



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Comprehending queries over finite maps

Citation for published version:

Ricciotti, W 2023, Comprehending queries over finite maps. in *PPDP'23: Proceedings of the 25th International Symposium on Principles and Practice of Declarative Programming.*, 8, ACM International Conference Proceedings Series, ACM Association for Computing Machinery, pp. 1-12, 25th International Symposium on Principles and Practice of Declarative Programming, Lisbon, Portugal, 22/10/23.
<https://doi.org/10.1145/3610612>

Digital Object Identifier (DOI):

[10.1145/3610612](https://doi.org/10.1145/3610612)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

PPDP'23: Proceedings of the 25th International Symposium on Principles and Practice of Declarative Programming

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Comprehending queries over finite maps

Wilmer Ricciotti

research@wilmer-ricciotti.net

Laboratory for Foundations of Computer Science, University of Edinburgh
Edinburgh, Scotland, UK

Abstract

Recent programming languages research has developed *language-integrated query*, a convenient technique to seamlessly embed a domain-specific database query language into a general-purpose host programming language; such queries are then automatically converted to the language understood by the target DBMS (e.g. SQL) while at the same time taking advantage of the host language’s type-checker to prevent failure at run-time. The embedded query language is often equipped with a rewrite system which normalizes queries to a form that can be directly translated to the DBMS query language.

However, the theoretical foundations of such rewrite systems have not been explored to their full extent, particularly when constructs like grouping and aggregation, which are ubiquitous in real-world database queries, are involved. In this work, we propose an extension of the nested relational calculus with grouping and aggregation which can provide such foundations. Along with strong normalization and translatability to SQL we show that, remarkably, this extension can also blend with *shredding* techniques proposed in the literature to allow queries with a nested relational type to be executed on the DBMS.

CCS Concepts: • Information systems → Query languages; • Software and its engineering → Formal language definitions.

Keywords: language-integrated query, nested relations, finite maps, multisets

ACM Reference Format:

Wilmer Ricciotti. 2018. Comprehending queries over finite maps. *J. ACM* 37, 4, Article 111 (August 2018), 20 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Author’s address: Wilmer Ricciotti, research@wilmer-ricciotti.net, Laboratory for Foundations of Computer Science, University of Edinburgh, 10 Crichton St, Edinburgh, Scotland, UK, EH8 9AB.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

0004-5411/2018/8-ART111 \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Many, if not all, real-world applications involve interaction with data, often stored and managed in a database system with its own domain-specific language for queries and updates, with relational database management systems (RDBMS) using the SQL standard among the most popular. This architecture has many benefits, particularly separation of concerns: programmers (supposedly) need only express the declarative needs of the application through queries or updates, leaving database implementors free to choose efficient implementation strategies. Nevertheless, database programming can be an unpleasant chore since SQL (for example) presents a rather different interface and abstractions to programmers than most general-purpose languages in which the main application logic is implemented do. Addressing this difficulty has been a major subject of research.

Previous work dating to the late 1980s [31] has established the value of viewing database collection types such as sets and bags as *monads* equipped with operations such as “return” (singleton) and “bind” (concat-map, flat-map, or one-generator comprehension), usually equipped with additional operations such as a zero (empty collection) and plus (union). From a programming languages point of view, this perspective has the distinct advantage of taming queries by making them an instance of a well-understood and explored interface (particularly in languages such as Haskell that directly support monads via type classes). And indeed, there has now been considerable work exploiting this correspondence, which began in earnest with Wong’s Kleisli system, and has subsequently informed systems such as Microsoft’s LINQ for SQL (particularly the F# implementation), the Links cross-tier web programming language, and libraries such as Quill for Scala.

Despite the success of this general approach, dark corners remain, due to the fact that query languages such as SQL are considerably more restricted than general-purpose programming languages, particularly those those with higher-order functions such as F#, Scala or Links. The free combination of type constructors such as collection and record types is a given in most programming languages, but such *nested collections* are not normally supported in RDBMSs (though this facility is supported by the SQL:2003 standard, it is not widely adopted). Similarly, it seems natural to use higher-order functions to decompose query expressions into smaller, reusable and more readable, chunks, but SQL does not support higher-order functions (and even using user-defined

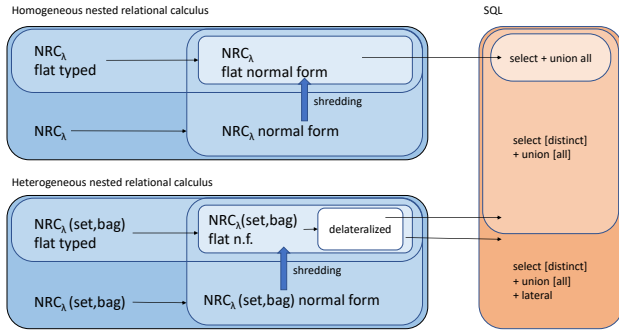


Figure 1. Relationship between fragments of \mathcal{NRC} and of SQL.

first-order functions inside queries can have a high cost). As a result, some previous work has considered the following problems: How can we identify which host language query expressions can actually be performed on the database? When features that seem natural in the host language are not directly supported by the database, can these nevertheless be simulated somehow? Proposals addressing the first question include type-and-effect systems (as in Links [15]) or staged calculi (for F# [3]) to identify query expressions that can safely be performed on the database, while the second has been addressed in part by developing query normalization techniques that simplify queries that employ nonrecursive higher-order functions and *shredding* transformations that simulate queries over nested data using queries over flat data (that is, queries that manipulate only collections of records of primitive values, as SQL requires.)

Core calculi like the nested relational calculus (\mathcal{NRC}) [2] offer an elegant way to describe the monadic sublanguage that is mutually understandable to a general purpose language and SQL and have been employed to study query normalization and shredding in the presence of non-recursive higher-order functions [4–6]. Such techniques were initially investigated for bag semantics only, but have been since extended to work with queries mixing set and bag semantics [21, 22, 24].

Figure 1 shows the relationship between fragments of two variants of \mathcal{NRC} and fragments of SQL. In the homogeneous \mathcal{NRC} , using bag semantics, terms whose type is a flat relation can be normalized and directly translated to SQL queries employing `SELECT` and `UNION ALL`. Terms with a nested relational type eventually rewrite to nested relational forms that can be converted to multiple flat relational normal forms by shredding, which can be translated to SQL as described above.

The heterogeneous calculus (also termed $\mathcal{NRC}_\lambda(Set, Bag)$) adds the ability to mix sets and bags in the same query. Just like in the homogeneous variant, every term can be normalized and normal forms with a nested relational type can be




Figure 2. Extending \mathcal{NRC} to target SQL with grouping.

shredded to flat normal forms. Predictably, these can be translated to a larger SQL fragment including `SELECT DISTINCT` and `UNION` operations; however, generally we will also need to use `LATERAL` joins (available since SQL:1999); an optional delateralization step converts $\mathcal{NRC}_\lambda(Set, Bag)$ normal forms to a form that can be translated to SQL without using `LATERAL`. Notice that the heterogeneous calculus, including its rewrite system, is a superset of the homogeneous one, therefore terms of this calculus only employing bag semantics are guaranteed to be translated to the same SQL queries produced by the homogeneous calculus.

Another area where improvement is needed, and which is the focus of this paper, is the coverage of the most important SQL features. In particular, so-called OLAP or *on-line analytic processing* queries rely on *grouping and aggregation* capabilities of SQL. The importance of these capabilities is highlighted by their use in 16 out of 22 of the TPC-H benchmark queries ([30]), which are by far the most widely used benchmark for SQL query processing and optimization. We want to eventually make language-integrated query “TPC-H complete”, or in other words capable of handling (in a principled way) all of the features needed for the TPC-H queries. As a simple concrete example, suppose we have an employee database and we wish to calculate the average salary in each department. In SQL this can be done as follows:

```
q_group := SELECT d.name, AVERAGE(e.salary)
           FROM department d, employee e
           WHERE d.id = e.dept
           GROUP BY d.name
```

We intend to support grouping and aggregation in a functional programming language by means of further extensions to \mathcal{NRC} : the intended operation of such an extension is described by Figure 2. We start with minimal expectations concerning the operations that should be supported; since our goal is to allow idiomatic functional programs to be converted to SQL queries, these operations need not mimic SQL queries closely, but it is important for them to provide a natural interface for the functional programmer. Our initial proposal is based on the understanding that grouping and aggregation are informally seen as separate operations, and the reason why SQL requires grouping to be associated with aggregation (aside from pathological cases) is based on its limited type system, where all queries must evaluate to finite collections of records of scalars (also termed (*flat*) *relations*).

Operationally, grouping takes a relation and partitions it into a number of subrelations, each of which is indexed by a key for that group. While a way to represent such an indexing would be by means of a function type, general functions cannot be represented as relations, meaning such types would not be a suitable source for a translation to SQL. However, the maps resulting from grouping associate the elements of a *finite* domain set to finite relations (i.e. finite collections): we call these maps *finite maps*. If we represent multiset collections with element type T as $[T]$ and finite maps from records of type ρ to multiset collections with element type T as $[T]_\rho$, we can see that finite maps can be represented as SQL relations by employing the isomorphism:

$$[T]_\rho \simeq [\rho \times T]$$

Now we tentatively list the operations that a language with finite maps, grouping, and aggregation should support:

- given a collection l of type $[T]$ and an indexing function f assigning a key of record type ρ to any value of type T , $groupBy\ f\ l$ should return a finite map associating each $k : \rho$ to the collection of those values v in l for which k is equal to the key computed from v , i.e. $f\ v$:

$$groupBy : (T \rightarrow \rho) \rightarrow [T] \rightarrow [T]_\rho$$

- a lookup operation which, given a finite map m of type $[T]_\rho$ and a record k of type ρ , returns the collection associated by m to the key k :

$$lookup : [T]_\rho \rightarrow \rho \rightarrow [T]$$

- an aggregation function α takes a collection of type $[S]$ and returns a value of type T (the language will enforce the property that such aggregation functions should be expressible in SQL); but for the very common case in which aggregation is performed on the result of grouping (i.e. on a finite map), we should be able to express an operation $aggBy$ to lift aggregations from collections to indexed collections:

$$aggBy : ([T] \rightarrow T) \rightarrow [S]_\rho \rightarrow [\rho \times T]$$

- Crucially, we will also need to extend comprehension from collections to finite maps; we can come up with several options, which we distinguish by annotating the comprehension generator by different superscripts.

Perhaps the most obvious solution is to provide comprehension for finite maps by implicitly converting them to their isomorphic collections: we denote this operation with $\overset{\mathcal{M}}{\leftarrow}$ generators

$$\left[R \mid (k, x) \overset{\mathcal{M}}{\leftarrow} M \right]$$

This should take the disjoint unions of all the R evaluated for each key k of M and each value x in the collection returned by $lookup\ M\ k$. Since the values that M maps to k are considered separately rather than as a collection, this

form of comprehension is not very natural. Alternatively, we can consider a generator returning keys k together with the collection v of all values associated to k by M : we use $\overset{\mathcal{G}}{\leftarrow}$ generators to denote this kind of comprehension:

$$\left[R \mid (k, v) \overset{\mathcal{G}}{\leftarrow} M \right]$$

Here, each key k will be returned only once, together with a v which will equal $lookup\ M\ k$; since v can be computed from M and k , we can even drop it and consider generators returning only the keys of a given finite map, which we denote by $\overset{\mathcal{K}}{\leftarrow}$:

$$\left[R \mid k \overset{\mathcal{K}}{\leftarrow} M \right]$$

While we could consider a language employing any of these comprehensions (potentially expressing the other two as derived operators), the queries we express will have to be converted to SQL. In order for this to happen we will need to provide a rewrite system to convert the more liberal functional queries into normal forms that can be easily translated to SQL.

2 A nested relational calculus with finite maps

2.1 Syntax

We introduce an extension of the heterogeneous nested relational calculus $\mathcal{NRC}_\lambda(Set, Bag)$ [21] with finite maps: we will call the new calculus $\mathcal{NRC}_{\lambda\gamma}$. Like its predecessor, $\mathcal{NRC}_{\lambda\gamma}$ allows the mixing of two different kinds of collections: sets and bags (also known as multisets). The main extension of the new calculus consists in a *grouping* operator γ which converts collections into maps assigning collections to *keys* in a certain finite domain; these finite maps can be seen as a generalization of collections and thus collection operators such as unions and comprehensions are extended to finite maps. The calculus also supports aggregation operators from bags of values to simple values.

In our previous work we provided two different versions for most collection operators depending on whether they applied to sets or bags (for instance: \cup for set union, and \uplus for disjoint bag union): this created a large syntactic overhead by requiring us to consider all of these operators twice. In the syntax below, instead, we factorize such operator pairs by using the same symbol (say \uplus for union over any collection), and discriminate between the two cases by annotating said operator with its collection kind $\mathcal{S} = \text{set}$ or bag (i.e. $\uplus^{\text{set}}, \uplus^{\text{bag}}$). We will often omit the annotation when it is irrelevant or can be easily inferred from the context.

Types:

$$\begin{aligned} \mathcal{S} &::= \text{set} \mid \text{bag} \\ S, T &::= b \mid S \rightarrow T \mid \langle P \rangle \mid [T]_P^\delta \\ P &::= \overrightarrow{\ell : T} \end{aligned}$$
Terms:

$$\begin{aligned} L, M, N &::= x \mid t \mid c(\overrightarrow{M}) \mid \langle \overrightarrow{\rho} \rangle \mid M.\ell \\ &\mid \lambda x.M \mid MN \mid \alpha(M) \mid M \textbf{where}^\delta N \\ &\mid []^\delta \mid [\rho \triangleright M]^\delta \mid M \#^\delta N \mid [M \mid x \leftarrow N]^\delta \\ &\mid \gamma_{x,\rho}^\delta(M) \mid M \otimes^\delta N \mid \left[M \mid k \xleftarrow{P} N \right]^\delta \\ &\mid \delta M \mid \iota M \mid \text{aggBy}_{z.\ell=\alpha(z.\ell')}^\delta(M) \\ \rho &::= \overrightarrow{\ell = M} \\ \alpha &::= \text{count} \mid \text{sum} \mid \text{min} \mid \text{max} \mid \text{avg} \end{aligned}$$

$\mathcal{NRC}_{\lambda\gamma}$ includes refined types $[T]_P^{\text{set}}$ and $[T]_P^{\text{bag}}$ representing respectively set-valued and bag-valued finite maps: terms of these types behave as functions whose domain is described by a *row type* P (where a row type is defined as a collection of items in the form $\ell : T'$, such that ℓ is a label and T' a regular type) and whose codomain is either a set or a bag of values of the target type T ; such maps are *finite* because they will return a non-empty collection only for a finite subset of their domain. When P is an empty row, we obtain the traditional types $[T]^{\text{set}}$ and $[T]^{\text{bag}}$ of sets and bags over an object type T . Row types are also used to express record types $\langle P \rangle$. The grammar of types is completed by atomic types b (which must include Booleans \mathbf{B} and numbers \mathbf{N}) and function types $S \rightarrow T$.

The terms allow many standard forms including variables x , applied constants $c(\overrightarrow{M})$, function abstraction and application ($\lambda x.M$ and $M N$). Given a row ρ consisting of items in the form $\ell = M$ associating a term to a label, $\langle \rho \rangle$ represents the record containing that row. As usual, $M.\ell$ is used to access field ℓ in a record M .

Maps include terms that are similar to the collection terms in $\mathcal{NRC}_\lambda(\text{Set}, \text{Bag})$, but with some differences: while the empty map $[]^\delta$ and union map $M \#^\delta N$ are syntactically indistinguishable from the corresponding concepts of $\mathcal{NRC}_\lambda(\text{Set}, \text{Bag})$, the singleton map $[\rho \triangleright M]^\delta$ represents an expression associating the record $\langle \rho \rangle$ to the singleton collection $[M]^\delta$ – similarly to what we did with map types, we identify $[M]^\delta$ with the singleton map $[\cdot \triangleright M]^\delta$ having the empty row \cdot as its source. The operation $[M \mid x \leftarrow N]^\delta$ is also syntactically identical to comprehension in previous versions of \mathcal{NRC} , but as we will see later its typing rule is relaxed; we will call this operation *value comprehension map* to distinguish it from a *key comprehension operation* that we will introduce soon. The one armed conditional $M \textbf{where}^\delta N$ evaluating to M or $[]^\delta$ depending on whether the Boolean N is true or false is also unsurprising.

All of the aforementioned term forms can be applied to either collection kind; two operations that defy this scheme are deduplication δM (which converts a bag-valued map to a set-valued one by deduplicating all of its outputs) and promotion ιM (performing the inverse operation by converting the set outputs of the map M to bags containing the exact same elements, with multiplicity equal to one).

So far, the terms of $\mathcal{NRC}_{\lambda\gamma}$ are very similar to those of $\mathcal{NRC}_\lambda(\text{Set}, \text{Bag})$, although with refined typing and semantics; however, $\mathcal{NRC}_{\lambda\gamma}$ also introduces new operations to manipulate maps that have no counterpart in $\mathcal{NRC}_\lambda(\text{Set}, \text{Bag})$: the simplest of these is the lookup operation $M \otimes^\delta N$ which, given a map M and a key N , returns the value of M associated to N . Since a finite map M only associates an output to a finite number of keys N , but $\mathcal{NRC}_{\lambda\gamma}$ allows its types to have infinite inhabitants, it may happen that N is not in the domain of M : in this case, the semantics of $M \otimes^\delta N$ is the empty collection.

The grouping operation $\gamma_{x,\rho}^\delta(M)$ binds the variable x in the row ρ : if M is a collection (a pure set or bag), for each occurrence of an element N in M , the grouping operation produces a singleton map sending the row $\rho [N/x]$ to N ; if M is itself a proper map rather than a pure collection, the behaviour is similar, but the rows resulting from ρ are combined with those already present in M .

The groupwise aggregation $\text{aggBy}_{z.\ell=\alpha(z.\ell')}^\delta(M)$ takes as its main input M , which must be a finite map whose target type is a collection of records of base type (or, for short, a *flat relational map*); the secondary input is an aggregation clause using a bound variable z (used to reference the target of the map M), arbitrary output attributes $\overrightarrow{\ell}$, and input attributes $\overrightarrow{\ell'}$ matching attributes of the relation M .

For example, assume $M_{\text{factors}} : [\langle f : \mathbf{N} \rangle]_{\text{num}:\mathbf{N}}^{\text{bag}}$ is a bag-valued relational map associating integers smaller than 100 to the bag of their prime factors, with their multiplicity: so for example $\langle \text{num} = 2 \rangle$ is mapped to the bag $[\langle f = 2 \rangle]^{\text{bag}}$, $\langle \text{num} = 12 \rangle$ to the bag $[\langle f = 2 \rangle, \langle f = 2 \rangle, \langle f = 3 \rangle]^{\text{bag}}$, and $\langle \text{num} = 99 \rangle$ to the bag $[\langle f = 3 \rangle, \langle f = 3 \rangle, \langle f = 11 \rangle]^{\text{bag}}$. Then, by using the row $(\text{small} = x \leq 10)$ (which evaluates to true for integers x less than 10, and to false otherwise), we can construct a map $\gamma_{x.\text{small}=x \leq 10}(M_{\text{factors}})$ which maps $\langle \text{num} = 12, \text{small} = \text{true} \rangle$ to $[\langle f = 2 \rangle, \langle f = 2 \rangle, \langle f = 3 \rangle]^{\text{bag}}$, $\langle \text{num} = 12, \text{small} = \text{false} \rangle$ to $[]^{\text{bag}}$, $\langle \text{num} = 99, \text{small} = \text{true} \rangle$ to $[\langle f = 3 \rangle, \langle f = 3 \rangle]^{\text{bag}}$ and $\langle \text{num} = 99, \text{small} = \text{false} \rangle$ to $[\langle f = 11 \rangle]^{\text{bag}}$. Groupwise aggregation can be used to count the number of prime factors of each number in the map: the map resulting from $\text{aggBy}_{z.c=\text{count}(z.f)}(M_{\text{factors}})$ maps 2 to the singleton $[\langle c = 1 \rangle]^{\text{bag}}$, and 12 to $[\langle c = 3 \rangle]^{\text{bag}}$. Notice that since the output type of a finite map is always a collection, the use of aggBy results in a map returning singletons rather than naked values.

The final operation is $\left[M \mid k \stackrel{\mathcal{K}}{\leftarrow} N \right]^{\mathcal{S}}$, which we will call *key comprehension*: it is semantically equivalent to performing a set comprehension over the keys of the map N returning, for each key, a map M where the actual key has been substituted for the variable k . Notice that even in the bag case $\left[M \mid k \stackrel{\mathcal{K}}{\leftarrow} N \right]^{\text{bag}}$, each key of N is only considered once, since there does not seem to be much use to the concept of multiplicity of the input of a finite map.

To make the syntax of comprehensions less pedantic, we will allow ourselves to use placeholder symbols \diamond and \clubsuit in comprehension generators like $\left[M \mid x \stackrel{\diamond}{\leftarrow} N \right]^{\text{set}}$ or $\left[M \mid x \stackrel{\clubsuit}{\leftarrow} N \right]^{\text{bag}}$ to mean either value or key comprehension. Finally, we will use the syntactic sugar $\left[M \mid x_1 \stackrel{\diamond}{\leftarrow} N_1, \dots, x_n \stackrel{\clubsuit}{\leftarrow} N_m \right]$, defined as $\left[\dots \left[M \mid x_n \stackrel{\clubsuit}{\leftarrow} N_m \right] \dots \mid x_1 \stackrel{\diamond}{\leftarrow} N_1 \right]$ (and similarly for bags).

2.2 Typing rules

The type system for $\mathcal{NRC}_{\lambda\gamma}$ is given in Figure 3. The rules for variables, records, projections, functions, and applications are standard. A fixed signature Σ assigns to constants and aggregate operators their type: the former take a sequence of arguments of scalar type and return a scalar; the latter take as input a single bag of scalar-typed items and return a scalar.

The typing rules for finite maps are largely based on those of $\mathcal{NRC}_{\lambda}(\text{Set}, \text{Bag})$ collections, but with important differences. Empty maps $[]$ are allowed at any finite map type, whereas singletons have a map type corresponding to their arguments: if ρ has a row type P and M has type T , then $[\rho \triangleright M]^{\mathcal{S}}$ has type $[T]_P^{\mathcal{S}}$. Unions combine collections preserving the type of their arguments (which must be the same).

The typing rules for value comprehension are more interesting: to construct the term $\left[M \mid x \leftarrow N \right]^{\mathcal{S}}$, we will require the generator N to be a pure collection of type $[T]^{\mathcal{S}}$; then the output term M , existing in an extended context where x gets values of type T , can have any set-valued map type $[S]_P^{\text{set}}$; the result will have the same type as M – it must be stressed that this is a significant extension over $\mathcal{NRC}_{\lambda}(\text{Set}, \text{Bag})$, as in that calculus M (and the comprehension output) would need to be pure collections. Deduplication and promotion convert a bag-valued map type into the corresponding set-valued map type and vice-versa; one-armed conditionals are similar as their \mathcal{NRC} version, except for the fact that their output can be a map rather than a pure collection.

As for key comprehensions $\left[M \mid k \stackrel{\mathcal{K}}{\leftarrow} N \right]^{\mathcal{S}}$, the generator N can of course be a proper map of type $[T]_{P'}^{\mathcal{S}}$ (as it would otherwise be pointless to perform comprehension over its keys), and the output M once again must have a \mathcal{S} -valued map type when evaluated in an extended context: however,

this time the type of the bound variable k must be a tuple over the row type P .

Grouping $\gamma_{x,\rho}(M)$ adds new keys to the map M : if M has type $[T]_P$ and ρ has row type P' in the extended context where $x : T$, then the whole term has type $[T]_{P \oplus P'}$, where $P \oplus P'$, only defined when P and P' have disjoint domains, contains exactly those label-type associations that appear in either P or P' ; should P and P' have overlapping label sets, the typechecker will reject the term. The lookup operation $M \otimes N$ requires M to have a map type $[T]_P$ and N to be a tuple over the row type P (matching the expected input of M), and returns the pure set of type $[T]$ associated to that input. Notice that, in order to syntactically avoid unnecessary “detour” lookup operations, we require that the lookup should only happen on proper finite maps rather than pure collections: therefore, if M is a pure collection and N has an empty tuple type, $M \otimes N$ is not syntactically well-formed; this is not a limitation because the expected semantics of that term would be the same as that of M .

Groupwise aggregation $\text{aggBy}_{z.\ell_n = \alpha_n(z.\ell'_n)}(M)$ assumes that M has type $[\langle P' \rangle]_P$: for each key k of M and for each tuple z in $M \otimes k$, aggBy will produce a key-value pair where the key is k (as it was in the input) and the value is a singleton tuple containing fields $\ell_i = \alpha_i(\ell'_i)$ for $i = 1, \dots, n$. The types of the fields in the codomain of map M must match the input of the aggregate functions using them (meaning that, for each $\alpha_i(z.\ell'_i)$, if the type of α_i is $[b_i]^{\text{bag}} \Rightarrow b'_i$, then the row type P' must associate the label ℓ'_i to the b_i matching the input of α_i , where $i = 1, \dots, n$). The resulting type is again a finite map with the same domain as that of its input M ; however each output of the map will be a singleton tuple (due to the use of aggregation) whose fields are $\ell_i : b'_i$ determined by the aggregation clause.

While value comprehension, grouping, lookup and groupwise aggregation have similar typing rules for both collection kinds, key comprehension is less obvious. The reason for this is that in a bag-valued finite map, the output provides multiplicity information, but there is no such information for its input. Consequently, the typing rule for key comprehensions returning bag-valued maps still requires the generator to be a set-valued map; however it is always possible to deduplicate a bag-valued map and perform key comprehension on the resulting set-valued map.

Example 2.1. The query `q_group` from the introduction can be expressed in $\mathcal{NRC}_{\lambda\gamma}$ by a combination of comprehension, singleton finite maps, and groupwise aggregation:

$$M_{\text{group}} := \text{aggBy}_{z.\text{salary} = \text{avg}(z.\text{salary})} \left(\left[\left[\text{dept} = d.\text{name} \right] \triangleright \langle \text{salary} = e.\text{salary} \rangle \right] \text{where } (d.\text{id} = e.\text{dept}) \mid d \leftarrow \text{department}, e \leftarrow \text{employee} \right]$$

$$\begin{array}{c}
551 \\
552 \\
553 \\
554 \\
555 \\
556 \\
557 \\
558 \\
559 \\
560 \\
561 \\
562 \\
563 \\
564 \\
565 \\
566 \\
567 \\
568 \\
569 \\
570 \\
571 \\
572 \\
573 \\
574 \\
575 \\
576 \\
577 \\
578 \\
579 \\
580 \\
581 \\
582 \\
583 \\
584 \\
585 \\
586 \\
587 \\
588 \\
589 \\
590 \\
591 \\
592 \\
593 \\
594 \\
595 \\
596 \\
597 \\
598 \\
599 \\
600 \\
601 \\
602 \\
603 \\
604 \\
605
\end{array}$$

$$\begin{array}{c}
\frac{x : T \in \Gamma}{\Gamma \vdash x : T} \quad \frac{\Sigma(c) = \vec{b}_n \Rightarrow b'}{(\Gamma \vdash M_i : b_i)_{i=1,\dots,n}} \quad \frac{\Sigma(\alpha) = [b]^{\text{bag}} \Rightarrow b'}{\Gamma \vdash M : [b]^{\text{bag}}} \\
\frac{\vec{\ell} \text{ is duplicate-free}}{\Gamma \vdash \overrightarrow{\ell_n = M_n} : \overrightarrow{\ell_n : T_n}} \quad \frac{\Gamma \vdash \rho : P}{\Gamma \vdash \langle \rho \rangle : \langle P \rangle} \quad \frac{\Gamma \vdash M : \langle P \rangle \quad \ell \in \text{dom}(P)}{\Gamma \vdash M.\ell : P(\ell)} \\
\frac{\Gamma, x : S \vdash M : T}{\Gamma \vdash \lambda x.M : S \rightarrow T} \quad \frac{\Gamma \vdash M : S \rightarrow T \quad \Gamma \vdash N : S}{\Gamma \vdash (M N) : T} \quad \frac{\Gamma \vdash M : [T]_P^\delta}{\Gamma \vdash N : \mathbf{B}} \\
\frac{\Gamma \vdash \rho : P}{\Gamma \vdash [\rho]^\delta : [T]_P^\delta} \quad \frac{\Gamma \vdash M : T}{\Gamma \vdash [\rho \triangleright M]^\delta : [T]_P^\delta} \quad \frac{\Gamma \vdash M : [T]_P^\delta \quad \Gamma \vdash N : [T]_P^\delta}{\Gamma \vdash M \#^\delta N : [T]_P^\delta} \quad \frac{\Gamma, x : T \vdash M : [S]_P^\delta}{\Gamma \vdash N : [T]^\delta} \\
\frac{\Gamma \vdash M : [T]_P^{\text{bag}}}{\Gamma \vdash \delta M : [T]_P^{\text{set}}} \quad \frac{\Gamma \vdash M : [T]_P^{\text{set}}}{\Gamma \vdash \iota M : [T]_P^{\text{bag}}} \quad \frac{\Gamma \vdash M : [T]_P^\delta \quad \Gamma, x : T \vdash \rho : P' \quad \rho \text{ not empty}}{\Gamma \vdash \gamma_{x,\rho}^\delta(M) : [T]_{P \oplus P'}^\delta} \quad \frac{\Gamma \vdash M : [T]_P^\delta \quad \Gamma \vdash N : \langle P \rangle \quad P \text{ not empty}}{\Gamma \vdash M \otimes^\delta N : [T]^\delta} \\
\frac{\Gamma, x : \langle P' \rangle \vdash M : [S]_P^\delta}{\Gamma \vdash N : [T]_{P'}^{\text{set}}} \quad \frac{\Gamma \vdash M : \left[\langle P' \rangle \right]_P^\delta}{(\Sigma(\alpha_i) = [b_i]^{\text{bag}} \Rightarrow b'_i)_{i=1,\dots,n} \quad (P(\ell'_i) = b_i)_{i=1,\dots,n}} \\
\Gamma \vdash \left[M \mid k \xrightarrow{\mathcal{K}} N \right]^\delta : [S]_P^\delta \quad \Gamma \vdash \text{aggBy}_{z.\ell_n = \alpha_n(z.\ell'_n)}^\delta(M) : \left[\overrightarrow{\langle \ell_n : b'_n \rangle} \right]_P^\delta
\end{array}$$

Figure 3. Type system of $\mathcal{NRC}_{\lambda\gamma}$.

2.3 Rewrite system

In Figure 4 we give a rewrite system to normalize $\mathcal{NRC}_{\lambda\gamma}$ queries. These include beta reduction for applied functions and a similar rule for projection on a record literal. We import from $\mathcal{NRC}_{\lambda}(Set, Bag)$ all of the normalization rules involving collections (although in this calculus they will be applied to finite maps); among those rules, we note the beta-like rule for contracting value comprehensions whose generators are singletons and the “associativity” rules for unnesting value comprehensions:

$$\begin{array}{c}
592 \\
593 \\
594 \\
595
\end{array}$$

$$\begin{array}{c}
[M \mid x \leftarrow [N]^\delta]^\delta \rightsquigarrow M [N/x] \\
[L \mid y \leftarrow [N \mid x \leftarrow M]^\delta]^\delta \rightsquigarrow [L \mid x \leftarrow M, y \leftarrow N]^\delta
\end{array}$$

with the side condition that, to prevent variable capture, the second rule can only be applied if x does not appear free in L .

We now wonder whether similar rules would work for key comprehension: the beta rule needs to be adapted to use the key in the singleton (rather than the value) as the variable replacement:

$$\begin{array}{c}
603 \\
604 \\
605
\end{array}$$

$$[M \mid k \xrightarrow{\mathcal{K}} [\rho \triangleright N]] \rightsquigarrow M [\langle \rho \rangle / k]$$

As for unnesting, we have an interesting situation: for set-valued maps, any combination of value and key comprehension can be unnested:

$$\begin{array}{c}
639 \\
640 \\
641 \\
642
\end{array}$$

$$[L \mid y \leftarrow \clubsuit [N \mid x \leftarrow \diamond M]^{\text{set}}]^{\text{set}} \rightsquigarrow [L \mid x \leftarrow \diamond M, y \leftarrow \clubsuit N]^{\text{set}}$$

where \diamond and \clubsuit can stand for value or key comprehension, indifferently; on the other hand, when we use bag-valued maps, we are unable to unnest a comprehension out of a key comprehension

$$\begin{array}{c}
643 \\
644 \\
645 \\
646 \\
647 \\
648 \\
649 \\
650
\end{array}$$

$$[L \mid y \xrightarrow{\mathcal{K}} [N \mid x \leftarrow \diamond M]^{\text{set}}]^{\text{bag}} \not\rightsquigarrow [L \mid x \leftarrow \diamond M, y \xrightarrow{\mathcal{K}} N]^{\text{bag}}$$

Recall that the generator of a key comprehension returning a bag-valued map must be a set-valued map: semantically there is an implicit promotion from set to bag, preventing unnesting from being sound. So, a key-comprehension returning a bag-valued map behaves similarly to a value-comprehension whose generator is in the form ιN : even if N is a comprehension, it is wrapped in a ι operation which blocks unnesting.

On the other hand, unnesting a comprehension of either kind out of a value comprehension is semantically sound,

$$\begin{aligned}
& (\lambda x.M) N \rightsquigarrow M[N/x] & \langle \dots, \ell = M, \dots \rangle . \ell \rightsquigarrow M \\
& \gamma_{x,\rho}(\square) \rightsquigarrow \square & \gamma_{x,\rho}([\rho' \triangleright M]) \rightsquigarrow [\rho [M/x] \oplus \rho' \triangleright M] \\
& \gamma_{x,\rho}(M \# N) \rightsquigarrow \gamma_{x,\rho}(M) \# \gamma_{x,\rho}(N) & \gamma_{x,\rho}(\gamma_{x,\rho'}(M)) \rightsquigarrow \gamma_{x,\rho \oplus \rho'}(M) \\
& \gamma_{x,\rho}([L \mid y \stackrel{\diamond}{\leftarrow} M]) \rightsquigarrow [\gamma_{x,\rho}(L) \mid y \stackrel{\diamond}{\leftarrow} M] & \text{(if } y \notin \text{FV}(x.\rho)) \\
& \gamma_{x,\rho}(M \text{ where } N) \rightsquigarrow \gamma_{x,\rho}(M) \text{ where } N \\
& \square \otimes N \rightsquigarrow \square & (L \# M) \otimes N \rightsquigarrow (L \otimes N) \# (M \otimes N) \\
& [L \mid x \stackrel{\diamond}{\leftarrow} M] \otimes N \rightsquigarrow [L \otimes N \mid x \stackrel{\diamond}{\leftarrow} M] & \text{(if } x \notin \text{FV}(N)) \\
& (L \text{ where } M) \otimes N \rightsquigarrow (L \otimes N) \text{ where } M \\
& [\rho \triangleright M] \otimes N \rightsquigarrow [M] \text{ where } (\langle \rho \rangle = N) \\
& \square \# M \rightsquigarrow M & M \# \square \rightsquigarrow M & [\square \mid x \stackrel{\diamond}{\leftarrow} M] \rightsquigarrow \square & [M \mid x \stackrel{\diamond}{\leftarrow} \square] \rightsquigarrow \square \\
& [M \mid x \leftarrow [N]] \rightsquigarrow M[N/x] & [M \mid k \stackrel{\mathcal{K}}{\leftarrow} [\rho \triangleright N]^{\text{set}}] \rightsquigarrow M[\langle \rho \rangle / k] \\
& [M \# N \mid x \stackrel{\diamond}{\leftarrow} R] \rightsquigarrow [M \mid x \stackrel{\diamond}{\leftarrow} R] \# [N \mid x \stackrel{\diamond}{\leftarrow} R] \\
& [M \mid x \stackrel{\diamond}{\leftarrow} N \# R]^{\text{set}} \rightsquigarrow [M \mid x \stackrel{\diamond}{\leftarrow} N]^{\text{set}} \# [M \mid x \stackrel{\diamond}{\leftarrow} R]^{\text{set}} \\
& [M \mid x \leftarrow N \# R]^{\text{bag}} \rightsquigarrow [M \mid x \leftarrow N]^{\text{bag}} \# [M \mid x \leftarrow R]^{\text{bag}} \\
& [L \mid y \stackrel{\star}{\leftarrow} [N \mid x \stackrel{\diamond}{\leftarrow} M]^{\text{set}}]^{\text{set}} \rightsquigarrow [L \mid x \stackrel{\diamond}{\leftarrow} M, y \stackrel{\star}{\leftarrow} N]^{\text{set}} & \text{(if } x \notin \text{FV}(L)) \\
& [L \mid y \leftarrow [N \mid x \stackrel{\diamond}{\leftarrow} M]^{\text{bag}}]^{\text{bag}} \rightsquigarrow [L \mid x \stackrel{\diamond}{\leftarrow} M, y \leftarrow N]^{\text{bag}} & \text{(if } x \notin \text{FV}(L)) \\
& [M \mid x \stackrel{\diamond}{\leftarrow} N \text{ where } R] \rightsquigarrow [M \mid x \stackrel{\diamond}{\leftarrow} N] \text{ where } R \\
& \delta[\square]^{\text{bag}} \rightsquigarrow [\square]^{\text{set}} & \delta[\rho \triangleright M]^{\text{bag}} \rightsquigarrow [\rho \triangleright M]^{\text{set}} & \delta(M \# N) \rightsquigarrow \delta M \# \delta N & \delta \gamma_{x,\rho}(M) \rightsquigarrow \gamma_{x,\rho}(\delta M) \\
& \delta[M \mid x \stackrel{\diamond}{\leftarrow} N]^{\text{bag}} \rightsquigarrow [\delta M \mid x \stackrel{\diamond}{\leftarrow} \delta N]^{\text{set}} & \delta(M \text{ where }^{\text{bag}} N) \rightsquigarrow \delta M \text{ where }^{\text{set}} N \\
& \iota[\square]^{\text{set}} \rightsquigarrow [\square]^{\text{bag}} & \iota[\rho \triangleright M]^{\text{set}} \rightsquigarrow [\rho \triangleright M]^{\text{bag}} & \iota(M \text{ where }^{\text{set}} N) \rightsquigarrow \iota M \text{ where }^{\text{bag}} N & \delta \iota M \rightsquigarrow M \\
& M \text{ where true} \rightsquigarrow M & M \text{ where false} \rightsquigarrow \square \\
& \square \text{ where } M \rightsquigarrow \square \\
& (N \# R) \text{ where } M \rightsquigarrow (N \text{ where } M) \# (R \text{ where } M) \\
& [N \mid x \stackrel{\diamond}{\leftarrow} R] \text{ where } M \rightsquigarrow [N \text{ where } M \mid x \stackrel{\diamond}{\leftarrow} R] & \text{(if } x \notin \text{FV}(M)) \\
& (R \text{ where } M) \text{ where } N \rightsquigarrow R \text{ where } (M \wedge N)
\end{aligned}$$

Figure 4. Normalization rules for $\mathcal{NRC}_{\lambda\gamma}$

hence we can state the rewrite rule:

$$[L \mid y \leftarrow [N \mid x \stackrel{\diamond}{\leftarrow} M]^{\text{bag}}]^{\text{bag}} \rightsquigarrow [L \mid x \stackrel{\diamond}{\leftarrow} M, y \leftarrow N]^{\text{bag}}$$

For similar reasons, in the set-valued case, key comprehension distributes over unions both in head and in generator position, just like value comprehension, but in the bag-valued case, key comprehension only distributes over disjoint unions in head position:

$$\begin{aligned}
& [L \# M \mid x \stackrel{\diamond}{\leftarrow} N]^{\delta} \rightsquigarrow [L \mid x \stackrel{\diamond}{\leftarrow} N]^{\delta} \# [M \mid x \stackrel{\diamond}{\leftarrow} N]^{\text{set}} \\
& [L \mid x \stackrel{\diamond}{\leftarrow} M \# N]^{\text{set}} \rightsquigarrow [L \mid x \stackrel{\diamond}{\leftarrow} M]^{\text{set}} \# [L \mid x \stackrel{\diamond}{\leftarrow} N]^{\text{set}} \\
& [L \mid x \leftarrow M \# N]^{\text{bag}} \rightsquigarrow [L \mid x \leftarrow M]^{\text{bag}} \# [L \mid x \leftarrow N]^{\text{bag}} \\
& [L \mid k \stackrel{\mathcal{K}}{\leftarrow} M \# N]^{\text{bag}} \rightsquigarrow [L \mid k \stackrel{\mathcal{K}}{\leftarrow} M]^{\text{bag}} \# [L \mid k \stackrel{\mathcal{K}}{\leftarrow} N]^{\text{bag}}
\end{aligned}$$

Finally, we have a look at the reduction rules involving grouping and lookup. They state that:

- empty collections are absorbent elements for both operations;
- grouping over a singleton maps $(\gamma_{x,\rho}([\rho' \triangleright M]))$ extends the key ρ' of the singleton with new fields obtained from the grouping criterion ρ , instantiated on the value M ;
- lookup on a singleton map $([\rho \triangleright M] \otimes N)$ reduces to a conditional statement evaluating to the singleton collection containing M if and only if $\langle \rho \rangle = N$, and to an empty collection otherwise;
- grouping and lookup commute with unions and disjoint unions, with the head of a comprehension, and with the left-hand argument of a **where**;

- nested γ s can be merged by merging their grouping criteria;

Notice that there are no rewrite rules involving groupwise aggregation. Remember that the main purpose of our rewrite system is to simplify away unnecessarily nested intermediate collections: since the input and output types of aggBy are already flat, no rewrite rules are needed for our purposes (however, its argument M may still be involved in rewrites).

Remark. We can easily express the \mathcal{G} -comprehension from the introduction using key comprehension and finite map application

$$\left[R \mid (k, v) \stackrel{\mathcal{G}}{\leftarrow} M \right] := \left[R \mid R \otimes k/v \mid k \stackrel{\mathcal{K}}{\leftarrow} M \right]$$

The reason why we adopt $\stackrel{\mathcal{K}}{\leftarrow}$ as primitive instead of $\stackrel{\mathcal{G}}{\leftarrow}$ is that it has better algebraic properties. It can be shown that $\stackrel{\mathcal{G}}{\leftarrow}$ does not distribute over unions:

$$\left[R \mid (k, v) \stackrel{\mathcal{G}}{\leftarrow} M \uplus N \right] \not\sim \left[R \mid (k, v) \stackrel{\mathcal{G}}{\leftarrow} M \right] \uplus \left[R \mid (k, v) \stackrel{\mathcal{G}}{\leftarrow} N \right]$$

By unfolding the definition of $\stackrel{\mathcal{G}}{\leftarrow}$, we see that, in order for such a rule to be sound, we would need substitution to commute with unions, but that is generally not the case ($R \mid M \uplus N/v \not\equiv R \mid M/v \uplus R \mid N/v$).

3 Strong normalization

We now move to proving an important termination property of the rewrite system presented in the previous section: all well-typed $\mathcal{NRC}_{\lambda\gamma}$ terms will eventually be reduced to their normal form in a finite number of steps, regardless of the reduction strategy. This property is known as *strong normalization*.

Our proof is derived as a corollary of the strong normalization proof for the smaller calculus $\mathcal{NRC}_{\lambda}(Set, Bag)$, via a translation. We prove that every well-typed $\mathcal{NRC}_{\lambda\gamma}$ term can be translated to a well-typed $\mathcal{NRC}_{\lambda}(Set, Bag)$ term, and to every reduction step on well-typed $\mathcal{NRC}_{\lambda\gamma}$ terms there correspond one or more reduction steps on $\mathcal{NRC}_{\lambda}(Set, Bag)$ via the same translation procedure. Then, since there are no infinite reduction sequences in $\mathcal{NRC}_{\lambda}(Set, Bag)$, there can be no infinite reduction sequences in $\mathcal{NRC}_{\lambda\gamma}$.

The embedding of $\mathcal{NRC}_{\lambda\gamma}$ into $\mathcal{NRC}_{\lambda}(Set, Bag)$ is defined in Figure 5. The rule for grouping is type directed because we need to know the row associated to the grouping input (or in other words, the labels ℓ within the tuple x .1). A minor complication is the translation of aggregations (both simple and groupwise), since they have no counterpart in $\mathcal{NRC}_{\lambda}(Set, Bag)$; we map both of them to a distinguished unary constant \bullet – this is sufficient because aggregations do not participate in reductions.

Lemma 3.1. *If $\Gamma \vdash M : T$ in $\mathcal{NRC}_{\lambda\gamma}$, then $[\Gamma] \vdash [M] : [T]$ in $\mathcal{NRC}_{\lambda}(Set, Bag)$.*

Theorem 3.2. *Whenever $M \rightsquigarrow N$ in $\mathcal{NRC}_{\lambda\gamma}$, we have that $[M] \overset{\dagger}{\rightsquigarrow} [N]$ in $\mathcal{NRC}_{\lambda}(Set, Bag)$.*

The proof is by induction and case analysis on the reduction rule, unfolding the definition of the embedding as needed. Some care is needed when handling reduction rules involving $\gamma_{x,\rho}((M))$, due to the more involved definition of the corresponding embedding. More details of the proof are given in the appendix.

Theorem 3.3. *If M is a well-typed $\mathcal{NRC}_{\lambda\gamma}$ term, there is no infinite reduction sequence starting with it.*

Proof. By Theorem 3.2, every infinite $\mathcal{NRC}_{\lambda\gamma}$ reduction sequence starting with M can be simulated by an infinite $\mathcal{NRC}_{\lambda}(Set, Bag)$ reduction sequence starting with $[M]$. By Lemma 3.1, since M is well-typed in $\mathcal{NRC}_{\lambda\gamma}$, $[M]$ is well-typed in $\mathcal{NRC}_{\lambda}(Set, Bag)$. But given that well-typed $\mathcal{NRC}_{\lambda}(Set, Bag)$ terms are strongly normalizing, there can be no infinite reduction starting with $[M]$. Consequently, M must be strongly normalizing in $\mathcal{NRC}_{\lambda\gamma}$. \square

4 Grammar of normal forms

The remaining results of this paper rely on the structure of $\mathcal{NRC}_{\lambda\gamma}$ normal forms. To ease the reasoning on such terms, it will be useful to derive a grammar of normal forms, which can be used to perform case analysis and recursion without considering all of the cases of the general grammar of terms, many of which are precluded by the fact that the term has been normalized and must therefore be pruned.

We actually go a step further and add to normalization a few “administrative” rules and assumptions: in so doing, we effectively present “standardized” normal forms for which the following conditions are met:

- n -ary unions $C_1 \uplus^{\text{set}} \dots \uplus^{\text{set}} C_n$ and $D_1 \uplus^{\text{bag}} \dots \uplus^{\text{bag}} D_n$ are represented as $\bigcup \vec{C}$ and $\biguplus \vec{D}$ respectively; for empty \vec{C} or \vec{D} , their grand unions stand for empty collections $[\]^{\text{set}}$, $[\]^{\text{bag}}$;
- variables with record type and comprehensions whose head is not a singleton are kept in eta-long form using the following rules:

$$\frac{x : \langle \vec{\ell} : \vec{T} \rangle}{x \rightsquigarrow \langle \vec{\ell} = x.\vec{\ell} \rangle}$$

$$\frac{M : [T]_{\ell:S}^S \quad N = \left[\vec{\ell} = k.\vec{\ell} \triangleright v \right]}{\left[M \text{ where } X \mid \vec{F} \right]^S \rightsquigarrow \left[N \text{ where } X \mid \vec{F}, k \stackrel{\mathcal{K}}{\leftarrow} M, v \leftarrow M \otimes k \right]^S}$$

- comprehensions without a guard are considered syntactically equal to those with a trivial guard:

$$\left[[\rho \triangleright M] \mid x \stackrel{\diamond}{\leftarrow} N \right] = \left[[\rho \triangleright M] \text{ where true } \mid x \stackrel{\diamond}{\leftarrow} N \right]$$

$$\begin{aligned}
881 \quad [b] &= b & [S \rightarrow T] &= [S] \rightarrow [T] & [\langle P \rangle] &= \langle [P] \rangle & 936 \\
882 \quad [[T]_P^\delta] &= [\langle \langle [P] \rangle, [T] \rangle]^\delta & [\overrightarrow{\ell : T}] &= \overrightarrow{\ell : [T]} & 937 \\
883 & & & & 938 \\
884 & & & & 939 \\
885 \quad [c(\overrightarrow{M})] &= c(\overrightarrow{M}) & [\langle \rho \rangle] &= \langle [\rho] \rangle & [M.\ell] &= [M] . \ell & [\overrightarrow{\ell = M}] &= \overrightarrow{\ell = [M]} & 940 \\
886 \quad [\lambda x^T.M] &= \lambda x^{[T]}. [M] & [MN] &= [M] [N] & [M \text{ where }^\delta N] &= [M] \text{ where }^\delta [N] & 941 \\
887 \quad [[]^\delta] &= []^\delta & [[\rho \triangleright M]^\delta] &= [\langle \langle [\rho] \rangle, [M] \rangle]^\delta & [M \text{ ++ }^\delta N] &= [M] \text{ ++ }^\delta [N] & 942 \\
888 \quad [\delta M] &= \delta [M] & [\iota M] &= \iota [M] & [[M \mid x \leftarrow N]^\delta] &= [[M] [x.2/x] \mid x \leftarrow [N]]^\delta & 943 \\
889 \quad [[M \mid k \xleftarrow{\mathcal{K}} N]^\text{set}] &= [[M] \mid k \leftarrow \text{dom}([N])]^\text{set} & [[M \mid k \xleftarrow{\mathcal{K}} N]^\text{bag}] &= [[M] \mid k \leftarrow \text{idom}([N])]^\text{bag} & 944 \\
890 & & & & 945 \\
891 & & & & 946 \\
892 \quad [\gamma_{x,\rho}^\delta(M)] &= [[\langle \overrightarrow{\ell = x.1.\ell} \oplus [\rho] [x.2/x], x.2 \rangle]^\delta \mid x \leftarrow [M]]^\delta & 947 \\
893 & & & & 948 \\
894 \quad [M \otimes^\delta N] &= [[z.2]^\delta \text{ where }^\delta z.1 = [N] \mid z \leftarrow [M]]^\delta & (z \notin \text{FV}(N)) & 949 \\
895 \quad [\alpha(M)] &= [\text{aggBy}_{z.\ell_n = \alpha_n(z.\ell'_n)}^\delta (M)] = \bullet([M]) & \text{dom}(M) &= [[z.1]^\text{set} \mid z \leftarrow M]^\text{set} & 950 \\
896 & & & & 951 \\
897 & & & & 952 \\
898 & & & & 953 \\
899 & & & & 954 \\
900 & & & & 955
\end{aligned}$$

Figure 5. Embedding of $\mathcal{NRC}_{\lambda_Y}$ into $\mathcal{NRC}_{\lambda}(\text{Set}, \text{Bag})$.

- singletons that do not appear as the head of a comprehension are represented as trivial comprehensions

$$[\rho \triangleright M] = [[\rho \triangleright M] \mid]$$

Similar rules have been considered in previous versions of \mathcal{NRC} . Once the above additional rules are considered, the normal forms of the resulting rewrite system are described by the grammar in Figure 6, as stated by the following result.

Theorem 4.1. *Every well-typed term M of $\mathcal{NRC}_{\lambda_Y}$ in normal form generated by the grammar in Figure 6.*

The proof is by induction on the structure of M , noticing that the induction hypothesis states that the subterms of the M for which we are proving this theorem are generated by the grammar.

4.1 Normal forms of relational maps

For the purpose of generating SQL queries, we are not interested in all the terms of $\mathcal{NRC}_{\lambda_Y}$ in their full generality, but we will instead focus on particular types of terms that we intend to translate to SQL: we call these terms nested and flat relational maps.

A nested relational map (**nrm**) is a term of finite map type whose keys are rows of basic type, and whose values are, inductively, tuples of nested relational maps:

$$\frac{\rho = \overrightarrow{k : b} \quad (T_i \text{ nrm})}{\left[\langle \overrightarrow{\ell : T} \rangle \right]_\rho \text{ nrm}} \quad \frac{\rho = \overrightarrow{k : b}}{\left[\langle \overrightarrow{\ell : b'} \rangle \right]_\rho \text{ nrm}}$$

Under the additional condition that function typed terms should not appear as the argument of deduplication δ and promotion ι (as in $\mathcal{NRC}_{\lambda}(\text{Set}, \text{Bag})$), we can easily see that all function terms are normalized away and all the variables bound by comprehensions (equivalently, all bound variables *tout court*, since there can be no lambda abstractions in the

General NF	$M ::= B \mid U \mid W \mid Q \mid R$	956
	$\rho ::= \overrightarrow{\ell = M}$	957
Indeterminate NF	$X ::= x \mid X.\ell \mid XM$	958
Base-typed NF	$B ::= X \mid c(\overrightarrow{B}) \mid \alpha(R)$	959
Tuple-typed NF	$U ::= X \mid \langle \rho \rangle$	960
Function-typed NF	$W ::= X \mid \lambda x.M$	961
Set-typed NF	$Q ::= []^\text{set} \mid C \mid Q \text{ ++ }^\text{set} Q$	962
	$C ::= H \mid [C \mid F]^\text{set}$	963
	$H ::= I \mid I \text{ where }^\text{set} B \mid Y$	964
	$I ::= N \mid [\rho \triangleright M]^\text{set}$	965
	$N ::= X \mid \gamma_{x,\rho}^\text{set}(N) \mid N \otimes^\text{set} U$	966
	$\mid \delta X \mid \delta t \mid Y$	967
	$F ::= x \leftarrow N \mid k \xleftarrow{\mathcal{K}} N$	968
	$Y ::= \text{aggBy}_{z.\ell = \alpha(z.\ell')}^\text{set}(Q)$	969
Bag-typed NF	$R ::= []^\text{bag} \mid D \mid R \text{ ++ }^\text{bag} R$	970
	$D ::= J \mid [D \mid G]^\text{bag} \mid Z$	971
	$J ::= L \mid L \text{ where }^\text{bag} B$	972
	$L ::= O \mid [\rho \triangleright M]^\text{bag}$	973
	$O ::= X \mid \gamma_{x,\rho}^\text{bag}(O) \mid O \otimes^\text{bag} U$	974
	$\mid t \mid \iota Q \mid Z$	975
	$G ::= x \leftarrow O \mid k \xleftarrow{\mathcal{K}} Q$	976
	$Z ::= \text{aggBy}_{z.\ell = \alpha(z.\ell')}^\text{bag}(R)$	977
		978
		979
		980
		981
		982
		983
		984
		985
		986
		987
		988
		989
		990

Figure 6. Normal forms of $\mathcal{NRC}_{\lambda_Y}$

normal forms of relational maps) have a tuple type containing either base types or, inductively, nested relational maps.

Thanks to these considerations, we can derive a further simplified grammar of nested relational maps, described by Figure 7. This grammar is made slightly more compact by

991 $M ::= B \mid Q \mid R$
 992 $\rho ::= \overrightarrow{\ell = M}$
 993 $\rho^* ::= \overrightarrow{\ell = \vec{B}}$
 994 $B ::= x.\ell \mid c(\vec{B}) \mid \alpha(R)$
 995 $Q ::= \bigcup \vec{C}$
 996 $C ::= \left[[\rho^* \triangleright \langle \rho \rangle]^{\text{set}} \text{ where }^{\text{set}} B \mid \vec{F} \right] \mid Y$
 997 $N ::= x.\ell \mid \gamma_{x,\rho}^{\text{set}}(N) \mid N \otimes^{\text{set}} \langle \rho^* \rangle \mid \delta(x.\ell) \mid \delta t \mid Y$
 998 $Y ::= \text{aggBy}_{z.\ell=\alpha(z.\ell')}^{\text{set}}(Q^*)$
 999 $F ::= x \leftarrow N \mid k \xleftarrow{\mathcal{K}} N$
 1000 $R ::= \biguplus \vec{D}$
 1001 $D ::= \left[[\rho^* \triangleright \langle \rho \rangle]^{\text{bag}} \text{ where }^{\text{bag}} B \mid \vec{G} \right]^{\text{bag}} \mid Z$
 1002 $O ::= x.\ell \mid \gamma_{x,\rho}^{\text{set}}(O) \mid O \otimes^{\text{bag}} \langle \rho^* \rangle \mid t \mid \iota Q^* \mid Z$
 1003 $G ::= x \leftarrow O \mid k \xleftarrow{\mathcal{K}} Q^*$
 1004 $Z ::= \text{aggBy}_{z.\ell=\alpha(z.\ell')}^{\text{bag}}(R^*)$

1005 (Q^* and R^* are defined similarly to Q and R , but replacing
 1006 deep references to ρ (i.e. those in C and D) with ρ^* .)

1007 **Figure 7.** Nested relational map normal forms of $\mathcal{NRCE}_{\lambda\gamma}$

1008 means of syntactic sugar for multi-generator comprehensions, representing as usual head-nested comprehensions.

1009 A final simplification consists in limiting relational maps to the flat case, in which no collection nesting on the values of a map is allowed: in other words, flat relational maps are terms whose type is in the form $\left[\overrightarrow{\langle \ell : b' \rangle} \right]_{k:b}$.

1010 Applying this constraint to the grammar for nested relational normal forms we obtain the grammar in Figure 8.

1011 5 Translation to SQL

1012 The grammar of normal forms is amenable to be used to define algorithmic procedures, including a translation of $\mathcal{NRCE}_{\lambda\gamma}$ normal forms to SQL; clearly, only flat normal forms need to be translated, as nested queries do not have a direct representation in SQL due to typing limitations. We devised one such translation and used it to implement language-integrated database queries with grouping and aggregation in the Links programming language [7]: due to space constraints, we give here a high level description of the translation of comprehension forms, referring the reader to the appendix for further details.

1013 To translate comprehensions, our procedure operates by returning a complete **SELECT ... FROM ... WHERE ...** statement: for instance

$$\begin{aligned}
 & \left(\left[[\rho \triangleright \langle \rho' \rangle]^{\text{bag}} \text{ where }^{\text{bag}} X \mid \vec{G} \right]^{\text{bag}} \right)^{\text{sql}} \\
 & = \text{SELECT } (\rho)_{\mathcal{K}}^{\text{sql}}, (\rho')_{\mathcal{V}}^{\text{sql}} \text{ FROM } (G)^{\text{sql}} \text{ WHERE } (X)^{\text{sql}}
 \end{aligned}$$

1046 $M ::= B \mid Q^* \mid R^*$
 1047 $\rho^* ::= \overrightarrow{\ell = \vec{B}}$
 1048 $B ::= x.\ell \mid c(\vec{B}) \mid \alpha(R^*)$
 1049 $Q^* ::= \bigcup \vec{C}^*$
 1050 $C^* ::= \left[[\rho^* \triangleright \langle \rho^* \rangle]^{\text{set}} \text{ where }^{\text{set}} B \mid \vec{F}^* \right] \mid Y$
 1051 $N^* ::= \gamma_{x,\rho^*}^{\text{set}}(N^*) \mid N^* \otimes^{\text{set}} \langle \rho^* \rangle \mid \delta t \mid Y$
 1052 $F^* ::= x \leftarrow N^* \mid k \xleftarrow{\mathcal{K}} N^*$
 1053 $Y ::= \text{aggBy}_{z.\ell=\alpha(z.\ell')}^{\text{set}}(Q^*)$
 1054 $R^* ::= \biguplus \vec{D}^*$
 1055 $D^* ::= \left[[\rho^* \triangleright \langle \rho^* \rangle]^{\text{bag}} \text{ where }^{\text{bag}} B \mid \vec{G}^* \right]^{\text{bag}} \mid Z$
 1056 $O^* ::= \gamma_{x,\rho^*}^{\text{set}}(O^*) \mid O^* \otimes^{\text{bag}} \langle \rho^* \rangle \mid t \mid \iota Q^* \mid Z$
 1057 $G^* ::= x \leftarrow O^* \mid k \xleftarrow{\mathcal{K}} Q^*$
 1058 $Z ::= \text{aggBy}_{z.\ell=\alpha(z.\ell')}^{\text{bag}}(R^*)$

1059 **Figure 8.** Flat relational map normal forms of $\mathcal{NRCE}_{\lambda\gamma}$

1060 (with or without **DISTINCT**, depending on whether we are performing a set or a bag comprehension); the **where** clause is translated to the **WHERE** statement, and the comprehension generators are translated to the **FROM** clause; that leaves the singleton key-value pair – namely, $[\rho \triangleright \langle \rho' \rangle]$; both ρ and ρ' will go to the **SELECT** statement, but ρ shall use the auxiliary $(\rho)_{\mathcal{K}}^{\text{sql}}$ translation, marking that row as a key, while ρ' shall use the $(\rho')_{\mathcal{V}}^{\text{sql}}$ translation, marking it as a value row. Concretely, the two different translations rename the attributes using prefixes that univocally identify them as keys or values. Generators G are translated differently depending on whether they represent value comprehensions or key comprehensions: value comprehensions use a simple recursive call $(G)^{\text{sql}}$; key comprehensions use a specialized auxiliary translation that removes the prefix from the attributes of $(G)^{\text{sql}}$ that are marked as keys, and drops the attributes marked as values entirely.

1061 5.1 Query shredding and tabular functions

1062 Previous work [4] introduced *query shredding*, a technique which allows database queries with a nested collection type to be decomposed (i.e., “shredded”) into multiple flat queries which can be expressed in SQL, thus run by a typical SQL-based DBMS, yielding partial results that are then *stitched* back together into the desired final value having a nested relational type. More recently, [22] proposed $\mathcal{NRCE}_{\mathcal{G}}$, an extension of \mathcal{NRCE} with tabular functions which greatly simplifies the study of query shredding. $\mathcal{NRCE}_{\mathcal{G}}$ shares some similarities with the finite maps resulting from the grouping operators of our $\mathcal{NRCE}_{\lambda\gamma}$. In $\mathcal{NRCE}_{\mathcal{G}}$, tabular functions are created using the graph operators $\mathcal{G}^{\text{set}}(-; -)$ and $\mathcal{G}^{\text{bag}}(-; -)$. We compare the graph operator with the grouping operator

from $\mathcal{NRC}_{\lambda\gamma}$ in the set case (note that we have taken the liberty of adapting the types the subterms of $\mathcal{G}^{\text{set}}(-; -)$ to match them to the style used in this paper):

$$\frac{\Gamma \vdash L : [\vec{P}]^{\text{set}} \quad \Gamma \vdash M : [T]^{\text{set}}}{\Gamma, x : \vec{P} \vdash M : [T]^{\text{set}}} \quad \frac{\Gamma, x : \vec{T} \vdash \rho : P}{\Gamma \vdash \gamma_{x,\rho}^{\text{set}}(M) : [T]_P^{\text{set}}}$$

On the right-hand side, the grouping constructor $\gamma_{x,\rho}^{\text{set}}(M)$ works by producing a different key ρ for each element x of the input collection M : this effectively groups the elements of M sharing the same key ρ ; on the left-hand side, the graph constructor $\mathcal{G}^{\text{set}}(x \leftarrow L; M)$ takes the elements x of domain collection L and returns for each of them a codomain collection M which may depend on x : we may view this finite map as a “grouping”, where the grouping keys are the elements of L .

The tabular functions of $\mathcal{NRC}_{\mathcal{G}}$ cannot be used in aggregations, but they do allow access to “groups” using the lookup syntax \otimes (which we borrowed for $\mathcal{NRC}_{\lambda\gamma}$).

While the lookup operation plays a very similar role in the two calculi, the introduction rules for tabular functions/finite maps are rather dissimilar. Despite the superficial differences, we can show that each operation can be expressed in terms of the other:

$$\mathcal{G}^{\text{set}}(x \leftarrow L; M) ::= \left[\gamma_{x,\rho_x}^{\text{set}}(M) \mid x \leftarrow L \right]$$

$$\gamma_{x,\rho}^{\text{set}}(M) ::= \left[\mathcal{G}^{\text{set}}(_ \leftarrow [\rho]^{\text{set}}; [x]^{\text{set}}) \mid x \leftarrow M \right]$$

(where $_$ is any variable name chosen to be fresh with respect to its scope, and ρ_x is the row corresponding to the eta-expansion of the tuple-typed x , defined as $\vec{\ell} = x.\vec{\ell}$, whenever x has type $\langle \vec{\ell} : \vec{S} \rangle$; additionally, in the second equation above we have extended the $\mathcal{NRC}_{\mathcal{G}}$ notion of comprehension to allow it to return a tabular function, which was not allowed in the original formulation but is a straightforward extension nonetheless).

Therefore, $\mathcal{NRC}_{\lambda\gamma}$ is very similar to $\mathcal{NRC}_{\mathcal{G}}$; however, $\mathcal{NRC}_{\mathcal{G}}$ did not provide any rewrite system and was largely intended as syntactic sugar that could be translated to plain \mathcal{NRC}_{λ} . On the contrary, $\mathcal{NRC}_{\lambda\gamma}$ provides its own notion of reduction and its normal forms can be directly used to implement language-integrated database queries.

5.2 A query shredding judgment for $\mathcal{NRC}_{\lambda\gamma}$

Since $\mathcal{G}^{\text{set}}(-; -)$ can be intended as a defined operator within $\mathcal{NRC}_{\lambda\gamma}$, we wonder whether it is possible to extend query shredding to $\mathcal{NRC}_{\lambda\gamma}$, thus allowing one to run nested relational queries involving grouping operations. The answer is affirmative and it involves, as it turns out, a limited redesign of the shredding judgment from [22], of which we show some key rules in Figure 9. As in the case of $\mathcal{NRC}_{\lambda}(\text{Set}, \text{Bag})$, the shredding algorithm performs structural recursion over the grammar of nested relational normal

$$\frac{\varphi \notin \text{dom}(\Psi) \quad \Phi; \Theta, \vec{F} \vdash \rho \Rightarrow \vec{\rho} \mid \Psi}{\Phi; \Theta \vdash \left[[\rho^* \triangleright \rho]^{\text{set}} \text{ where } B \mid \vec{F} \right]^{\text{set}} \Rightarrow \varphi \otimes^{\text{set}} \langle \text{dom}(\Theta) \rangle \mid \Psi[\varphi \mapsto \mathcal{G}(\Theta; [\rho^* \triangleright \vec{\rho}]^{\text{set}} \text{ where } B \mid \vec{F}^{\text{set}})]}$$

$$\frac{\varphi \notin \text{dom}(\Psi) \quad \Phi; \Theta, \vec{G} \vdash \rho \Rightarrow \vec{\rho} \mid \Psi}{\Phi_0; \Theta \vdash \left[[\rho^* \triangleright \rho]^{\text{bag}} \text{ where } B \mid \vec{G} \right]^{\text{bag}} \Rightarrow \varphi \otimes^{\text{bag}} \langle \text{dom}(\Theta) \rangle \mid \Psi[\varphi \mapsto \mathcal{G}(\Theta; [\rho^* \triangleright \vec{\rho}]^{\text{bag}} \text{ where } B \mid \vec{G}^{\text{bag}})]}$$

$$(x \leftarrow O)^\delta \triangleq \begin{cases} x \leftarrow Q^* & \text{if } O = {}_t Q^* \\ x \leftarrow \delta O & \text{else} \end{cases} \quad \Phi \setminus \vec{\psi} \triangleq [(\varphi \mapsto N) \in \Phi \mid \varphi \notin \vec{\psi}]$$

$$(k \xleftarrow{\mathcal{X}} Q^*)^\delta \triangleq Q^*$$

Figure 9. Shredding rules.

forms; in $\mathcal{NRC}_{\lambda\gamma}$, however, we will need to consider some more cases for nested relational map normal forms that are not pure collections.

The *shredding judgment* describes the process by which, given a normalized $\mathcal{NRC}_{\lambda\gamma}$ query, each of its subqueries having a nested finite map type is lifted (in a manner analogous to lambda-lifting [12]) to an independent finite map query: more specifically, shredding will produce a *shredding environment* (denoted by Φ, Ψ, \dots), which is a finite map associating special *graph variables* φ, ψ to $\mathcal{NRC}_{\lambda\gamma}$ terms:

$$\Phi, \Psi, \dots ::= [\overrightarrow{\varphi \mapsto M}]$$

The shredding judgment has the following form:

$$\Phi; \Theta \vdash M \Rightarrow \check{M} \mid \Psi$$

where the \Rightarrow symbol separates the input (to the left) from the output (to the right). The normalized $\mathcal{NRC}_{\lambda\gamma}$ term M is the query that is being considered for shredding; M may contain free variables declared in Θ , which must be a sequence of $\mathcal{NRC}_{\lambda}(\text{Set}, \text{Bag})$ set comprehension bindings. Θ is initially empty, but during shredding it is extended with parts of the input that have already been processed. Similarly, the input shredding environment Φ is initially empty, but will grow during shredding to collect shredded queries that have already been generated.

The output of shredding consists of a shredded term \check{M} and an output shredding environment Ψ . Ψ extends Φ with the new queries obtained by shredding M ; \check{M} is an output $\mathcal{NRC}_{\lambda\gamma}$ query obtained from M by lifting its collection typed subqueries to independent flat queries defined in Ψ .

The shredding of finite maps in normal form (i.e. unions, comprehensions, and groupwise aggregations) is performed by means of *query lifting*: we turn the collection into a globally defined (graph) query, which will be associated to a fresh name φ and instantiated to the local comprehension context by graph application. This operation converts local subterms

into global graphs: thus, when shredding a map, besides processing its subterms recursively, we will need to extend the output shredding environment with a definition for the new global graph φ . In the interesting case of comprehensions, φ is defined by graph-abstracting over the comprehension context Θ ; notice that, since we are only shredding normalized terms, we are allowed to reason on the limited number of cases allowed by the grammar of normal forms and, in particular, it is easy to see that the judgment for bag comprehensions must ensure that generators \vec{G} be converted into sets using the operation G^δ .

6 Related work

This paper follows a line of research at the intersection of programming languages and databases encompassing over three decades. Inspired by Trinder and Wadler's work on understanding database queries as a form of monadic comprehension syntax [31], Buneman et al. developed the nested relational calculus [1, 2, 29] with nested collections and first-order functions. Wong [32] proved the *conservativity theorem* stating that any query computable in that language can be expressed without resorting to subcomputations with a greater degree of nesting than that of its input or output (whichever is greater), and provided a strongly normalizing rewrite system to simplify queries involving subcomputations of an unnecessarily nested type; this also implied that flat-flat \mathcal{NRC} queries could be translated to SQL, allowing a practical implementation [33].

By adding higher-order functions, Cooper extended \mathcal{NRC} to a superset of the simply-typed lambda calculus [6]. Support for *heterogeneous* queries mixing collection kinds such as sets and bags, initially advocated by Grust and Scholl [10], was added by Ricciotti and Cheney who also gave a proof of strong normalization of their rewrite system [21, 24]. Advancements in the formal semantics of SQL by Guagliardo and Libkin [11] also made it possible to give a mechanized proof that \mathcal{NRC} queries using sets and bags can be translated to SQL with **LATERAL** joins [23].

Further research developed language-integrated query, i.e. techniques to integrate a domain-specific database query sublanguage into a general-purpose host language. Microsoft gave a commercial implementation, termed LINQ [16, 28], and Cheney et al. showed how to use normalization techniques to improve the reliability and performance of LINQ queries [3]. Language-integrated query was also implemented, in a form closely based on \mathcal{NRC} , in the Links programming language [15], which uses an effect system to identify computations that can be run as database queries. Furthermore, language-integrated query is available in Scala and Haskell through libraries such as Quill [20] and DSH [9].

Particularly relevant to our paper is the research devoted to translating nested relational queries to multiple flat SQL queries. A shredding technique was devised by Cheney et

al. [4] to accomplish this goal in \mathcal{NRC} with bag semantics; shredding has also proved useful to achieve greater parallelism in large-scale distributed query processing [26]. A rationalization of shredding by Ricciotti and Cheney called *query lifting* [22] using ideas from the well-understood concept of lambda-lifting makes it possible to process queries mixing sets and bags; the definition of query lifting employs finite maps (partially inspired by [8]), which are used as the basis of this paper.

Various other works have proposed query calculi or language-integrated query facilities offering grouping and aggregation, including Libkin and Wong's \mathcal{BQL} [14], Suzuki et al.'s \mathcal{QUEL} [13, 27], and Okura and Kameyama's \mathcal{QELG} [18, 19]. Unlike $\mathcal{NRC}_{\lambda\gamma}$, none of these works supports nested collections or first-class grouping independent of aggregation, but \mathcal{QELG} provides optimization techniques to produce more efficient SQL queries, which we view as complementary to our proposal. To enable the optimization of hybrid database and linear algebra workload, Shaikha et al. developed *semi-ring dictionaries* [25], which are conceptually similar to our finite maps; their work even proposes a variant of \mathcal{NRC} allowing independent grouping, but crucially does not provide a rewrite system or normalization.

7 Conclusion

Our calculus $\mathcal{NRC}_{\lambda\gamma}$ provides a formalism which can be used to develop a principled implementation of language-integrated query with grouping and aggregation. Unlike other proposals, we allow grouping to happen independently of aggregation: this allows queries to be expressed in a way that is more natural in a general-purpose programming language, while at the same time making it possible to convert such queries to idiomatic SQL queries. Commercial implementations such as Microsoft's LINQ are known to fail on grouping queries due to an imperfect understanding of their theory [17]: we believe that $\mathcal{NRC}_{\lambda\gamma}$ provides as a rational basis to reason on such queries and are using it in an extension the Links programming language.

At the moment, $\mathcal{NRC}_{\lambda\gamma}$ does not address the efficient execution of queries: its normal forms are considerably more involved than those of \mathcal{NRC}_λ or even $\mathcal{NRC}_\lambda(\text{Set}, \text{Bag})$, suggesting that there is probably space for optimizations, particularly in the case of complex queries obtained compositionally from simpler ones, as shown by Okura and Kameyama in the context of the \mathcal{QELG} language [19]: we do expect to be able to translate their optimizations to $\mathcal{NRC}_{\lambda\gamma}$, and the experimental evaluation of this (possibly together with more optimizations) will be the subject of our future work.

References

- [1] Buneman, P., Libkin, L., Suciu, D., Tannen, V., Wong, L.: Comprehension syntax. *SIGMOD Record* 23 (1994)
- [2] Buneman, P., Naqvi, S., Tannen, V., Wong, L.: Principles of programming with complex objects and collection types. *Theor. Comput. Sci.*

- 1321 149(1) (1995). [https://doi.org/10.1016/0304-3975\(95\)00024-Q](https://doi.org/10.1016/0304-3975(95)00024-Q)
- 1322 [3] Cheney, J., Lindley, S., Wadler, P.: A practical theory of language-integrated query. In: ICFP (2013). <https://doi.org/10.1145/2500365.2500586>
- 1323 [4] Cheney, J., Lindley, S., Wadler, P.: Query shredding: efficient relational evaluation of queries over nested multisets. In: SIGMOD. pp. 1027–1038. ACM (2014). <https://doi.org/10.1145/2588555.2612186>
- 1324 [5] Cooper, E.: Programming language features for web application development. Ph.D. thesis, University of Edinburgh (2009)
- 1325 [6] Cooper, E.: The script-writer’s dream: How to write great SQL in your own language, and be sure it will succeed. In: DBPL (2009). https://doi.org/10.1007/978-3-642-03793-1_3
- 1326 [7] Cooper, E., Lindley, S., Wadler, P., Yallop, J.: Links: web programming without tiers. In: FMCO (2007). https://doi.org/10.1007/978-3-540-74792-5_12
- 1327 [8] Gibbons, J., Henglein, F., Hinze, R., Wu, N.: Relational algebra by way of adjunctions. Proc. ACM Program. Lang. 2(ICFP) (Jul 2018). <https://doi.org/10.1145/3236781>
- 1328 [9] Giorgidze, G., Grust, T., Schreiber, T., Weijers, J.: Haskell boards the Ferry - database-supported program execution for Haskell. In: IFL. pp. 1–18. No. 6647 in LNCS, Springer-Verlag (2010)
- 1329 [10] Grust, T., Scholl, M.H.: How to comprehend queries functionally. J. Intell. Inf. Syst. 12(2-3), 191–218 (1999). <https://doi.org/10.1023/A:1008705026446>
- 1330 [11] Guagliardo, P., Libkin, L.: A formal semantics of SQL queries, its validation, and applications. PVLDB (2017). <https://doi.org/10.14778/3151113.3151116>
- 1331 [12] Johnsson, T.: Lambda lifting: Treansforming programs to recursive equations. In: FPCA. pp. 190–203 (1985). https://doi.org/10.1007/3-540-15975-4_37
- 1332 [13] Kiselyov, O., Katsushima, T.: Sound and efficient language-integrated query - maintaining the ORDER. In: APLAS 2017. pp. 364–383 (2017). https://doi.org/10.1007/978-3-319-71237-6_18
- 1333 [14] Libkin, L., Wong, L.: Query languages for bags and aggregate functions. J. Comput. Syst. Sci. 55(2) (1997)
- 1334 [15] Lindley, S., Cheney, J.: Row-based effect types for database integration. In: TLDI (2012). <https://doi.org/10.1145/2103786.2103798>
- 1335 [16] Meijer, E., Beckman, B., Bierman, G.M.: LINQ: reconciling object, relations and XML in the .NET framework. In: SIGMOD (2006). <https://doi.org/10.1145/1142473.1142552>
- 1336 [17] .NET EF Core: Issue #30173: Improve GroupBy support, <https://github.com/dotnet/efcore/issues/30173>
- 1337 [18] Okura, R., Kameyama, Y.: Language-integrated query with nested data structures and grouping. In: FLOPS. pp. 139–158 (2020). https://doi.org/10.1007/978-3-030-59025-3_9
- 1338 [19] Okura, R., Kameyama, Y.: Reorganizing queries with grouping. In: Proceedings of the 19th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences. p. 50–62. GPCE 2020, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3425898.3426960>, <https://doi.org/10.1145/3425898.3426960>
- 1339 [20] Quill: Compile-time language integrated queries for Scala. Open source project, <https://github.com/getquill/quill>
- 1340 [21] Ricciotti, W., Cheney, J.: Mixing set and bag semantics. In: DBPL. pp. 70–73 (2019). <https://doi.org/10.1145/3315507.3330202>
- 1341 [22] Ricciotti, W., Cheney, J.: Query lifting: Language-integrated query for heterogeneous nested collections. In: ESOP. pp. 579–606 (2021)
- 1342 [23] Ricciotti, W., Cheney, J.: A formalization of sql with nulls. J. Autom. Reasoning 66, 989–1030 (2022). <https://doi.org/10.1007/s10817-022-09632-4>, <https://doi.org/10.1007/s10817-022-09632-4>
- 1343 [24] Ricciotti, W., Cheney, J.: Strongly-normalizing higher-order relational queries. LMCS 18(3) (2022)
- 1344 [25] Shaikhha, A., Huot, M., Smith, J., Olteanu, D.: Functional collection programming with semi-ring dictionaries. Proc. ACM Program. Lang. 6(OOPSLA1) (apr 2022). <https://doi.org/10.1145/3527333>, <https://doi.org/10.1145/3527333>
- 1345 [26] Smith, J., Benedikt, M., Nikolic, M., Shaikhha, A.: Scalable querying of nested data. Proc. VLDB Endow. 14(3), 445–457 (nov 2020). <https://doi.org/10.14778/3430915.3430933>, <https://doi.org/10.14778/3430915.3430933>
- 1346 [27] Suzuki, K., Kiselyov, O., Kameyama, Y.: Finally, safely-extensible and efficient language-integrated query. In: PEPM. pp. 37–48 (2016). <https://doi.org/10.1145/2847538.2847542>
- 1347 [28] Syme, D.: Leveraging .NET meta-programming components from F#: integrated queries and interoperable heterogeneous execution. In: ML Workshop (2006)
- 1348 [29] Tannen, V., Buneman, P., Wong, L.: Naturally embedded query languages. In: ICDT. LNCS, vol. 646. Springer (1992)
- 1349 [30] Transaction processing performance council: TPC-H benchmark (2022), <http://www.tpc.org>
- 1350 [31] Trinder, P., Wadler, P.: Improving list comprehension database queries. In: TENCEN ’89. (1989). <https://doi.org/10.1109/TENCEN.1989.176921>
- 1351 [32] Wong, L.: Normal forms and conservative extension properties for query languages over collection types. J. Comput. Syst. Sci. 52(3) (1996). <https://doi.org/10.1006/jcss.1996.0037>
- 1352 [33] Wong, L.: Kleisli, a functional query system. J. Funct. Programming 10(1) (2000). <https://doi.org/10.1017/S0956796899003585>

A Proofs from Section 3

Lemma A.1. For all M, N , and x , we have $\llbracket M \llbracket N/x \rrbracket \rrbracket = \llbracket M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$.

Proof. By structural induction on M : we comment on a few relevant cases:

- $\llbracket (\lambda y.M) \llbracket N/x \rrbracket \rrbracket = \llbracket \lambda y.M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$: we assume without loss of generality that y is fresh with respect to x and N ; by definition, the lhs is equal to $\lambda y. \llbracket M \llbracket N/x \rrbracket \rrbracket$, which can be rewritten to $\lambda y. \llbracket M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$ by the IH; this is equal to the rhs by definition.

- $\llbracket \left[L \mid k \xleftarrow{\mathcal{K}} M \right]^{\text{set}} \llbracket N/x \rrbracket \rrbracket = \llbracket \left[L \mid k \xleftarrow{\mathcal{K}} M \right]^{\text{set}} \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$: we assume without loss of generality that k is fresh with respect to x and N ; by definition, the lhs is equal to

$$\llbracket \left[L \llbracket N/x \rrbracket \right] \mid k \leftarrow \text{dom}(\llbracket M \llbracket N/x \rrbracket \rrbracket) \rrbracket^{\text{set}}$$

By the IH, we rewrite $\llbracket L \llbracket N/x \rrbracket \rrbracket$ to $\llbracket L \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$ and $\llbracket M \llbracket N/x \rrbracket \rrbracket$ to $\llbracket M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$; then we can see that the resulting term is equal to the rhs by unfolding the definitions.

- $\llbracket \gamma_{y,\rho}(M) \llbracket N/x \rrbracket \rrbracket = \llbracket \gamma_{y,\rho}(M) \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$: we assume without loss of generality that y is fresh with respect to x and N ; by definition, the lhs is equal to

$$\llbracket \left[\langle \ell = y.1.\dot{\ell} \oplus [\rho \llbracket N/x \rrbracket] [y.2/y], y.2 \rangle \mid y \leftarrow \llbracket M \llbracket N/x \rrbracket \rrbracket \right] \rrbracket$$

By the IH, we rewrite $\llbracket M \llbracket N/x \rrbracket \rrbracket$ to $\llbracket M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$, and $[\rho \llbracket N/x \rrbracket]$ to $[\rho] \llbracket \llbracket N \rrbracket/x \rrbracket$; by the substitution lemma, we prove $[\rho] \llbracket \llbracket N \rrbracket/x \rrbracket [y.2/y] = [\rho] [y.2/y] \llbracket \llbracket N \rrbracket/x \rrbracket$; by this, we can prove that the term is equal to the rhs.

- $\llbracket (L \otimes M) \llbracket N/x \rrbracket \rrbracket = \llbracket L \otimes M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$: we choose z fresh with respect to x, L, M, N ; then by definition, the lhs is equal to $\llbracket [z.2] \text{ where } z.1 = \llbracket M \llbracket N/x \rrbracket \rrbracket \mid z \leftarrow \llbracket L \llbracket N/x \rrbracket \rrbracket \rrbracket$; by the IH, this is equal to

$$\llbracket [z.2] \text{ where } z.1 = \llbracket M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket \mid z \leftarrow \llbracket L \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket \rrbracket$$

which by definition is equal to the rhs. \square

Theorem 3.2. Whenever $M \rightsquigarrow N$ in $\mathcal{NR}\mathcal{C}_{\lambda\gamma}$, we have that $\llbracket M \rrbracket \rightsquigarrow \llbracket N \rrbracket$ in $\mathcal{NR}\mathcal{C}_{\lambda}(\text{Set}, \text{Bag})$.

Proof. By induction and case analysis on the reduction rule. We consider here some key cases:

- $\llbracket (\lambda x.M) N \rrbracket \rightsquigarrow \llbracket M \llbracket N/x \rrbracket \rrbracket$: we prove that $\llbracket M \llbracket N/x \rrbracket \rrbracket = \llbracket M \rrbracket \llbracket \llbracket N \rrbracket/x \rrbracket$. By Lemma A.1.
- $\llbracket \gamma_{x,\rho}(\square) \rrbracket = \llbracket \left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x.2/x], x.2 \rangle \mid x \leftarrow \square \right] \rrbracket \rightsquigarrow \llbracket \square \rrbracket$: trivial.

- $\llbracket \gamma_{x,\rho}([\rho' \triangleright M]) \rrbracket = \llbracket \left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x.2/x], x.2 \rangle \mid x \leftarrow \llbracket \llbracket \llbracket \rho' \rrbracket \rrbracket, \llbracket M \rrbracket \rrbracket \right] \rightsquigarrow \left[\langle \ell = \llbracket \rho' \rrbracket \oplus [\rho] \llbracket \llbracket M \rrbracket/x \rrbracket \rrbracket, \llbracket M \rrbracket \right] \rrbracket = \llbracket [\rho] \llbracket M/x \rrbracket \oplus \rho' \triangleright M \rrbracket$: we reduce the comprehension in the lhs, and then perform as many reductions on the resulting projection redexes as needed to obtain the rhs. This also uses Lemma A.1 just like the lambda-application case.
- $\llbracket \gamma_{x,\rho}(\gamma_{x,\rho'}(M)) \rrbracket = \llbracket \left[\left[\langle \ell = y.1.\dot{\ell} \oplus [\rho] [y.2/y], y.2 \rangle \mid y \leftarrow \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho'] [x.2/x], x.2 \rangle \mid x \leftarrow \llbracket M \rrbracket \right] \right] \rightsquigarrow \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x.2/x] \oplus [\rho'] [x.2/x], x.2 \rangle \mid x \leftarrow \llbracket M \rrbracket \right] \right] \right] \rrbracket = \llbracket \gamma_{x,\rho \oplus \rho'}(M) \rrbracket$: we perform unnesting on the lhs, then reduce the comprehension binding y to a singleton; finally, we reduce the projections as many times as needed to obtain the rhs.
- $\llbracket \gamma_{x,\rho}(M \# N) \rrbracket = \llbracket \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x/x.2], x.2 \rangle \mid x \leftarrow \llbracket M \rrbracket \# \llbracket N \rrbracket \right] \rightsquigarrow \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x/x.2], x.2 \rangle \mid x \leftarrow \llbracket M \rrbracket \right] \# \left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x/x.2], x.2 \rangle \right]^{\text{set}} \mid x \leftarrow \llbracket N \rrbracket \right] \rrbracket = \llbracket \gamma_{x,\rho}(M) \# \gamma_{x,\rho}(N) \rrbracket$: trivial.
- $\llbracket \gamma_{x,\rho}([L \mid y \leftarrow M]) \rrbracket = \llbracket \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x/x.2], x.2 \rangle \mid x \leftarrow \llbracket [L] [y.2/y] \mid y \leftarrow \llbracket M \rrbracket \rrbracket \right] \rightsquigarrow \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x.2/x], x.2 \rangle \right] [y.2/y] \mid y \leftarrow \llbracket M \rrbracket, x \leftarrow \llbracket [L] [y.2/y] \rrbracket \right] \rrbracket = \llbracket \left[\gamma_{x,\rho}(L) \mid y \leftarrow M \right] \rrbracket$: we assume, without loss of generality, that $y \neq x$ and $y \notin \text{FV}(M, \rho)$; we perform unnesting on the lhs and, thanks to the freshness conditions, that the resulting term is equal to the rhs.
- $\llbracket \gamma_{x,\rho}([L \mid k \xleftarrow{\mathcal{K}} M]^{\text{set}}) \rrbracket = \llbracket \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x/x.2], x.2 \rangle \right]^{\text{set}} \mid x \leftarrow \llbracket [L] [k.1/k] \mid k \leftarrow \llbracket M \rrbracket \rrbracket \right]^{\text{set}} \rightsquigarrow \left[\left[\langle \ell = x.1.\dot{\ell} \oplus [\rho] [x.2/x], x.2 \rangle \right]^{\text{set}} [k.1/k] \mid y \leftarrow \llbracket M \rrbracket, x \leftarrow \llbracket [L] [y.1/k] \rrbracket \right]^{\text{set}} \rrbracket$

1541 $= \left[\left[\gamma_{x,\rho}(L) \mid k \stackrel{\mathcal{K}}{\leftarrow} M \right] \right]$: we assume, without loss of gen-
 1542 erality, that $k \neq x$ and $k \notin \text{FV}(M, \rho)$; we perform
 1543 unnesting on the lhs and, thanks to the freshness con-
 1544 ditions, that the resulting term is equal to the rhs.

- 1545 • $\left[\left[\right] \otimes N \right] = \left[\left[\langle \rangle, x.2 \rangle \right] \text{ where } x.1 = \lceil N \mid x \leftarrow \left[\right] \right]$
 1546 $\stackrel{+}{\sim} \left[\left[\right] \right]$: trivial.
- 1547 • $\left[(L \# M) \otimes N \right]$
 1548 $= \left[\left[\langle \rangle, x.2 \rangle \right] \text{ where } x.1 = \lceil N \mid x \leftarrow \lceil L \rceil \# \lceil M \rceil \right]$
 1549 $\stackrel{+}{\sim} \left[\left[\langle \rangle, x.2 \rangle \right] \text{ where } x.1 = \lceil N \mid x \leftarrow \lceil L \rceil \# \right.$
 1550 $\left. \left[\langle \rangle, x.2 \rangle \right] \text{ where } x.1 = \lceil N \mid x \leftarrow \lceil M \rceil \right] \right]$
 1551 $= \left[(L \otimes N) \# (M \otimes N) \right]$: trivial.
- 1552 • $\left[\lceil L \mid x \leftarrow M \right] \otimes N \right]$
 1553 $= \left[\left[\langle \rangle, y.2 \rangle \right] \text{ where } y.1 = \lceil N \right]$
 1554 $\left[\mid y \leftarrow \lceil \lceil L \rceil [x.2/x] \mid x \leftarrow \lceil M \rceil \right] \right]$
 1555 $\stackrel{+}{\sim} \left[\left[\langle \rangle, y.2 \rangle \right] \text{ where } y.1 = \lceil N \right]$
 1556 $\left[\mid x \leftarrow \lceil M \rceil, y \leftarrow \lceil L \rceil [x.2/x] \right]$
 1557 $= \left[\lceil L \otimes N \mid x \leftarrow M \right] \right]$: we assume, without loss of
 1558 generality, that $y \neq x$, and $y \notin \text{FV}(N)$; we perform
 1559 unnesting on the lhs and, by the freshness conditions,
 1560 show that the resulting term is equal to the rhs.
- 1561 • $\left[\left[\lceil L \mid k \stackrel{\mathcal{K}}{\leftarrow} M \right]^{\text{set}} \otimes N \right]$
 1562 $= \left[\left[\langle \rangle, y.2 \rangle \right] \text{ where } y.1 = \lceil N \right]$
 1563 $\left[\mid y \leftarrow \left[\left[\lceil L \rceil [k.1/k] \mid k \leftarrow \lceil M \rceil \right]^{\text{set}} \right]^{\text{set}} \right]$
 1564 $\stackrel{+}{\sim} \left[\left[\langle \rangle, y.2 \rangle \right] \text{ where } y.1 = \lceil N \right]$
 1565 $\left[\mid k \leftarrow \lceil M \rceil, y \leftarrow \lceil L \rceil [k.1/k] \right]^{\text{set}}$
 1566 $= \left[\left[\lceil L \otimes N \mid k \stackrel{\mathcal{K}}{\leftarrow} M \right]^{\text{set}} \right]$: we assume, without loss of
 1567 generality, that $y \neq k$, and $y \notin \text{FV}(N)$; we perform
 1568 unnesting on the lhs and, by the freshness conditions,
 1569 show that the resulting term is equal to the rhs.
- 1570 • For all reductions happening within a context, the
 1571 thesis is obtained by an application of the induction
 1572 hypothesis. \square

1581 B Proofs from Section 4

1582 **Theorem 4.1.** Every well-typed term M of $\mathcal{NRC}_{\lambda\gamma}$ in
 1583 normal form generated by the grammar in Figure 6.

1584 *Proof.* By induction on the structure of M . Notice that the
 1585 induction hypothesis states that the subterms of the M for
 1586 which we are proving this theorem are generated by the
 1587 grammar. By cases:

- 1588 • Terms in the form $x, t, \langle \rho \rangle, \lambda x.M, \left[\right]^{\mathcal{S}}, \left[\rho \triangleright M' \right]^{\mathcal{S}}$: the
 1589 proof is trivial, or a direct consequence of the induction
 1590 hypothesis.
- 1591 • $c(\vec{M})$: by IH, \vec{M}' is generated by the grammar; we
 1592 reason by cases on the possible productions of \vec{M}' and

1596 see that, in order for the term to be well-typed, we
 1597 must have $\vec{M}' = \vec{B}: c(\vec{B})$ is generated by the grammar
 1598 for M , which proves the thesis.

- 1599 • $\alpha(M')$: by IH, M' is generated by the grammar; we
 1600 reason by cases on the possible productions of M' and
 1601 see that, in order for the term to be well-typed, we
 1602 must have $M' = R: \alpha(R)$ is generated by the grammar
 1603 for M , which proves the thesis.
- 1604 • $M'.\ell$: by IH, M' is generated by the grammar; we reason
 1605 by cases on the possible productions of M' and see
 1606 that, in order for the term to be well-typed, we must
 1607 have $M' = U: U$ can be either X or $\langle \rho \rangle$, however $\langle \rho \rangle.\ell$
 1608 is not in normal form; that only leaves $X.\ell$, which is
 1609 generated by the grammar for M , as required.
- 1610 • $M' M''$: by IH, M' and M'' are generated by the gram-
 1611 mar; we reason by cases on the possible productions of
 1612 M' and see that, in order for the term to be well-typed,
 1613 we must have $M' = W: W$ can be either X or $\lambda x.M'''$,
 1614 however $(\lambda x.M''') M''$ is not in normal form; that only
 1615 leaves $X M''$, which is generated by the grammar for
 1616 M , as required.
- 1617 • $M' \#^{\text{set}} M''$: by IH, M' and M'' are generated by the
 1618 grammar; we reason by cases on the possible produc-
 1619 tions of M' and M'' and see that, in order for the term
 1620 to be well-typed, we must have $M' = Q', M'' = Q''$;
 1621 $Q' \#^{\text{set}} Q''$ is generated by the grammar for M , as
 1622 required.
- 1623 • $M' \text{ where}^{\text{set}} M''$: by IH, M' and M'' are generated
 1624 by the grammar; we reason by cases on the possible
 1625 productions of M' and M'' and see that, in order for
 1626 the term to be well-typed, we must have $M' = Q$ and
 1627 $M'' = B$; furthermore, in order for the term to be in
 1628 normal form, we prove by deep case analysis on the
 1629 possible productions of Q that we must have $Q = I$;
 1630 then, $I \text{ where}^{\text{set}} B$ is generated by the grammar for M ,
 1631 as required.
- 1632 • $\delta M'$: by IH, M' is generated by the grammar; we reason
 1633 by cases on the possible productions of M' and see that,
 1634 in order for the term to be well-typed, we must have
 1635 $M' = R$; by deep cases analysis on the productions
 1636 starting at R , we see that, in order for δR to be in
 1637 normal form, we must have $R = X$ or $R = t$; both
 1638 δX and δt are generated by the grammar for M , as
 1639 required.
- 1640 • $\iota M'$: by IH, M' is generated by the grammar; we reason
 1641 by cases on the possible productions of M' and see that,
 1642 in order for the term to be well-typed, we must have
 1643 $M' = Q: \iota Q$ is generated by the grammar for M , as
 1644 required.
- 1645 • $\left[M' \mid x \leftarrow M'' \right]^{\text{set}}$: by IH, M' and M'' are generated
 1646 by the grammar; we reason by cases on the possible
 1647 productions of M' and M'' and see that, in order for
 1648 the term to be well-typed, we must have $M' = Q'$,

$M'' = Q''$; by more case analysis, we see that $Q' = C$ (otherwise, the term would not be in normal form); by a similar reasoning, Q'' must be in the form N : we thus easily show that $[C \mid x \leftarrow N]^{\text{set}}$ is generated by the grammar for M , as required.

- $\left[M' \mid k \xleftarrow{\mathcal{K}} M'' \right]^{\text{set}}$: similar to the value comprehension case above.
- $\gamma_{x,\rho}^{\text{set}}(M')$: by IH, M' is generated by the grammar; we reason by cases on the possible productions of M' and see that, in order for the term to be well-typed and in normal form, we must have $M' = N$; $\gamma_{x,\rho}^{\text{set}}(N)$ is generated by the grammar for M , as required.
- $M' \otimes^{\text{set}} M''$: by IH, M' and M'' are generated by the grammar; we reason by cases on the possible productions of M' and M'' and see that, in order for the term to be well-typed, we must have $M' = Q$, $M'' = U$; furthermore, by cases on the productions for Q , in order for the term to be in normal form we will need $Q = N$; then we see that the term $N \otimes^{\text{set}} U$ is generated by the grammar, as required.
- $\text{aggBy}_{\ell=\alpha}^{\text{set}}(M')$: by IH, M' is generated by the grammar; we reason by cases on the possible productions of M' and see that, in order for the term to be well-typed, we must have $M' = Q$; then we see that $\text{aggBy}_{\ell=\alpha}^{\text{set}}(Q)$ is generated by the grammar, as required.
- $M' \uplus^{\text{bag}} M''$: by IH, M' and M'' are generated by the grammar; we reason by cases on the possible productions of M' and M'' and see that, in order for the term to be well-typed, we must have $M' = R'$, $M'' = R''$; $R' \uplus^{\text{bag}} R''$ is generated by the grammar for M , as required.
- $[M' \mid x \leftarrow M'']^{\text{bag}}$: by IH, M' and M'' are generated by the grammar; we reason by cases on the possible productions of M' and M'' and see that, in order for the term to be well-typed, we must have $M' = R'$, $M'' = R''$; by more case analysis, we see that $R' = D$ (otherwise, the term would not be in normal form); by a similar reasoning, R'' must be in the form O : we thus easily show that $[D \mid x \leftarrow O]^{\text{bag}}$ is generated by the grammar for M , as required.
- $\left[M' \mid k \xleftarrow{\mathcal{K}} M'' \right]^{\text{bag}}$: by IH, M' and M'' are generated by the grammar; we reason by cases on the possible productions of M' and M'' and see that, in order for the term to be well-typed, we must have $M' = R'$, $M'' = Q''$; by more case analysis, we see that $R' = D$ (otherwise, the term would not be in normal form); we thus easily show that $\left[D \mid k \xleftarrow{\mathcal{K}} Q'' \right]^{\text{bag}}$ is generated by the grammar for M , as required.
- $M' \text{ where }^{\text{bag}} M''$: by IH, M' and M'' are generated by the grammar; we reason by cases on the possible productions of M' and M'' and see that, in order for

the term to be well-typed, we must have $M' = R$ and $M'' = B$; furthermore, in order for the term to be in normal form, we prove by deep case analysis on the possible productions of R that we must have $R = L$; then, $L \text{ where }^{\text{bag}} B$ is generated by the grammar for M , as required.

- $\gamma_{x,\rho}^{\text{bag}}(M')$: by IH, M' is generated by the grammar; we reason by cases on the possible productions of M' and see that, in order for the term to be well-typed and in normal form, we must have $M' = O$; $\gamma_{x,\rho}^{\text{bag}}(O)$ is generated by the grammar for M , as required.
- $M' \otimes^{\text{bag}} M''$: by IH, M' and M'' are generated by the grammar; we reason by cases on the possible productions of M' and M'' and see that, in order for the term to be well-typed, we must have $M' = R$, $M'' = U$; furthermore, by cases on the productions for R , in order for the term to be in normal form we will need $R = O$; then we see that the term $O \otimes^{\text{bag}} U$ is generated by the grammar, as required.
- $\text{aggBy}_{\ell=\alpha}^{\text{bag}}(M')$: by IH, M' is generated by the grammar; we reason by cases on the possible productions of M' and see that, in order for the term to be well-typed, we must have $M' = R$; then we see that $\text{aggBy}_{\ell=\alpha}^{\text{bag}}(R)$ is generated by the grammar, as required. \square

C Details of the translation to SQL (Section 5)

The grammar of normal forms is amenable to be used to define algorithmic procedures, including a translation of $\mathcal{NRC}_{\lambda_Y}$ normal forms to SQL; clearly, only flat normal forms need to be translated, as nested queries do not have a direct representation in SQL due to typing limitations. We give one such translation in Figure 10: we define a main translation operation $(\cdot)^{\text{sql}}$ along with auxiliary translations $(\cdot)_{K'}^{\text{sql}}$, $(\cdot)_{V'}^{\text{sql}}$, $(\cdot)_{GK'}^{\text{sql}}$, and operations *keys*, *vals*, *attrs*.

The main translation $(\cdot)^{\text{sql}}$ assigns trivially empty queries to $[\]$ and, by recursion, union/disjoint union queries to $\mathcal{NRC}_{\lambda_Y}$ queries employing \uplus . In the more involved comprehensions cases, the translation operates by returning a complete `SELECT ... FROM ... WHERE ...` statement (with or without `DISTINCT`, depending on whether we are performing a set or a bag comprehension); the `where` clause is translated to the `WHERE` statement, and the comprehension generators are translated to the `FROM` clause; that leaves the singleton key-value pair – namely, $[\rho \triangleright \langle \rho' \rangle]$; both ρ and ρ' will go to the `SELECT` statement, but ρ shall use the auxiliary $(\rho)_{K'}^{\text{sql}}$ translation, marking that row as a key, while ρ' shall use the $(\rho')_{V'}^{\text{sql}}$ translation, marking it as a value row. Concretely, the two different translations tag the attribute names so that attributes used as grouping keys (in the form

1761 1@ ℓ) will always be distinguished from attributes used as
 1762 values (in the form 2@ ℓ).

1763 The rest of the $\mathcal{NR}\mathcal{C}_{\lambda\gamma}$ normal forms can only appear as
 1764 comprehension generators: these include table references t
 1765 (translated using `SELECT` with the value-tagging translation
 1766 $(\cdot)_V^{\text{sql}}$), deduplicated tables (which use `SELECT DISTINCT` *,
 1767 promoted set queries (translated by recursion: from an SQL
 1768 point of view, promotion does not change the semantics of a
 1769 query).

1770 The translation of grouping ($\gamma_{x,\rho}^{\text{set}}(N)$ or $\gamma_{x,\rho}^{\text{bag}}(O)$) uses a
 1771 `SELECT (DISTINCT)` statement, in which the values from the
 1772 table resulting from the evaluation of N (or O) are linked with
 1773 keys $(\rho)_K^{\text{sql}}$; notice that, although we know by reasoning on
 1774 typing that N (or O) must be a pure collection (rather than a
 1775 finite map), the recursive translation will return a table with
 1776 value-tagged attributes: the $\text{vals}(\cdot)$ operation returns the
 1777 list of value-tagged attributes. Dually, the lookup operation
 1778 $(N \oplus^{\text{set}} \langle \rho \rangle$ or $O \oplus^{\text{bag}} \langle \rho \rangle)$ also produces a `SELECT (DISTINCT)`
 1779 statement using the $\text{vals}(\cdot)$ operation, but the input relation
 1780 N or O must be a proper finite map, thus we need to filter
 1781 those keys matching the lookup row ρ . Groupwise aggrega-
 1782 tion is translated by means of a `SELECT` statement using
 1783 the appropriate aggregates, along with a `GROUP BY` clause
 1784 (notice that since we know that each key will be mapped
 1785 to a singleton, we can avoid using `DISTINCT` even in the set
 1786 case).

1787 The main translation is used in this form to process value
 1788 comprehensions; for key comprehensions, we provide a spe-
 1789 cific translation $(\cdot)_{GK}^{\text{sql}}$ whose only cases are grouping, with
 1790 set or bag semantics, and (implicitly promoted) set sub-
 1791 queries Q ; the main purpose of this particular translation is
 1792 to collect only the key attributes from the input map (using
 1793 $\text{keys}(\cdot)$ in the inner set query case) and then retag them as
 1794 value attributes; since we never want duplicate keys, not
 1795 even in bags, all of the cases are translated using `DISTINCT`.
 1796 In both kinds of comprehension, our translation can gener-
 1797 ate lateral joins (as it is normal for queries mixing sets an
 1798 bags [21]) which require the SQL:1999 keyword `LATERAL`.
 1799 Lateral joins can always be removed using the technique
 1800 shown in [22] (in exchange for additional query complexity).

1801 Aggregation can only be performed on pure collections
 1802 of basic values: we convert these to collections of trivial
 1803 unary tuples, translate them recursively, and finally apply
 1804 the corresponding SQL aggregation operation to obtain the
 1805 result we want.

1806 As a concluding remark, we note that this translation
 1807 procedure handles the grouping operator *without* employing
 1808 SQL `GROUP BY` statements, save when it is the argument of
 1809 a groupwise aggregation: this is because `GROUP BY` clauses
 1810 can only be used when attributes other than the grouping
 1811 keys have been aggregated.

1813 **Example C.1.** The $\mathcal{NR}\mathcal{C}_{\lambda\gamma}$ query M_{group} from Example 2.1
 1814 is translated to SQL as follows;

1816 $(M_{\text{group}})^{\text{sql}} :=$
 1817 `SELECT x.1@dept AS 1@dept,`
 1818 `AVERAGE(x.2@salary) as 2@salary`
 1819 `FROM (SELECT d.name AS 1@dept,`
 1820 `e.salary AS 2@salary`
 1821 `FROM department d, employee e`
 1822 `WHERE d.id = e.dept)`
 1823 `GROUP BY 1@dept`
 1824

1825 This is slightly more complicated than the original query
 1826 `q_group` from the introduction due to the separation of
 1827 grouping and groupwise aggregation in the $\mathcal{NR}\mathcal{C}_{\lambda\gamma}$ term;
 1828 however, this kind of nesting is easily optimized in most
 1829 DBMS; alternatively, we can get rid of such artifacts in a
 1830 postprocessing step.
 1831

1832 D Full definition of query shredding

1833 (Section 5.2)

1834 Figure 11 shows the full definition of the shredding judgment
 1835 for $\mathcal{NR}\mathcal{C}_{\lambda\gamma}$. We comment the rules that were omitted from
 1836 the main body of this article.
 1837

1838 The rules for the shredding judgment operate as follows:
 1839 the first rule expresses the fact that a normalized term of
 1840 base type B does not contain subexpressions with nested
 1841 collection type, therefore it can be shredded to itself, leaving
 1842 the shredding environment Φ unchanged. This rule is also
 1843 applied for plain aggregation, despite the fact that they have
 1844 a collection subterm: from the grammar of flat relational
 1845 normal forms, we know that such aggregatins must be in
 1846 the form $\alpha(R^*)$, where R^* is a flat collection. Since R^* is flat,
 1847 it does not need to be shredded recursively.
 1848

1849 In the case of rows, we perform shredding pointwise on
 1850 each field, connecting the input and output shredding envi-
 1851 ronments in a pipeline, and finally combining together the
 1852 shredded subterms in the obvious way.
 1853

1854 The shredding of set and bag unions uses the query lifting
 1855 technique and is performed by recursion on the subterms,
 1856 using the same plumbing technique we employed for tuples;
 1857 additionally, we optimize the output shredding environment
 1858 by removing the graph queries $\vec{\psi}$ resulting from recursion,
 1859 since they are absorbed into the new graph φ .
 1860

1861 Finally, groupwise aggregation employs query lifting as
 1862 well, since it returns a finite map. Since we know that the
 1863 input is a flat relation, we do not need to shred the input
 1864 query recursively.
 1865

1871	$(\emptyset)^{\text{sql}}$	=	SELECT 42 WHERE 0 = 1	$(x.\ell)^{\text{sql}}$	=	$x.\ell$	1926
1872	$(c(\vec{X}))^{\text{sql}}$	=	$(c)^{\text{sql}} (\vec{X})^{\text{sql}}$	$(\alpha(M))^{\text{sql}}$	=	SELECT $(\alpha)^{\text{sql}}(x)$ FROM $\left(\left[[(\bullet = z)]^{\text{bag}} \mid z \leftarrow M \right]^{\text{bag}} \right)^{\text{sql}}$	1927
1873	$(\vec{\ell} = \vec{X})^{\text{sql}}$	=	$(X_1)^{\text{sql}} \text{ AS } \ell_1, \dots, (X_n)^{\text{sql}} \text{ AS } \ell_n$				1928
1874	$(\vec{\ell} = \vec{X})^{\text{sql}}$	=	$(X_1)^{\text{sql}} \text{ AS } 1@ \ell_1, \dots, (X_n)^{\text{sql}} \text{ AS } 1@ \ell_n$	$(\vec{\ell} = \vec{X})_V^{\text{sql}}$	=	$(X_1)^{\text{sql}} \text{ AS } 2@ \ell_1, \dots, (X_n)^{\text{sql}} \text{ AS } 2@ \ell_n$	1929
1875	$(\vec{\ell} = \vec{X})_K^{\text{sql}}$	=	$(C_1)^{\text{sql}} \text{ UNION } \dots \text{ UNION } (C_n)^{\text{sql}}$	$(\vec{\ell} = \vec{X})_D^{\text{sql}}$	=	$(D_1)^{\text{sql}} \text{ UNION ALL } \dots \text{ UNION ALL } (D_n)^{\text{sql}}$	1930
1876	$(\bigcup \vec{C})^{\text{sql}}$	=	$(C_1)^{\text{sql}} \text{ UNION } \dots \text{ UNION } (C_n)^{\text{sql}}$				1931
1877	$(t)^{\text{sql}}$	=	SELECT $(\vec{\ell} = x.\ell)_V^{\text{sql}}$ FROM t (if $t : [\langle \vec{\ell} : T \rangle]^{\text{bag}}$)	$(\iota(Q))^{\text{sql}}$	=	$(Q)^{\text{sql}}$	1932
1878	$(\delta t)^{\text{sql}}$	=	SELECT DISTINCT * FROM $(t)^{\text{sql}}$				1933
1879	$(Q)^{\text{sql}}$	=	SELECT DISTINCT $keys(Q, z)$ FROM $(Q)^{\text{sql}}$ z				1934
1880	$(\gamma_{x,\rho}^{\text{set}}(N))^{\text{sql}}$	=	SELECT DISTINCT $(\rho)_K^{\text{sql}}, vals(N, x)$ FROM $(N)^{\text{sql}}$ x				1935
1881	$(\gamma_{x,\rho}^{\text{set}}(N))_{GK}^{\text{sql}}$	=	SELECT DISTINCT $(\rho)_V^{\text{sql}}$ FROM $(N)^{\text{sql}}$ x				1936
1882	$(\gamma_{x,\rho}^{\text{bag}}(O))^{\text{sql}}$	=	SELECT $(\rho)_K^{\text{sql}}, vals(O, x)$ FROM $(O)^{\text{sql}}$ x (if $O : [\langle \vec{\ell} : T \rangle]^{\text{bag}}$)				1937
1883	$(\gamma_{x,\rho}^{\text{bag}}(O))_{GK}^{\text{sql}}$	=	SELECT DISTINCT $(\rho)_V^{\text{sql}}$ FROM $(O)^{\text{sql}}$ x				1938
1884	$(N \otimes^{\text{set}} \langle \vec{k}_i = M_i \rangle)^{\text{sql}}$	=	SELECT DISTINCT $vals(N, x)$ FROM $(N)^{\text{sql}}$ x WHERE $(x.1@ \vec{k}_i = (M_i)^{\text{sql}})$				1939
1885	$(O \otimes^{\text{bag}} \langle \vec{k}_i = M_i \rangle)^{\text{sql}}$	=	SELECT $vals(O, x)$ FROM $(O)^{\text{sql}}$ x WHERE $(x.1@ \vec{k}_i = (M_i)^{\text{sql}})$				1940
1886	$(aggBy^{\text{set}}_{z.\ell=\alpha(z.\ell')} (Q))^{\text{sql}}$	=	SELECT $keys(Q, x), (\alpha)^{\text{sql}}(x.2@ \ell')$ AS $2@ \ell$ FROM $(Q)^{\text{sql}}$ x GROUP BY $keys(Q, x)$				1941
1887	$(aggBy^{\text{bag}}_{z.\ell=\alpha(z.\ell')} (R))^{\text{sql}}$	=	SELECT $keys(R, x), (\alpha)^{\text{sql}}(x.2@ \ell')$ AS $2@ \ell$ FROM $(R)^{\text{sql}}$ x GROUP BY $keys(R, x)$				1942
1888	$(\vec{F}, x \leftarrow N)^{\text{sql}}$	=	$\begin{cases} (\vec{F})^{\text{sql}}, ((N)^{\text{sql}}) x & (N \text{ closed wrt. } vars(F)) \\ (\vec{F})^{\text{sql}}, \text{LATERAL } ((N)^{\text{sql}}) x & (\text{otherwise}) \end{cases}$				1943
1889	$(\vec{F}, x \xleftarrow{\mathcal{K}} N)^{\text{sql}}$	=	$\begin{cases} (F)^{\text{sql}}, ((N)_{GK}^{\text{sql}}) x & (N \text{ closed wrt. } vars(F)) \\ (F)^{\text{sql}}, \text{LATERAL } ((F)_{GK}^{\text{sql}}) x & (\text{otherwise}) \end{cases}$				1944
1890	$(\vec{G}, x \leftarrow O)^{\text{sql}}$	=	$\begin{cases} (\vec{G})^{\text{sql}}, ((O)^{\text{sql}}) x & (O \text{ closed wrt. } vars(G)) \\ (\vec{G})^{\text{sql}}, \text{LATERAL } ((O)^{\text{sql}}) x & (\text{otherwise}) \end{cases}$				1945
1891	$(\vec{G}, x \xleftarrow{\mathcal{K}} Q^*)^{\text{sql}}$	=	$\begin{cases} (\vec{G})^{\text{sql}}, ((Q^*)_{GK}^{\text{sql}}) x & (Q^* \text{ closed wrt. } vars(G)) \\ (\vec{G})^{\text{sql}}, \text{LATERAL } ((G)_{GK}^{\text{sql}}) x & (\text{otherwise}) \end{cases}$				1946
1892	$(\left[[\rho \triangleright \langle \rho' \rangle]^{\text{set}} \text{ where }^{\text{set}} X \mid \vec{F} \right]^{\text{set}})^{\text{sql}}$	=	SELECT DISTINCT $(\rho)_K^{\text{sql}}, (\rho')_V^{\text{sql}}$ FROM $(\vec{F})^{\text{sql}}$ WHERE $(X)^{\text{sql}}$				1947
1893	$(\left[[\rho \triangleright \langle \rho' \rangle]^{\text{bag}} \text{ where }^{\text{bag}} X \mid \vec{G} \right]^{\text{bag}})^{\text{sql}}$	=	SELECT $(\rho)_K^{\text{sql}}, (\rho')_V^{\text{sql}}$ FROM $(\vec{G})^{\text{sql}}$ WHERE $(X)^{\text{sql}}$				1948
1894			$keys(Q, z) = z.1@ \ell_1 \text{ AS } 2@ \ell_1, \dots, z.1@ \ell_n \text{ AS } 2@ \ell_n$ (if $Q : [T]_{\vec{k},S}^{\text{set}}$ or $[T]_{\vec{k},S}^{\text{bag}}$)				1949
1895			$vals(Q, z) = z.2@ \ell_1, \dots, z.2@ \ell_n$ (if $Q : [\langle \vec{\ell} : T \rangle]_P^{\text{set}}$ or $[\langle \vec{\ell} : T \rangle]_P^{\text{bag}}$)				1950
1896							1951
1897							1952
1898							1953
1899							1954
1900							1955
1901							1956
1902							1957
1903							1958
1904							1959
1905							1960
1906							1961
1907							1962
1908							1963
1909							1964
1910							1965
1911							1966
1912							1967
1913							1968
1914							1969
1915							1970
1916							1971
1917							1972
1918							1973
1919							1974
1920							1975
1921							1976
1922							1977
1923							1978
1924							1979
1925							1980

Figure 10. Translation to SQL

1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035

2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090

$$\begin{array}{c}
\frac{}{\Phi; \Theta \vdash B \Rightarrow B \mid \Phi} \\
\\
\frac{(\Phi_{i-1}; \Theta \vdash M_i \Rightarrow \check{M}_i \mid \Phi_i)_{i=1, \dots, n}}{\Phi_0; \Theta \vdash \vec{\ell} = \vec{M} \Rightarrow \vec{\ell} = \vec{M} \mid \Phi_n} \\
\\
\frac{\varphi \notin \text{dom}(\Phi_n) \quad (\Phi_{i-1}; \Theta \vdash C_i \Rightarrow \psi_i \otimes^{\text{set}} \langle \text{dom}(\Theta) \rangle \mid \Phi_i)_{i=1, \dots, n}}{\Phi_0; \Theta \vdash \bigcup \vec{C} \Rightarrow \varphi \otimes^{\text{set}} \langle \text{dom}(\Theta) \rangle \mid (\Phi_n \setminus \vec{\psi})[\varphi \mapsto \bigcup \Phi_n(\vec{\psi})]} \\
\\
\frac{\varphi \notin \text{dom}(\Phi_n) \quad (\Phi_{i-1}; \Theta \vdash D_i \Rightarrow \psi_i \otimes^{\text{bag}} \langle \text{dom}(\Theta) \rangle \mid \Phi_i)_{i=1, \dots, n}}{\Phi_0; \Theta \vdash \biguplus \vec{D} \Rightarrow \varphi \otimes^{\text{bag}} \langle \text{dom}(\Theta) \rangle \mid (\Phi_n \setminus \vec{\psi})[\varphi \mapsto \biguplus \Phi_n(\vec{\psi})]} \\
\\
\frac{\varphi \notin \text{dom}(\Psi) \quad \Phi; \Theta, \vec{F} \vdash \rho \Rightarrow \check{\rho} \mid \Psi}{\Phi; \Theta \vdash \left[[\rho^* \triangleright \langle \rho \rangle]^{\text{set}} \text{ where } B \mid \vec{F} \right]^{\text{set}} \Rightarrow \varphi \otimes^{\text{set}} \langle \text{dom}(\Theta) \rangle \mid \Psi[\varphi \mapsto \mathcal{G}(\Theta; \left[[\rho^* \triangleright \check{\rho}]^{\text{set}} \text{ where } B \mid \vec{F} \right]^{\text{set}})]} \\
\\
\frac{\varphi \notin \text{dom}(\Psi) \quad \Phi; \Theta, \vec{G} \vdash \rho \Rightarrow \check{\rho} \mid \Psi}{\Phi; \Theta \vdash \left[[\rho^* \triangleright \rho]^{\text{bag}} \text{ where } B \mid \vec{G} \right]^{\text{bag}} \Rightarrow \varphi \otimes^{\text{bag}} \langle \text{dom}(\Theta) \rangle \mid \Psi[\varphi \mapsto \mathcal{G}(\Theta; \left[[\rho^* \triangleright \check{\rho}]^{\text{bag}} \text{ where } B \mid \vec{G} \right]^{\text{bag}})]} \\
\\
\frac{\varphi \notin \text{dom}(\Phi) \quad \Phi; \Theta \vdash \text{aggBy}_{z, \ell = \alpha(z, \ell')}^{\text{set}} \longrightarrow (Q^*) \Rightarrow \varphi \otimes^{\text{set}} \langle \text{dom}(\Theta) \rangle \mid \Phi[\varphi \mapsto \mathcal{G}(\Theta; \text{aggBy}_{z, \ell = \alpha(z, \ell')}^{\text{set}} \longrightarrow (Q^*))]} \\
\\
\frac{\varphi \notin \text{dom}(\Phi) \quad \Phi; \Theta \vdash \text{aggBy}_{z, \ell = \alpha(z, \ell')}^{\text{bag}} \longrightarrow (R^*) \Rightarrow \varphi \otimes^{\text{bag}} \langle \text{dom}(\Theta) \rangle \mid \Phi[\varphi \mapsto \mathcal{G}(\Theta; \text{aggBy}_{z, \ell = \alpha(z, \ell')}^{\text{bag}} \longrightarrow (R^*))]} \\
\\
(x \leftarrow O)^\delta \triangleq \begin{cases} x \leftarrow Q^* & \text{if } O = \iota Q^* \\ x \leftarrow \delta O & \text{else} \end{cases} \quad \Phi \setminus \vec{\psi} \triangleq [(\varphi \mapsto N) \in \Phi \mid \varphi \notin \vec{\psi}] \\
(k \xrightarrow{\mathcal{X}} Q^*)^\delta \triangleq Q^*
\end{array}$$

Figure 11. Shredding rules.