
Static Checking of Dynamically Generated Queries in Database Applications¹

Carl Gould

Abstract

Many data-intensive applications dynamically construct and execute queries in response to client requests. Java servlets, for example, can create string representations of SQL queries and then, using JDBC, send the queries to a database server for execution. The servlet programmer enjoys static checking via Java's strong type system. However, the Java type system does little to check for possible errors in the dynamically generated SQL query strings. Thus, a type error in a generated selection query (e.g., comparing a string attribute with an integer) can result in an SQL runtime exception. Currently, such defects must be rooted out through careful testing; worse yet, if they are not, the customer might discover them at runtime. In this paper, a sound, static, program analysis technique to verify the correctness of dynamically generated query strings is presented. The analysis technique will be described, and soundness results for the static analysis algorithm will be provided. The paper also describes the details of a prototype tool based on the algorithm, and includes several illustrative defects found in senior software-engineering student-team projects, online tutorial examples, and a real world purchase order system written by the author.

Introduction

Data-intensive applications often dynamically construct database query strings and execute them. For example, a typical Java servlet web service constructs SQL query strings and dispatches them over a JDBC connector to an SQL-compliant database. In this example scenario, the Java servlet program generates and manipulates SQL queries as string data. Here, I refer to Java as the meta-language used to manipulate object-language programs in SQL. I use a concrete example throughout the paper to explain my analysis technique.

Consider a front-end Java servlet for a grocery store, with an SQL-driven database back-end. The database has a table `INVENTORY`, containing a list of all items in the store. This table includes three columns, among others: `RETAIL`, `WHOLESALE`, and `TYPE`. The `RETAIL` and `WHOLESALE` columns are both of type integer, indicating their respective costs in cents. The `TYPE` column is an integer, representing the product type-codes of the items in the table. In the grocery store database, there is another table, `TYPES`, used to look up type-codes. This table contains the columns `TYPECODE`, `TYPEDESC`, and `NAME`, which, respectively, are of the types integer, varchar (a string), and varchar. The following example code fragment illustrates some common errors that programmers might make when programming Java servlet applications:

```
ResultSet getPerishablePrices(String lowerBound) {
    String query = "SELECT '$' "
        + "(RETAIL/100) FROM INVENTORY "
        + "WHERE ";
    if (lowerBound != null) {
        query += "WHOLESALE > " + lowerBound + " AND ";
    }
    query += "TYPE IN (" + getPerishableTypeCode() + ")";
    return statement.executeQuery(query);
}
String getPerishableTypeCode() {
    return "SELECT TYPECODE, TYPEDESC FROM TYPES "
        + "WHERE NAME = 'fish' OR NAME = 'meat'";
}
```

The method `getPerishablePrices` constructs the string `query` to hold an SQL `SELECT` statement to return the prices of all the perishable items, and executes the query. It uses the string returned by the method `getPerishableTypeCode` as a sub-query. In the code, `+` is the concatenation operator, and the clause `TYPE IN (...)` checks whether the type-code `TYPE` matches any of the type-codes of the perishable items. If `lowerBound` is "595," for example, then the query to be executed is:

```
SELECT '$' || (RETAIL/100) FROM INVENTORY
WHERE WHOLESALE > 595 AND TYPE IN
(SELECT TYPECODE, TYPEDESC FROM TYPES
WHERE NAME = 'fish' OR NAME = 'meat');
```

Several different runtime errors can arise with this example. I list them below. Note that *none of these* would be caught by Java's type system:

Error (1). The expression `'$' || (RETAIL/100)` concatenates the character `'$'` with the result of the numeric expression `RETAIL/100`. While some database systems will implicitly type-cast the numeric result to a string, many do not, and will issue a runtime error.

Error (2). Consider the expression `WHOLESALE > lowerBound`. The variable `lowerBound` is declared as a string, and the `WHOLESALE` column is of type integer. As long as `lowerBound` is indeed a string representing a number, there are no type errors. However, this is risky: nothing (certainly not the Java type system itself) keeps the string variable `lowerBound` from containing non-numeric characters.

Error (3). The string returned by the method `getPerishableTypeCode()` constitutes a sub-

query that selects two columns from the table `TYPES`. Because the `IN` clause of SQL supports only sub-queries returning a single column (in this context), a runtime error would arise. This can happen if the method `getPerishableType-Code()` did return a single column before, but was inadvertently changed to return two columns.

This specific combination of Java as the meta-language and SQL as the object-language is widely used today. The databases receiving these constructed SQL queries certainly perform syntax and semantic checking of the queries. But because these queries are dynamically generated, errors are only discovered at runtime. It would be desirable to catch these errors statically in the source code.

In this paper, I present a static analysis to flag potential errors or guarantee their absence in dynamically generated SQL queries. My approach is based on a combination of automata-theoretic techniques [9], and a variant of the context-free language (CFL) reachability problem [14, 15].

As a first step, my analysis builds upon a static string analysis to build a *conservative* representation of the generated query strings as a finite-state automaton. Then, I statically check the finite-state automaton with a modified version of the context-free language reachability algorithm. My analysis is sound in the sense that if it does not find any errors, then such errors do not occur at runtime. I have implemented the analysis and tested the tool on realistic programs using JDBC, including senior software-engineering student-team projects, online tutorial examples, and a real-world purchase order system. The tool is able to detect some known and some unknown errors in these programs. Although it has not been tuned for performance, it runs its operations in only a few minutes on all test programs. Furthermore, the analysis empirically appears quite precise, with a low false-positive error rate.


The rest of the paper begins with background on the string analysis and context-free language reachability and a brief overview of my analysis. Then I present the analysis in more detail and discuss my experimental setup and results. Finally, I survey related work and conclude with a discussion of possible future work.

Background and Analysis Overview

STATIC STRING ANALYSIS OF JAVA PROGRAMS

The analysis makes use of the string analysis of Java programs reported in [6]. Essentially, it is an interprocedural data-flow analysis [10, 12] to approximate the semantics of a program's string manipulation expressions. The analysis is similar to a pointer analysis [3] for imperative languages or a control-flow analysis (0-CFA) [17] for functional languages. It approximates the set of possible strings that the program might generate for a particular string variable at a particular program location of interest; these locations are called hotspots. The string analysis produces a finite state automaton (FSA) that conservatively approximates the set of possible strings for each hotspot specified; that is, for each hotspot, the automaton accepts a larger set of strings than is actually produced by the program. In my earlier example, the statement `return statement.executeQuery(query);` is a hotspot for that program.

The string analysis works on Java bytecode. It starts by finding the hotspots in the Java program. I simply mark every location with a call to the method `executeQuery` (such as `return statement.executeQuery(query)` in my example) as a hotspot. Then the analysis abstracts away the control flow of the program, and creates a *flow graph* representing the possible string expressions. The graph captures the flow of strings and string operations in a program; all else is abstracted away. The nodes in a flow graph correspond to variables or expressions in the program, and the edges represent directed *def-use relationships* for the possible data-flow. For example, for the statement `query = str1 + str2;` the following graph nodes and edges are created:



In the graph, the node labeled *concat* represents the concatenation expression `str1 + str2`, with the edges labeled 1 and 2 being the corresponding first and second arguments. The edge between nodes labeled *concat.* and *query* indicates the assignment. The other expressions and operators are treated similarly; full details can be found in [6]. Next, this flow graph is reduced to an extended context-free grammar² by treating the nodes of the flow graph as terminals and non-terminals of the grammar. For example, the flow graph given earlier in this section yields the following grammar rules:


$$\begin{aligned} C &::= S1 S2 \\ Q &::= C \end{aligned}$$

where *S1*, *S2*, *C*, and *Q* correspond to the respective nodes for `str1`, `str2`, *concat*, and *query*. In general, the grammar generated from a flow graph is not regular (not even context-free as mentioned earlier). To make further analysis computationally tractable, I widen this grammar to a regular language. The widening step allows syntax checking of the generated strings against a grammar.³ Since the (fairly technical) details of this step are not the main focus of this work, I omit them here and refer interested readers to [6].


CONTEXTFREE LANGUAGE REACHABILITY

In the next step, the FSA is processed by a context-free language (CFL) reachability algorithm, which forms the foundation of my analysis. I give a brief description of the problem and the algorithm here (cf. [14,15] for further details). The CFL-reachability problem takes as inputs a context-free grammar *G* with terminals *T* and non-terminals *N*, and a directed graph *D* with edges labeled with symbols from $T \cup N$. Let *S* be the start symbol of *G*, and $\Sigma = T \cup N$. A path in the graph is called an *S*-path if its word is derived from the start symbol *S*. The CFL-reachability problem is to find all pairs of vertices *s* and *t* such that there is an *S*-path between *s* and *t*.

The algorithm to solve the CFL-reachability problem uses dynamic programming, and also relates to dynamic transitive closure [23],⁴ which underlies many standard program analysis algorithms, such as type systems based on subtyping, alias analysis, and control-flow analysis [2,3,17]. The algorithm first normalizes the grammar *G* such that each production's right-hand side contains no more than two symbols. This is easily done by introducing new non-terminal symbols. Then new derived edges are added to *D* based on the productions of *G*. For example, suppose *G* has the production $X ::= A B$, and *A* contains the following edges:



The algorithm adds a dotted edge:



The algorithm repeatedly applies the above transformation to the graph *D* until no more new edges can be added. Any pair of nodes *s* and *t* with an edge labeled *X* in the final graph has an *X*-path from *s* to *t* in

my transformation of the automaton, I have the FSA shown in Figure 2b.⁵ To achieve this, I use a depth-first traversal of the original automaton, which groups letters into tokens (in the same sense as those in the lexical analysis phase of a compiler [1]). I then use these tokens to create an equivalent FSA with transitions over the keywords, literals, and delimiters of my object-language. In addition, white-spaces are removed from the automaton in this step.

RECONSTRUCTION OF TYPE ENVIRONMENTS

For an SQL query, the declared types of various columns are given in a database schema. This is similar to the notion of a type environment in standard type systems for language such as C, Java, and ML to look up types of variables. To illustrate, consider the sample SQL query `SELECT NAME FROM EMPLOYEE WHERE SALARY > 20,000`. The information that `NAME` is of type `varchar`, and `SALARY` is of type `integer` is not explicit in the above query expression, but is stored separately in the schema of the table `EMPLOYEE`. For the same reason, my generated FSA does not have this information either. I need to reconstruct it from the database schemas. I now describe how I use the CFL-reachability analysis to obtain the column-name to type mapping from the schema.

In this step, I assume that the generated FSA is syntactically correct, i.e., that the query strings produced by the FSA are all of valid SQL syntax. This assumption is reinforced by the string analysis [6], because it performs syntax-checking of the generated automaton.




Figure 3. An FSA with two table contexts.

appear in any of several different contexts. For example, suppose I have the automaton shown in Figure 3. The type of the column `NAME` can be different depending on which of the two paths in the automaton is taken. In one path, its type is determined by the schema's definition of `TABLE1`; in the other, it is determined by that of `TABLE2`. My solution to this problem is based on a variant of the CFL-reachability algorithm. I apply the algorithm with the context-free grammar for SQL queries and with the transformed automaton as inputs. In essence, I use the CFL reachability algorithm to parse the automaton. This is very similar to general context-free parsing (which is also of cubic time complexity). However, instead of parsing a particular query, I work with an automaton which produces a potentially infinite number of query strings.

Nonterminal	Productions
<i>select_stmt</i>	::= <i>select ;</i>
<i>select</i>	::= <i>select_part1 select_part2</i>
<i>select_part1</i>	::= <i>SELECT column_list</i>
<i>select_part2</i>	::= <i>FROM table_list</i>
	<i>FROM table_list where_clause</i>
<i>where_clause</i>	::= <i>WHERE condition</i>
<i>condition</i>	::= <i>logical_term</i>
	<i>NOT logical_term</i>
	<i>condition OR logical_term</i>
	<i>exp_simple IN subquery</i>
<i>... Additional rules elided</i>	

Table 1. SQL `SELECT` statement grammar.

In Table 1, I show an excerpt of the grammar I use for SQL's `SELECT` statement [7]. Non-terminals are in italics, and terminals are in capital letters. My grammar is not yet complete, but could be easily made so by adding more rules. The CFL-reachability algorithm described in [14] requires a normalized grammar such that the right-hand side of any production has no more than two symbols.

My implementation of CFL-reachability has been extended so that it works with productions with no more than three symbols on their right-hand sides. This extension allows me to use more naturally written grammars.

The modified CFL-reachability algorithm enables me to find the type context (i.e., type environment) of each path through the automaton. I can then use this information to match every column with all of its possible types. The type contexts are discovered by annotating the automaton with the derivation in use while running the CFL-reachability algorithm. In particular, whenever the CFL-reachability process adds a non-terminal edge to the automaton, I store references in that edge to the edges making up this derivation. This is similar to actions in syntax-directed translation [1]; I build a collection of parse trees for the automaton. In fact, the step in which I reconstruct a complete type environment is similar to attribute grammars in syntax directed translation [1].

To illustrate this process, I will run through a few steps using an example. First, let me return to the simple example in Figure 3. Here are a few steps (shown in Figure 4) of running my algorithm for discovering the type environments:

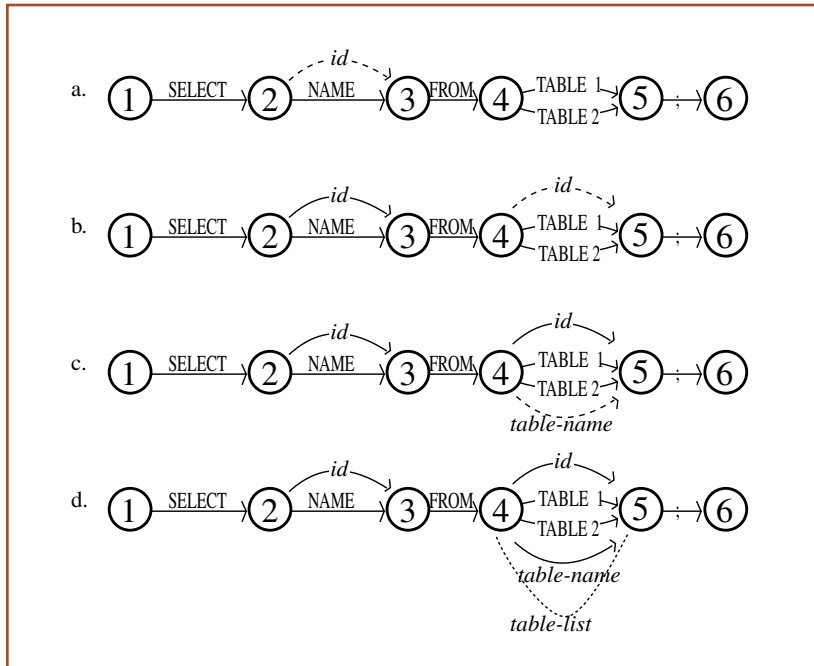


Figure 4. Sample steps in discovering type environments.

However, the edge is already present in the automaton. Instead of simply stopping (as is done in the standard CFL reachability algorithm), my algorithm adds a second derivation reference to the already-present *id* edge. This allows each context to be discovered when searching through the derivation edges. Notice that I can handle loops in the automaton by exploiting its simple looping structure. Due to space limitations, I omit the details here.

Figure 4c. After these *id* edges are added, an edge labeled *table_name* is added from node 4 to node 5, which has a single derivation—the *id* edge from node 4 to node 5.

Figure 4d. Then, an edge labeled *table_list* will be added from node 4 to node 5. Next, an edge labeled *select_part2* will be added from node 3 to node 5. This process continues, and eventually an edge labeled *select_stmt* is added from node 1 to node 6 (the final graph is not shown due to space limitations).

Now let me go back to my earlier example. After running my type environment reconstruction algorithm on this example, there is an edge from the start state of the automaton to its final state, labeled *select_stmt*. The edge has a single derivation with two parts: a *select* edge and an edge labeled with the

Figure 4a. Because **NAME** is not a keyword, it must be an identifier, so an edge labeled *id* is added between nodes 2 and 3.

Figure 4b. The same goes for the two edges between nodes 4 and 5. However, the *id* edge from node 4 to 5 has two distinct derivations. This is different from the standard CFL-reachability algorithm. With the standard algorithm, if an edge to be added is already present, then nothing needs to be done for that edge. In my example, suppose there is already an *id* edge from 4 to 5, which was added through the edge labeled **TABLE 1**. When the edge labeled **TABLE 2** is to be processed, another *id* edge from 4 to 5 must be added.

Nonterminal	Productions
<i>integer</i>	$::=$ <i>integer</i> + <i>integer</i> <i>integer</i> - <i>integer</i> <i>integer</i> * <i>integer</i> (<i>integer</i>) ABS <i>integer</i>
<i>decimal</i>	$::=$ (<i>decimal</i>) <i>integer</i> / <i>integer</i>
<i>varchar</i>	$::=$ (<i>varchar</i>) UPPER <i>varchar</i> <i>varchar</i> <i>varchar</i>
<i>boolean</i>	$::=$ <i>integer compare_op integer</i> <i>integer</i> IN (<i>integer</i>)

Table 2. SQL SELECT statement type-system grammar

a column name, and must be a literal. At this point, I determine the literal's type and add it to the graph as an edge.

In type environment reconstruction, some errors can be discovered. For instance, I can determine whether there is an invalid column that does not exist in any of the applicable tables for that column. If this happens, an error can be reported as either an improperly-quoted literal or a non-existent column. Other errors are also detected in this step, including duplicate table references (the same table appears more than once in the **FROM** clause), duplicate uses of the same table alias (two tables are assigned the same alias), and non-existent tables (the schema does not have a table referenced in the **FROM** clause).

TYPECHECKING

In the final step of the analysis, I perform type-checking on the automaton produced in the previous step, as described in the previous section. At this stage, the automaton has been annotated to show the types of column names and literals. SQL's simple type system lets me treat the type system as a context-free grammar. For example, the type rule for additions over integers looks like this:

The above rule can be viewed as equivalent to the grammar rule: $integer ::= integer + integer$, which states that an integer plus an integer is again an integer. The other rules for type-checking SQL expressions can be handled in a similar manner. This is possible due to SQL's simple type language—a collection of atomic types. This is in contrast to general purpose programming languages that have more complicated type structures. Table 2 shows a small subset of the context-free grammar for the type-checking rules of SQL's **SELECT** statement. All the rules are straightforward, except the last one, which merits some explanation. The rule says that the conditional expression **IN** is well-typed if the sub-query inside the parentheses (here required to be of type *integer*) reduces to a single column of type *integer*, when the outer expression is an integer. Thus, this rule requires the sub-query to be of the correct type. A similar rule, not shown here, exists for the *varchar* type. Notice that I did not specify all the rules—for example the rules for the **SELECT** statement. These rules are obvious, and I omit them in this paper due to space limitations.

As in a standard type system, if none of the rules applies for a language construct, then a type error is discovered.

delimiter ';'. The select edge also has a single derivation: a link to a *select_part1* edge and a link to a *select_part2* edge. These links form different parse trees for the automaton. By a top-down traversal of parse trees, I can determine which tables apply to which column lists, and with the schema for these tables, I can determine the possible types of each column. Then, I add type edges to the graph, replacing column names with their corresponding types. Notice that, because of the possibility of multiple contexts, a column may have more than one type.

In my example, I know that **RETAIL** is a column in the table **INVENTORY**, so I add an edge labeled *integer* to the graph. When I come across a primitive such as the edge labeled **100**, I determine that it is not

I apply the CFL-reachability algorithm using the grammar in Table 2 to propagate type information. If during the process, there is an expression that does not match any one of the right-hand sides of the rules, then an error is discovered. In some sense, there are implicit error rules in the grammar, such as $errortype ::= integer + varchar$.

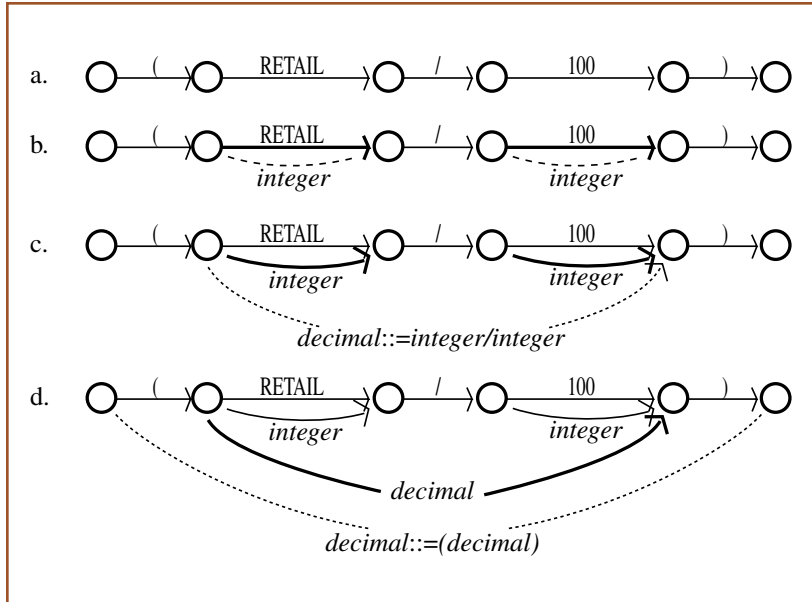


Figure 5. Sample steps of running CFL reachability.

object-programs specified by the automaton are type-correct. On the other hand, if type-checking does produce an *errortype* edge, the analysis reports potential errors and displays a sample derivation that causes the type error. Note, however, that, due to imprecision in the automaton characterization of the object-programs, a reported error may not be an actual error in the original Java program. In the case of SQL, my analysis is precise under the assumption that all the object-programs specified by the automaton are feasible in the original Java source program.




Figure 6. Discovering a type error.

I illustrate type-checking with an example. Consider the small snippet shown in Figure 5a, which is taken from my working example. Before type-checking begins, the automaton is annotated with type information for the column names and literals, as shown in Figure 5b. Figures 5c-d show a few steps of type propagation using my grammar rules.

If my type-checking step does not produce any edges labeled *errortype*, then all the object-programs specified by the automaton are type-correct. On the other hand, if type-checking does produce an *errortype* edge, the analysis reports potential errors and displays a sample derivation that causes the type error. Note, however, that, due to imprecision in the automaton characterization of the object-programs, a reported error may not be an actual error in the original Java program. In the case of SQL, my analysis is precise under the assumption that all the object-programs specified by the automaton are feasible in the original Java source program.

Figure 6 shows the edge *errortype* being added to a snippet of my example program, corresponding to the concatenation error between the character '\$' and the numeric result of the division. Note that many irrelevant edges are omitted in the figure. The two other errors present in my example can also be discovered in a similar manner.

Figure 6 shows the edge *errortype* being added to a snippet of my example program, corresponding to the concatenation error between the character '\$' and the numeric result of the division. Note that many irrelevant edges are omitted in the figure. The two other errors present in my example can also be discovered in a similar manner.

CORRECTNESS OF THE ANALYSIS

Due to conservative approximations, my analysis may report a spurious (infeasible) error. In the next section, I present experimental data to support the claim that the analysis is rather precise and has low false-positive rates.

Theorem 3.1 (Soundness). *My analysis is sound. In other words, if the analysis does not report any errors, then the generated object programs are type-safe.*

Proof. [Sketch] I give a brief justification of the soundness theorem. I assume that the automaton that I operate on is a conservative approximation of the set of possible SQL query strings for a particular hotspot. This is guaranteed by the correctness of the string analysis in [6]. I also make the assumption that the query strings produced by the automaton are syntactically correct,⁶ which is crucial for the soundness of my analysis. Next, I can show that CFL-reachability considers all possible derivations of the query strings because the query strings are all of the correct syntax. Thus, the type-environment reconstruction step is correct, meaning that a column name is considered in all possible tables. For example, if I annotate an edge labeled by a column with the type integer, there must be a path in the automaton containing that edge such that the particular column is found to be of type integer. At the end of this step, all columns are labeled with a superset of actual types. Finally, if the resulting automaton contains a path with a type error, it will be detected because CFL-reachability considers all possible edge combinations.

Experimental Evaluation

I have tested a prototype tool that embodies my approach in its ability to detect programming errors in Java/JDBC applications. As any SQL developer will attest, every database vendor implements a different version of SQL; thus, checkers such as this one require some porting effort for each different database. I have implemented my analysis for the **SELECT** statement specified by the grammar for Oracle version 7 [7]. This grammar is a subset of what is specified in the SQL-92 standard. Adding support for other statements or different vendors is not difficult, because I have separated the type environment reconstruction and type-checking steps. In most cases, I would simply need to modify my syntax grammar and/or type-system rules (specified as input files in my analysis) and the recursive-descent code to traverse the parse trees and map column names with their possible types. In order to have a sound analysis, I have built a strict semantics into my tool: if a program is deemed type-safe by my analysis, it should be type-safe on any database system.

Because the semantics of many database systems is not as strict as the one enforced by my tool, the tool may report an error that some database systems consider legitimate. My tool is implemented in Java and uses the string analysis in [6] for computing the FSA, which in turn uses the Soot framework [20] to parse class files and compute interprocedural control-flow graphs. I have tested my tool on various test programs, including student team projects from an undergraduate software engineering class, sample code from online tutorials found on the web, and code from other projects made available to me. Table 3 lists the test programs and summarizes my results. For each test program, I list the Java source code size (number of lines of source code), number of hotspots in the program, number of columns in the database schema, generated automaton size (number of edges and nodes), analysis time (split into automaton generation and semantic analysis), and the numbers of various warnings and errors found (Table 4). Note

Test Programs (S) Student (W) Web Download (I) Industrial	Java Program (Lines)	Size			Analysis Time (sec)		Errors Found			
		Hotspots	Schema (Columns)	Total Automaton Edges/Nodes	Automaton Generation	Semantic Analysis	Warnings	Total Errors	Confirmed Errors	False Errors
Smi (S)	1559	1	19	35 / 27	1.5	7.1	0	0	0	0
CFWorkshop (S)	36	5	13	47 / 52	0.6	1.3	0	0	0	0
TicTacToe (S)	2888	2	26	134 / 121	6.4	144.8	3	1	1	0
WebBureau (W)	50	10	21	152 / 162	0.5	2.5	0	1	1	0
Checkers (S)	6615	4	36	181 / 138	11.8	97.8	0	15	15	0
JuegoParadis (S)	6135	13	29	259 / 206	27.0	45.0	0	9	0	9
Reservations (S)	2385	22	54	368 / 383	1.7	29.1	0	0	0	0
OfficeTalk (S)	5812	29	14	655 / 525	7.0	120.8	0	2	2	0
PurchaseOrders (I)	642	51	82	1324 / 1373	1.3	173.3	41	10	9	1

Table 3. Experimental Results

Error Kind	Description	test programs
Type Mismatch	Concatenation of fields with wrong types.	PurchaseOrders
Type Mismatch	Possibly unquoted string compared with a varchar column.	PurchaseOrders
Type Mismatch	Quoting a numerical value and treating it numerical.	Checkers, Office-, Talk, TicTacToe
Semantic Error	Ambiguous column selection.	WebBureau
Semantic Error	Column not found.	Checkers, PurchaseOrders
Spurious Error	Column not found (due to imprecision of the string analysis).	PurchaseOrders, JuegoParadis
Warning	Comparing a numerical value with a possibly non-numerical value.	PurchaseOrders, TicTacToe

Table 4. Breakdown of errors and warnings.

minutes. I expect my analysis to scale to large systems, because analyses based on the same underlying algorithms have been shown to scale to millions of source lines of C [8, 18]. Further experiments are needed to verify this claim.

Table 4 shows a breakdown of the kinds of errors found in the test programs, which I now explain in more detail:

Concatenation of fields with wrong types. This is the same error as in the concatenation ‘\$’ || (RETAIL/100). After discovering this error in porting a program to a different (more strict) database, my tool has been used to find all instances of the error in the “PurchaseOrders” program.

Possibly unquoted string. Assume we have a comparison such as `NAME = α` , where α represents an unknown string. If there are no quotes in a string that α possibly represents, then such an error occurs.

Quoting a numerical value. This error happens when a numerical value is quoted but still treated as a numerical literal. This is a common error in student projects using MySQL, which permits numerical literals to be quoted. Many other database systems also consider this an error, because quoted numerical literals are of type `varchar`.

Ambiguous column selection. My tool detected such an error in some sample code from a tutorial website (<http://web-bureau.com/modules/sql.php>). This error is quite subtle, and it appears to be unknown. The particular statement is:

```
SELECT customer id FROM customers c, orders o
WHERE c.customer id = o.customer id;
```

The error is that the database does not know from which table it should choose the `customer_id`. Certainly, in this particular statement, it seems not to matter which `customer_id` it selects, but in general, the semantics of the column list should not depend on the outcome of the `WHERE` clause.

Column not found. This error happens when a column name does not exist in any of the tables in the `FROM` clause. I found two distinct causes of this error—one a real error and the other a spurious error:

Real error. The schema of the database does not include this column. This can be caused by either selecting a nonexistent column or missing the quotes around a literal, which as a result is treated as a column.

that the test programs are sorted by automaton size, since for my analysis it is a good measure of the complexity of the programs.

All experiments were done on a machine with a 2GHz Intel Xeon processor and 1 GB RAM, running Linux kernel 2.4.20. The results indicate that my analysis is rather precise, i.e., that it has a low false-positive rate. Because my analysis is sound, if the tool does not report any error on a program, then I have verified that the program is type-correct. In addition, although I have not tuned the performance of my implementation, the analysis is still quite efficient; it was able to analyze each of my test programs within a matter of

Spurious error. This error is due to imprecision in the string analysis. Consider the following example, where `makeQuery` is a public method taking a string parameter `tables` to construct a `FROM` clause:

```
public String makeQuery(String tables) f
    return "SELECT name FROM " + tables;g
```

The string analysis adds an α edge to the table list of the `FROM` clause (meaning any table is possible), in addition to the concrete table list it finds through analyzing other input classes. The reason is that this method is public, and the string analysis expects the possibility of other calls to the method and thus views the input classes as being incomplete. The presence of this α edge causes the analysis to search for the column in an empty table list, which certainly fails. This is the only kind of spurious error I found in the test programs. These errors can easily be filtered out by modifying the string analysis to consider its input as a complete set of classes.

Warning. I found one type of warning in the test programs. It is the same as the one illustrated in the earlier running example with which I began, in which, at runtime, a numerical column is compared with what is *possibly* a non-numerical value.

I have shown that my tool can detect errors in non-trivial programs. To further evaluate the tool and to provide a programming aid to students, I plan to experiment by using the tool in an undergraduate software engineering class.

Related Work

Perhaps the most closely related work is the string analysis of Christensen, Møller, and Schwartzbach [6], which forms the basis of my analysis. Their string analysis ensures that the generated object-programs are syntactically correct. However, it does not provide any guarantee of semantic correctness for the object-programs. Many domain-specific languages have been proposed for ensuring correctness of dynamically generated web documents. Most of these are language extensions enhanced with tree manipulation capabilities, instead of the string manipulations that I deal with in this work. In addition, they usually guarantee only that there will be *syntactic* correctness, not *semantic* correctness. There are two research efforts in static validation of dynamically generated web documents such as HTML that are worth noting here. In [16], the authors propose a typed, higher-order template language that provides safety of the dynamically generated web documents within the `<bigwig>` project [5], an extension to Java for high-level web service development. Their type system is based on standard data-flow analysis techniques [10, 12]. Another work along the same lines is the work by Braband et al [4] to statically validate dynamically generated HTML documents against the official DTD for XHTML. Their work also is based on a data-flow analysis that computes a summary of all possible documents at a particular point in the program. The summary graph is then validated against the official DTD for XHTML. This work is also done in the context of the `<bigwig>` language. In [11], a test adequacy criterion for data-intensive applications is presented. My approach is complementary; static analysis can save testing time, but testing can discover logical defects not related to SQL query construction.

In a broader context, my research can be viewed as an instance of providing static safety guarantee for meta-programming [19]. Macros were the earliest meta-programming technique in which the issue of correctness of generated code first arose. Using powerful macro programming languages, macro programmers clearly needed to worry about the correctness of the generated code. How could such macro meta-programs be statically checked for correctness? The widely used `cpp` macro pre-processor does little checking, and allows one to write arbitrary macros without regard to correctness. The program-

mable syntax macros of Weise & Crew [22], which work at the level of correct abstract-syntax tree (AST) fragments, guarantee that generated code is syntactically correct with respect (specifically) to the C language. Static type-checking is used to guarantee that AST fragments (e.g., Expressions, Statements, etc.) are assembled correctly by macro meta-programs. Another issue is scoping of generated names: macro expansion should not “capture” variable names in an unexpected manner. Hygienic macro expansion algorithms, beginning with Kohlbecker et al [13], provide these guarantees.

More recent work seeks to extend the guarantees for syntax and scoping to the semantics of the generated code. The work of Taha & Sheard [19] and others is concerned with guaranteeing (in a functional programming setting) that generated code is type-safe. I do not introduce a new macro language, like [22], nor work in a uniform functional setting, like [19], with functional languages both at meta- and target-levels. My goal is simply to ensure that strings passed into a database from an arbitrary Java program are type-safe SQL queries from the perspective of a given database schema. I expect that the general technique outlined in this paper can be extended to apply in other settings as well.

Conclusions and Future Work

I have presented a sound, static analysis technique for verifying the correctness of dynamically generated SQL query strings in database applications. My technique is based on applications of a string analysis for Java programs and a variant of the context-free language reachability algorithm. I have implemented this technique and have performed extensive testing of the tool on realistic programs. The tool has detected known and unknown errors in these programs, and its low false-positive rates on my test programs show that it is rather precise.

Future work could go in several promising directions. To increase the usability of my tool for debugging, it would be interesting to map an error path that I find in the automaton to the original Java source. One possible approach is to carry line numbers of the flow-graph nodes in the source code to the automaton, so that a path in the automaton can be associated with a set of source lines in the original Java program. A related problem is to check the correct uses of the query results in the source program. Consider the following example:

```
query = ``SELECT NAME, SALARY FROM EMPLOYEE``;
rset = statement.executeQuery(query);
...
salary = rset.getInt(1); // should be getInt(2)
name = rset.getString(2); // should be getString(1)
```

The result set `rset` is a table of pairs of type `String` (the `NAME` column) and `int` (the `SALARY` column). The statement `salary = rset.getInt(1)` attempts to read the first field of the pair, a string, and treat it as an integer. The last statement has a similar error. By mapping my analysis results back to the original program, I should be able to detect this class of errors.

The class of embedded SQL command injection problems [21] in web and database applications is also interesting to consider. The problem is that, without proper input validation, an attacker can supply arbitrary code to be executed by a web or database server, an extremely dangerous situation. I plan to extend my approach to deal with this class of errors. Finally, thus far I have considered only SQL. I believe my technique can be generalized and plan to investigate how to extend it to analyze dynamically generated programs in other languages.

Acknowledgments

I wish to thank my mentors and collaborators, Dr. Zhendong Su and Dr. Premkumar Devanbu, for their support of this research. I also wish to thank Anders Møller and Aske Simon Christensen for useful discussions about this research, as well as several students in the Software Engineering class at UC Davis for providing me with the source codes of their projects so that I could evaluate my tool. Finally, I thank the anonymous reviewers of the International Conference on Software Engineering 2004 for their valuable comments.

Notes

1. A Startup Fund from the University of California, Davis, awarded to Zhendong Su, supported this research. The research of a third collaborator, Premkumar Devanbu, was supported by NSF (both CISE & ITR programs).
2. This extension handles operators or functions, if any.
3. By contrast, the grammar is narrowed to a finite state automaton in [6] for syntax checking. This is done because, in general, checking the containment of a regular language by a context-free language is undecidable [9]. However, it is possible to work directly with the grammar from the flow graph or with one widened to a context-free grammar because the intersection of a context-free grammar with a regular expression is still context-free, and furthermore, it is easy to check the emptiness of a context-free grammar.
4. The problem is to maintain transitive closure of a graph while new basic graph edges can be added during graph closure.
5. In the figure, α denotes an unknown string.
6. This can be enforced by the string analysis in [6].

References

- [1] A. Aho, R. Sethi, and J. Ullman. *Compilers, Principles, Techniques and Tools*. Addison-Wesley, 1986.
- [2] R. Amadio and L. Cardelli. Subtyping recursive types. In *Proceedings of the 18th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 1991; 104-118.
- [3] L. Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU, University of Copenhagen, May 1994. DIKU report 94/19.
- [4] C. Braband, A. Møller, and M. Schwartzbach. Static validation of dynamically generated HTML. In *ACM SIGPLAN. SIGSOFT Workshop on Program Analysis for Software Tools and Engineering*, June 2001; 38-45.
- [5] C. Brabrand, A. Møller, and M. Schwartzbach. The <bigwig> project. *ACM Transactions on Internet Technology*, 2002; 2(2):79-114.
- [6] A. Christensen, A. Møller, and M. Schwartzbach. Precise analysis of string expressions. In *Proceedings of the 10th International Static Analysis Symposium*, 2003; 1-18.
- [7] J. Guyot. BNF index of SQL for Oracle 7. Available at <http://cui.unige.ch/db-research/Enseignement/analyseinfo/SQL7/>.
- [8] N. Heintze and O. Tardieu. Ultra-fast aliasing analysis using CLA: A million lines of C code in a second. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2001; 254-263.

- [9] J. Hopcroft and J. Ullman. *Introduction to Automata Theory, Language, and Computation*. Addison-Wesley, Reading, MA, 1979.
- [10] J. Kam and J. Ullman. Global data flow analysis and iterative algorithms. *Journal of the ACM*, Jan. 1976; 23(1): 158-171.
- [11] G. M. Kapfhammer and M. L. Soffa. A family of test adequacy criteria for database-driven applications. In *Proceedings of the 9th European Software Engineering Conference and the 11th ACM SIGSOFT Symposium on Foundations of Software Engineering*, 2003; 98-107.
- [12] G. Kildall. A unified approach to global program optimization. In *Conference Record of the ACM Symposium on Principles of Programming Languages, SIGACT and SIGPLAN*, 1973; 194-206.
- [13] E. Kohlbecker, D. Friedman, M. Felleisen, and B. Duba. Hygenic macro expansion. In *Proceedings of the 1986 ACM Conference on Lisp and Functional Programming*, 1986; 151-159.
- [14] D. Melski and T. Reps. Interconvertibility of set constraints and context-free language reachability. In *Proceedings of the 1997 ACM Symposium on Partial Evaluation and Semantics-Based Program Manipulation, PEPM '97*, 1997; 74-89.
- [15] T. Reps, S. Horwitz, and M. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Proceedings of the 22nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 1995; 49-61.
- [16] A. Sandholm and M. Schwartzbach. A type system for dynamic web documents. In *Proceedings of the 27th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2000; 290-301.
- [17] O. Shivers. Control flow analysis in Scheme. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, June 1988; 164-174.
- [18] Z. Su, M. Fähndrich, and A. Aiken. Projection merging: Reducing redundancies in inclusion constraint graphs. In *Proceedings of the 27th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2000; 81-95.
- [19] W. Taha and T. Sheard. Multi-stage programming with explicit annotations. In *Proceedings of the ACM-SIGPLAN Symposium on Partial Evaluation and semantic based program manipulations*, 1997; 203-217.
- [20] R. Vallee-Rai, L. Hendren, V. Sundaresan, P. Lam, E. Gagnon, and P. Co. Soot. A Java optimization framework. In *Proc. IBM Centre for Advanced Studies Conference, CASCON' 99*. IBM, Nov. 1999.
- [21] J. Viega and G. McGraw. *Building Secure Software: How to Avoid Security Problems the Right Way*. Addison Wesley, 2001.
- [22] D. Weise and R. Crew. Programmable syntax macros. In *SIGPLAN Conference on Programming Language Design and Implementation*, pages 156-165, 1993.
- [23] D. Yellin. Speeding up dynamic transitive closure for bounded degree graphs. *Acta Informatica*, July 1993; 30(4): 369-384.

