional to rotate the $\dashv$ symbol by 90 degrees.

Adjunctions really do model natural correspondences between categories, generalizing Galois connections from lattice theory; they are generally embedding–retraction pairs. For example, the product category Set^2 has pairs (A, B) of Set -objects A and B as objects and pairs $(f, g) : (A, B) \rightarrow (C, D)$ of Set -arrows $f : A \rightarrow C$ and $g : B \rightarrow D$ as arrows, and the diagonal functor $\Delta : \text{Set} \rightarrow \text{Set}^2$ takes object A and arrow f in Set to (A, A) and (f, f) in Set^2 ; then the adjunction $\Delta \dashv \times$ between Set^2 and Set in Section 3 captures the equivalence between arrows $\Delta A \rightarrow (B, C)$ in Set^2 (that is, pairs of functions with a common domain) and arrows $A \rightarrow B \times C$ in Set (that is, functions constructing pairs)—an embedding of the pair-forming functions within the function pairs. A more important example for us is the adjunction $\text{Free} \dashv U$ between CMon and Set in Section 3, capturing the embedding of a set A as the ‘free’ commutative monoid $\text{Free } A = (\text{Bag } A, \emptyset, \uplus)$, and the retraction $U (M, \epsilon, \otimes) = M$ of a commutative monoid to its carrier. The 1-to-1 correspondence in this case is between bag aggregations $\text{Free } A \rightarrow (M, \epsilon, \otimes)$ which are bag homomorphisms to some commutative monoid, and plain functions $A \rightarrow U M$ from bag elements into the commutative monoid. The triangle identities specialize to give all the important properties of aggregations.

2.5 Monads

A *monad* (T, η, μ) consists of an endofunctor T equipped with two natural transformations, unit $\eta : \text{Id} \rightarrow T$ and multiplication $\mu : T \circ T \rightarrow T$, satisfying three identities: $\mu \cdot T \eta = \text{id}$, $\mu \cdot \eta T = \text{id}$, and $\mu \cdot T \mu = \mu \cdot \mu T$. The composition of adjoint functors $R \circ L$ always forms a monad; the unit η of the monad is the unit of the adjunction, and the multiplication is constructed from the counit as $\mu = R \epsilon L$.

Bulk types such as `Bag` are monads; η constructs a singleton bag, and μ flattens a bag of bags to a bag. The `Bag` monad arises from the adjunction $\text{Free} \dashv U$ discussed above.

2.6 Algebras

An *algebra* for endofunctor F is a pair $(A, f : F A \rightarrow A)$, with A called the *carrier* of the algebra and f the *action*. A *homomorphism* h between F -algebras (A, f) and (B, g) is an arrow $h : A \rightarrow B$ in the underlying category that is coherent with the actions: $h \cdot f = g \cdot F h$ (see Figure 1(b)). The identity arrow is a homomorphism, and homomorphisms compose; so F -algebras and their homomorphisms form a category $F\text{-Alg}(\mathcal{C})$.

There is a *forgetful* functor $U^F : F\text{-Alg}(\mathcal{C}) \rightarrow \mathcal{C}$ discarding the algebraic structure: $U^F (A, f) = A$ and $U^F (h : (A, f) \rightarrow (B, g)) = h : A \rightarrow B$. Under mild conditions on F and $\mathcal{C}$ (for example, for polynomial functors F on `Set`, that is, sums of products), the forgetful functor U^F has a left adjoint $\text{Free}^F \dashv U^F$ (see Figure 2(h)). The functor Free^F maps an object A to the *free* F -algebra $\text{Free}^F A$, which can be seen as the terms built from signature F on variables A . Similar to the closely related adjunction $\text{Free} \dashv U$ between `CMon` and `Set` above, the adjunction $\text{Free}^F \dashv U^F$ states the 1-to-1 correspondence between the homomorphisms $\text{Free}^F A \rightarrow (B, g)$ on the free algebra and the mappings $A \rightarrow U^F (B, g) = A \rightarrow B$ on their variables. In the special case that $A = 0$, the initial object in the underlying category, Free^F yields the initial F -algebra $(\mu F, \text{In})$, and the unique homomorphism to a target algebra is called a *fold*. The adjunction is crucial for programming with algebraic datatypes: these are interpreted as initial algebras, so there is a unique (the empty) mapping on the set of variables, and therefore a homomorphism is completely determined by specifying a target algebra.

An *algebra for a monad* (T, η, μ) (or ‘Eilenberg–Moore algebra’) is an algebra $(A, f : T A \rightarrow A)$ for T considered as an endofunctor, which additionally respects the monadic structure: $f \cdot \eta A = \text{id} A$ and $f \cdot \mu A = f \cdot T f$. Monad algebras model well-behaved mappings from structured objects; in particular, algebras for a bulk type such as `Bag` are the aggregation operations, such as *sum* and *maximum*.

3 RELATIONAL ALGEBRA

By exploiting a small number of fundamental adjunctions between familiar categories of structured objects, we can build up a collection of mathematical operators that form a basis for relational algebra. We start in this section with *finite bags*; in Section 4, we introduce the *pointed sets* we need for *finite maps*; in Section 5, we combine these to yield *indexed tables*; in Section 6, we define grouping to indexed tables, and prove that equijoin implemented using grouping agrees with the naive nested-loop version; finally, in Section 7, we present a modest generalisation of monads called *grading* in order to accommodate finiteness. It turns out that a little category theory goes a surprisingly long way: everything we need arises from the handful of adjunctions in Figure 2. This is intellectually satisfying. But more than that: the adjunctions not only give us the appropriate structures, they also provide precisely the equational laws needed for typical query optimization, as we discuss in Section 8.

$$\begin{array}{ccc}
\text{Set} \xleftarrow[\Delta]{\begin{smallmatrix} 0 \\ \perp \end{smallmatrix}} \text{Set}^0 \xleftarrow[\perp]{\Delta} \text{Set} & \text{Set} \xleftarrow[\Delta]{+} \text{Set}^2 \xleftarrow[\times]{\Delta} \text{Set} & \text{Set}_* \xleftarrow[\cup]{\begin{smallmatrix} L \\ \perp \end{smallmatrix}} \text{Set} \\
\Delta : A \mapsto * & \Delta : A \mapsto (A, A) & L : A \mapsto (A \cup \{*\}, *) \\
0 : * \mapsto \{ \} & + : (A, B) \mapsto A + B & \cup : (A, a) \mapsto A \\
1 : * \mapsto \{*\} & \times : (A, B) \mapsto A \times B & \\
\text{(a)} & \text{(b)} & \text{(c)}
\end{array}$$

$$\begin{array}{ccc}
\text{Set} \xleftarrow[\perp]{\begin{smallmatrix} -\times P \\ \perp \end{smallmatrix}} \text{Set} & \text{Rel} \xleftarrow[\perp]{\begin{smallmatrix} -\times P \\ \perp \end{smallmatrix}} \text{Rel} \\
(-)^P : (f : A \rightarrow B) \mapsto (f \times P : A \times P \rightarrow B \times P) & (- \times P) : (R : A \leftrightarrow B) \mapsto (R \times P : A \times P \leftrightarrow B \times P) \\
(-)^P : (f : A \rightarrow B) \mapsto ((f \cdot) : A^P \rightarrow B^P) & (P \times -) : (R : A \leftrightarrow B) \mapsto (P \times R : P \times A \leftrightarrow P \times B) \\
\text{(d)} & \text{(e)}
\end{array}$$

$$\begin{array}{ccc}
\text{CMon} \xleftarrow[\cup]{\text{Free}} \text{Set} & \text{Rel} \xleftarrow[\text{E}]{\text{J}} \text{Set} & \text{F-Alg}(\mathcal{C}) \xleftarrow[\cup^F]{\text{Free}^F} \mathcal{C} \\
\text{Free} : A \mapsto (\text{Bag } A, \emptyset, \uplus) & \text{J} : (f : A \rightarrow B) \mapsto \{(a, f \ a) \mid a \in A\} & \text{Free}^F : A \mapsto \text{"terms of sign. F with vars A"} \\
\cup : (M, \epsilon, \otimes) \mapsto M & \text{E} : (R : A \leftrightarrow B) \mapsto \lambda X \subseteq A. \{b \in B \mid \exists a \in X. a \ R \ b\} & \cup^F : (A, f) \mapsto A \\
\text{(f)} & \text{(g)} & \text{(h)}
\end{array}$$

Fig. 2. The adjunctions underlying relational algebra

As a first approximation, a database table is a finite bag of values—a bag rather than a set or a list, as we wish to preserve the multiplicity of elements but ignore any ordering. Preserving multiplicity is important for non-idempotent aggregations, such as summing a collection of numbers:

sum (fmap amount (filter ($\lambda i \rightarrow i.\text{due} < \text{today}$) invoices))

With sets rather than bags as the bulk datatype, the query simply doesn't work: coincidentally duplicate amounts are incorrectly discarded.

The values contained in a relational database are typically records, but we do not wish to limit attention to this particular class of types. The reason is simple: intermediate tables often contain non-record values; for instance, in the example above, the argument to *sum* is a bag of monetary values; and the running example in the introduction yielded a bag of pairs of bags as the final result. The operations of the relational algebra enjoy straightforward implementations in terms of bags—Figure 3 summarizes the notation, with full definitions to follow.

The function *single* takes a single element to a singleton bag, and $\uplus$ is bag union. We use bag brackets as a shorthand to define bags by listing their elements, instead of using *single* and $\uplus$; so $\{a, b, b\} = \text{single } a \uplus \text{single } b \uplus \text{single } b$ is the bag that contains *a* once and *b* twice. Even though the bag model of database tables is rather naive, exploring the theory of bags in more depth is still useful: it will pay dividends later when we discuss more sophisticated models.

Most of the operations and their accompanying properties arise out of the adjunction

$$\text{CMon} \xleftarrow[\cup]{\text{Free}} \text{Set}$$

table of V values	Bag V	projection π_f	Bag f
empty table	$\emptyset$	selection σ_p	<i>filter</i> p
singleton table	<i>single</i>	aggregation in monoid M	<i>reduce</i> M
union of tables	$\uplus$		
cartesian product of tables	$\times$		
neutral element	$\wr()$		

Fig. 3. Relational algebra operators, implemented in terms of bags

between the category **CMon** of commutative monoids and their homomorphisms and the category **Set** of sets and total functions (Figure 2(f)). Very briefly, the forgetful functor U maps a commutative monoid $M = (M, \epsilon, \otimes)$ to its carrier set M , and a homomorphism $h: (M_1, \epsilon_1, \otimes_1) \rightarrow (M_2, \epsilon_2, \otimes_2)$ to its underlying function $h: M_1 \rightarrow M_2$ on the carriers; the algebraic structure is forgotten. Its left adjoint Free takes a set A to $\text{Free } A = (\text{Bag } A, \emptyset, \uplus)$, the so-called *free commutative monoid* on A , and lifts a function f on elements to homomorphism $\text{Bag } f$ on bags. Adjointness means that homomorphisms of type $\text{Free } A \rightarrow M$ are in 1-to-1 correspondence with functions of type $A \rightarrow U M$, where A is a set and M is a commutative monoid. (To be precise, there is a 1-to-1 correspondence classically; in a constructive setting, one would have to take the usual care over what is meant by ‘equals’, perhaps assuming decidable equality on A [Carette 2018].) This bijection is natural in A and M . Let us now investigate what this tells us about the operations of relational algebra.

Monoid operations. First of all, we record that $(\text{Bag } A, \emptyset, \uplus)$ is a commutative monoid. The empty bag and bag union

$$\begin{aligned} \emptyset &: \text{Bag } A \\ (\uplus) &: \text{Bag } A \times \text{Bag } A \rightarrow \text{Bag } A \end{aligned}$$

satisfy the required laws:

$$\emptyset \uplus s = s = s \uplus \emptyset \quad (1a)$$

$$(s \uplus t) \uplus u = s \uplus (t \uplus u) \quad (1b)$$

$$s \uplus t = t \uplus s. \quad (1c)$$

Projections. The type of finite bags is functorial in the elements, as it can be written as a composition of two functors, $\text{Bag} = U \circ \text{Free}$; so mapping over bags satisfies the two functor laws.

$$\text{Bag } id = id \quad (2a)$$

$$\text{Bag } (g \cdot f) = \text{Bag } g \cdot \text{Bag } f \quad (2b)$$

The functor Free maps a function to a homomorphism; Bag inherits this property (since U acts as the identity on arrows).

$$\text{Bag } f \emptyset = \emptyset \quad (3a)$$

$$\text{Bag } f \wr a \wr = \wr f a \wr \quad (3b)$$

$$\text{Bag } f (s \uplus t) = \text{Bag } f s \uplus \text{Bag } f t \quad (3c)$$

These three equations fully determine $\text{Bag } f$. In the special case that the elements of the bag are records, and function f is a record projection (preserving some fields and discarding others), then $\text{Bag } f$ can be seen as projecting some columns out of a database table. But the functorial action is more general than this: the bag elements need not be records, and the function f may perform some non-trivial computation on each element rather than merely a projection.

Aggregations. The adjunction $\text{Free} \dashv \text{U}$ captures a 1-to-1 correspondence between homomorphisms of type $\text{Free } A \rightarrow M$ in **CMon** and functions of type $A \rightarrow \text{U } M$ in **Set**. Setting $M := \text{Free } A$, the correspondent in **Set** of the identity homomorphism $\text{id} (\text{Free } A)$ is

$$\text{single } A : A \rightarrow \text{U } (\text{Free } A) ,$$

which maps an element to a singleton bag; *single* is the unit of the adjunction. We also write $\text{single } A \ a$ as $\lceil a \rceil$. Conversely, setting $A := \text{U } M$, the correspondent in **CMon** of the identity function $\text{id} (\text{U } M)$ is

$$\text{reduce } M : \text{Free } (\text{U } M) \rightarrow M ,$$

which is the unique homomorphism from the free commutative monoid on the carrier of M to the commutative monoid M itself that takes the singleton $\lceil a \rceil$ to a ; this is the counit of the adjunction. Using the units, we can capture the 1-to-1 correspondence as an equivalence:

$$h = \text{reduce } M \cdot \text{Free } f \iff \text{U } h \cdot \text{single } A = f , \quad (4)$$

for all homomorphisms $h : \text{Free } A \rightarrow M$ and functions $f : A \rightarrow \text{U } M$. In other words, a homomorphism from the free commutative monoid is uniquely determined by its behaviour on singleton bags. The homomorphism h is sometimes called the extension of f .

Turning to the properties, $\text{reduce } M$ for monoid $M = (M, \epsilon, \otimes)$ is a homomorphism:

$$\text{reduce } M \ \emptyset = \epsilon \quad (5a)$$

$$\text{reduce } M \ \lceil a \rceil = a \quad (5b)$$

$$\text{reduce } M \ (s \uplus t) = \text{reduce } M \ s \otimes \text{reduce } M \ t . \quad (5c)$$

Furthermore, *single* and *reduce* are natural transformations:

$$\text{U } (\text{Free } f) \cdot \text{single } A = \text{single } B \cdot f \quad (6a)$$

$$h \cdot \text{reduce } M = \text{reduce } N \cdot \text{Free } (\text{U } h) , \quad (6b)$$

for all functions $f : A \rightarrow B$ and homomorphisms $h : M \rightarrow N$. Equation (6a) is the point-free presentation of (3b).

Simple examples of homomorphisms from the free commutative monoid are aggregations such as counting and conjoining.

aggregation	monoid	on singletons
<i>count</i>	$(\mathbb{N}, 0, +)$	$\lceil a \rceil \mapsto 1$
<i>sum</i>	$(\mathbb{Z}, 0, +)$	$\lceil a \rceil \mapsto a$
<i>max</i>	$(\mathbb{Z}, \text{minBound}, \text{max})$	$\lceil a \rceil \mapsto a$
<i>min</i>	$(\mathbb{Z}, \text{maxBound}, \text{min})$	$\lceil a \rceil \mapsto a$
<i>all</i>	$(\mathbb{B}, \text{True}, \wedge)$	$\lceil a \rceil \mapsto a$
<i>any</i>	$(\mathbb{B}, \text{False}, \vee)$	$\lceil a \rceil \mapsto a$

The last four operations target monoids that are also idempotent, so these would continue to work if bags were replaced by sets—unlike the first two operations.

Strictly speaking, a homomorphism such as $h = \text{reduce } (\mathbb{N}, 0, +)$ is an arrow in **CMon**, and we ought to write ‘ $\text{U } h$ ’ for the corresponding function in **Set**; but in the remainder of the paper we will often abuse notation by omitting the U .

Selection. A restriction or selection, here called *filter*, is a further example of a homomorphism from the free commutative monoid. A filter can be seen as the intersection of a finite bag with an arbitrary set represented by its $\mathbb{B}$ -valued characteristic function.

$$\begin{aligned} \text{filter} &: (A \rightarrow \mathbb{B}) \rightarrow (\text{Bag } A \rightarrow \text{Bag } A) \\ \text{filter } p &= \text{reduce } (\text{Free } A) \cdot \text{Bag } (\text{guard } p) \\ \text{guard} &: (A \rightarrow \mathbb{B}) \rightarrow (A \rightarrow \text{Bag } A) \\ \text{guard } p \ a &= \text{if } p \ a \ \text{then } \lfloor a \rfloor \ \text{else } \emptyset \end{aligned}$$

In other words, *filter* p is the extension of *guard* p . The latter function can be neatly written using a guarded bag comprehension: *guard* $p = \lambda a \rightarrow \lfloor a \mid p \ a \rfloor$, which motivates the choice of name. Turning to the properties, *filter* p is a homomorphism and a dinatural transformation (7d).

$$\text{filter } p \ \emptyset = \emptyset \quad (7a)$$

$$\text{filter } p \ \lfloor a \rfloor = \text{guard } p \ a \quad (7b)$$

$$\text{filter } p \ (s \uplus t) = \text{filter } p \ s \uplus \text{filter } p \ t \quad (7c)$$

$$\text{filter } p \cdot \text{Bag } f = \text{Bag } f \cdot \text{filter } (p \cdot f) \quad (7d)$$

Monad operations. The functor $\text{Bag} = \text{U} \circ \text{Free}$ is furthermore a monad: the composition of a right adjoint with its left adjoint always yields a monad. The unit η of the monad is *single*; the multiplication μ , which takes a bag of bags to its union, is constructed from the counit by

$$\begin{aligned} \text{union } A &: \text{Bag } (\text{Bag } A) \rightarrow \text{Bag } A \\ \text{union } A &= \text{U } (\text{reduce } (\text{Free } A)) \end{aligned}$$

By construction, *single* and *union* satisfy the monad laws:

$$\text{union } A \cdot \text{Bag } (\text{single } A) = \text{id } (\text{Bag } A) \quad (8a)$$

$$\text{union } A \cdot \text{single } (\text{Bag } A) = \text{id } (\text{Bag } A) \quad (8b)$$

$$\text{union } A \cdot \text{Bag } (\text{union } A) = \text{union } A \cdot \text{union } (\text{Bag } A) \quad (8c)$$

The fact that Bag is a monad justifies the use of bag comprehensions, as these can be desugared into applications of the monad operations (together with a ‘zero’) [Wadler 1992]:

$$\lfloor e \mid \rfloor = \eta \ e \quad (9a)$$

$$\lfloor e \mid b \rfloor = \text{if } b \ \text{then } \eta \ e \ \text{else } \emptyset \quad (9b)$$

$$\lfloor e \mid p \leftarrow m \rfloor = \text{Bag } (\lambda p \rightarrow e) \ m \quad (9c)$$

$$\lfloor e \mid q, q' \rfloor = \mu \ \lfloor \lfloor e \mid q' \rfloor \mid q \rfloor \quad (9d)$$

In particular, one can show by induction that the functorial action of Bag distributes over the comprehension syntax:

$$\text{Bag } f \ \lfloor e \mid q \rfloor = \lfloor f \ e \mid q \rfloor \quad (10)$$

Cartesian product. The unit type 1 and the product bifunctor $\times$ turn the category of sets into a so-called symmetric monoidal category: 1 is the neutral element of $\times$ up to a natural isomorphism, and $\times$ is associative and commutative up to natural isomorphisms. These two structures arise from the two adjunctions (Figure 2(a)(b)).

$$\text{Set}^0 \xrightleftharpoons[\text{1}]{\Delta} \text{Set} \qquad \text{Set}^2 \xrightleftharpoons[\times]{\Delta} \text{Set}$$

Very briefly, Δ is overloaded to denote the *diagonal* functor $\text{Set} \rightarrow \text{Set}^k$ for any k , taking each object A and arrow f to k copies of the same; arrows in the one-object, one-arrow category Set^0

are in 1-to-1 correspondence with functions to a singleton set (there is precisely one of each), and pairs of functions with a common domain are in 1-to-1 correspondence with functions constructing pairs.

The functor Bag has *strength*, a broadcast operation of type $A \times \text{Bag } B \rightarrow \text{Bag } (A \times B)$. Moreover, Bag can be extended to a *lax monoidal* functor (‘lax’ because *unit* and $\times$ are one-way mappings rather than isomorphisms):

$$\begin{aligned} \text{unit} &: \text{Bag } 1 \dashrightarrow \text{really } 1 \rightarrow \text{Bag } 1 \\ \text{unit} &= \ulcorner () \urcorner \\ (\times) &: \text{Bag } A \times \text{Bag } B \rightarrow \text{Bag } (A \times B) \\ s \times t &= \ulcorner (a, b) \mid a \leftarrow s, b \leftarrow t \urcorner \end{aligned}$$

The unit bag and cartesian product satisfy:

$$s \times \text{unit} \cong s \cong \text{unit} \times s \quad (11a)$$

$$s \times (t \times u) \cong (s \times t) \times u \quad (11b)$$

$$s \times t \cong t \times s \quad (11c)$$

(the notation is a bit loose: $a \cong b$ denotes equality modulo the respective canonical isomorphism α , that is, $\text{Bag } \alpha a = b$), so we also say that Bag is a *commutative* or *symmetric* monoidal monad; in particular,

$$\ulcorner (a, b) \mid a \leftarrow s, b \leftarrow t \urcorner = \ulcorner (a, b) \mid b \leftarrow t, a \leftarrow s \urcorner \quad (12)$$

A lax-monoidal functor with strength is also known as an *applicative functor* or *idiom* [McBride and Paterson 2008].

We have introduced the cartesian product of bags mainly for completeness. It is expensive to compute, so should generally be avoided in favour of joins. To do so is difficult with mere bags of values, which are a rather naive model of database tables; to implement joins efficiently, we require *indexed* tables. These are modelled by finite maps, which are based on the notion of a pointed set, to which we turn next.

4 POINTED SETS

A *pointed set* (A, a) is a set A with a distinguished point or base element $a \in A$. We may write null_A or simply *null* for the point of A . Let (A, null_A) and (B, null_B) be pointed sets; a total function $f : A \rightarrow B$ is *point-preserving* if $f \text{ null}_A = \text{null}_B$. Pointed sets and point-preserving functions form the category Set_* , which is equivalent to the category of sets and partial functions.

Let A be a set, and let (B, null_B) be a pointed set. The *domain* of the function $f : A \rightarrow (B, \text{null}_B)$ is the set of all arguments that are mapped to non- null_B values:

$$\text{dom } f = \{a \in A \mid f a \neq \text{null}_B\} \quad (13)$$

(In general, a function between pointed sets will have an infinite domain; in Section 4 we focus on finite maps, which by definition have finite domain.)

To add a base point to an arbitrary set, we can use lifting (called *Maybe* in Haskell, and *option* in ML). Indeed, the *Maybe* monad $(1 + -)$ arises as the composition $U \circ L$, where

$$\text{Set}_* \xrightleftharpoons[U]{L} \text{Set}$$

is the adjunction between the category of pointed sets and point-preserving functions and the category of sets and total functions (Figure 2(c)). Very briefly, the forgetful functor U maps (A, a) to A and a point-preserving function to its underlying function. Its left adjoint L lifts a set A

to $(\text{Maybe } A, \text{Nothing})$ and lifts a function on A to the obvious point-preserving function. The adjunction $L \dashv U$ tells us a lot about the structure of $\mathbf{Set}_*$: since left adjoints preserve initial objects, $0 \cong L 0 = (\text{Maybe } 0, \text{Nothing}) \cong (1, ())$, the pointed set $(1, ())$ is initial in $\mathbf{Set}_*$; it is also final, as right adjoints preserve final objects. Right adjoints furthermore preserve products, so

$$\begin{aligned} 0 &\cong (1, ()) \\ 1 &\cong (1, ()) \\ (A, \text{null}_A) \times (B, \text{null}_B) &\cong (A \times B, (\text{null}_A, \text{null}_B)) \end{aligned}$$

Some types have an obvious point. For example, the empty bag is a canonical choice for bags, and more generally the neutral element for any monoid. Thus we can turn Bag into a functor $\text{Bag}_* : \mathbf{Set} \rightarrow \mathbf{Set}_*$:

$$\begin{aligned} \text{Bag}_* A &= (\text{Bag } A, \emptyset) \\ \text{Bag}_* f &= \text{Bag } f \end{aligned}$$

Property (3a) shows that $\text{Bag } f$ is point-preserving. Indeed, almost all the bag operations preserve points: union $\uplus$ (1a); projection $\text{Bag } f$ (3a); selection *filter* p (7a); aggregation *reduce* M (5a), assuming that the neutral element of a monoid is its point; and cartesian product $\times$, since $\emptyset \times \emptyset = \emptyset$. The only notable exceptions are *single* and hence *unit*. This implies that Bag_* is not a monoidal functor, only a so-called *semi-monoidal* functor between the semi-monoidal categories $(\mathbf{Set}, \times)$ and $(\mathbf{Set}_*, \times)$ ('semi-monoidal' meaning lacking the unit)—this will be important later.

Finite Maps. Let K be a set, and let V be a pointed set. A *map* of type $\text{Map } K \ V$ is a total function from K to V ; a *finite map* of type $\text{Map}_* K \ V$ is a map that yields null_V for all but finitely many arguments. Finite maps are also known as key-value maps, association lists, and dictionaries; they are more appropriate for modelling databases than possibly infinite maps are, because the latter cannot in general be aggregated.

$$\text{Map}_* K \ V = \{ s : K \rightarrow V \mid \text{dom } s \text{ is finite} \}$$

We can turn $\text{Map}_* K$ into an endofunctor over the category $\mathbf{Set}_*$ of pointed sets and point-preserving functions.

$$\begin{aligned} \text{Map}_* K \ V &= (\text{Map } K \ V, \lambda k \rightarrow \text{null}_V) \\ \text{Map}_* K \ f &= \lambda s \rightarrow f \cdot s \end{aligned}$$

The point is the empty map, which sends all its arguments to null_V . The functorial action on arrows is just post-composition, which is point-preserving when f is, as it sends $\lambda k \rightarrow \text{null}_V$ to $\lambda k \rightarrow \text{null}_W$ for $f : V \rightarrow W$. We record that $\text{Map}_* K$ satisfies the functor laws:

$$\text{Map}_* K \ \text{id} = \text{id} \tag{14a}$$

$$\text{Map}_* K \ (g \cdot f) = \text{Map}_* K \ g \cdot \text{Map}_* K \ f \tag{14b}$$

It is important to note that $\text{Map}_* K \ f$ does not necessarily preserve the domain of its argument map: it acts as a *filter* if f sends any non-*null* values to *null*. For example,

$$\text{Map}_* K \ (\lambda v \rightarrow \text{if bad } v \text{ then null else } v) \ s$$

filters out the 'bad' values in s . The codomain of a finite map is defined

$$\begin{aligned} \text{cod} : \text{Map}_* K \ V &\rightarrow \text{Bag}_* V \\ \text{cod } s &= \{ s \ k \mid k \leftarrow \text{dom } s \} \end{aligned}$$

(Note that *cod* yields a finite bag, whereas *dom* returns a set.)

Laws of exponents. The finite map functor satisfies a number of properties which are commonly known as the ‘laws of exponents’ (write $\text{Map}_* K V$ as V^K to see why).

$$\text{Map}_* 0 V \cong 1 \quad (15a)$$

$$\text{Map}_* 1 V \cong V \quad (15b)$$

$$\text{Map}_* (K_1 + K_2) V \cong \text{Map}_* K_1 V \times \text{Map}_* K_2 V : (\nabla) \quad (15c)$$

$$\text{curry} : \text{Map}_* (K_1 \times K_2) V \cong \text{Map}_* K_1 (\text{Map}_* K_2 V) : \text{curry}^\circ \quad (15d)$$

The isomorphisms are natural in the type V of values. For unrestricted functions, they are consequences of the currying adjunction $(- \times P) \dashv (-)^P$ (Figure 2(d)). Property (15d) follows directly from this adjunction, and Property (15b) follows from currying and $A \times 1 \cong A$, the fact that product is monoidal.

Properties (15a) and (15c) are also a consequence of currying, but indirectly: this adjunction can be used to form yet another one, namely that $(X^{(-)})^{\text{op}} \dashv X^{(-)}$, which swaps the arguments of a binary function. Since $X^{(-)}$ is right adjoint and contravariant, it takes 0 to 1 and $+$ to $\times$, and the properties follow.

All of these properties inform the way we can use maps in practice. For example, the penultimate isomorphism (15c) expresses that two functions with the same target type can be represented by a single function from the sum (coproduct) of their source types (‘case analysis’).

$$s \nabla t = \lambda k \rightarrow \text{case } k \text{ of } \{ \text{Inl } a \rightarrow s \ a; \text{Inr } b \rightarrow t \ b \}$$

The last isomorphism (15d) captures ‘currying’: a function of two arguments can be treated as a function of the first argument whose values are functions of the second argument.

$$\text{curry}^\circ s = \lambda(k_1, k_2) \rightarrow (s \ k_1) \ k_2$$

We have to show that the isomorphisms are still valid when restricted to finite maps. For (15a) and (15b) this is trivial. The domain of $s \nabla t$ and $\text{curry}^\circ s$ are given by

$$\text{dom } (s \nabla t) = \{ \text{Inl } a \mid a \leftarrow \text{dom } s \} \cup \{ \text{Inr } b \mid b \leftarrow \text{dom } t \}$$

$$\text{dom } (\text{curry}^\circ s) = \{ (k_1, k_2) \mid k_1 \leftarrow \text{dom } s, k_2 \leftarrow \text{dom } (s \ k_1) \}$$

Clearly, $\text{dom } (s \nabla t)$ is finite if and only if both $\text{dom } s$ and $\text{dom } t$ are [Connelly and Morris 1995, Proposition 1.1]. Likewise, $\text{curry}^\circ s$ is a finite map if and only if s and each partial application of s is finite [Connelly and Morris 1995, Proposition 1.3]. It remains to verify that the isomorphisms are also point-preserving. For (15a) and (15b) this is again immediate. For (15c) this follows from $\text{null} \cdot \text{Inl} = \text{null}$ and $\text{null} \cdot \text{Inr} = \text{null}$. Finally, for (15d) we calculate

$$\begin{aligned} & \text{curry}^\circ \text{null} \\ &= \{ \text{definition of } \text{curry}^\circ \} \\ & \lambda(k_1, k_2) \rightarrow (\text{null } k_1) \ k_2 \\ &= \{ \text{definition of } \text{null} \text{ for finite maps} \} \\ & \lambda(k_1, k_2) \rightarrow \text{null } k_2 \\ &= \{ \text{definition of } \text{null} \text{ for finite maps} \} \\ & \lambda(k_1, k_2) \rightarrow \text{null} \\ &= \{ \text{definition of } \text{null} \text{ for finite maps} \} \\ & \text{null} \end{aligned}$$

These isomorphisms are important, as they can be used to implement finite maps generically for finite sums and finite products [Hinze 2000]. Since the isomorphisms are effectively turned into equalities, the mappings back and forth are constant-time operations.

Union. The ‘laws of exponents’ come in two flavours: the laws above vary the key type and keep the value type fixed; the laws below fix the key type and vary the value type—the unit type 1 is clearly pointed.

$$\text{empty} : 1 \cong \text{Map}_* K 1 \quad (15e)$$

$$\text{merge} : \text{Map}_* K V_1 \times \text{Map}_* K V_2 \cong \text{Map}_* K (V_1 \times V_2) \quad (15f)$$

The last isomorphism is natural in V_1 and V_2 . It implements the merge or union of two finite maps.

$$\begin{aligned} \text{empty} &: \text{Map}_* K 1 \quad \text{-- really } 1 \rightarrow \text{Map}_* K 1 \\ \text{empty} &= \lambda k \rightarrow () \\ \text{merge} &: \text{Map}_* K V_1 \times \text{Map}_* K V_2 \rightarrow \text{Map}_* K (V_1 \times V_2) \\ \text{merge } (s, t) &= \lambda k \rightarrow (s \, k, t \, k) \end{aligned}$$

Again we have to check that the isomorphisms are still valid when restricted to finite maps.

$$\begin{aligned} \text{dom empty} &= \{ \} \\ \text{dom } (\text{merge } (s, t)) &= \text{dom } s \cup \text{dom } t \end{aligned}$$

The unique arrow to 1 is certainly finite ($\text{empty} = \text{null}$). Clearly, $\text{merge } (s, t)$ is a finite map if (and indeed, only if) both s and t are. Furthermore, the isomorphisms are point-preserving. For (15e) this is trivial. For (15f) this follows from $\text{fst} \cdot \text{null} = \text{null}$ and $\text{snd} \cdot \text{null} = \text{null}$.

These two definitions turn $\text{Map}_* K$ into a strong monoidal functor over Set_* . (We use the ‘strong’ in ‘strong monoidal functor’ to indicate that the maps $1 \rightarrow F \, 1$ for unit and $F \, A \times F \, B \rightarrow F \, (A \times B)$ for product are isomorphisms. To minimize confusion, we use ‘lax monoidal functor with strength’ to indicate the existence of tensorial strength, a broadcast operation of type $A \times F \, B \rightarrow F \, (A \times B)$, and avoid the ambiguous term ‘strong lax monoidal functor’.) The empty finite map and the merge operation satisfy:

$$\text{merge } (s, \text{empty}) \cong s \cong \text{merge } (\text{empty}, s) \quad (16a)$$

$$\text{merge } (s, \text{merge } (t, u)) \cong \text{merge } (\text{merge } (s, t), u) \quad (16b)$$

$$\text{merge } (s, t) \cong \text{merge } (t, s) \quad (16c)$$

(again, $\cong$ means equality modulo the respective canonical isomorphism). However, $\text{Map}_* K$ is *not* an applicative functor, because it does not support a strength operation of type $A \times \text{Map}_* K \, B \rightarrow \text{Map}_* K \, (A \times B)$: the resulting map would in general not be finite. In particular, we cannot define the *pure* operation of applicative functors in terms of the *unit*, ie $\text{pure } a = \text{Map}_* K \, (\lambda() \rightarrow a) \, \text{unit}$, because $\lambda() \rightarrow a$ is not point-preserving.

5 INDEXED TABLES

An indexed table is a finite map from keys (indices) to finite bags of values.

$$\text{Table } K \, V = \text{Map}_* K (\text{Bag}_* V)$$

Note that $\text{Table } K$ is just the composition of $\text{Map}_* K$ and Bag_* . Since these are both semi-monoidal functors, and since semi-monoidal functors compose, $\text{Table } K : \text{Set}_* \rightarrow \text{Set}_*$ is a semi-monoidal functor too.

Figure 4 summarizes the operations on tables; full definitions follow. We use the monoidal structure of $\text{Map}_* K$ to lift operations on bags to operations on tables: the lifted version of the unary bag function f is $\text{Map}_* K \, f$; the lifted version of the binary bag function g is $\text{Map}_* K \, g \cdot \text{merge}$.

K-indexed table of V values	Table $K \ V$
empty table	<i>empty</i>
singleton table (k, v)	$k \mapsto \{v\}$
union of tables	$\text{Map}_* K (\uplus) \cdot \text{merge}$
projection π_f	$\text{Map}_* K (\text{Bag}_* f)$
selection σ_p	$\text{Map}_* K (\text{filter } p)$
aggregation in monoid M	$\text{Map}_* K (\text{reduce } M)$
natural join	$\text{Map}_* K (\times) \cdot \text{merge}$

Fig. 4. Relational algebra operators, implemented in terms of indexed tables

Indexing. Relations are in 1-to-1 correspondence with set-valued functions. This correspondence is captured by the adjunction

$$\mathbf{Rel} \xleftarrow[\text{E}]{\text{J}} \mathbf{Set}$$

between the category of sets and relations and the category of sets and total functions (Figure 2(g)). Very briefly, the functor J embeds $\mathbf{Set}$ into $\mathbf{Rel}$, viewing functions as relations. Its right adjoint is the existential image functor E , which sends a set to its powerset and a relation R to the function $\lambda X \rightarrow \{b \in B \mid \exists a \in X. a R b\}$. Now, the 1-to-1 correspondence $\mathbf{Set} (K \times V) \cong \mathbf{Map} K (\mathbf{Set} V)$ remains valid if restricted to finite sets and finite maps, and furthermore if finite sets are replaced by finite bags:

$$ix : \text{Bag}_* (K \times V) \cong \text{Map}_* K (\text{Bag}_* V) : ix^\circ \quad (17)$$

The isomorphism is natural in V . The operation ix can be seen as providing a view or an indexing structure on a table (given as a bag of key-value pairs). It is defined

$$ix \ s = \lambda k \rightarrow \{v \mid (k', v) \leftarrow s, k = k'\}$$

The domain of $ix \ s$ is given by

$$\text{dom} (ix \ s) = \{k \mid (k, v) \leftarrow s\}$$

Clearly, ix sends a finite bag to a finite map, which in turn maps each key to a finite bag. The isomorphism is point-preserving, as the empty bag is sent to the empty map. With care and an appropriate choice of representation, ix can be made to run in linear time [Henglein and Hinze 2013].

A useful derived isomorphism is relational currying:

$$\text{Table} (K_1 \times K_2) \ V \cong \text{Table} K_1 (K_2 \times V) \quad (18)$$

which allows us to shift a key into and out of an index. The proof is straightforward:

$$\begin{aligned} & \text{Table} (K_1 \times K_2) \ V \\ = & \{ \text{definition of Table} \} \\ & \text{Map}_* (K_1 \times K_2) (\text{Bag}_* V) \\ \cong & \{ \text{currying (15d)} \} \\ & \text{Map}_* K_1 (\text{Map}_* K_2 (\text{Bag}_* V)) \\ \cong & \{ \text{indexing (17)} \} \\ & \text{Map}_* K_1 (\text{Bag}_* (K_2 \times V)) \\ = & \{ \text{definition of Table} \} \\ & \text{Table} K_1 (K_2 \times V) \end{aligned}$$

Consequently, the isomorphism is given by the composition $\text{Map}_* K_1 ix^\circ \cdot \text{curry}$. As an aside, equation (18) is related to the adjunction

$$\mathbf{Rel} \begin{array}{c} \xleftarrow{-\times P} \\ \perp \\ \xrightarrow{P \times -} \end{array} \mathbf{Rel}$$

(Figure 2(e)), capturing the fact that the relation types $(A \times P) \leftrightarrow B$ and $A \leftrightarrow (P \times B)$ are both essentially the powerset of $A \times P \times B$; this shows that $\mathbf{Rel}$ is a closed monoidal category—but note that $\times$ is not the categorical product in $\mathbf{Rel}$, and that $\mathbf{Rel}$ is not cartesian closed.

6 EQUIJOIN BY INDEXING

The ‘index by’ or ‘group by’ operation

$$\begin{aligned} \text{indexBy} &: \text{Bag}_* V \times (V \rightarrow K) \rightarrow \text{Map}_* K (\text{Bag}_* V) \\ s \text{ 'indexBy' } f &= ix (\text{Bag}_* (f \triangle id) s) \end{aligned}$$

is the crux to an efficient implementation of joins: the equijoin

$$\begin{aligned} (\bowtie) &:: Eq\ k \Rightarrow (a \rightarrow k) \rightarrow (b \rightarrow k) \rightarrow \text{Bag}\ a \rightarrow \text{Bag}\ b \rightarrow \text{Bag}\ (a, b) \\ x \text{ } f \bowtie_g y &= \wr (a, b) \mid a \leftarrow x, b \leftarrow y, f\ a == g\ b \wr \end{aligned}$$

can be computed by indexing x and y by their keys, merging these two indexed tables, performing small cartesian products at each key—thus far, a full outer-join—then extracting the elements from the table:

$$x \text{ } f \bowtie_g y = \text{elems} (\text{Map}_* (\times) (\text{merge} (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g)))$$

where

$$\begin{aligned} \text{elems} &: \text{Table}\ K\ V \rightarrow \text{Bag}_* V \\ \text{elems} &= \mu \cdot \text{cod} = \text{Bag}_* \text{snd} \cdot ix^\circ \end{aligned}$$

Of course, a properly engineered database engine should not recompute an index each time a table is accessed: it should memoise the index, and maintain the cached index as the table changes.

In this section, we prove the equality

$$\begin{aligned} &\wr (a, b) \mid a \leftarrow x, b \leftarrow y, f\ a == g\ b \wr \\ &= \text{elems} (\text{Map}_* K (\times) (\text{merge} (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g))) \end{aligned}$$

demonstrating the correctness of the grouping implementation of the equijoin $x \text{ } f \bowtie_g y$. For convenience, we recall the types

$$\begin{aligned} \text{Map}_* K\ V &= K \rightarrow V \\ \text{Table}\ K\ V &= \text{Map}_* K (\text{Bag}_* V) \end{aligned}$$

(Bag_* and Map_* specifically denote *finite* bags and maps, but finiteness is not relevant for this proof), and functions

$$\begin{aligned} (\times) &: \text{Bag}_* A \times \text{Bag}_* B \rightarrow \text{Bag}_* (A \times B) \\ \text{filter} &: (A \rightarrow \mathbb{B}) \rightarrow (\text{Bag}_* A \rightarrow \text{Bag}_* A) \\ \text{indexBy} &: \text{Bag}_* V \times (V \rightarrow K) \rightarrow \text{Table}\ K\ V \\ \text{merge} &: \text{Map}_* K\ V_1 \times \text{Map}_* K\ V_2 \rightarrow \text{Map}_* K\ (V_1 \times V_2) \\ \text{elems} &: \text{Table}\ K\ V \rightarrow \text{Bag}_* V \end{aligned}$$

We use the following lemmas, each of which is easy to prove:

- (1) elems is a post-inverse of indexBy :

$$\text{elems} (x \text{ 'indexBy' } f) = x$$

(2) lookup in an indexed table is a filter:

$$(x \text{ 'indexBy' } f) k = \text{filter } (\lambda a \rightarrow k == f a) x$$

(3) independent filters promote through cartesian product:

$$\text{filter } (\lambda(a, b) \rightarrow p a \wedge q b) (x \times y) = \text{filter } p x \times \text{filter } q y$$

(4) lookup in merged indices is a pair of lookups:

$$\text{merge } (x, y) k = (x k, y k)$$

(5) functoriality of $\text{Map}_* K$:

$$\text{Map}_* K g f = g \cdot f$$

We have

$$\begin{aligned} & \{ (a, b) \mid a \leftarrow x, b \leftarrow y, f a == g b \} \\ &= \text{elems } (\text{Map}_* K (\times) (\text{merge } (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g))) \\ \iff & \{ \text{elems and indexBy (1)} \} \\ & \text{elems } (\{ (a, b) \mid a \leftarrow x, b \leftarrow y, f a == g b \} \text{ 'indexBy' } (f \cdot \text{fst})) \\ &= \text{elems } (\text{Map}_* K (\times) (\text{merge } (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g))) \\ \iff & \{ \text{Leibniz} \} \\ & \{ (a, b) \mid a \leftarrow x, b \leftarrow y, f a == g b \} \text{ 'indexBy' } (f \cdot \text{fst}) \\ &= \text{Map}_* K (\times) (\text{merge } (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g)) \\ \iff & \{ \text{function equality: } (f = g) \iff (\forall a. f a = g a) \} \\ & \forall k. (\{ (a, b) \mid a \leftarrow x, b \leftarrow y, f a == g b \} \text{ 'indexBy' } (f \cdot \text{fst})) k \\ &= (\text{Map}_* K (\times) (\text{merge } (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g))) k \end{aligned}$$

so it suffices to prove the latter equation. We have, for arbitrary k :

$$\begin{aligned} & (\{ (a, b) \mid a \leftarrow x, b \leftarrow y, f a == g b \} \text{ 'indexBy' } (f \cdot \text{fst})) k \\ &= \{ \text{lookup in table (2)} \} \\ & \text{filter } (\lambda(a, b) \rightarrow f a == k) \{ (a, b) \mid a \leftarrow x, b \leftarrow y, f a == g b \} \\ &= \{ \text{comprehensions} \} \\ & \text{filter } (\lambda(a, b) \rightarrow f a == k) (\text{filter } (\lambda(a, b) \rightarrow f a == g b) (x \times y)) \\ &= \{ \text{filter composition} \} \\ & \text{filter } (\lambda(a, b) \rightarrow (f a == k) \wedge (f a == g b)) (x \times y) \\ &= \{ \text{equality an equivalence relation} \} \\ & \text{filter } (\lambda(a, b) \rightarrow (f a == k) \wedge (g b == k)) (x \times y) \\ &= \{ \text{independent filters (3)} \} \\ & \text{filter } (\lambda a \rightarrow f a == k) x \times \text{filter } (\lambda b \rightarrow g b == k) y \\ &= \{ \text{lookup in table (2)} \} \\ & (x \text{ 'indexBy' } f) k \times (y \text{ 'indexBy' } g) k \\ &= \{ \text{lookup in merged indices (4)} \} \\ & (\times) (\text{merge } (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g) k) \\ &= \{ \text{functoriality of maps (5)} \} \\ & (\text{Map}_* K (\times) (\text{merge } (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g))) k \end{aligned}$$

As it stands, this implementation is efficient up to a point. Provided that the predicates used for selection, the functions used for projection, and the equality comparisons used for grouping and joining all take constant time, it takes time linear in the number of tuples in the input to compute a query involving selection, projection, and join. The result is a bag of pairs of bags, as in Section 1; it is important to note that one cannot in general multiply out the pairs of bags to make one big bag of pairs in linear time, because the resulting bag may have more than linearly many tuples.

When the predicates and comparisons are more expensive, such as on unbounded strings, then one needs something like discrimination-based techniques [Henglein and Larsen 2010] in order to retain the linear time guarantee.

7 GRADED MONADS

Let us take stock of where we have got to. We model database tables by some collection type. It is important for that collection type to form a monad, because only monads support the convenient comprehension syntax for queries; we started off simply with $\text{Bag } V$. We elaborated this to tables $\text{Map } K (\text{Bag}_* V)$, in order to support grouping and indexing for efficient joins. We had to restrict these to finite tables $\text{Table } K V = \text{Map}_* K (\text{Bag}_* V)$, in order also to accommodate aggregation. But sadly, although maps and tables form monads, *finite* maps and tables do not, for the same reason that they are only semi-monoidal and not fully monoidal: the units for the map and table monads, of type $V \rightarrow \text{Map } K V$ and $V \rightarrow \text{Map } K (\text{Bag}_* V)$ respectively, yield constant functions that represent *infinite* tables (at least, when the key type K is unbounded).

Fortunately, there is a way to enjoy both finiteness and the monadic structure. This is to use *graded monads* [Fujii et al. 2016], also sometimes called ‘indexed monads’ [Orchard et al. 2014] or ‘parametric monads’ [Katsumata 2014]—that is, to consider $\text{Table } K$ as a family of endofunctors for various K that collectively act like a monad, rather than actually being a single monad for all or for a fixed K .

Formally, an M -graded monad (T, η, μ) over a monoid $(M, \epsilon, \otimes)$ is an indexed family of endofunctors T_m for $m : M$, with natural transformations $\eta : \text{Id} \rightarrow T_\epsilon$ and $\mu_{m,n} : T_m \circ T_n \rightarrow T_{m \otimes n}$ satisfying indexed versions of the three coherence equations of ordinary monads:

$$\mu_{m,\epsilon} \cdot T_m \eta = \text{id} \quad (19a)$$

$$\mu_{\epsilon,n} \cdot \eta T_n = \text{id} \quad (19b)$$

$$\mu_{m,n \otimes p} \cdot T_m \mu_{n,p} = \mu_{m \otimes n,p} \cdot \mu_{m,n} T_p \quad (19c)$$

A familiar example involves the monoid $(\mathbb{N}, 1, \times)$ of naturals under multiplication, and the family Vector_n of functors representing vectors of length n ; the unit $\eta : A \rightarrow \text{Vector}_1 A$ constructs a singleton vector, and the multiplication $\mu_{m,n} : \text{Vector}_m (\text{Vector}_n A) \rightarrow \text{Vector}_{m \times n} A$ flattens a rectangular vector of vectors into one long vector of elements.

We want the related monoid $(\mathbb{K}^*, \langle \rangle, ;)$ of finite sequences of finite key types $\mathbb{K}$, under the empty sequence $\langle \rangle$ and sequence concatenation $(;)$. (We could instead take finite bags of key types, if we wanted to consider the column order of a table not to be significant.) To be concrete, let $\Pi : \mathbb{K}^* \rightarrow \mathbb{K}$ denote the cartesian product ΠKs of a finite sequence $Ks : \mathbb{K}^*$ of finite key types. Then (T, η, μ) is a monad graded over $(\mathbb{K}^*, \langle \rangle, ;)$, where the carrier $T_{Ks} V = \text{Table } (\Pi Ks) V$ is finite maps from Ks to V , the unit $\eta : V \rightarrow T_{\langle \rangle} V$ constructs a singleton finite map on the empty sequence of key types, and $\mu_{Ks, Ks'} : T_{Ks} (T_{Ks'} V) \rightarrow T_{Ks; Ks'} V$ is relational uncurrying (18).

How do graded monads fit into the adjunction story? At first sight, it does not look promising. An ordinary monad is necessarily an endofunctor on a category $\mathcal{C}$, arising as it does from an adjunction between two categories. An M -graded monad, on the other hand, is a functor $T : \mathcal{M} \times \mathcal{C} \rightarrow \mathcal{C}$ (viewing monoid M as a discrete category $\mathcal{M}$), or equivalently $T : \mathcal{M} \rightarrow \mathcal{C}^{\mathcal{C}}$; neither of these is an endofunctor, so neither can arise directly from an adjunction.

Nevertheless, there is an adjunction story to graded monads, albeit under a layer of disguise. A *strict action* of monoid $M = (M, \epsilon, \otimes)$ on category $\mathcal{D}$ is a family A_m of endofunctors on $\mathcal{D}$ for $m : M$, such that $A_\epsilon = \text{Id}$ and $A_{m \otimes n} = A_m \circ A_n$; it is a special case of an M -graded monad in which the η and μ are identities. Then an M -graded monad in category $\mathcal{C}$ factorises into an adjunction $L \dashv R : \mathcal{C} \rightarrow \mathcal{D}$ and a strict M -action on $\mathcal{D}$ [Fujii et al. 2016; Street 1972]. Indeed, each graded

```

[e | ]           = return e
[e | b]          = if b then return e else mzero
[e | p ← m]      = liftM (λp → e) m
[e | q, r]        = mult [[e | r] | q]
[e | (q | r), s] = liftM (λ(vq, vr) → [e | s]) (mzip [vq | q] [vr | r])
[e | q, then group by b using f, r]
                 = liftM (λys → case (fmap v1q ys, ..., fmap vnq ys) of vq → [e | r])
                   (f (λvq → b) [vq | q])

```

Fig. 5. Definition by translation of Haskell's generalised comprehensions

monad gives rise to a whole family of such factorisations, with two canonical such—a ‘least’ one and a ‘greatest’ one. Conversely, any strict M -action on a category $\mathcal{D}$ can be transferred over an adjunction $L \dashv R : \mathcal{C} \rightarrow \mathcal{D}$ to yield an M -graded monad on $\mathcal{C}$. An ordinary monad T arising from adjunction $L \dashv R$ is a graded monad over the trivial singleton monoid; in that case, the canonical factorisations reduce to the Kleisli and Eilenberg–Moore adjunctions for the monad, and conversely, T coincides with the transfer of the identity monad (the trivial strict action of the trivial singleton monoid) over $L \dashv R$. The constructions are given in full by [Fujii et al. \[2016\]](#); the details do not add any more insight beyond the above, so we skip them here, although for completeness we summarize them in the online appendix.

In Section 6 we had the characterisation

$$x \text{ f}\bowtie_g y = \{ (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a = g \ b \} \\ = \text{elems} (\text{Map}_* (\times) (\text{merge} (x \text{ 'indexBy' } f, y \text{ 'indexBy' } g)))$$

of equijoin in terms of indexed tables. The linear execution time of the latter combinator version is appealing; but it is a shame to have had to give up on the comprehension notation of the former in order to obtain it. On the face of it, this sacrifice seems inevitable: the combinator version is heterogeneous, involving different monads Bag_* and Table K as the input and output of *indexBy*, whereas traditional comprehension notation is homogeneous.

Nevertheless, it is possible to generalise the comprehension notation to support heterogeneity with grouping [[Peyton Jones and Wadler 2007](#)] and merging [[Plasmeijer and van Eekelen 1995](#)], as has been done in the Glasgow Haskell Compiler [[Giorgidze et al. 2011b](#); [Glasgow 2015](#)]. Haskell's generalised comprehension for a monad M is defined by translation, as shown in Figure 5. In a nutshell, the translation is as follows. The first four clauses precisely match our properties of bag comprehensions (9a)–(9d) in Section 3: *mult* and *return* are the multiplication μ and unit η of the monad M , *liftM* its functorial action, and *mzero* the neutral element of the monoid—in particular, *mzero* = $\emptyset$ for $M = \text{Bag}_*$.

In the fifth and sixth clauses, v^q denotes the tuple of variables bound by qualifiers q , and v_i^q projects out the i th component from this tuple; and similarly for v^r and r . In the fifth clause, the monad M is assumed to support an appropriate ‘zip’ function

$$\text{mzip} :: M \ a \rightarrow M \ b \rightarrow M \ (a, b)$$

For example, with $M = \text{List}$, a suitable instantiation is *mzip* = *zip*; so the comprehension

$$[a + b + c \mid (a, b) \leftarrow [(1, 3), (2, 4)] \mid c \leftarrow [5, 6]]$$

desugars to

$$\text{liftM} (\lambda((a, b), c) \rightarrow a + b + c) (\text{zip} [(a, b) \mid (a, b) \leftarrow [(1, 3), (2, 4)]] [c \mid c \leftarrow [5, 6]])$$

and hence evaluates to $[9, 12]$. More interestingly, with $M = \text{Table K}$, an appropriate instantiation is $mzip = \text{merge}$.

In the sixth clause, the argument f is required to have type $(a \rightarrow K) \rightarrow M a \rightarrow N (F a)$, where N is a (possibly different) monad, and F some functor; note that this introduces heterogeneity into the comprehension. This translation [Peyton Jones and Wadler 2007] is particularly ingenious, in that the variables v^q bound by qualifiers q and used in key expression b are rebound by the **case** expression for the scope containing e and r . For example, GHC provides a function

$$\text{groupWith} :: \text{Ord } b \Rightarrow (a \rightarrow b) \rightarrow [a] \rightarrow [[a]]$$

such that $\text{groupWith } k \text{ } xs$ which categorizes the elements of a list xs according to k , groups them by category, and sorts the groups; then the comprehension

$$[(\text{head } c, \text{sum } n) \mid (c, n) \leftarrow [(\text{'a'}, 3), (\text{'b'}, 4), (\text{'a'}, 5)], \\ \text{then group by } c \text{ using groupWith}]$$

desugars to

$$\text{liftM } (\lambda ys \rightarrow \text{case } (\text{fmap fst } ys, \text{fmap snd } ys) \text{ of } (c, n) \rightarrow [(\text{head } c, \text{sum } n) \mid]) \\ (\text{groupWith } (\lambda(c, n) \rightarrow c) [(c, n) \mid (c, n) \leftarrow [(\text{'a'}, 3), (\text{'b'}, 4), (\text{'a'}, 5)]])$$

and hence evaluates to $[(\text{'a'}, 8), (\text{'b'}, 4)]$ (note that c and n are bound initially to characters and integers, and then rebound to lists of such). For the purposes of this paper, we will use indexBy for f , thereby instantiating both M and F to Bag_* and N to Table K .

(In GHC [Glasgow 2015], the choice of $mzip$ function is specified indirectly by way of a subclass MonadZip of Monad , whereas the grouping function f is specified directly.) The full translation of the extended comprehension notation is discussed in more detail by its designers [Giorgidze et al. 2011a,b; Peyton Jones and Wadler 2007] and, in this particular context, in a related paper [Gibbons 2016]. Others have designed embedded query DSLs with grouping and aggregation [Suzuki et al. 2016], but not as far as we know in terms of comprehensions.

With these extensions, we can write the equijoin $x \bowtie_g y$ in Haskell as follows:

$$\text{elems } [\text{cp } a \text{ } b \mid a \leftarrow x, \text{then group by } f \text{ using indexBy} \\ \mid b \leftarrow y, \text{then group by } g \text{ using indexBy}]$$

—note how the variable a is bound first to an element of x by the qualifier $a \leftarrow x$, and used as such in the key expression f ; but rebound to a finite bag of such elements by the grouping function indexBy , and used as such in the term $\text{cp } a \text{ } b$. The example query from Section 1 turns out as follows:

$$\text{reduce} \\ [\text{cp name amount} \\ \mid (\text{cid}, \text{name}) \leftarrow \text{customers}, \text{then group by cid using indexBy} \\ \mid (\text{iid}, \text{cust}, \text{due}, \text{amount}) \leftarrow \text{invoices}, \text{due} < \text{today}, \text{then group by cust using indexBy} \\]$$

So graded monads are really a small generalisation of ordinary monads, and enjoy generalisations of the same properties. Crucially, they still support the same comprehension notation, satisfying the same equational properties. In particular, the translation of grouping and zipping comprehensions in Figure 5 works as well for graded monads as for ungraded ones, and for heterogeneous comprehensions (where grouping takes an $M A$ -collection to a nested $N (F A)$ collection) as for homogeneous ones (where M , N , and F coincide). We do not repeat and renumber all the previous comprehension properties here; but it is important to note that subsequent appeals to properties

about ‘comprehensions’ (for example, in Section 8 below) depend implicitly on their definition for graded monads.

8 QUERY TRANSFORMATION

The classical System R approach to database query optimization [Selinger et al. 1979] is based on enumerating a number of extensionally equivalent expressions for an input query, estimating the execution cost of each of them, and then generating code for the best one. Many of the equivalences between query expressions amount to algebraic properties of the various operators, which in turn arise from the adjunctions we have explored above. We discuss some of the most important such equivalences in what follows. As a reminder, the relational algebra operators have the following specifications:

$$\begin{aligned}\pi_f &= \text{Bag } f \\ \sigma_p &= \text{filter } p \\ x \bowtie_g y &= \{ (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a = g \ b \ \}\end{aligned}$$

(although equijoin has the more efficient implementation discussed in Section 6).

We do not cover equivalences that depend on data invariants, such as functional dependencies, nor how to use statistics on the actual data in the tables in order to choose among equivalent queries; both aspects depend on information outside our model, and “whereof one cannot speak, thereof one must be silent”.

8.1 High-School Algebra

Tarski’s *high-school algebra identities* [Burris and Lee 1993; Fiore et al. 2006] are the eleven basic laws of the positive naturals that one typically learns at school: addition is associative and commutative; multiplication is associative and commutative with neutral element 1; multiplication distributes over addition; exponentiation promotes through addition ($x \uparrow (y + z) = (x \uparrow y) \times (x \uparrow z)$) and distributes over multiplication ($((x \times y) \uparrow z) = (x \uparrow z) \times (y \uparrow z)$); exponentiation is a right action of the monoid of multiplication ($(x \uparrow (y \times z)) = (x \uparrow y) \uparrow z$ and $x \uparrow 1 = x$); and 1 is a left zero of exponentiation.

But the positive naturals are not the only model of this equational theory; another model is given by $(\text{Set}, \uplus, \times, \rightarrow, \{*\})$, which arises from sets under disjoint union, cartesian product, function spaces, and singletons. So too are finite non-empty sets under the same operations; and set cardinality is a theory morphism from finite non-empty sets to the positive naturals (that is, it preserves all the laws). These properties of sets fall out of the adjunctions in Figure 2(a),(b) yielding sums, products, and singletons.

To show that multiplication distributes through addition, we use the *Yoneda Embedding*, which states that

$$B \cong A \iff \mathcal{C}(A, -) \cong \mathcal{C}(B, -)$$

One can see the Yoneda Embedding as a high-level reasoning principle of *indirect proof*; in the special case of ordered sets, it becomes

$$(b = a) \iff (\forall c. (a \leq c) \Leftrightarrow (b \leq c))$$

We can use the Yoneda Embedding to show that $L(A_1 + A_2) \cong L A_1 + L A_2$ whenever we have an adjunction $L \dashv R$ —it suffices to show $\text{Set}(L(A_1 + A_2), X) \cong \text{Set}(L A_1 + L A_2, X)$ for arbitrary X ,

and so we calculate:

$$\begin{aligned}
 & \text{Set } (L (A_1 + A_2), X) \\
 \cong & \quad \{ L \dashv R \} \\
 & \text{Set } (A_1 + A_2, R X) \\
 \cong & \quad \{ + \dashv \Delta \} \\
 & \text{Set } (A_1, R X) \times \text{Set } (A_2, R X) \\
 \cong & \quad \{ L \dashv R \} \\
 & \text{Set } (L A_1, X) \times \text{Set } (L A_2, X) \\
 \cong & \quad \{ + \dashv \Delta \} \\
 & \text{Set } (L A_1 + L A_2, X)
 \end{aligned}$$

Since $- \times B$ is a left adjoint, this gives us the special case $(A_1 + A_2) \times B \cong A_1 \times B + A_2 \times B$ of distributing a product over a sum.

8.2 Projections

Relational projections π_i are a special case of the functorial action $F i$ of the functor F representing database tables (whether that functor be Bag , Bag_* , or Table K), reading i as both the names of some columns and the function that extracts those fields from a single record. Equivalences for projections are therefore special cases of the functor laws. For example, assuming columns i are a subset of columns j , then field extraction i subsumes j (that is, $i \cdot j = i$), and so

$$\begin{aligned}
 & \pi_i \cdot \pi_j \\
 = & \quad \{ \text{projections as functorial action} \} \\
 & \text{Bag } i \cdot \text{Bag } j \\
 = & \quad \{ \text{functors preserve composition (2b)} \} \\
 & \text{Bag } (i \cdot j) \\
 = & \quad \{ \text{subsumption: } i \cdot j = i \} \\
 & \text{Bag } i \\
 = & \quad \{ \text{projection as functorial action} \} \\
 & \pi_i
 \end{aligned}$$

8.3 Selections

We represent the relational selection σ_p as *filter* p , reading p as a predicate, that is, a $\mathbb{B}$ -valued function. We can simplify and combine filters as follows, where *false*, *true*, and conjunction are lifted to predicates.

$$\text{filter false } s = \emptyset \tag{20a}$$

$$\text{filter true } s = s \tag{20b}$$

$$\text{filter } (p \wedge q) s = \text{filter } p (\text{filter } q s) \tag{20c}$$

All of these equalities are easy to show, since *filter* p , constant functions and the identity are all homomorphisms, and two homomorphisms from $\text{Free } A$ are equal iff they coincide on singleton bags (4). In particular, the last equality shows that consecutive selections can be fused.

The property that *filter* p is a dinatural transformation (7d) shows how to commute a selection σ_p and a projection π_i , when predicate p depends only on columns i : then as boolean predicates,

$p \cdot i = p$, and so

$$\begin{aligned}
 & \sigma_p \cdot \pi_i \\
 = & \{ \text{selections and projections} \} \\
 & \text{filter } p \cdot \text{Bag } i \\
 = & \{ \text{dinaturality (7d)} \} \\
 & \text{Bag } i \cdot \text{filter } (p \cdot i) \\
 = & \{ \text{subsumption: } p \cdot i = p \} \\
 & \text{Bag } i \cdot \text{filter } p \\
 = & \{ \text{selections and projections} \} \\
 & \pi_i \cdot \sigma_p
 \end{aligned}$$

8.4 Aggregations

Selection $\sigma_p = \text{reduce } (\text{Free } A) \cdot \text{Bag } (\text{guard } p)$ is a particular homomorphism; aggregations more generally are also of this form, as we saw in Section 3, and enjoy similar properties, for similar reasons. For example, a homomorphism $h = \text{reduce } M \cdot \text{Bag } f$ for monoid $M = (M, \epsilon, \otimes)$ can always absorb a projection:

$$\begin{aligned}
 & h \cdot \pi_i \\
 = & \{ \text{homomorphism, projection} \} \\
 & \text{reduce } M \cdot \text{Bag } f \cdot \text{Bag } i \\
 = & \{ \text{functor composition (2b)} \} \\
 & \text{reduce } M \cdot \text{Bag } (f \cdot i)
 \end{aligned}$$

and a selection:

$$\begin{aligned}
 & h \cdot \sigma_p \\
 = & \{ \text{selection} \} \\
 & h \cdot \text{reduce } (\text{Free } A) \cdot \text{Bag } (\text{guard } p) \\
 = & \{ h : \text{Free } A \rightarrow M, \text{ naturality (6b)} \} \\
 & \text{reduce } M \cdot \text{Bag } h \cdot \text{Bag } (\text{guard } p) \\
 = & \{ \text{functor composition (2b)} \} \\
 & \text{reduce } M \cdot \text{Bag } (h \cdot \text{guard } p) \\
 = & \{ \text{guard} \} \\
 & \text{reduce } M \cdot \text{Bag } (\lambda a \rightarrow h \text{ (if } p \text{ } a \text{ then } \lambda a \text{ } \text{ else } \emptyset)) \\
 = & \{ \text{homomorphisms: } h \text{ } \lambda a \text{ } = f \text{ } a \text{ (5b), } h \text{ } \emptyset = \epsilon \text{ (5a)} \} \\
 & \text{reduce } M \cdot \text{Bag } (\lambda a \rightarrow \text{if } p \text{ } a \text{ then } f \text{ } a \text{ else } \epsilon)
 \end{aligned}$$

8.5 Comprehensions

The monad laws justify certain manipulations of monad comprehensions. For example, associativity of qualifier sequencing follows from associativity of monad multiplication μ :

$$\begin{aligned}
 & \lambda e \mid (q, q'), q'' \int \\
 = & \{ \text{sequencing of qualifiers (9d)} \} \\
 & \mu \lambda e \mid q'' \int \mid q, q' \int \\
 = & \{ \text{sequencing of qualifiers (9d) again} \} \\
 & \mu (\mu \lambda \lambda e \mid q'' \int \mid q' \int \mid q \int) \\
 = & \{ \text{associativity (8c): } \mu \cdot \mu = \mu \cdot \text{Bag } \mu \} \\
 & \mu (\text{Bag } \mu \lambda \lambda e \mid q'' \int \mid q' \int \mid q \int) \\
 = & \{ \text{functors over comprehension (10)} \} \\
 & \mu \lambda \mu \lambda e \mid q'' \int \mid q' \int \mid q \int \\
 = & \{ \text{sequencing of qualifiers (9d)} \} \\
 & \mu \lambda \lambda e \mid q', q'' \int \mid q \int \\
 = & \{ \text{sequencing of qualifiers (9d) again} \} \\
 & \lambda e \mid q, (q', q'') \int
 \end{aligned}$$

(Indeed, this implication is an equivalence [Wadler 1992].) The above shows how to flatten an inner comprehension nested in the term part of an outer comprehension; for one nested in the qualifier part, we have:

$$\begin{aligned}
 & \lambda e \mid p \leftarrow \lambda e' \mid q \int \int \\
 = & \{ \text{generators (9c)} \} \\
 & \text{Bag } (\lambda p \rightarrow e) \lambda e' \mid q \int \\
 = & \{ \text{functors over comprehension (10)} \} \\
 & \lambda (\lambda p \rightarrow e) e' \mid q \int
 \end{aligned}$$

8.6 Joins

Join order optimization is one of the most powerful optimization techniques, even though it is known to be insufficient for generating worst-case optimal join implementations [Ngo et al. 2012]. It can be shown that extending *merge* from binary to n -ary functions, combined with the generic trie implementation of *ix* based on (15a)–(15d) [Henglein and Hinze 2013] and symbolic Cartesian product constructors, yields worst-case optimal joins that dramatically outperform conventional query processors on cyclic join queries, specifically triangle queries, also in practice [Henglein and Larsen 2010].

Cartesian product and equijoin are associative and commutative (up to canonical isomorphisms), and so queries involving two or more joins offer many opportunities for rearranging the query in order to improve performance. The pivotal observation is the efficient implementation of equijoin:

$$\begin{aligned}
 x \text{ } f \bowtie_g y &= \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \int \\
 &= \text{elems } (\text{Map}_* \ K \ (\times) \ (\text{merge } (x \text{ } \text{indexBy}' f, y \text{ } \text{indexBy}' g)))
 \end{aligned}$$

as discussed in Section 6.

The comprehension version is much easier to manipulate than the combinator-style one. For example, here is the proof that join is commutative, up to the isomorphism $\text{swap } (a, b) = (b, a)$ on

pairs:

$$\begin{aligned}
& x \text{ } f \bowtie_g \text{ } y \\
&= \{ \text{equijoin as comprehension} \} \\
&\quad \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \] \\
&= \{ \text{Bag is a commutative monad (12)} \} \\
&\quad \lambda (a, b) \mid b \leftarrow y, a \leftarrow x, f \ a == g \ b \] \\
&= \{ \text{swap; == is symmetric} \} \\
&\quad \lambda \text{ swap } (b, a) \mid b \leftarrow y, a \leftarrow x, g \ b == f \ a \] \\
&= \{ \text{functors over comprehension (10)} \} \\
&\quad \text{Bag swap } \lambda (b, a) \mid b \leftarrow y, a \leftarrow x, g \ b == f \ a \] \\
&= \{ \text{equijoin as comprehension} \} \\
&\quad \text{Bag swap } (y \text{ } g \bowtie_f \text{ } x)
\end{aligned}$$

Similarly, join is associative, up to the nested pair isomorphism $\text{assoc } ((a, b), c) = (a, (b, c))$:

$$\begin{aligned}
& (x \text{ } f \bowtie_g \text{ } y) \text{ } g \cdot \text{snd} \bowtie_h \text{ } z \\
&= \{ \text{equijoin as comprehension} \} \\
&\quad \lambda ((a, b), c) \mid (a, b) \leftarrow \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \], c \leftarrow z, g \ b == h \ c \] \\
&= \{ \text{comprehensions} \} \\
&\quad \lambda ((a, b), c) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b, c \leftarrow z, g \ b == h \ c \] \\
&= \{ \text{assoc; } c \text{ not free in } f, g \} \\
&\quad \lambda \text{ assoc } (a, (b, c)) \mid a \leftarrow x, b \leftarrow y, c \leftarrow z, g \ b == h \ c, f \ a == g \ b \] \\
&= \{ \text{functors over comprehension (10)} \} \\
&\quad \text{Bag assoc } \lambda (a, (b, c)) \mid a \leftarrow x, b \leftarrow y, c \leftarrow z, g \ b == h \ c, f \ a == g \ b \] \\
&= \{ \text{equijoin as comprehension} \} \\
&\quad \text{Bag assoc } (x \text{ } f \bowtie_{g \cdot \text{fst}} \text{ } (y \text{ } g \bowtie_h \text{ } z))
\end{aligned}$$

8.7 Combining Joins with Other Operations

A projection $\pi_{i \times j}$ can be promoted through an equijoin $f \bowtie_g$ if it is compatible with the join criterion, that is, if there exist f', g' such that $(f \ a == g \ b) \Leftrightarrow (f' \ (i \ a) == g' \ (j \ b))$:

$$\begin{aligned}
& \pi_{i \times j} (x \text{ } f \bowtie_g \text{ } y) \\
&= \{ \text{projection, equijoin} \} \\
&\quad \text{Bag } (i \times j) \ \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \] \\
&= \{ \text{functors over comprehension (10)} \} \\
&\quad \lambda (i \times j) \ (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \] \\
&= \{ \text{product} \} \\
&\quad \lambda (i \ a, j \ b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \] \\
&= \{ \text{compatibility} \} \\
&\quad \lambda (i \ a, j \ b) \mid a \leftarrow x, b \leftarrow y, f' \ (i \ a) == g' \ (j \ b) \] \\
&= \{ \text{functors} \} \\
&\quad \lambda (a', b') \mid a' \leftarrow \text{Bag } i \ x, b' \leftarrow \text{Bag } j \ y, f' \ a' == g' \ b' \] \\
&= \{ \text{projection, equijoin} \} \\
&\quad (\pi_i \ x) \text{ } f' \bowtie_{g'} \text{ } (\pi_j \ y)
\end{aligned}$$

Similarly, if a selection σ_p after a join uses a predicate p that can be factorised to act independently on the components of the pair—that is, there exist q, r with $p \ (a, b) = q \ a \wedge r \ b$ with b not free in

q a and a not free in r b —then the selection can be promoted through the join:

$$\begin{aligned}
 & \sigma_p (x \bowtie_g y) \\
 = & \{ \text{selection, equijoin} \} \\
 & \lambda (a, b) \mid (a, b) \leftarrow \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \}, p \ (a, b) \} \\
 = & \{ \text{factorising predicate} \} \\
 & \lambda (a, b) \mid (a, b) \leftarrow \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b \}, q \ a \wedge r \ b \} \\
 = & \{ \text{comprehensions} \} \\
 & \lambda (a, b) \mid a \leftarrow x, b \leftarrow y, f \ a == g \ b, q \ a, r \ b \} \\
 = & \{ b \text{ is not free in } q \ a \} \\
 & \lambda (a, b) \mid a \leftarrow x, q \ a, b \leftarrow y, r \ b, f \ a == g \ b \} \\
 = & \{ \text{comprehensions} \} \\
 & \lambda (a, b) \mid a \leftarrow \lambda a \mid a \leftarrow x, q \ a \}, b \leftarrow \lambda b \mid b \leftarrow y, r \ b \}, f \ a == g \ b \} \\
 = & \{ \text{selection, equijoin} \} \\
 & \sigma_q \ x \bowtie_g \sigma_r \ y
 \end{aligned}$$

To sum up, not only do adjunctions provide the structures required for database operations, but they also give rise to the properties we need for optimizations.

9 CONCLUSION

Comprehensions as a notation for queries, generalizable to various kinds of collection monad, have for decades been a staple of programming interfaces to data. Monads enjoy rich mathematical structure, deriving from category theory, and manifesting as a body of equivalences that provides a basis for many query optimizations. However, the monadic structure alone does not provide a good story about joins in queries; and reducing joins to cartesian products leads to extremely inefficient implementations. Efficient implementations—taking time linear in the size of the input—require a different explanation. We have presented such an explanation; it is just as rich and elegant as the naive monadic story, being similarly based on a few simple adjunctions.

Our construction depends crucially on the interplay between multiple monads. Bags are a sweet spot for expressivity of queries: lists would preserve too many distinctions, in terms of ordering of results—so denying some opportunities for query optimization, in particular for joins. On the other hand, sets collapse too many distinctions, in terms of multiplicity of elements—thereby denying non-idempotent aggregations such as count and sum. But bags alone are insufficient for representing the indexing and merging that is necessary for an efficient implementation of joins; for that, we also need maps. And for maps to be aggregatable, they need to be finite. And finally, since finite maps are not a plain monad, we had to introduce the refined notion of graded monads.

This necessary interplay between heterogeneous monads has the unfortunate consequence that we can no longer write a query involving a join as a traditional comprehension, because these can only be homogeneous. Nevertheless, as we showed in Section 7, recent extensions to Haskell’s comprehension notation to support grouping [Peyton Jones and Wadler 2007] and merging [Plasmeijer and van Eekelen 1995] support the necessary heterogeneity that allows us to express equijoins efficiently in the comprehension notation.

This means that we do not have to resort to the combinator style just to get efficient implementation of joins; comprehensions still suffice as a surface syntax. Nevertheless, the combinator style is arguably more appropriate as an intermediate representation; database management systems traditionally translate input queries from the comprehension style (‘relational calculus’) to combinator style (‘relational algebra’) in preparation for query optimization (although it is possible to carry out the query transformations directly in the comprehension style, as for example Cheney et al.

[2013] do). And conversely, it is possible to reconstruct joins from queries written in the ordinary comprehension style, albeit at the cost of considerable analysis.

The investigations presented here are a first step towards a greater goal. We have shown that adjunctions give rise to all the combinators needed to give a rigorous explanation of the selections, projections, and joins of relational algebra, providing both the basis for an efficient implementation and the justification for the body of transformations needed in query optimization. We are optimistic that these constructions might form a firm foundation for an intermediate representation for implementing database connection frameworks such as Microsoft’s LINQ [Meijer 2007; Syme 2006]. In general, this work lends evidence to the argument that expressive type systems (here, notably the use of type classes for datatype-generic programming) are important for accommodating embedded domain-specific languages.

ACKNOWLEDGEMENTS

We are grateful to the members of IFIP Working Groups 2.1 on *Algorithmic Languages and Calculi* and 2.11 on *Program Generation* for valuable discussions about this work at meetings; special thanks are due to Sam Lindley, Conor McBride, and Jacques Carette. We also thank the anonymous reviewers of earlier versions for their helpful comments. This work has been funded by UK EPSRC project *Unifying Theories of Generic Programming* (grant number EP/J010995/1) and Innovation Fund Denmark project *HIPERFIT* (grant number 10-092299).

REFERENCES

- Alexander Alexandrov, Andreas Künft, Asterios Katsifodimos, Felix Schüler, Lauritz Thamsen, Odej Kao, Tobias Herb, and Volker Markl. 2015. Implicit Parallelism through Deep Language Embedding. In *SIGMOD record*. ACM, 47–61. <https://doi.org/10.1145/2723372.2750543>
- Steve Awodey. 2006. *Category Theory*. Oxford University Press.
- Peter Buneman, Leonid Libkin, Dan Suciu, Val Tannen, and Limsoon Wong. 1994. Comprehension Syntax. *SIGMOD Record* 23, 1 (1994), 87–96. <https://doi.org/10.1145/181550.181564>
- Stanley Burris and Simon Lee. 1993. Tarski’s High School Identities. *Amer. Math. Monthly* 100, 3 (March 1993), 231–236.
- Jacques Carette. 2018. Email correspondence. (May 2018). Personal communication.
- James Cheney, Sam Lindley, and Philip Wadler. 2013. A Practical Theory of Language-Integrated Query. In *International Conference on Functional Programming*, Greg Morrisett and Tarmo Uustalu (Eds.). ACM, 403–416. <https://doi.org/10.1145/2500365.2500586>
- Eugenia Cheng. 2015. *Cakes, Custard, and Category Theory*. Profile Books.
- Richard H. Connelly and F. Lockwood Morris. 1995. A Generalization of the Trie Data Structure. *Mathematical Structures in Computer Science* 5, 3 (September 1995), 381–418. <https://doi.org/10.1017/S096012950000803>
- Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. 2006. Links: Web Programming without Tiers. In *Formal Methods for Components and Objects (Lecture Notes in Computer Science)*, Vol. 4709. Springer, 266–296. https://doi.org/10.1007/978-3-540-74792-5_12
- John Darlington. 1975. Application of Program Transformation to Program Synthesis. In *IRIA Symposium on Proving and Improving Programs*. Arc-et-Senans, France, 133–144.
- Christopher J. Date. 2004. *An Introduction to Database Systems* (8th ed.). Pearson.
- Mary Fernandez, Jerome Simeon, and Philip Wadler. 2001. A Semi-Monad for Semi-Structured Data. In *International Conference on Database Theory (Lecture Notes in Computer Science)*, Jan Van den Bussche and Victor Vianu (Eds.), Vol. 1973. Springer, 263–300. https://doi.org/10.1007/3-540-44503-X_18
- Marcelo Fiore, Roberto Di Cosmo, and Vincent Balat. 2006. Remarks on Isomorphisms in Typed Lambda Calculi with Empty and Sum Types. *Annals of Pure and Applied Logic* 141, 1-2 (2006), 35–50. <https://doi.org/10.1016/j.apal.2005.09.001>
- Soichiro Fujii, Shin-ya Katsumata, and Paul-André Mellies. 2016. Towards a Formal Theory of Graded Monads. In *Foundations of Software Science and Computation Structures (Lecture Notes in Computer Science)*. Springer-Verlag, 513–530. https://doi.org/10.1007/978-3-662-49630-5_30
- Jeremy Gibbons. 2016. Comprehending Ringads. In *A List of Successes that can Change the World (Lecture Notes in Computer Science)*, Sam Lindley, Conor McBride, Don Sannella, and Phil Trinder (Eds.), Vol. 9600. Springer, 132–151. https://doi.org/10.1007/978-3-319-30936-1_7

- George Giorgidze, Torsten Grust, Tom Schreiber, and Jeroen Weijers. 2011a. Haskell Boards the Ferry: Database-Supported Program Execution for Haskell. In *Implementation and Application of Functional Languages (Lecture Notes in Computer Science)*, Jurriaan Hage and Marco T. Morazán (Eds.), Vol. 6647. Springer, 1–18. https://doi.org/10.1007/978-3-642-24276-2_1
- George Giorgidze, Torsten Grust, Nils Schweinsberg, and Jeroen Weijers. 2011b. Bringing Back Monad Comprehensions. In *Haskell Symposium*. ACM, 13–22. <https://doi.org/10.1145/2034675.2034678>
- Glasgow 2015. *Glasgow Haskell Compiler Users' Guide, Version 7.10.1*. https://downloads.haskell.org/~ghc/7.10.1/docs/html/users_guide/index.html.
- Torsten Grust and Marc H. Scholl. 1999. How to Comprehend Queries Functionally. *Journal of Intelligent Information Systems* 12, 2-3 (1999), 191–218. <https://doi.org/10.1023/A:1008705026446>
- Fritz Henglein and Ralf Hinze. 2013. Distributive Sorting and Searching: From Generic Discrimination to Generic Tries. In *Asian Symposium on Programming Languages and Systems (Lecture Notes in Computer Science)*, Chung-chieh Shan (Ed.), Vol. 8301. Springer, 315–332. https://doi.org/10.1007/978-3-319-03542-0_23
- Fritz Henglein and Ken Friis Larsen. 2010. Generic Multiset Programming with Discrimination-Based Joins and Symbolic Cartesian Products. *Higher-Order and Symbolic Computation* 23, 3 (2010), 337–370. <https://doi.org/10.1007/s10990-011-9078-8>
- Ralf Hinze. 2000. Generalizing Generalized Tries. *Journal of Functional Programming* 10, 4 (2000), 327–351. <https://doi.org/10.1017/S0956796800003713>
- Ralf Hinze. 2003. Fun with Phantom Types. In *The Fun of Programming*, Jeremy Gibbons and Oege de Moor (Eds.). Palgrave Macmillan, 245–262.
- Shin-ya Katsumata. 2014. Parametric Effect Monads and Semantics of Effect Systems. In *Principles of Programming Languages*. ACM, 633–645. <https://doi.org/10.1145/2535838.2535846>
- Leonid Libkin and Limsoon Wong. 1997. Query Languages for Bags and Aggregate Functions. *J. Comput. System Sci.* 55, 2 (1997), 241–272. <https://doi.org/10.1006/jcss.1997.1523>
- Conor McBride and Ross Paterson. 2008. Functional Pearl: Applicative Programming with Effects. *Journal of Functional Programming* 18, 1 (2008), 1–13. <https://doi.org/10.1017/S0956796807006326>
- Erik Meijer. 2007. Confessions of a Used Programming Language Salesman. In *Object-Oriented Programming: Systems, Languages and Applications*. ACM, 677–694. <https://doi.org/10.1145/1297027.1297078>
- Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. 2012. Worst-Case Optimal Join Algorithms. In *Principles of Database Systems*. ACM, 37–48. <https://doi.org/10.1145/2213556.2213565>
- Dominic A. Orchard, Tomas Petricek, and Alan Mycroft. 2014. The Semantic Marriage of Monads and Effects. *CoRR* abs/1401.5391 (2014).
- Simon Peyton Jones and Philip Wadler. 2007. Comprehensive Comprehensions. In *Haskell Workshop*. ACM, 61–72. <https://doi.org/10.1145/1291201.1291209>
- Benjamin C. Pierce. 1991. *Basic Category Theory for Computer Scientists*. MIT Press.
- Rinus Plasmeijer and Marko van Eekelen. 1995. *Concurrent Clean Language Report (Version 1.0)*. Technical Report. University of Nijmegen. <ftp://ftp.science.ru.nl/pub/Clean/old/Clean10/doc/refman.ps.gz>.
- Jacob T. Schwartz, Robert B. K. Dewar, Ed Dubinsky, and Edmond Schonberg. 1986. *Programming with Sets: An Introduction to SETL*. Springer, New York.
- Patricia Griffiths Selinger, Morton M. Astrahan, Donald D. Chamberlin, Raymond A. Lorie, and Thomas G. Price. 1979. Access Path Selection in a Relational Database Management System. In *SIGMOD record*. ACM, 23–34. <https://doi.org/10.1145/582095.582099>
- Ross Street. 1972. Two Constructions on Lax Functors. *Cahiers de Topologie et Géométrie Différentielle Catégoriques* 13, 3 (1972), 217–264. <http://eudml.org/doc/91107>
- Kenichi Suzuki, Oleg Kiselyov, and Yuki Yoshi Kameyama. 2016. Finally, Safely-Extensible and Efficient Language-Integrated Query. In *Partial Evaluation and Program Manipulation*, Martin Erwig and Tiark Rompf (Eds.). ACM, 37–48. <https://doi.org/10.1145/2847538.2847542>
- Don Syme. 2006. Leveraging .NET Meta-programming Components from F#: Integrated Queries and Interoperable Heterogeneous Execution. In *ML Workshop*. ACM, New York, NY, USA, 43–54. <https://doi.org/10.1145/1159876.1159884>
- Philip W. Trinder. 1991. Comprehensions, a Query Notation for DBPLs. In *Database Programming Languages*. 55–68.
- Philip Wadler. 1992. Comprehending Monads. *Mathematical Structures in Computer Science* 2, 4 (1992), 461–493. <https://doi.org/10.1017/S0960129500001560>
- Limsoon Wong. 2000. Kleisli, a Functional Query System. *Journal of Functional Programming* 10, 1 (2000), 19–56.