

# Effective Quotation

## Relating approaches to language-integrated query

James Cheney Sam Lindley

The University of Edinburgh  
jcheney@inf.ed.ac.uk,  
Sam.Lindley@ed.ac.uk

Gabriel Radanne

ENS Cachan  
gabriel.radanne@zoho.com

Philip Wadler

The University of Edinburgh  
wadler@inf.ed.ac.uk

### Abstract

Language-integrated query techniques have been explored in a number of different language designs. We consider two different, type-safe approaches employed by Links and F#. Both approaches provide rich dynamic query generation capabilities, and thus amount to a form of heterogeneous staged computation, but to date there has been no formal investigation of their relative expressiveness. We present two core calculi Eff and Quot, respectively capturing the essential aspects of language-integrated querying using effects in Links and quotation in LINQ. We show via translations from Eff to Quot and back that the two approaches are equivalent in expressiveness. Based on the translation from Eff to Quot, we extend a simple Links compiler to handle queries.

**Categories and Subject Descriptors** D.3.1 [Formal Definitions and Theory]; D.3.2 [Language Classifications]: Applicative (functional) languages; H.2.3 [Languages]: Query languages

**Keywords** language-integrated query; effects; quotation

### 1. Introduction

Increasingly, programming involves coordinating data and computation among several layers, such as server-side, client-side and database layers of a typical three-tier Web application. The interaction between the host programming language (e.g. Java, C#, F#, Haskell or some other general-purpose language) running on the server and the query language (e.g. SQL) running on the database is particularly important, because the relational model and query language provided by the database differ from the data structures of most host languages. Conventional approaches to embedding database queries within a general-purpose language, such as Java's JDBC, provide the programmer with precise control over performance but are subject to typing errors and security vulnerabilities such as SQL injection attacks [35]. Object-relational mapping (ORM) tools and libraries, such as Java's Hibernate, provide a popular alternative by wrapping database access and update in type-safe object-oriented interfaces, but this leads to a loss of control over the structure of generated queries, which makes it difficult to understand and improve performance [14].

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 the author(s) 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.

PEPM '14, January 20–21, 2014, San Diego, CA, USA.  
Copyright is held by the owner/author(s). Publication rights licensed to ACM.  
ACM 978-1-4503-2619-3/14/01...\$15.00.  
http://dx.doi.org/10.1145/2543728.2543738

| employees  |        |         | tasks  |             |
|------------|--------|---------|--------|-------------|
| dpt        | name   | salary  | emp    | tsk         |
| "Product"  | "Alex" | 40,000  | "Alex" | "build"     |
| "Product"  | "Bert" | 60,000  | "Bert" | "build"     |
| "Research" | "Cora" | 50,000  | "Cora" | "abstract"  |
| "Research" | "Drew" | 70,000  | "Cora" | "build"     |
| "Sales"    | "Erik" | 200,000 | "Cora" | "call"      |
| "Sales"    | "Fred" | 95,000  | "Cora" | "dissemble" |
| "Sales"    | "Gina" | 155,000 | "Cora" | "enthuse"   |
|            |        |         | "Drew" | "abstract"  |
|            |        |         | "Drew" | "enthuse"   |
|            |        |         | "Erik" | "call"      |
|            |        |         | "Erik" | "enthuse"   |
|            |        |         | "Fred" | "call"      |
|            |        |         | "Gina" | "call"      |
|            |        |         | "Gina" | "dissemble" |

Figure 1. Sample Data

To avoid these so-called *impedance mismatch* problems, a number of *language-integrated query* techniques for embedding queries into general-purpose programming languages have emerged, which seek to reconcile the goals of type-safety and programmer control. Two distinctive styles of language-integrated query have emerged:

- Employ some form of static analysis or type system to identify parts of programs that can be turned into queries (e.g. Kleisli [38], Links [8], Batches for Java [36]).
- Extend a conventional language with explicit facilities for quotation or manipulation of query code (e.g. LINQ [21], Ur/Web [5], Database-Supported Haskell [13]).

Links is an example of the first approach. It uses a type-and-effect system [32] to classify parts of programs as executable only on the database, executable only on the host programming language, or executable anywhere. For example, consider the employee and task data in tables in Figure 1. The following code

```
for (x <- employees)
where(x.salary > 50000)
[(name=x.name)]
```

retrieves the names of employees earning over \$50,000, specifically ["Bert", "Drew", "Erik", "Fred", "Gina"]. In Links, the same code can be run either on the database (if `employees` and `tasks` are tables) or in the host language. If executed as a query, the interpreter generates a single (statically defined) SQL query that can take advantage of the database's indexing or other query optimisation; if executed in-memory, the expression will by default be interpreted as a quadratic nested loop. (Efficient in-memory implementations of query expressions are also possible [16].)

In contrast, in Microsoft’s LINQ (supported in C#, F#, and some other .NET languages), the programming language is extended with query-like syntax. For example, the same query as above can be written in F# as:

```
query { for x in employees
  where (x.salary > 50000)
  yield {name=x.name} }
```

This is just syntactic sugar for code that builds and manipulates quotations. In F#, this facility is built explicitly on top of language support for quotation [22, 29] and its *computation expression* syntax [26]. The above F# query expression is implemented by quoting the code inside the `query{ ... }` brackets and translating it (at run time) to C# values of type `Expression<T>`, which are converted to SQL by the .NET LINQ to SQL library.

The above example is rather simplistic: the query is *static*, that is, does not depend on any run-time data. Static queries can be handled easily even by libraries such as JDBC, and systems such as Links and LINQ provide the added benefit of type-safety. However, most queries are generated *dynamically*, depending on some run-time data. The ability to generate dynamic queries is essential for database programming. Libraries such as JDBC allow queries to be parameterized over base type values such as strings or integers, ensuring that values are correctly escaped to prevent SQL injection attacks. Both Links and LINQ go significantly further: they allow constructing dynamic queries using  $\lambda$ -abstraction and run-time normalisation, while retaining type safety and preventing SQL injection. However, this capability comes with its own pitfalls: it can be difficult to predict when an expression can be turned into a single query.

To address this problem, Cooper [7] showed how to extend Links so that performance-critical code can be highlighted with the `query` keyword. Links will statically check that the enclosed expression will definitely translate to a single query (neither failing at run-time, nor generating multiple queries). We refer to this as the *single-query guarantee*. In database theory, conservativity results due to Wong [37] and others provided a single-query guarantee in the case of first-order queries: any query expression having flat input and output types can be turned into an SQL query. This idea provided the basis for the Kleisli system [38], which was a source of inspiration to Links; the single-query guarantee was generalised to the higher-order case by Cooper [7], who also gave a static type-and-effect system that showed how to embed queries in a higher-order general-purpose language. Subsequent work on Links [19] generalised this to use row typing and effect polymorphism.

The possibility of generating LINQ queries dynamically in ad hoc cases was discussed by Syme [29] and Petricek [24, 25]. The F# and LINQ to SQL libraries in Microsoft .NET do not provide a single-query guarantee for dynamic queries; instead, they attempt to generate a single query but sometimes fail or generate multiple queries. In our recent paper [4] we showed that Cooper’s approach to normalisation for Links can be transferred to provide systematic support for abstraction in LINQ in F#, providing a single-query guarantee. In the rest of this paper, we consider the F# LINQ approach with this extension.

Nevertheless, there are still apparent differences between the approaches. For example, in LINQ, a query expression cannot be (easily) reused as ordinary code. This potentially leads to the need to write (essentially) the same code twice, once for ordinary use and once for use on the database. Code duplication can interfere with the use of functional abstraction to construct queries. For example, the following Links code

```
fun elem(x,xs) {
  not(empty(for (y <- xs) where (x == y) [()]))
}
```

```
fun canDo(name,tsk) {
  elem("build", for (t <- tasks)
    where (t.emp == name)
    [t.tsk])
}
query { for (x <- employees)
  where (canDo(x.name,"build"))
  [(name=x.name)] }
```

defines functions `elem` and `canDo` that test respectively whether a value is an element of a collection and whether an employee can do a certain task. The Links effect system correctly determines that `elem` can be run anywhere, and that `canDo` can be run on the database. When the query is to be executed, Links normalises the query by inlining `elem` and `canDo` and performing other transformations to generate a single SQL query [7]. In contrast, naively executing this code might involve loading all of the data from the `employees` table, and running one subquery to compute `canDo` for each `employees` row in-memory.

In F#, it is possible to do something similar, but only by explicitly quoting `elem` and `canDo`.

```
let elem = <@ fun x xs ->
  query { for y in xs
    exists(y = x) } @>
let canDo = <@ fun name tsk ->
  (%elem) tsk (for t in tasks
    where (t.emp = name)
    yield t.tsk) @>
query { for x in employees
  where ((%canDo) x.name "build")
  yield {name=x.name} }
```

The quoted version of `elem` is spliced into the query using antiquotation (`%elem`). If we need the `elem` function in both query and non-query code, its code must be duplicated, or we need to evaluate or generate compiled code for it at runtime. (F#’s quotation library does include `Eval` and `Compile` functions that can be used for this purpose, but it is not clear that these actually generate efficient code at runtime, nor is it convenient to write this boilerplate code.)

It is important to note that SQL does not natively support general recursion or first-class functional abstraction (although there are recent proposals to support the latter [15]). Nonrecursive lambda-abstraction is supported in query expressions in Links and F#, but it is eliminated in the process of generating an SQL query. Recursive functions can also be used to construct queries from data in the host language, in both Links and LINQ, but care is needed to make the staging explicit. For example, in F# we can define a predicate that tests whether an employee can do all tasks in some list as follows:

```
let rec canDoAll(tsks) =
  match tsks with
  [] -> <@ fun name -> true @>
  | tsk::tsks' -> <@ fun name ->
    (%canDo) name tsk && (%canDoAll tsks') name @>
query {
  for x in employees
  where ((%canDoAll ["build","call"]) x.name)
  yield {name=x.name} }
```

This is also possible in Links, but we need to use function abstraction and hoist subcomputations to satisfy the effect type system:

```
fun canDoAll(tsks) {
  switch (tsks) {
  case [] -> fun (name) {true}
```

```

case (tsk::tsks') ->
  var p = canDoAll(tsk);
  fun (name) { canDo(name,tsk) && p(name) } } }
query {
  for (x <- employees)
  where (canDoAll(["build","call"])(x.name))
  [(name=x.name)] }

```

We have to hoist the recursive call to `canDoAll` and name it so that it is clear to the type system that the recursive computation does not depend on values in the database that are not directly available to the host language interpreter. Arguably, in this case F#'s explicit quotation and antiquote annotations clarify the distinction between staging and functional abstraction, whereas in Links the distinction is not as explicit in the program syntax.

Both techniques are essentially heterogeneous forms of staged computation, based on a common foundation of manipulating partially-evaluated query expressions (or query fragments) at run time in order to construct SQL queries. The single-query property [4, 7] guarantees each query expression succeeds in generating one and only one query, even if lambda-abstraction or recursion is used to construct the queries. However, several natural questions about the relative strengths of the two approaches remain unanswered: Can we translate the Links effect-based approach to the (seemingly lower-level) LINQ quotation-based approach? If so, this might suggest a fruitful implementation strategy. Conversely, do we lose any expressiveness by providing the (seemingly higher-level) effect-based Links approach? Or can we always (in principle) translate LINQ-style quotation-based code to Links-style effect-based code?

In this paper, we consider the problem of relating the *expressiveness* of the two approaches to language-integrated query represented by Links and LINQ. Database query languages are often limited in expressiveness: for example, plain SQL conjunctive queries cannot express recursive properties such as transitive closure. Thus, understanding the relative expressiveness of different (Turing-incomplete) query languages, and the tradeoff with complexity of query evaluation or optimisation, are important issues in database theory [1]. We are interested in a dual question: what is the expressiveness of a programming language that generates queries? Given that both Links and LINQ approaches provide a measure of support for dynamic queries, can they express the same classes of dynamic queries?

To make this question precise, we introduce two core calculi:  $\text{Eff}_{\leq}$ , representing the effect-based approach supported by Links, and Quot, representing the quotation-based approach adopted in LINQ in F#. The former is similar to Cooper's core language [7]; the latter is essentially the same as the T-LINQ core language [4]. Both core languages make simplifying assumptions compared to Links and F# respectively, but we argue that they capture what is essential about the two approaches, making them suitable for a formal comparison that avoids preoccupation with other distracting details (e.g. Links's support for client-side programming [9] or F#'s support for objects [30], or different facilities for polymorphism in both languages.)

In database theory, the expressiveness of a language is usually measured by the set of functions definable in it, according to a conventional denotational semantics of database queries. However, this notion of expressiveness is not very interesting for general-purpose languages: two Turing-complete programming languages are always (by definition) expressively equivalent in this sense. Felleisen [12] and Mitchell [23] proposed notions of expressiveness based on restricted forms of translation among different languages. However, neither of these notions seems appropriate for relating language-integrated query formalisms.

Programs interact with a database that may be concurrently updated by other programs, and we want a notion of equivalence that takes query behaviour into account while abstracting over the possible concurrent behaviours of the database. For example, we want to consider a program that issues query  $Q$  to the database *inequivalent* to another program that reads all the database tables into memory and executes  $Q$  in-memory. In addition, we wish to abstract as much as possible over the possible behaviours of the database: databases are typically concurrently accessed and updated by many applications, and we want our notion of expressiveness to minimise assumptions about the behaviour of the database. Thus, we define the semantics of both languages as labeled transitions, where labels are either silent transitions or pairs  $(q, V)$  consisting of database queries and responses. We consider two programs *query-equivalent* if, given the same input, they have the same possible (finite and infinite) query/response traces  $(q_1, V_1), \dots, (q_n, V_n), \dots$ .

We show that  $\text{Eff}_{\leq}$  programs can be translated to query-equivalent Quot programs via a two-stage translation: first we eliminate subeffecting by duplicating code in the *doubling* translation, then we introduce explicit quotation and antiquote in the *splicing* translation. Perhaps more surprisingly, we can also give a converse translation from Quot to  $\text{Eff}_{\leq}$ , which translates quoted code to thunks (functions with unit domain): thus, the two approaches are expressively equivalent up to query-equivalence.

The current version of Links is interpreted, and query normalisation depends on being able to inspect code at run time. The translation from  $\text{Eff}_{\leq}$  to Quot suggests a compilation strategy by translating Links-style code to explicitly quoted code.

In the rest of this paper, we present the following contributions:

- We propose an appropriate notion of dynamic query behaviour suitable for comparing the expressiveness of different language-integrated query techniques.
- We provide a detailed exploration of the relationship between implicit, effect-based (Links/ $\text{Eff}_{\leq}$ ) and explicit, quotation-based (LINQ/Quot) approaches, giving type- and semantics-preserving translations in each direction.
- We discuss an application of the translation from  $\text{Eff}_{\leq}$  to Quot to support compilation of Links programs with embedded queries, along with preliminary experimental results.

The rest of this paper is structured as follows. Section 2 presents necessary background material from prior work, and defines the desired notion of equivalence of programs with respect to observable query behaviour. Section 3 presents  $\text{Eff}_{\leq}$  and Quot, giving their syntax, type systems, and operational semantics. Section 4 presents translations between  $\text{Eff}_{\leq}$  and Quot. Section 5 presents a practical application of the translation from  $\text{Eff}_{\leq}$  to Quot, which serves as the basis for a prototype compiler for Links that supports run-time dynamic query generation. Section 6 provide additional discussion of related work and Section 7 concludes.

## 2. Background

The Nested Relational Calculus (NRC) is a widely-studied core language for database queries corresponding closely to monadic comprehension syntax [2, 3]. Previous work [4, 7, 19, 37] has shown how first- and higher-order variants of NRC can be used for language-integrated query. We give the syntax of first-order NRC in Figure 2. Extended examples of the use of NRC are presented in prior work [2, 4].

We let  $x$  range over variables,  $c$  range over constants, and  $op$  range over primitive operators. Records  $(\bar{\ell} = q)$  and field projections  $q.\ell$  are standard. We write  $[]$  for an empty bag,  $[q]$  for the singleton bag containing the element  $q$ , and  $q + q'$  for the union of bags  $q$  and  $q'$ . We write for  $(x^A \leftarrow q)$   $q'$  for a bag comprehen-

|             |        |       |   |
|-------------|--------|-------|---|
| (Query)     | $q$    | $::=$ | $x \mid c \mid op(\bar{q}) \mid \text{if } q \ q' \mid \overline{\langle \ell = q \rangle} \mid q.\ell$ |
|             |        |       | $\mid \square \mid [q] \mid q_1 \ ++ \ q_2 \mid \text{for } (x \leftarrow q) \ q'$                      |
|             |        |       | $\mid \text{table } t$  |
| (Base type) | $O$    | $::=$ | $Int \mid Bool \mid String \mid \dots$  |
| (Type)      | $A, B$ | $::=$ | $O \mid \langle \ell = A \rangle \mid [A]$  |
| (Row type)  | $R$    | $::=$ | $\langle \ell : O \rangle$  |

**Figure 2.** Syntax of NRC

sion, which for each element  $x$  in  $q$  evaluates  $q'$ , then computes the union of the resulting bags. We write *table*  $t$  for the relational database table  $t$ . In order to keep normalisation as simple as possible we restrict ourselves to one sided conditionals over collections (equivalent to SQL where clauses). The expression *if*  $q \ q'$  evaluates to  $q'$  if  $q$  evaluates to true and  $\square$  if  $q$  evaluates to false. Lindley and Cheney [19] describe how to normalise in the presence of general conditionals; briefly, the idea is to push conditionals inside records and translate *if*  $q \ q' \ q''$  to *(if*  $q \ q') \ ++ \ (\text{if } (-q) \ q'')$  when  $q', q''$  are of bag type.

The NRC types include base types (*Int*, *Bool*, *String*, etc.), record types  $\langle \ell : A \rangle$ , and bag types  $[A]$ . Row types are *flat* record types restricted to contain base types (just like rows in SQL queries).

Conservativity results (see e.g. Wong [37]) ensure that any NRC expression  $M$  having a flat return type and flat inputs can be normalised to a form that corresponds directly to SQL.

We assume a fixed signature  $\Sigma$  mapping constants  $c$  to base types, operators  $op$  to functions on base types, and table references to flat bag types  $[R]$ . We omit typing or evaluation rules for queries; these are standard and implicit in the typing and operational semantics rules of  $\text{Eff}_{\leq}$  and  $\text{Quot}$  given later.

We will model the behaviour of a database server nondeterministically: whenever a query is posed, the response may be any value of the appropriate type. Thus, we fix a set  $\Omega$  of all pairs  $(q, V)$  such that whenever  $\vdash q : A$  we have  $\vdash V : A$ . Further constraints, reflecting the semantics of the query language or integrity constraints on the database tables, could be imposed. Our results concerning expressiveness are parametric in  $\Omega$  (provided it is at least type-safe and respects query equivalence).

**DEFINITION 1.** *Let  $L$  be a set of actions, including a “silent” action  $\tau$ , and let  $\mu$  range over elements of  $L$ . A labeled transition system over some set of labels  $L$  is a structure  $(X, \longrightarrow)$  where  $\longrightarrow \subseteq X \times L \times X$ . We write  $x \xrightarrow{\mu} y$  when  $(x, \mu, y) \in \longrightarrow$  for some  $\mu \in L$  and  $x \xrightarrow{\tau} y$  when  $x \xrightarrow{\tau} y$ . We write  $\Longrightarrow$  for  $\longrightarrow^*$  and write  $\xrightarrow{\mu} \text{for } \Longrightarrow \circ \xrightarrow{\mu} \circ \Longrightarrow$ . Suppose that  $x_1 \xrightarrow{\mu_1} x_2 \xrightarrow{\mu_2} \dots x_n \xrightarrow{\mu_n} \dots$  is a (finite or infinite) run starting from  $x_1$ . We say that the (finite or infinite) word  $\mu_1 \dots \mu_n \dots$  obtained from concatenating the non- $\tau$  labels is the trace of the run, and we define  $\text{Tr}(x)$  to be the set of all traces starting from  $x$ .*

**DEFINITION 2.** *Let  $(X, \longrightarrow_X)$  and  $(Y, \longrightarrow_Y)$  be labeled transition systems over the same label set  $L$ . We say that  $x \in X$  trace-simulates  $y \in Y$  (or  $x \lesssim y$ ) if  $\text{Tr}(x) \subseteq \text{Tr}(y)$ , and  $x \in X$  and  $y \in Y$  are trace-equivalent ( $x \approx y$ ) if  $\text{Tr}(x) = \text{Tr}(y)$ .*

In the rest of this paper, we are concerned with instances of trace-equivalence where the labels are either  $\tau$  or pairs  $(q, V) \in \Omega$ .

### 3. Formalisation

#### 3.1 Source Language: $\text{Eff}_{\leq}$

$\text{Eff}_{\leq}$  is a higher-order nested relational calculus over bags augmented with an effect type system for issuing flat NRC relational

queries. It is similar to calculi considered in previous work on query compilation [7, 19]. A difference is that  $\text{Eff}_{\leq}$  models two ground effects (**pl** and **db**), whereas the other calculi model just one (indicating code that cannot be run in the database). Like Cooper’s system,  $\text{Eff}_{\leq}$  provides subeffecting. Our previous work [19], and its implementation in the Links web programming language [8] uses row-based effect polymorphism instead; while effect polymorphism is more flexible, we focus on the simpler subtype-based approach here.

The effects of  $\text{Eff}_{\leq}$  are given by the following grammar.

|                  |        |       |                       |
|------------------|--------|-------|-----------------------|
| (ground effects) | $X, Y$ | $::=$ | <b>pl</b>   <b>db</b> |
| (effects)        | $E$    | $::=$ | $X \mid \text{any}$   |

Code that runs in the programming language has the **pl** effect. Code that runs in the database has the **db** effect. Code that can be run in either place, i.e. anywhere, has the **any** effect. The typing rules will enforce a *subeffecting* discipline, allowing a function that has the **any** effect to be applied with either the **pl** or the **db** effect.

The types of  $\text{Eff}_{\leq}$  are given by the following grammar.

|         |        |       |   |
|---------|--------|-------|---|
| (types) | $A, B$ | $::=$ | $O \mid \langle \ell : A \rangle \mid [A] \mid A \rightarrow^E B$ |
|---------|--------|-------|---|

Types in  $\text{Eff}_{\leq}$  extend the NRC base types, records, and collections with function types  $A \rightarrow^E B$  annotated by an effect  $E$ .

The terms of  $\text{Eff}_{\leq}$  are given by the following grammar.

|           |       |   |
|-----------|-------|---|
| $L, M, N$ | $::=$ | $x \mid c \mid op(\overline{M}) \mid \lambda^E x^A. M \mid M \ N$               |
|           |       | $\mid \text{if } L \ M \mid \langle \ell = M \rangle \mid M.\ell$               |
|           |       | $\mid \square \mid [M] \mid M \ ++ \ N \mid \text{for } (x^A \leftarrow M) \ N$ |
|           |       | $\mid \text{rec } f^{A \rightarrow P^1 B} x^A. M \mid \text{fold } L \ M \ N$   |
|           |       | $\mid \text{query } M \mid \text{table } t$                                     |

We will often omit type and effect annotations on bindings. We include them so that later we can define translations on terms rather than on judgements (Section 4).

The syntax is in most cases similar to NRC or standard. Lambda-abstractions  $\lambda^E x^A. M$  are annotated with an effect  $E$ . Recursive functions  $\text{rec } f \ x.M$  and folds are only available in the programming language. The fold operation is the only elimination form for bag values. The special form *query*  $M$  denotes a query returning results of type  $[R]$ , where  $M$  has type  $[R]$ .

A type environment ascribes types to variables.

|                  |       |                            |
|------------------|-------|----------------------------|
| $\Gamma, \Delta$ | $::=$ | $\cdot \mid \Gamma, x : A$ |
|------------------|-------|----------------------------|

The typing rules are given in Figure 3. In this figure and elsewhere, notations such as  $\overline{M}$  or  $\overline{\Gamma} \vdash \overline{M} : A!E$  abbreviate lists of terms, judgments, etc. The typing judgment is of the form  $\Gamma \vdash M : A!E$ , indicating that in context  $\Gamma$ , term  $M$  has type  $A$  and effect  $E$ . The effect  $E$  can be either **pl** or **db**, indicating that  $M$  can only be executed by the host programming language or by the database respectively, or it can be **any**, indicating that  $M$  can be executed in either context.

We again assume a signature  $\Sigma$  that maps each constant  $c$  to its underlying type, each primitive operator  $op$  to its type (e.g.  $\Sigma(\wedge) = (Bool, Bool) \rightarrow Bool$  and  $\Sigma(+)= (Int, Int) \rightarrow Int$ ), and each table  $t$  to the type of its rows. (For simplicity we assume that the same base types, constants and primitive operations are available to both the host and database.) The rules are quite standard; most of them are parametric in the effect. The most interesting rules are those that do something non-trivial with effects, such as **TABLE**, **QUERY**, **REC**, **FOLD** and **APP**. The **TABLE** rule ensures that table references *table*  $t$  can only be used in a database context. The **QUERY** rule requires that the body of a query expression *query*  $M$  must be run in the database, and the query expression itself must be invoked from the programming language. This rule also implicitly requires that a query result type  $[R]$  is a bag of records (note that  $R$  is a row type). Recursive functions (rule **REC**) and folds (rule



|   |  |   |   |  |  |  |  |  |  |
|---|--|---|---|--|--|--|--|--|--|
| $\Gamma \vdash M : A ! E$   |  | $E_1 \leq E_2$  |   |  |  |  |  |  |  |
| <b>VAR</b><br>$\frac{}{\Gamma, x : A \vdash x : A ! E}$   | <b>CONST</b><br>$\frac{\Sigma(c) = A}{\Gamma \vdash c : A ! E}$  | <b>OP</b><br>$\frac{\Sigma(op) = (\overline{O}) \rightarrow O \quad \overline{\Gamma \vdash M : O ! E}}{\Gamma \vdash op(\overline{M}) : O ! E}$          | <b>LAM</b><br>$\frac{\Gamma, x : A \vdash M : B ! E}{\Gamma \vdash \lambda x^A . M : A \rightarrow^E B ! E'}$   |  |  |  |  |  |  |
| <b>APP</b><br>$\frac{\Gamma \vdash M : A \rightarrow^{E'} B ! E \quad \Gamma \vdash N : A ! E \quad E' \leq E}{\Gamma \vdash M N : B ! E}$  | <b>IF</b><br>$\frac{\Gamma \vdash L : Bool ! E \quad \Gamma \vdash M : [A] ! E}{\Gamma \vdash \text{if } L M : [A] ! E}$ | <b>RECORD</b><br>$\frac{\overline{\Gamma \vdash M : A ! E}}{\Gamma \vdash \langle \ell = \overline{M} \rangle : \langle \ell : \overline{A} \rangle ! E}$ |   |  |  |  |  |  |  |
| <b>PROJECT</b><br>$\frac{}{\Gamma \vdash M : \langle \ell : \overline{A} \rangle ! E}$  | <b>EMPTY</b><br>$\frac{}{\Gamma \vdash [] : [A] ! E}$  | <b>SINGLETON</b><br>$\frac{}{\Gamma \vdash M : [A] ! E}$  | <b>UNION</b><br>$\frac{\Gamma \vdash M : [A] ! E \quad \Gamma \vdash N : [A] ! E}{\Gamma \vdash M ++ N : [A] ! E}$  |  |  |  |  |  |  |
| <b>FOR</b><br>$\frac{\Gamma \vdash M : [A] ! E \quad \Gamma, x : A \vdash N : [B] ! E}{\Gamma \vdash \text{for } (x^A \leftarrow M) N : [B] ! E}$   | <b>TABLE</b><br>$\frac{\Sigma(t) = [A]}{\Gamma \vdash \text{table } t : [A] ! \text{db}}$                                | <b>QUERY</b><br>$\frac{\Gamma \vdash M : [R] ! \text{db}}{\Gamma \vdash \text{query } M : [R] ! \text{pl}}$   | <b>REC</b><br>$\frac{\Gamma, f : A \rightarrow^{\text{pl}} B, x : A \vdash M : B ! \text{pl}}{\Gamma \vdash \text{rec } f^{A \rightarrow^{\text{pl}} B} x^A . M : A \rightarrow^{\text{pl}} B ! E}$ |  |  |  |  |  |  |
| <b>FOLD</b><br>$\frac{\Gamma \vdash L : [A] ! \text{pl} \quad \Gamma \vdash M : B ! \text{pl} \quad \Gamma \vdash N : A \rightarrow^{\text{pl}} B \rightarrow^{\text{pl}} B ! \text{pl}}{\Gamma \vdash \text{fold } L M N : B ! \text{pl}}$ |  |   | <b>REFLEXIVITY</b><br>$\frac{}{E \leq E}$   | <b>ANYPL</b><br>$\frac{}{\text{any} \leq \text{pl}}$ | <b>ANYDB</b><br>$\frac{}{\text{any} \leq \text{db}}$ |  |  |  |  |

**Figure 3.** Typing and Subeffecting Rules for  $\text{Eff}_{\leq}$

|   |   |   |   |  |  |  |  |  |  |  |
|---|---|---|---|--|--|--|--|--|--|--|
| $\Gamma \vdash M : A$   |   |   |   |  |  |  |  |  |  |  |
| <b>VAR</b><br>$\frac{}{\Gamma, x : A \vdash x : A}$   | <b>CONST</b><br>$\frac{\Sigma(c) = A}{\Gamma \vdash c : A}$   | <b>OP</b><br>$\frac{\Sigma(op) = (\overline{O}) \rightarrow O \quad \overline{\Gamma \vdash M_i : O_i}}{\Gamma \vdash op(\overline{M}) : O}$                        | <b>LIFT</b><br>$\frac{}{\Gamma \vdash \text{lift } M : \text{Expr} < O >}$                            | <b>LAM</b><br>$\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x^A . M : A \rightarrow B}$ |  |  |  |  |  |  |
| <b>APP</b><br>$\frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash M N : B}$ | <b>IF</b><br>$\frac{\Gamma \vdash L : Bool \quad \Gamma \vdash M : [A]}{\Gamma \vdash \text{if } L M : [A]}$                                    | <b>RECORD</b><br>$\frac{}{\Gamma \vdash \langle \ell = \overline{M} \rangle : \langle \ell : \overline{A} \rangle}$   | <b>PROJECT</b><br>$\frac{}{\Gamma \vdash M . \ell_i : A_i}$   | <b>EMPTY</b><br>$\frac{}{\Gamma \vdash [] : [A]}$  |  |  |  |  |  |  |
| <b>SINGLETON</b><br>$\frac{}{\Gamma \vdash M : [A]}$  | <b>UNION</b><br>$\frac{\Gamma \vdash M : [A] \quad \Gamma \vdash N : [A]}{\Gamma \vdash M ++ N : [A]}$  | <b>FOR</b><br>$\frac{\Gamma \vdash M : [A] \quad \Gamma, x : A \vdash N : [B]}{\Gamma \vdash \text{for } (x \leftarrow M) N : [B]}$                                 | <b>QUOTE</b><br>$\frac{\Gamma \vdash M : A}{\Gamma \vdash \langle @ M @ \rangle : \text{Expr} < A >}$ |  |  |  |  |  |  |  |
| <b>QUERY</b><br>$\frac{}{\Gamma \vdash \text{query } M : [R]}$  | <b>REC</b><br>$\frac{\Gamma, f : A \rightarrow B, x : A \vdash M : B}{\Gamma \vdash \text{rec } f^{A \rightarrow B} x^A . M : A \rightarrow B}$ | <b>FOLD</b><br>$\frac{\Gamma \vdash L : [A] \quad \Gamma \vdash M : B \quad \Gamma \vdash N : A \rightarrow B \rightarrow B}{\Gamma \vdash \text{fold } L M N : B}$ |   |  |  |  |  |  |  |  |

|  |   |   |  |  |  |  |  |  |  |  |
|--|---|---|--|--|--|--|--|--|--|--|
| $\Gamma; \Delta \vdash M : A$  |   |   |  |  |  |  |  |  |  |  |
| <b>VARQ</b><br>$\frac{}{\Gamma; \Delta, x : A \vdash x : A}$   | <b>CONSTQ</b><br>$\frac{\Sigma(c) = A}{\Gamma; \Delta \vdash c : A}$  | <b>OPQ</b><br>$\frac{\Sigma(op) = (\overline{O}) \rightarrow O \quad \overline{\Gamma; \Delta \vdash M : O}}{\Gamma; \Delta \vdash op(\overline{M}) : O}$ | <b>LAMQ</b><br>$\frac{\Gamma; \Delta, x : A \vdash M : B}{\Gamma; \Delta \vdash \lambda x^A . M : A \rightarrow B}$  |  |  |  |  |  |  |  |
| <b>APPQ</b><br>$\frac{\Gamma; \Delta \vdash M : A \rightarrow B \quad \Gamma; \Delta \vdash N : A}{\Gamma; \Delta \vdash M N : B}$ | <b>IFQ</b><br>$\frac{\Gamma; \Delta \vdash L : Bool \quad \Gamma; \Delta \vdash M : A}{\Gamma; \Delta \vdash \text{if } L M : A}$ | <b>RECORDQ</b><br>$\frac{}{\Gamma; \Delta \vdash \langle \ell = \overline{M} \rangle : \langle \ell : \overline{A} \rangle}$                              | <b>PROJECTQ</b><br>$\frac{}{\Gamma; \Delta \vdash M . \ell_i : A_i}$   |  |  |  |  |  |  |  |
| <b>EMPTYQ</b><br>$\frac{}{\Gamma; \Delta \vdash []^A : [A]}$   | <b>SINGLETONQ</b><br>$\frac{}{\Gamma; \Delta \vdash M : [A]}$   | <b>UNIONQ</b><br>$\frac{\Gamma; \Delta \vdash M : [A] \quad \Gamma; \Delta \vdash N : [A]}{\Gamma; \Delta \vdash M ++ N : [A]}$                           | <b>FORQ</b><br>$\frac{\Gamma; \Delta \vdash M : [A] \quad \Gamma; \Delta, x : A \vdash N : [B]}{\Gamma; \Delta \vdash \text{for } (x \leftarrow M) N : [B]}$ |  |  |  |  |  |  |  |
| <b>TABLE</b><br>$\frac{\Sigma(t) = [R]}{\Gamma; \Delta \vdash \text{table } t : [R]}$  |   | <b>ANTIQUOTE</b><br>$\frac{\Gamma \vdash M : \text{Expr} < A >}{\Gamma; \Delta \vdash (\%M) : A}$   |  |  |  |  |  |  |  |  |

**Figure 4.** Typing Rules for Quot

(value)  
 $V, W ::= c \mid \lambda x.M \mid \text{rec } f x.M \mid \overline{\langle \ell = \bar{V} \rangle} \mid \overline{[V]}$   
(evaluation context)  
 $\mathcal{E} ::= [\ ] \mid \text{op}(\overline{V}, \mathcal{E}, \overline{M}) \mid \mathcal{E} M \mid V \mathcal{E}$   
 $\mid \overline{\langle \ell = \bar{V}, \ell' = \mathcal{E}, \ell'' = \overline{M} \rangle} \mid \mathcal{E}.\ell \mid [\mathcal{E}]$   
 $\mid \mathcal{E} \text{ ++ } M \mid [V] \text{ ++ } \mathcal{E} \mid \text{for } (x \leftarrow \mathcal{E}) N$   
 $\mid \text{if } \mathcal{E} M$

**Figure 5.** Values and Evaluation Contexts for Eff

$$\begin{aligned} \text{op}(\overline{V}) &\longrightarrow \delta(\text{op}, \overline{V}) \\ (\lambda x.M) V &\longrightarrow M[x := V] \\ (\text{rec } f x.M) V &\longrightarrow M[f := \text{rec } f x.M, x := V] \\ \overline{\langle \ell = \bar{V} \rangle}.\ell_i &\longrightarrow V_i \\ \text{if true } M &\longrightarrow M \\ \text{if false } M &\longrightarrow [\ ] \\ [\ ] \text{ ++ } V &\longrightarrow V \\ ([V] \text{ ++ } V') \text{ ++ } V'' &\longrightarrow [V] \text{ ++ } (V' \text{ ++ } V'') \\ \text{for } (x \leftarrow [\ ]) N &\longrightarrow [\ ] \\ \text{for } (x \leftarrow [V] \text{ ++ } W) N &\longrightarrow N[x := V] \text{ ++ } \\ &\quad (\text{for } (x \leftarrow W) N) \\ \text{fold } [\ ] M N &\longrightarrow M \\ \text{fold } (([V] \text{ ++ } W) M N) &\longrightarrow N V (\text{fold } W M N) \\ \text{query } M \xrightarrow{(|M|, V)} V &\quad ((|M|, V) \in \Omega) \\ \frac{M \xrightarrow{\mu} N}{\mathcal{E}[M] \xrightarrow{\mu} \mathcal{E}[N]} \end{aligned}$$

**Figure 6.** Operational Semantics for Eff<sub>z</sub>

FOLD) can only be applied in the programming language. The APP rule allows a function with the **any** effect to be applied in the programming language or the database.

We define a sub-language Eff by restriction of Eff<sub>z</sub> to programs in which the **any** effect is disallowed. The subeffecting constraint in the APP rule is superfluous in this sublanguage. We can translate any Eff<sub>z</sub> program to an equivalent, albeit longer, Eff program, by representing each **any** function as a pair of a **pl** function and a **db** function (Section 4.1).

### 3.2 Operational Semantics for Eff<sub>z</sub>

We now present small-step operational semantics for Eff<sub>z</sub>. The syntax of values and evaluation contexts is given in Figure 5. The values are standard. We write  $\overline{[V]}$  for  $[V_1] \text{ ++ } \dots \text{ ++ } [V_n] \text{ ++ } [\ ]$ . The operational semantics is defined by reduction relation  $M \xrightarrow{\mu} N$  as shown in Figure 6. It is parameterised by an interpretation  $\delta$  for each primitive operation  $\text{op}$ , and a set  $\Omega$  of possible query request and response pairs  $(q, V)$ , both of which respect types: if  $\Sigma(\text{op}) = \overline{O} \rightarrow O$  and  $\vdash \overline{V} : \overline{O}$  and  $V = \delta(\text{op}, \overline{V})$  then  $\vdash V : O$ , and if  $(q, V) \in \Omega$  and  $\vdash q : [R]$  then  $\vdash V : [R]$ .

The rules are standard apart from the one for query evaluation. Evaluation contexts  $\mathcal{E}$  enforce left-to-right call-by-value evaluation. Rule (query) evaluates a query  $M$  by first *normalising*  $M$  to yield an equivalent NRC query  $q = |M|$ , and then taking a transition labeled  $(q, V)$  yielding result value  $V$ . The normalisation function  $| - |$  is the same as that of Cheney et al. [4]. It first applies standard symbolic reduction rules (Figure 7) and then further ad hoc reduction rules (Figure 8). The former eliminate all nesting and abstraction from a closed term of flat bag type, while the latter account for the lack of uniformity in SQL (see [4] for further de-

$$\begin{aligned} (\lambda x.M) N &\rightsquigarrow M[x := N] \\ \overline{\langle \ell = \bar{M} \rangle}.\ell_i &\rightsquigarrow M_i \\ \text{for } (x \leftarrow [M]) N &\rightsquigarrow N[x := M] \\ \text{for } (y \leftarrow \text{for } (x \leftarrow L) M) N &\rightsquigarrow \text{for } (x \leftarrow L) \text{ for } (y \leftarrow M) N \\ \text{for } (x \leftarrow \text{if } L M) N &\rightsquigarrow \text{if } L (\text{for } (x \leftarrow M) N) \\ &\quad \text{for } (x \leftarrow [\ ]) N \rightsquigarrow [\ ] \\ \text{for } (x \leftarrow (L \text{ ++ } M)) N &\rightsquigarrow \\ &\quad (\text{for } (x \leftarrow L) N) \text{ ++ } (\text{for } (x \leftarrow M) N) \\ \text{if true } M &\rightsquigarrow M \\ \text{if false } M &\rightsquigarrow [\ ] \end{aligned}$$

**Figure 7.** Normalisation Stage 1: symbolic reduction

$$\begin{aligned} \text{for } (x \leftarrow L) (M \text{ ++ } N) &\hookrightarrow (\text{for } (x \leftarrow L) M) \text{ ++ } (\text{for } (x \leftarrow L) N) \\ \text{for } (x \leftarrow L) [\ ] &\hookrightarrow [\ ] \\ \text{if } L (M \text{ ++ } N) &\hookrightarrow (\text{if } L M) \text{ ++ } (\text{if } L N) \\ \text{if } L [\ ] &\hookrightarrow [\ ] \\ \text{if } L (\text{if } M N) &\hookrightarrow \text{if } (L \ \& \ M) N \\ \text{if } L (\text{for } (x \leftarrow M) N) &\hookrightarrow \text{for } (x \leftarrow M) (\text{if } L N) \end{aligned}$$

**Figure 8.** Normalisation Stage 2: ad hoc reduction

tails). Define  $|L| = N$  when  $L \rightsquigarrow^* M$  and  $M \hookrightarrow^* N$ , where  $M$  and  $N$  are in normal form with respect to  $\rightsquigarrow$  and  $\hookrightarrow$  respectively.

It is straightforward to show type soundness via the usual method of preservation and progress.

**PROPOSITION 3.** *If  $\Gamma \vdash M : A ! E$  and  $M \xrightarrow{\mu} N$  then  $\Gamma \vdash N : A ! E$ . If  $\Gamma \vdash M : A ! E$  then either  $M$  is a value or  $M \xrightarrow{\mu} N$  for some  $N$  and  $\mu$ .*

### 3.3 Target Language: Quot

Quot is a higher-order nested relational calculus over bags augmented with a quotation mechanism for constructing NRC queries using quotation.

Modulo superficial differences Quot is essentially the same as the T-LINQ core language [4]. Specifically, the differences are: the lexical syntax; Quot includes fold; Quot includes extra type annotations (to aid the translation to Eff); T-LINQ has a database construct whereas Quot has a table construct.

Where Eff<sub>z</sub> uses effect types to distinguish the *programming language* from the *database*, Quot uses quotation to distinguish the *host language* from the *query language*. Host language terms may build and evaluate quoted query language terms. For convenience (and ease of comparison with Eff<sub>z</sub>), we use the same syntax for the host and query languages, so that the query language is essentially a sublanguage of the host language (except for antiquotation and table  $t$ ). In general, the two languages may be different, as already discussed elsewhere [4].

The types of Quot are given by the following grammar.

$$\text{(types)} \quad A, B ::= O \mid \overline{\langle \ell : \bar{A} \rangle} \mid [A] \mid A \rightarrow B \mid \text{Expr} < A >$$

They are the same as for Eff<sub>z</sub>, except function types are not annotated with effects, and types are extended to include *quotation types*  $\text{Expr} < A >$ , which represent closed, quoted query terms of type  $A$ .

The terms of Quot are given by the following grammar.

$$\begin{aligned} L, M, N &::= x \mid c \mid \text{op}(\overline{M}) \mid \lambda x^A.M \mid M N \\ &\mid \text{if } L M \mid \overline{\langle \ell = \bar{M} \rangle} \mid M.\ell \\ &\mid [\ ] \mid [M] \mid M \text{ ++ } N \mid \text{for } (x^A \leftarrow M) N \\ &\mid \text{rec } f^{A \rightarrow B} x^A.M \mid \text{fold } L M N \\ &\mid \text{query } M \mid \text{table } t \\ &\mid \langle \mathcal{Q} M \mathcal{Q} \rangle \mid (\%M) \mid \text{lift } M \end{aligned}$$

The grammar is largely the same as that of  $\text{Eff}_\leq$ . The key difference is that effects are replaced by quotation constructs. Lambda-abstractions are no longer annotated with effects. More importantly, Quot includes quotation: we write  $\langle\!\langle M \rangle\!\rangle$  for the *quotation* operation, where  $M$  is a query term; we write  $\langle\!\langle M \rangle\!\rangle$  for the *antiquotation* operation, which splices a quoted term  $M$  into a query term; and we write  $\text{lift } M$  for the operation that coerces a value of base type to a quoted value. (This is a limited form of *cross-stage persistence* [31], which is otherwise unavailable because quotation expressions must be closed.)

Type environments  $\Gamma, \Delta$  are as follows.

$$\Gamma, \Delta ::= \cdot \mid \Gamma, x : A$$

The typing rules are given in Figure 4. There are two typing judgements: one for host terms, and the other for query terms. The judgement  $\Gamma \vdash M : A$  states that host term  $M$  has type  $A$  in type environment  $\Gamma$ . The judgement  $\Gamma; \Delta \vdash M : A$  states that query term  $M$  has type  $A$  in host type environment  $\Gamma$  and query type environment  $\Delta$ .

Most of the typing rules are standard and similar in both judgements. The variable typing rule (VARQ) for the query judgement does not allow the use of variables from the  $\Gamma$  environment; hence variables from  $\Gamma$  must be explicitly lifted (rule LIFT) or used within an antiquotation (rule ANTIQUOTE) in order to be used within a quotation.

The interesting rules are those that involve quotation, namely QUOTE, QUERY, ANTIQUOTE, and LIFT. Query terms can only be quoted in the host language (QUOTE). In the host language, a closed quoted term of flat bag type can be evaluated as a query (QUERY). A quoted term can be spliced into a query term (ANTIQUOTE). A host term of base type can be lifted to a query term (LIFT). Note that variables of bag, record, or function type in  $\Gamma$  cannot be lifted using LIFT. As in  $\text{Eff}_\leq$ , the  $\text{table } t$  construct is only available within a query; its use within queries is enabled by rule TABLE and its use elsewhere is forbidden because there is no similar rule for unquoted table references.

### 3.4 Operational Semantics for Quot

The operational semantics for Quot is standard, and is similar to that given in Cheney et al. [4] for the T-LINQ language. Values and evaluation contexts are given in Figure 9. The normalisation function  $\| - \|$  is the same as  $| - |$  but applied to Quot instead of Eff. Define  $\|L\| = N$  when  $L \rightsquigarrow^* M$  and  $M \hookrightarrow^* N$ , where  $M$  and  $N$  are in normal form with respect to  $\rightsquigarrow$  and  $\hookrightarrow$  respectively. The semantics is given in Figure 10; most rules are the same as in  $\text{Eff}_\leq$ , except those involving quotation and querying. The rules are parameterised by the same  $\delta$  and  $\Omega$  as the semantics of  $\text{Eff}_\leq$ .

**PROPOSITION 4.** *If  $\Gamma \vdash M : A$  and  $M \xrightarrow{\mu} N$  then  $\Gamma \vdash N : A$ . If  $\Gamma \vdash M : A$  then either  $M$  is a value or  $M \xrightarrow{\mu} N$  for some  $N$  and  $\mu$ .*

## 4. Translations

In this section we present translations that show that  $\text{Eff}_\leq$  and Quot can simulate one another. The translations are summarised in Figure 11. We first (Section 4.1) show how to compile away the **any**-effect, translating arbitrary  $\text{Eff}_\leq$  programs to **any**-free Eff programs. Next (Section 4.2) we show how to translate Eff programs to Quot programs. In the reverse direction, we first (Section 4.3) give a straightforward translation hoisting computations out of antiquotations in Quot, resulting in a normal form Quot' in which all antiquotations are of the form  $\langle\!\langle x \rangle\!\rangle$ . Finally (Section 4.4) we translate Quot' programs to Eff programs.

(value)  
 $V, W ::= c \mid \lambda x. M \mid \text{rec } f x. M \mid \langle\ell = V\rangle \mid [\overline{V}] \mid \langle\!\langle Q \rangle\!\rangle$   
 (quotation value)  
 $Q ::= x \mid c \mid \text{op}(\overline{Q}) \mid \lambda x. Q \mid Q Q' \mid$   
 $\quad \mid \text{if } Q Q' \mid \langle\ell = Q\rangle \mid Q.\ell$   
 $\quad \mid [] \mid [Q] \mid Q + Q' \mid \text{for } (x \leftarrow Q) Q' \mid \text{table } t$   
 (evaluation context)  
 $\mathcal{E} ::= [] \mid \text{op}(\overline{V}, \mathcal{E}, \overline{M}) \mid \text{lift } \mathcal{E} \mid \mathcal{E} M \mid V \mathcal{E}$   
 $\quad \mid \langle\ell = V, \ell' = \mathcal{E}, \ell'' = M\rangle \mid \mathcal{E}.\ell \mid [\mathcal{E}]$   
 $\quad \mid \mathcal{E} + M \mid [V] + \mathcal{E} \mid \text{for } (x \leftarrow \mathcal{E}) N$   
 $\quad \mid \text{if } \mathcal{E} M \mid \text{query } \mathcal{E} \mid \langle\!\langle Q \rangle\!\rangle \mid \langle\!\langle \langle\!\langle \mathcal{E} \rangle\!\rangle \rangle\!\rangle$   
 (quotation context)  
 $\mathcal{Q} ::= [] \mid \text{op}(\overline{Q}, \mathcal{Q}, \overline{M}) \mid \lambda x. \mathcal{Q}$   
 $\quad \mid \mathcal{Q} M \mid \mathcal{Q} \mathcal{Q} \mid \langle\ell = \mathcal{Q}, \ell' = \mathcal{Q}, \ell'' = M\rangle$   
 $\quad \mid \mathcal{Q}.\ell \mid [\mathcal{Q}] \mid \mathcal{Q} + M \mid \mathcal{Q} + \mathcal{Q}$   
 $\quad \mid \text{for } (x \leftarrow \mathcal{Q}) N \mid \text{for } (x \leftarrow \mathcal{Q}) \mathcal{Q}$   
 $\quad \mid \text{if } \mathcal{Q} M \mid \text{if } \mathcal{Q} \mathcal{Q}$

**Figure 9.** Values and Evaluation Contexts for Quot

$$\begin{array}{l} \text{op}(\overline{V}) \longrightarrow \delta(\text{op}, \overline{V}) \\ (\lambda x. M) V \longrightarrow M[x := V] \\ (\text{rec } f x. M) V \longrightarrow M[f := \text{rec } f x. M, x := V] \\ \langle\ell = V\rangle.\ell_i \longrightarrow V_i \\ \text{if true } M \longrightarrow M \\ \text{if false } M \longrightarrow [] \\ [] + V \longrightarrow V \\ ([V] + V') + V'' \longrightarrow [V] + (V' + V'') \\ \text{fold } [] M N \longrightarrow M \\ \text{fold } (([V] + W) M N \longrightarrow N V (\text{fold } W M N) \\ \text{for } (x \leftarrow []) N \longrightarrow [] \\ \text{for } (x \leftarrow [V] + W) N \longrightarrow N[x := V] + \\ \quad \quad \quad (\text{for } (x \leftarrow W) N) \\ \text{query } \langle\!\langle Q \rangle\!\rangle \xrightarrow{(\|Q\|, V)} V \quad ((\|Q\|, V) \in \Omega) \\ \text{lift } c \longrightarrow \langle\!\langle c \rangle\!\rangle \quad (\text{lift}) \\ \langle\!\langle Q \rangle\!\rangle \mid \langle\!\langle \langle\!\langle Q \rangle\!\rangle \rangle\!\rangle \longrightarrow \langle\!\langle Q \rangle\!\rangle \quad (\text{splice}) \end{array}$$

$$\frac{M \xrightarrow{\mu} N}{\mathcal{E}[M] \xrightarrow{\mu} \mathcal{E}[N]}$$

**Figure 10.** Operational Semantics for Quot

$$\begin{array}{c} \text{Eff}_\leq \xrightarrow{\langle\!\langle - \rangle\!\rangle (4.1)} \text{Eff} \xrightarrow{\| - \| (4.2)} \text{Quot} \\ \text{Eff}_\leq \xrightarrow{\supseteq} \text{Eff} \xrightarrow{\langle\!\langle - \rangle\!\rangle (4.4)} \text{Quot}' \xrightarrow{\subseteq (4.3)} \text{Quot} \end{array}$$

**Figure 11.** Summary of the Translations

### 4.1 From $\text{Eff}_\leq$ to Eff

As a first step, we can translate away all subeffecting through a global *doubling* translation that simulates each **any** function as a pair of a **pl** function and a **db** function. The type translation  $\langle A \rangle$  is shown in Figure 12. Type environments are translated pointwise:

$$\langle x_1 : A_1, \dots, x_n : A_n \rangle = x_1 : \langle A_1 \rangle, \dots, x_n : \langle A_n \rangle$$

Terms are translated as shown in Figure 13, where  $\langle M \rangle_X$  stands for the doubling translation of  $M$  with respect to target effect  $X$ . The type and term translations are structure-preserving except on **any**-function types, abstractions, and applications; these interest-

$$\begin{aligned}
\langle \text{Int} \rangle &= \text{Int} \\
\langle \text{Bool} \rangle &= \text{Bool} \\
\langle \text{String} \rangle &= \text{String} \\
\langle A \rightarrow^X B \rangle &= \langle A \rangle \rightarrow^X \langle B \rangle \quad X \in \{\text{db}, \text{pl}\} \\
\langle A \rightarrow^{\text{any}} B \rangle &= \langle \langle A \rangle \rightarrow^{\text{pl}} \langle B \rangle, \langle A \rangle \rightarrow^{\text{db}} \langle B \rangle \rangle \\
\langle \ell : A \rangle &= \langle \ell : \langle A \rangle \rangle \\
\langle [A] \rangle &= \langle [A] \rangle
\end{aligned}$$

**Figure 12.** Doubling Translation: types

$$\begin{aligned}
\langle x \rangle_X &= x \\
\langle c \rangle_X &= c \\
\langle \text{op}(\overline{M}) \rangle_X &= \text{op}(\langle \overline{M} \rangle_X) \\
\langle \lambda^Y x^A . M \rangle_X &= \lambda^Y x^{\langle A \rangle} . \langle M \rangle_Y \quad Y \in \{\text{db}, \text{pl}\} \\
\langle \lambda^{\text{any}} x^A . M \rangle_X &= \langle \lambda^{\text{pl}} x^{\langle A \rangle} . \langle M \rangle_{\text{pl}}, \lambda^{\text{db}} x^{\langle A \rangle} . \langle M \rangle_{\text{db}} \rangle \\
\langle M^{A \rightarrow^X B} N \rangle_X &= \langle M \rangle_X \langle N \rangle_X \\
\langle M^{A \rightarrow^{\text{any}} B} N \rangle_{\text{pl}} &= \langle M \rangle_{\text{pl}} . 1 \langle N \rangle_{\text{pl}} \\
\langle M^{A \rightarrow^{\text{any}} B} N \rangle_{\text{db}} &= \langle M \rangle_{\text{db}} . 2 \langle N \rangle_{\text{db}} \\
\langle \text{if } L M \rangle_X &= \text{if } \langle L \rangle_X \langle M \rangle_X \\
\langle \langle \ell = M \rangle \rangle_X &= \langle \ell = \langle M \rangle \rangle_X \\
\langle M . \ell \rangle_X &= \langle M \rangle_X . \ell \\
\langle \square \rangle_X &= \square \\
\langle [M] \rangle_X &= \langle [M] \rangle_X \\
\langle M + N \rangle_X &= \langle M \rangle_X + \langle N \rangle_X \\
\langle \text{for } (x^A \leftarrow M) N \rangle_X &= \text{for } (x^{\langle A \rangle} \leftarrow \langle M \rangle_X) \langle N \rangle_X \\
\langle \text{rec } f^{A \rightarrow \text{pl} B} x^A . M \rangle_{\text{pl}} &= \text{rec } f^{\langle A \rightarrow \text{pl} B \rangle} x^{\langle A \rangle} . \langle M \rangle_{\text{pl}} \\
\langle \text{fold } L M N \rangle_{\text{pl}} &= \text{fold } \langle L \rangle_{\text{pl}} \langle M \rangle_{\text{pl}} \langle N \rangle_{\text{pl}} \\
\langle \text{query } M \rangle_{\text{pl}} &= \text{query } \langle M \rangle_{\text{db}} \\
\langle \text{table } t \rangle_{\text{db}} &= \text{table } t
\end{aligned}$$

**Figure 13.** Doubling Translation: terms

ing cases are highlighted in grey boxes. Technically, this translation is defined by induction on the structure of typing derivations. However, the only cases where this matters are those for applications of **any**-functions. To avoid notational clutter, we write the translation in a syntax-directed style and only include type annotations in the cases for application.

The term translation is parameterised by the concrete effect: **pl** or **db**. It is a straightforward structure-preserving traversal, except on lambda-abstractions and applications of **any** functions. Each **any**-function is translated to a pair of a **pl**-function and a **db**-function. To translate an application of an **any**-function, the target effect parameter is used to determine whether to use the first or second element of the pair before applying the corresponding function.

The main correctness properties are as follows:

**THEOREM 5 (Type preservation).** *If  $\Gamma \vdash M : A ! E$  and  $E \leq X$ , then  $\langle \Gamma \rangle \vdash \langle M \rangle_X : \langle A \rangle ! X$ .*

**THEOREM 6 (any-freedom).** *If  $\Gamma \vdash M : A ! E$  and  $E \leq X$ , then judgement  $\langle \Gamma \rangle \vdash \langle M \rangle_X : \langle A \rangle ! X$  is derivable without **any**.*

The type preservation property is straightforward by structural induction on derivations. The **any**-freedom property clearly holds because the **any** effect appears nowhere on the right-hand-side of the definitions of the  $\langle - \rangle_X$  functions. An appropriate semantic correctness property can also be shown, relating the operational semantics of  $\text{Eff}_{\leq}$  programs and their Eff translations.

**THEOREM 7 (Semantics preservation).** *Assume  $\vdash M : A ! \text{pl}$ .*

1.  *$M$  is a value if and only if  $\langle M \rangle_{\text{pl}}$  is a value.*
2. *If  $M \xrightarrow{\mu} N$  then  $\langle M \rangle_{\text{pl}} \xrightarrow{\mu} \langle N \rangle_{\text{pl}}$ .*

3. *If  $\langle M \rangle_{\text{pl}}$  is reducible then there exists  $N$  such that  $\langle M \rangle_{\text{pl}} \xrightarrow{\mu} \langle N \rangle_{\text{pl}}$  and  $M \xrightarrow{\mu} N$ .*

Moreover, if  $\vdash M : A ! \text{db}$  then  $|\langle M \rangle_{\text{db}}| = |\langle M \rangle_{\text{pl}}|$ .

**COROLLARY 1.** *For any  $\vdash M : A ! \text{pl}$  in  $\text{Eff}_{\leq}$  we have  $M \approx \langle M \rangle_{\text{pl}}$ .*

## 4.2 From Eff to Quot

We give an effect-directed translation from Eff to Quot. The interpretations of types, terms and typing environments are parameterised by the concrete effect: **pl** or **db**. The type translation is written  $\llbracket A \rrbracket_X$ , where  $X$  is **pl** or **db** and  $A$  is an Eff-type, and is defined in Figure 14. In the programming language, **db** function types are interpreted as quoted functions. In the database, **pl** function types are interpreted as the unit type, which (as we will show) suffices because **pl** functions can never be called in the database.

One additional complication is how to deal with occurrences of variables bound within a **pl** context, that are also accessed by code within a **db** context. For example, consider the function  $\lambda^{\text{pl}} x . \langle \lambda^{\text{db}} y . (x, y), x \rangle$ : to translate this, we need to be able to convert the value of  $x$  to a quoted term, but Quot only allows lifting at base type. We deal with this by adding a special shadow variable  $x_{\text{db}}$  for each ordinary variable  $x$ , so that the value of  $x_{\text{db}}$  is a quoted version of  $x$ . In this example, the translation is

$$\lambda x . \text{let } x_{\text{db}} = \langle \otimes \downarrow_A(x) \otimes \rangle \text{ in } \langle \otimes \lambda y . \langle (\%x_{\text{db}}), y \rangle \otimes \rangle, x .$$

The special variables  $x_{\text{db}}$  and reification operation  $\downarrow_A(-)$  are explained in greater detail below.

To define the term translation, we need some auxiliary notation. The interpretations for terms and type environments are further parameterised by an effect environment  $\rho$ , which tracks the provenance of bound variables. This is necessary for interpreting a value bound in a programming language context and used in a database context, or vice-versa. The environment  $\rho$  is a finite map from variable names  $x$  to pairs  $A ! X$  denoting the type  $A$  and effect  $X$  at which  $x$  was bound. We write  $\varepsilon$  for the empty effect environment and  $\rho[x \mapsto A ! X]$  for the extension of effect environment  $\rho$  with the mapping  $x$  to  $A ! X$ .

We define a judgement  $\Gamma \vdash \rho$  which states that effect environment  $\rho$  is compatible with type environment  $\Gamma$ .

$$\begin{array}{c}
\text{EMPTYENV} \\
\hline
\vdash \varepsilon
\end{array}
\quad
\begin{array}{c}
\text{EXTENDENV} \\
\Gamma \vdash \rho \\
\hline
\Gamma, x : A \vdash \rho[x \mapsto A ! X]
\end{array}$$

The interpretation  $\llbracket \rho \rrbracket_X$  of an effect environment  $\rho$  at effect  $X$  is defined as:

$$\begin{aligned}
\llbracket \varepsilon \rrbracket_X &= \cdot \\
\llbracket \rho[x \mapsto A ! \text{db}] \rrbracket_{\text{db}} &= \llbracket \rho \rrbracket_{\text{db}}, x : \llbracket A \rrbracket_{\text{db}} \\
\llbracket \rho[x \mapsto A ! \text{pl}] \rrbracket_{\text{pl}} &= \llbracket \rho \rrbracket_{\text{pl}}, x : \llbracket A \rrbracket_{\text{pl}}, x_{\text{db}} : \llbracket A \rrbracket_{\text{db}} \\
\llbracket \rho[x \mapsto A ! Y] \rrbracket_X &= \llbracket \rho \rrbracket_X \quad \text{if } X \neq Y
\end{aligned}$$

which induces the interpretation of a type environment with respect to a compatible effect environment. Notice that the translation  $\llbracket \rho \rrbracket_{\text{pl}}$  introduces the shadow variables  $x_{\text{db}}$ . If  $\Gamma \vdash \rho$ , then we define  $\llbracket \Gamma \rrbracket_X^\rho$  as follows:

$$\llbracket \Gamma \rrbracket_X^\rho = \llbracket \rho \rrbracket_X$$

The term translation is defined as  $\llbracket M \rrbracket_X^\rho$  in Figure 16, where  $M$  is an Eff expression,  $X$  is an effect **db** or **pl** called the *target* of the translation, and  $\rho$  is an effect environment mapping variables to their types and effects (in Eff). The type translation is structure-preserving except at function types, and similarly the term translation is structure-preserving in most cases. The interesting cases (highlighted in grey) are those for function types, lambda-abstractions, or variables where the translation's target effect  $X$  does not match the actual effect. In addition, whenever a variable



$$\begin{aligned}
\llbracket \text{Int} \rrbracket_X &= \text{Int} \\
\llbracket \text{Bool} \rrbracket_X &= \text{Bool} \\
\llbracket \text{String} \rrbracket_X &= \text{String} \\
\llbracket A \rightarrow^X B \rrbracket_X &= \llbracket A \rrbracket_X \rightarrow \llbracket B \rrbracket_X \\
\llbracket A \rightarrow^{\text{db}} B \rrbracket_{\text{pl}} &= \text{Expr} \langle \llbracket A \rightarrow^{\text{db}} B \rrbracket_{\text{db}} \rangle \\
\llbracket A \rightarrow^{\text{pl}} B \rrbracket_{\text{db}} &= \langle \rangle \\
\llbracket \ell : A \rrbracket_X &= \langle \ell : \llbracket A \rrbracket_X \rangle \\
\llbracket [A] \rrbracket_X &= \llbracket [A] \rrbracket_X
\end{aligned}$$

**Figure 14.** Splicing Translation: types

$$\begin{aligned}
x @_{\rho} X &= x && \text{if } \rho(x) = A!X \\
x @_{\rho} \text{pl} &= \text{error} \llbracket A \rrbracket_{\text{pl}} && \text{if } \rho(x) = A! \text{db} \\
x @_{\rho} \text{db} &= \langle \%x_{\text{db}} \rangle && \text{if } \rho(x) = A! \text{pl} \\
\text{error}_A &= (\text{rec } f^{() \rightarrow A} x^{()} . f x) \langle \rangle \\
\downarrow_O(M) &= \langle \% \text{lift } M \rangle \\
\downarrow_{A \rightarrow \text{pl} B}(M) &= \langle \rangle \\
\downarrow_{A \rightarrow \text{db} B}(M) &= \langle \%M \rangle \\
\downarrow_{\langle \ell : A \rangle}(M) &= \langle \ell = \downarrow_A(M.\ell) \rangle \\
\downarrow_{[A]}(M) &= \langle \% \text{fold } M \langle \langle \rangle \rangle (\lambda x.\lambda y. \langle \langle \downarrow_A(x) \rangle \rangle + \langle \%y \rangle \langle \rangle) \rangle
\end{aligned}$$

**Figure 15.** Splicing Translation: auxiliary functions

$x$  is bound in the **pl** context, we bind an additional special variable  $x_{\text{db}}$  whose value is a quoted version of  $x$ ; intuitively, we need this quoted value to translate any occurrences of  $x$  within a **db** context. This affects all of the variable binding cases of  $\llbracket - \rrbracket_{\text{pl}}^{\rho}$  and is explained in more detail below. Variables are coerced to effect  $X$  using the  $x @_{\rho} X$  operation described below. Following the interpretation on types,  $\lambda^{\text{db}}$ -abstractions are coerced to a **pl** target by translating them to quoted lambda-abstractions, and  $\lambda^{\text{pl}}$ -abstractions are coerced to a **db** target by translating them to unit values.

The coercion operation, written  $x @_{\rho} Y$ , is defined in Figure 15. Given variable  $x$  of type  $\llbracket A \rrbracket_X$  (where  $\rho(x) = A!X$ ), it yields a corresponding term in Quot of type  $\llbracket A \rrbracket_Y$ . If  $X = Y$  then coercion leaves  $x$  unchanged. If  $X = \text{db}$  and  $Y = \text{pl}$ , then an error term of type  $\llbracket A \rrbracket_{\text{pl}}$  (implemented as a diverging term) results. This is sound (but not strictly necessary) because in a closed program a **db** variable can only be coerced to **pl** inside the body of a **pl** function bound in a **db** context, and such a function can never be applied. If  $X = \text{pl}$  and  $Y = \text{db}$ , then we splice in the value of  $x_{\text{db}}$ , the special variable bound to the *reified value* of  $x$  as a query term  $\langle \langle \downarrow_A(x) \rangle \rangle$ .

Reification is a type-directed operation that maps a term of type  $\llbracket A \rrbracket_{\text{pl}}$  to a corresponding term of type  $\llbracket A \rrbracket_{\text{db}}$ . A term of base type is reified by lifting and splicing. (Both are needed because lifting coerces a value of base type to its quotation.) A **pl** function is reified as unit. A **db** function is reified as an antiquotation. Records and lists are reified by structural recursion on the type (which is why we chose to include fold in the core calculi).

**THEOREM 8** (Type preservation). *Assume  $\Gamma \vdash \rho$ .*

1. If  $\Gamma \vdash M : A! \text{pl}$ , then  $\llbracket \Gamma \rrbracket_{\text{pl}}^{\rho}, \llbracket \Gamma \rrbracket_{\text{db}}^{\rho} \vdash \llbracket M \rrbracket_{\text{pl}}^{\rho} : \llbracket A \rrbracket_{\text{pl}}$ .
2. If  $\Gamma \vdash M : A! \text{db}$ , then  $\llbracket \Gamma \rrbracket_{\text{pl}}^{\rho}, \llbracket \Gamma \rrbracket_{\text{db}}^{\rho} \vdash \llbracket M \rrbracket_{\text{db}}^{\rho} : \llbracket A \rrbracket_{\text{db}}$ .

Notice that the translation of the database portion of the type environment  $\llbracket \Gamma \rrbracket_{\text{db}}^{\rho}$  appears in the programming language judgement  $\llbracket \Gamma \rrbracket_{\text{pl}}^{\rho}, \llbracket \Gamma \rrbracket_{\text{db}}^{\rho} \vdash \llbracket M \rrbracket_{\text{pl}}^{\rho} : \llbracket A \rrbracket_{\text{pl}}$  in Theorem 8(1). This is sound as any such database variable in  $\llbracket M \rrbracket_{\text{pl}}^{\rho}$  is interpreted as *error*. Database variables can never appear in a closed program.

**THEOREM 9** (Semantics preservation). *Assume  $\vdash M : A! \text{pl}$ .*

$$\begin{aligned}
\llbracket x \rrbracket_X^{\rho} &= x @_{\rho} X \\
\llbracket c \rrbracket_X^{\rho} &= c \\
\llbracket \text{op}(\overline{M}) \rrbracket_X^{\rho} &= \text{op}(\llbracket M \rrbracket_X^{\rho}) \\
\llbracket \lambda^{\text{pl}} x^A . M \rrbracket_{\text{pl}}^{\rho} &= \lambda x \llbracket A \rrbracket_{\text{pl}}^{\rho} . \\
&\quad \text{let } x_{\text{db}} = \langle \langle \downarrow_A(x) \rangle \rangle \text{ in} \\
&\quad \llbracket M \rrbracket_{\text{pl}}^{\rho[x \mapsto A! \text{pl}]} \\
\llbracket \lambda^{\text{db}} x^A . M \rrbracket_{\text{pl}}^{\rho} &= \langle \langle \lambda x \llbracket A \rrbracket_{\text{db}}^{\rho} . \llbracket M \rrbracket_{\text{db}}^{\rho[x \mapsto A! \text{db}]} \rangle \rangle \\
\llbracket \lambda^{\text{pl}} x^A . M \rrbracket_{\text{db}}^{\rho} &= \langle \rangle \\
\llbracket \lambda^{\text{db}} x^A . M \rrbracket_{\text{db}}^{\rho} &= \lambda x \llbracket A \rrbracket_{\text{db}}^{\rho} . \llbracket M \rrbracket_{\text{db}}^{\rho[x \mapsto A! \text{db}]} \\
\llbracket M N \rrbracket_X^{\rho} &= \llbracket M \rrbracket_X^{\rho} \llbracket N \rrbracket_X^{\rho} \\
\llbracket \text{if } L M \rrbracket_X^{\rho} &= \text{if } \llbracket L \rrbracket_X^{\rho} \llbracket M \rrbracket_X^{\rho} \\
\llbracket \langle \ell = M \rangle \rrbracket_X^{\rho} &= \langle \ell = \llbracket M \rrbracket_X^{\rho} \rangle \\
\llbracket M.\ell \rrbracket_X^{\rho} &= \llbracket M \rrbracket_X^{\rho} . \ell \\
\llbracket \square \rrbracket_X^{\rho} &= \square \\
\llbracket [M] \rrbracket_X^{\rho} &= \llbracket [M] \rrbracket_X^{\rho} \\
\llbracket M + N \rrbracket_X^{\rho} &= \llbracket M \rrbracket_X^{\rho} + \llbracket N \rrbracket_X^{\rho} \\
\llbracket \text{for } (x^A \leftarrow M) N \rrbracket_{\text{pl}}^{\rho} &= \text{for } (x \llbracket A \rrbracket_{\text{pl}}^{\rho} \leftarrow \llbracket M \rrbracket_{\text{pl}}^{\rho}) \\
&\quad (\text{let } x_{\text{db}} = \langle \langle \downarrow_A(x) \rangle \rangle \text{ in } \llbracket N \rrbracket_{\text{pl}}^{\rho[x \mapsto A! \text{pl}]} ) \\
\llbracket \text{for } (x^A \leftarrow M) N \rrbracket_{\text{db}}^{\rho} &= \text{for } (x \llbracket A \rrbracket_{\text{db}}^{\rho} \leftarrow \llbracket M \rrbracket_{\text{db}}^{\rho}) \llbracket N \rrbracket_{\text{db}}^{\rho[x \mapsto A! \text{db}]} \\
\llbracket \text{rec } f^{A \rightarrow \text{pl} B} x^A . M \rrbracket_{\text{pl}}^{\rho} &= \text{rec } f \llbracket A \rightarrow \text{pl} B \rrbracket x \llbracket A \rrbracket_{\text{pl}}^{\rho} . \\
&\quad \text{let } f_{\text{db}} = \langle \langle \downarrow_{A \rightarrow \text{pl} B}(f) \rangle \rangle \text{ in} \\
&\quad \text{let } x_{\text{db}} = \langle \langle \downarrow_A(x) \rangle \rangle \text{ in} \\
&\quad \llbracket M \rrbracket_{\text{pl}}^{\rho[f \mapsto (A \rightarrow \text{pl} B)! \text{pl}, x \mapsto A! \text{pl}]} \\
\llbracket \text{rec } f^{A \rightarrow \text{pl} B} x^A . M \rrbracket_{\text{db}}^{\rho} &= \langle \rangle \\
\llbracket \text{fold } L M N \rrbracket_{\text{pl}}^{\rho} &= \text{fold } \llbracket L \rrbracket_{\text{pl}}^{\rho} \llbracket M \rrbracket_{\text{pl}}^{\rho} \llbracket N \rrbracket_{\text{pl}}^{\rho} \\
\llbracket \text{query } M \rrbracket_{\text{pl}}^{\rho} &= \text{query } \llbracket M \rrbracket_{\text{db}}^{\rho} \\
\llbracket \text{table } t \rrbracket_{\text{db}}^{\rho} &= \text{table } t
\end{aligned}$$

**Figure 16.** Splicing Translation: terms

1.  $M$  is a value if and only if  $\llbracket M \rrbracket_X^{\varepsilon}$  is a value.
2. If  $M \xrightarrow{\mu} N$ , then  $\llbracket M \rrbracket_{\text{pl}}^{\varepsilon} \xrightarrow{\mu} \llbracket N \rrbracket_{\text{pl}}^{\varepsilon}$ .
3. If  $\llbracket M \rrbracket_{\text{pl}}^{\varepsilon}$  is reducible then there exist  $N$  and  $\mu$  such that  $\llbracket M \rrbracket_{\text{pl}}^{\varepsilon} \xrightarrow{\mu} \llbracket N \rrbracket_{\text{pl}}^{\varepsilon}$  and  $M \xrightarrow{\mu} N$ .

Moreover, if  $\vdash M : A! \text{db}$  then  $\llbracket \llbracket M \rrbracket_{\text{db}}^{\varepsilon} \rrbracket = \llbracket \llbracket M \rrbracket_{\text{db}}^{\varepsilon} \rrbracket$ .

**COROLLARY 2.** *If  $\vdash M : A! \text{pl}$  in Eff then  $M \approx \llbracket M \rrbracket_{\text{pl}}^{\varepsilon}$ .*

**REMARK 10.** *The translation could be simplified by optimising away or inlining unnecessary let-bindings of  $x_{\text{db}}$  variables. However, these simplifications complicate the correctness proof.*

### 4.3 From Quot to Quot'

In translating Quot to Eff, the translation of an antiquote presents a potential difficulty. It seems natural to translate an antiquoted Quot term  $\langle \%M \rangle$  into a corresponding Eff term with effect **db**. The problem is that  $M$  can perform arbitrary computation including recursion. The solution is to hoist any spliced computation out of the containing quotation. We can always soundly hoist antiquoted computations out of quotations as they never depend on the inner  $\Delta$  environment. Thus as a preprocessing step we replace all antiquoted terms with variables bound outside the scope of the containing quotation. For convenience, we also perform similar hoisting for lift and query expressions. The target language of this step, Quot', is the restriction of Quot such that each antiquotation, query, or lift may only be applied to a variable or value.

The hoisting translation uses the let form (as usual) as syntactic sugar for a lambda application:

$$\text{let } x^A = M \text{ in } N \equiv (\lambda x^A . N) M$$

$$\begin{aligned}
\langle \text{Int} \rangle_X &= \text{Int} \\
\langle \text{Bool} \rangle_X &= \text{Bool} \\
\langle A \rightarrow B \rangle_X &= \langle A \rangle_X \rightarrow^X \langle B \rangle_X \\
\langle \ell : A \rangle_X &= \langle \ell : \langle A \rangle_X \rangle \\
\langle [A] \rangle_X &= [ \langle A \rangle_X ] \\
\langle \text{Expr} \langle A \rangle \rangle_X &= \langle \rangle \rightarrow^{\text{db}} \langle A \rangle_{\text{db}}
\end{aligned}$$

Figure 17. Quot' to Eff: types

$$\begin{aligned}
\langle x \rangle_X &= x \\
\langle c \rangle_X &= c \\
\langle \text{op} \overline{M} \rangle_X &= \text{op}(\langle M \rangle_X) \\
\langle \lambda x^A. M \rangle_X &= \lambda x^{\langle A \rangle_X}. \langle M \rangle_X \\
\langle M N \rangle_X &= \langle M \rangle_X \langle N \rangle_X \\
\langle \text{if } L M \rangle_X &= \text{if} \langle L \rangle_X \langle M \rangle_X \\
\langle \ell = M \rangle_X &= \langle \ell = \langle M \rangle_X \rangle \\
\langle M. \ell \rangle_X &= \langle M \rangle_X. \ell \\
\langle [] \rangle_X &= [] \\
\langle [M] \rangle_X &= [ \langle M \rangle_X ] \\
\langle M ++ N \rangle_X &= \langle M \rangle_X ++ \langle N \rangle_X \\
\langle \text{for } (x^A \leftarrow M) N \rangle_X &= \text{for } (x^{\langle A \rangle_X} \leftarrow \langle M \rangle_X) \langle N \rangle_X \\
\langle \text{lift } M \rangle_{\text{pl}} &= \lambda^{\text{db}} x^{()}. \langle M \rangle_{\text{pl}} \\
\langle \langle \emptyset M \emptyset \rangle \rangle_{\text{pl}} &= \lambda^{\text{db}} x^{()}. \langle M \rangle_{\text{db}} \\
\langle \text{query } M \rangle_{\text{pl}} &= \text{query} (\langle M \rangle_{\text{pl}} \langle \rangle) \\
\langle \text{rec } f^{A \rightarrow B} x^A. M \rangle_{\text{pl}} &= \text{rec } f^{\langle A \rightarrow B \rangle_{\text{pl}}} x^{\langle A \rangle_{\text{pl}}} \langle M \rangle_{\text{pl}} \\
\langle \text{fold } L M N \rangle_{\text{pl}} &= \text{fold} \langle L \rangle_{\text{pl}} \langle M \rangle_{\text{pl}} \langle N \rangle_{\text{pl}} \\
\langle \text{table } t \rangle_{\text{db}} &= \text{table } t \\
\langle \langle \%M \rangle \rangle_{\text{db}} &= \langle M \rangle_{\text{pl}} \langle \rangle
\end{aligned}$$

Figure 18. Quot' to Eff: terms

Hoisting is defined by repeatedly applying the following rules

$$\begin{aligned}
\langle \emptyset Q[\langle \%M \rangle] \emptyset \rangle &\longrightarrow \text{let } x = M \text{ in } \langle \emptyset Q[\langle \%x \rangle] \emptyset \rangle \\
\text{lift } M &\longrightarrow \text{let } x = M \text{ in lift } x \\
\text{query } M &\longrightarrow \text{let } x = M \text{ in query } x
\end{aligned}$$

where  $M$  is not a variable or a value, and  $x$  is a fresh variable. Note that the structure of quotation contexts ensures that the terms hoisted out of a quotation are still evaluated left-to-right with respect to their original positions in the quotation.

We omit (routine but tedious) proofs of type-preservation and semantics-preservation for this transformation.

#### 4.4 From Quot' to Eff

We now give a translation from Quot' to Eff. The type translation is shown in Figure 17. It is structure-preserving except on closed quotation types, which are translated to **db** thunks.

Type environments are translated pointwise:

$$\langle x_1 : A_1, \dots, x_n : A_n \rangle_X = x_1 : \langle A_1 \rangle_X, \dots, x_n : \langle A_n \rangle_X$$

The term translation is shown in Figure 18. It is structure-preserving, except on lambda-abstractions, quotation, queries, anti-quotation and lifting. Lambda-abstractions are annotated with the appropriate effect. Quoted and lifted terms are translated to **db** thunks  $\lambda^{\text{db}} x^{()}. \langle M \rangle_{\text{db}}$  or  $\lambda^{\text{db}} x^{()}. \langle M \rangle_{\text{pl}}$  respectively. For lifting, note that it does not actually matter whether we use  $\langle - \rangle_{\text{pl}}$  or  $\langle - \rangle_{\text{db}}$  on  $M$ , since in a Quot' term,  $M$  will always be either a variable or constant of base type. Queries and antiquotations are translated to force the **db** thunks by applying to unit.

REMARK 11. *The  $\langle - \rangle_X$  translation uses general rules for translating quotation, lifting, and anti-quotation of arbitrary terms, but this is only correct for source expressions in Quot'. The Quot' terms resulting from the hoisting stage in the previous section will only have variables appearing as arguments to quotation, lifting, and anti-quotation. However, this invariant is not preserved by evalua-*

tion, because variables may be replaced by values during evaluation. This is why Quot' allows anti-quotation, lifting, and query operations to be applied to values as well as variables.

THEOREM 12 (Type preservation).

1. If  $\Gamma \vdash M : A$ , then  $(\Gamma)_{\text{pl}} \vdash (\langle M \rangle_{\text{pl}} : \langle A \rangle_{\text{pl}}) ! \text{pl}$ .
2. If  $\Gamma; \Delta \vdash M : A$ , then  $(\Gamma)_{\text{pl}}, (\Delta)_{\text{db}} \vdash (\langle M \rangle_{\text{db}} : \langle A \rangle_{\text{db}}) ! \text{db}$ .

THEOREM 13 (Semantics preservation). Assume  $\vdash M : A$ .

1. If  $M$  is a value then  $\langle M \rangle_{\text{pl}}$  is a value.
2. If  $\Gamma \vdash M : A$  and  $M \xrightarrow{\mu} N$ , then  $\langle M \rangle_{\text{pl}} (\xrightarrow{\mu} \cup \approx) \langle N \rangle_{\text{pl}}$ .
3. If  $\Gamma \vdash M : A$  and  $\langle M \rangle_{\text{pl}}$  is reducible then there exists  $N$  and  $\mu$  such that  $\langle M \rangle_{\text{pl}} (\xrightarrow{\mu} \cup \approx) \langle N \rangle_{\text{pl}}$  and  $M \xrightarrow{\mu} N$ .

Moreover, if  $\cdot \vdash Q : A$  then  $\langle \langle Q \rangle_{\text{db}} \rangle = \langle \langle Q \rangle_{\text{db}} \rangle_{\text{db}}$ .

REMARK 14. *Unlike the previous translations, these translations are not exact simulations of steps in Quot by steps in Eff. The reason is that the splicing rule in Figure 10 corresponds to  $\beta$ -value reduction under a  $\lambda$ -abstraction in Eff. Fortunately,  $\beta$ -value equivalence is valid for Eff modulo  $\approx$ .*

COROLLARY 3. *If  $\Gamma \vdash M : A$  in Quot then  $M \approx \langle M \rangle_{\text{pl}}$ .*

## 5. Application: Links query compilation

Holmes [17] previously developed a compiler for plain Links. Holmes' compiler translates Links programs (which include row polymorphism and other features not present in plain OCaml) to OCaml programs that manipulate tagged values. This typically improves the performance of computationally-intensive Links programs by 1-2 orders of magnitude. However, the compiler does not support Links' query or client-side Web programming features.

The translation from Eff to Quot shows how to compile effect-based Links code to quotation-based code. Moreover, quotation-based queries are relatively easy to translate to plain OCaml simply by translating quotation expressions to an explicit run-time abstract syntax tree representation, and extending the Links run-time library to include query normalisation. Accordingly, we have adapted Holmes' compiler to support queries, by first translating the Links intermediate representation (IR) code to eliminate effect-polymorphic functions via doubling, then inserting appropriate splicing annotations, and finally translating the resulting IR code to OCaml code that explicitly manipulates quoted terms.

The Links IR differs from  $\text{Eff}_{\leq}$  in some important respects: in particular, it employs row typing and effect polymorphism instead of subeffecting. In the current implementation, we handle a subset of Links and some polymorphic code is not handled. However, we believe the basic idea of the doubling translation can be adapted to handle polymorphism instead of subtyping. We view the current query compiler as a proof of concept demonstrating practical implications of the expressiveness results presented earlier; while it is not a mature compiler for Links, our experience with it reported here will help guide further development of such a compiler.

In the rest of this section we present some experiments showing the strengths and weaknesses of the doubling and splicing translations, which we hope will guide future work on compilation for language-integrated query. Figure 19 compares the interpreter (I), compiler without support for queries (C-Q) and query-enabled compiler (C+Q) on several examples.

We first consider programs that do not involve queries, and so can be handled by all three techniques. The `map-any` benchmark (Figure 19(a)) measures the time needed to map a simple `any`-function (which calls some other `any`-functions) over a list. The query-enabled compiler has a small, but measurable overhead compared to the plain compiler. The `sumlist` benchmark (Figure 19(b))

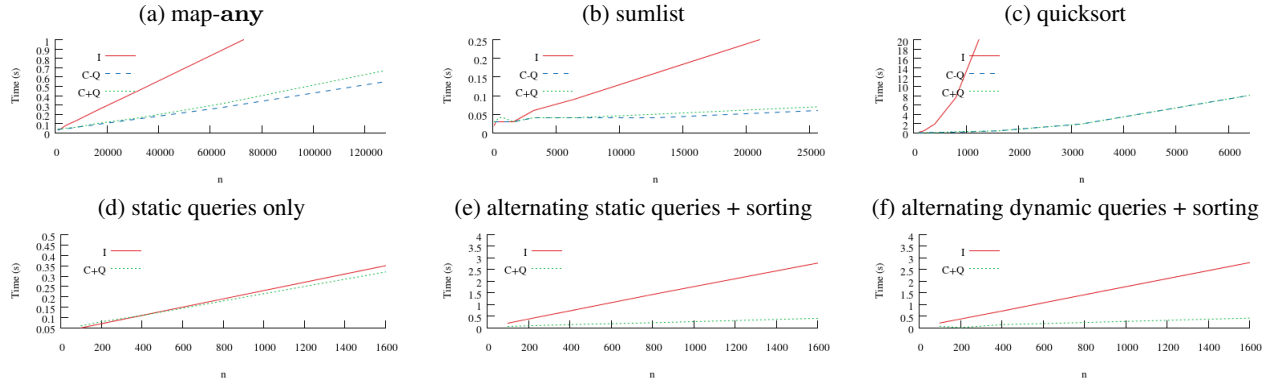


Figure 19. Experimental Results

measures the time to construct a list of the first  $n$  natural numbers, and compute its sum. Here, both compilers provide similar results, though again C+Q is slightly slower. The quicksort benchmark (Figure 19(c)) sorts a decreasing  $n$ -element list, exercising the quadratic worst-case behaviour of quicksort. In this case, both compilers provide similar speedup.

We next consider examples that involve a mix of queries and ordinary execution. These examples can only be run using the interpreter and query-enabled compiler. The first benchmark (Figure 19(d)) simply loops and generates  $n$  queries, with no other computation. There is no measurable difference between the two techniques. This is unsurprising since most of the time is spent in communicating with the database, and compilation does not affect this time, so speedup is bounded by Amdahl’s Law. Next, we consider a variant (Figure 19(e)) where in each iteration the program both issues a query and performs some computation, namely a quicksort of a 10-element list. As we saw before, the compiler yields a significant improvement on the quicksort code. Finally, we consider a similar benchmark (Figure 19(f)), where the query has a higher-order parameter, which takes one value for odd-numbered iterations and another for even-numbered iterations. The results of this benchmark show no appreciable difference from Figure 19(e).

These results show that while the translation from Eff to Quot introduces some overhead to ordinary code compared to the basic compiler, the query-enabled compiler still can realise significant gains for code that mixes queries and ordinary execution.

## 6. Related Work

There is a large and growing literature on approaches to language-integrated query, as well as language-based techniques for combining conventional execution with other execution models, such as MapReduce, GPU, and multicore-based data-parallelism. We discuss only work closely related to this paper; other recent papers [4, 19] compare our approach to language-integrated query with other work in greater depth.

Wadler [33] advocated monads as a technique for structuring programs, including list comprehensions and database queries. This approach was adopted in the nested relational calculus of Buneman et al. [2] and Wong [37] gave rewriting-based normalisation techniques for translating complex nested relational queries over flat data to SQL queries, implemented in the Kleisli system [38]. Links [7, 8, 19] built on this work in several ways, particularly in introducing the ability to compose queries using nonrecursive lambda-abstractions and recursion in the host language.

The LINQ approach adopted in C#, F# and other .NET languages also draws upon monadic comprehensions and nested re-

lational query languages [20, 21], but differs in its implementation strategy: query syntax in C# and F# is desugared to quoted abstract syntax trees which are manipulated and translated to SQL by a library. In F#, it is possible to write dynamic LINQ queries that fail at run time or generate unnecessarily large numbers of SQL queries. Our recent work [4] gave examples, and showed how Links’s normalisation algorithm can be adapted to F# to remedy this problem.

Our formalisation of LINQ-style Quot draws upon a long line of work on quotation and metaprogramming, starting with MetaML [27, 31]. Our approach is closest to the  $\lambda^\square$  calculus of Davies and Pfenning, which provides homogeneous closed quotation (the host and quoted languages coincide); for simplicity, we consider only one level of staging. As discussed elsewhere [4], open quotation can be simulated in Quot using lambda-abstraction, but better support for open quotation and multiple stages, possibly following the approach of Rhiger [27], may also be of interest. Our approach also has some similarities to Eckhardt et al.’s explicitly heterogeneous approach [11]. Reasoning about multi-stage programs is a well-known hard problem. Choi et al. [6] present translations from staged to unstaged programs that employ similar ideas to our translations, particularly hoisting pl-code out of db-code in the  $(-)$  translation. Inoue and Taha [18] present techniques for reasoning about call-by-value multi-stage programs.

Wadler and Thiemann demonstrate a close relationship between effect type systems and monads [34]. Our translation from Eff to Quot has some similarities to that work. However, the languages considered here are quite different from those in Wadler and Thiemann’s work; the latter employ reference types, effects that are sets of regions, and monads indexed by sets of regions, and there is nothing analogous to our doubling translation.

Felleisen [12] and Mitchell [23] presented different notions of expressiveness of programming languages, formulated in terms of different kinds of reductions preserving termination behaviour or observational equivalence. We adopt an ad hoc notion of equivalence based on preservation of query behaviour, which is inspired to some extent by the notion of (weak) bisimilarity familiar from concurrency theory [28]. However, in general nondeterministic labeled transition systems, bisimilarity is strictly stronger than trace equivalence, so it is possible that our translations do not preserve observable behaviour up to bisimulation. We intend to investigate whether our translations are (weak) bisimulations.

## 7. Conclusion

Combining database capabilities with general-purpose programming has been of interest for nearly thirty years [10]. Despite this long history, only within the last ten years have mature techniques

begun to appear in mainstream languages, with the chief example being Microsoft's LINQ, based on explicitly manipulating query code at run time using quotations. Over the same period, techniques developed in the Kleisli and Links languages have built on rigorous foundations of query rewriting to show how to type-safely embed nested relational queries in general-purpose languages, using type-and-effect systems.

In recent work [4], we started to bring these threads together, by showing that some techniques from Links, particularly query normalisation, can be adapted to LINQ in F# in order to improve the expressiveness of the latter. That work raised the question of the relative expressiveness of the two approaches with respect to dynamic query generation: Can Links express dynamic queries that LINQ in principle cannot, or vice versa?

In this paper, we provided a partial answer to this question: we proposed core languages  $\text{Eff}_{\leq}$  and  $\text{Quot}$  similar to those used in previous work on Links and LINQ respectively, and we gave semantics-preserving translations in both directions. This shows, surprisingly in our view, that the two approaches are equivalent in expressive power, at least relative to simple classes of queries: that is, while Links programs or LINQ programs may seem more convenient in different situations, in principle any program written using one approach can also be written using the other. In addition, we used one direction of the translation to extend a Links compiler with partial support for queries, demonstrating the effectiveness of quotation for compiling Links.

A number of areas for future work remain, including extending our translations to handle other query language features such as grouping, aggregation and nested results; extending the compiler to handle full Links including polymorphism; and completely eliminating the overhead of doubling, which we believe should be possible using closure conversion and storing each generated  $\text{db}$  function in a lookup table indexed by the code pointer of the corresponding  $\text{pl}$  function.

## Acknowledgments

This work is supported in part by a Google Research Award (Lindley), EPSRC grants EP/J014591/1 (Lindley) and EP/K034413/1 (Lindley, Wadler), and a Royal Society University Research Fellowship (Cheney).

## References

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] P. Buneman, L. Libkin, D. Suci, V. Tannen, and L. Wong. Comprehension syntax. *SIGMOD Record*, 23, 1994.
- [3] P. Buneman, S. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *Theor. Comput. Sci.*, 149(1), 1995.
- [4] J. Cheney, S. Lindley, and P. Wadler. A practical theory of language-integrated query. In *ICFP*, 2013.
- [5] A. J. Chlipala. Ur: statically-typed metaprogramming with type-level record computation. In *PLDI*, 2010.
- [6] W. Choi, B. Aktemur, K. Yi, and M. Tatsuta. Static analysis of multi-staged programs via unstaging translation. In *POPL*, pages 81–92. ACM, 2011.
- [7] E. Cooper. The script-writer's dream: How to write great SQL in your own language, and be sure it will succeed. In *DBPL*, 2009.
- [8] E. Cooper, S. Lindley, P. Wadler, and J. Yallop. Links: web programming without tiers. In *FMCO*, 2007.
- [9] E. Cooper and P. Wadler. The RPC calculus. In *PPDP*, 2009.
- [10] G. Copeland and D. Maier. Making Smalltalk a database system. *SIGMOD Rec.*, 14(2), 1984.
- [11] J. Eckhardt, R. Kaiabachev, E. Pasalic, K. N. Swadi, and W. Taha. Implicitly heterogeneous multi-stage programming. *New Generation Comput.*, 25(3):305–336, 2007.
- [12] M. Felleisen. On the expressive power of programming languages. *Sci. Comput. Programming*, 17:35–75, 1991.
- [13] G. Giorgidze, T. Grust, T. Schreiber, and J. Weijers. Haskell boards the Ferry - database-supported program execution for Haskell. In *IFL*, number 6647 in LNCS, pages 1–18. Springer-Verlag, 2010.
- [14] T. Goldschmidt, R. Reussner, and J. Winzen. A case study evaluation of maintainability and performance of persistency techniques. In *ICSE*, 2008.
- [15] T. Grust and A. Ulrich. First-class functions for first-order database engines. In *DBPL*, 2013. <http://arxiv.org/abs/1308.0158>.
- [16] F. Henglein and K. F. Larsen. Generic multiset programming with discrimination-based joins and symbolic cartesian products. *Higher-Order and Symbolic Computation*, 23(3):337–370, 2010.
- [17] S. Holmes. Compiling Links server-side code. Bachelor thesis, The University of Edinburgh, 2009.
- [18] J. Inoue and W. Taha. Reasoning about multi-stage programs. In *ESOP*, pages 357–376, 2012.
- [19] S. Lindley and J. Cheney. Row-based effect types for database integration. In *Proceedings of the 8th ACM SIGPLAN workshop on Types in language design and implementation, TLDI '12*, 2012.
- [20] E. Meijer. The world according to LINQ. *Commun. ACM*, 54(10):45–51, Oct. 2011.
- [21] E. Meijer, B. Beckman, and G. M. Bierman. LINQ: reconciling object, relations and XML in the .NET framework. In *SIGMOD*, 2006.
- [22] Microsoft. Query expressions (F# 3.0 documentation), 2013. <http://msdn.microsoft.com/en-us/library/vstudio/5hh225374.aspx>, accessed March 18, 2013.
- [23] J. Mitchell. On abstraction and the expressive power of programming languages. *Sci. Comput. Programming*, 21:141–163, 1993.
- [24] T. Petricek. Building LINQ queries at runtime in C#, 2007. <http://tomasp.net/blog/dynamic-linq-queries.aspx>.
- [25] T. Petricek. Building LINQ queries at runtime in F#, 2007. <http://tomasp.net/blog/dynamic-flinq.aspx>.
- [26] T. Petricek and D. Syme. The F# computation expression zoo. In *PADL*, 2014. To appear.
- [27] M. Rhiger. Staged computation with staged lexical scope. In *ESOP*, number 7211 in LNCS, pages 559–578. Springer-Verlag, 2012.
- [28] D. Sangiorgi. *Introduction to bisimulation and coinduction*. Cambridge University Press, 2012.
- [29] D. Syme. Leveraging .NET meta-programming components from F#: integrated queries and interoperable heterogeneous execution. In *ML*, 2006.
- [30] D. Syme, A. Granicz, and A. Cisternino. *Expert F# 3.0*. Apress, 2012.
- [31] W. Taha and T. Sheard. MetaML and multi-stage programming with explicit annotations. *Theor. Comput. Sci.*, 248(1-2):211–242, 2000.
- [32] J.-P. Talpin and P. Jouvelot. The type and effect discipline. *Inf. and Comput.*, 111(2), 1994.
- [33] P. Wadler. Comprehending monads. *Math. Struct. in Comp. Sci.*, 2(4), 1992.
- [34] P. Wadler and P. Thiemann. The marriage of effects and monads. *Transactions on Computational Logic*, 4(1), 2003.
- [35] G. Wassermann, C. Gould, Z. Su, and P. Devanbu. Static checking of dynamically generated queries in database applications. *ACM Trans. Softw. Eng. Methodol.*, 16, September 2007.
- [36] B. Wiedermann and W. R. Cook. Extracting queries by static analysis of transparent persistence. In *POPL*, 2007.
- [37] L. Wong. Normal forms and conservative extension properties for query languages over collection types. *J. Comput. Syst. Sci.*, 52(3), 1996.
- [38] L. Wong. Kleisli, a functional query system. *J. Funct. Program.*, 10(1), 2000.