

# Creating Reports with Query Objects

John Brant

Joseph Yoder

Department of Computer Science  
University of Illinois at Urbana-Champaign  
Urbana, IL 61801  
{brant, yoder}@cs.uiuc.edu

## Abstract

This paper contains a collection of patterns for creating database-reporting applications. While there are many different aspects of reporting, this paper focuses on the ability to dynamically create formulas and queries so that new reports can be generated at runtime. It does not discuss user interface issues or good database design. These patterns pull data from the database and manipulate data after it has been extracted from the database. This is accomplished by converting both queries and formulas into objects. These objects are then assembled into reports through the high-level *Report Objects* pattern.

## Introduction

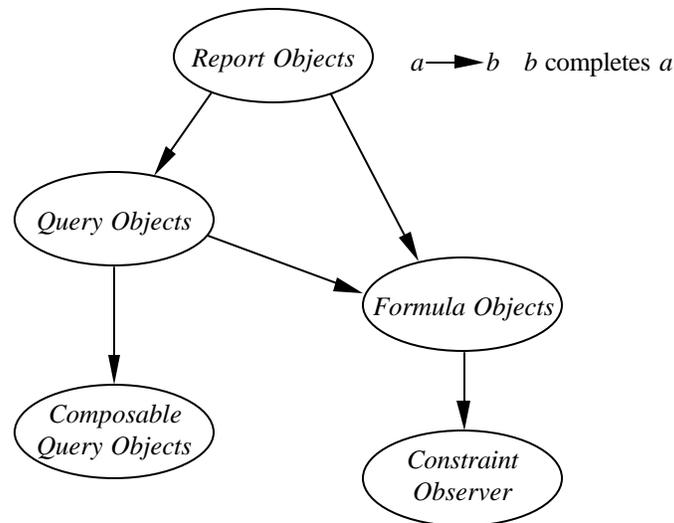
These database-reporting patterns are based upon our research using VisualWorks Smalltalk to dynamically generate reports that map to a relational database. We discovered these patterns while we were building a Business Modeling tool for Caterpillar. Although these patterns are described from a Smalltalk perspective, we feel that they are applicable to anyone building reporting applications. Furthermore, even though our examples apply to relational database queries, these ideas could be extended to other types of queries.

The *Caterpillar Business Modeling project* is a pilot to demonstrate how an appropriate tool can support financial analysis and business decision-making more effectively. This project translates legacy data into a relational database format where it is then mapped to and from objects. Reporting the financial aspects of the Business Model required building many reports based upon business logic and SQL queries.

Earlier versions used predefined reports, but this solution was not dynamic enough since users kept requesting new reports. Although predefined reports initially produce fewer bugs by statically checking the queries and values, they required programmatic changes whenever a new report or query is needed. Such changes can introduce bugs and are time consuming since developer must meet with the users to understand their request.

As a result, we researched making a flexible and user-configurable system. The difficulty here is to make a system that is both easy to understand and modify. If done incorrectly, you've just pushed off the task of programming to the end-users since they would have to specify everything the programmer would normally specify for them. Furthermore, such a system is more difficult to understand since everything is created at runtime and cannot be statically analyzed. This paper describes a collection of patterns, which avoids these two extremes by making objects such as query and formula objects that represent reports. These objects can be predefined so that they give the user reasonable defaults but can later be customized.

Our pattern language comprises five patterns (see Figure 1). *Report Objects* is the top-level pattern for defining reports. It consists of *Query Objects* and *Formula Objects*. *Query Objects* pull data from the database and can be composed into more complex queries by applying *Composable Query Object*. *Formula Objects* process data returned from *Query Objects* and are maintained by *Constraint Observers*.



**Figure 1: Reporting Patterns Overview**

---

## Report Objects

---

### **Motivation:**

You are creating a reporting application that, given some input from the user, must produce a report. While most reports are known during development, there are many that will need to be customized. For example, consider a financial application. The application allows users to view various aspects by selecting different input conditions. You may have designed a report to view sales by marketing areas, but the user might want to also view sales by model. If your reporting framework is not configurable, you will be required to program the desired changes, whereas with a configurable application, the user could define these reports at runtime. Most reporting applications will get their values from some form of a formula equation based upon results from a database query.

### **Problem:**

You need to create a configurable reporting application.

### **Forces:**

Most reports can be predefined with SQL code along with standard views, but these can't contain all of the possibilities that the user might want. Predefined reports are easier to understand since the code can be statically analyzed. However, new reports will require a programmer to modify the system and such modifications are prone to errors.

Alternatively, you could make everything configurable as many report writers already do, so that users could create their own reports. But giving the user the ability to define reports on-the-fly can be very difficult since you must present the options in an easy to understand format that is less complex than programming. Otherwise, you have just pushed the complexities of programming the system onto the users.

### **Solution:**

Define objects representing the report. These objects are made up of objects that process the data and those that view the data. A report is created by attaching processing objects to viewing objects. In this paper we only discuss the processing objects.

The processing objects can be further split into objects that fetch data from the database and objects that manipulate the data. Thus, you can build objects that represent the queries and objects that represent formulas needed in the reports. The *Query Objects* pattern allows you to select data from the database, while the *Formula Objects* pattern allows you to operate on data returned by *Query Objects*.

When reporting, you need to define queries and you need to apply functions on the results from queries. As noted in the forces, the solution can range from pre-defining the SQL queries to providing the user with a report-writer to dynamically generate any type of report. This pattern is more of a middle of the road solution that allows you to make objects that represent queries needed for the reports along with objects for constructing complex values by composing these results. This pattern provides the user with something that *Works Out of the Box*<sup>1</sup> [Foote & Yoder 96] while still allowing for customized reports.

### **Consequences:**

This pattern gives you more flexibility in creating reports. Since the way data is processed is an object, it can be manipulated at runtime. Unlike other implementations where methods define the processing, this technique provides more flexibility since objects are more easily manipulated than methods [Foote & Yoder 95].

Although this pattern provides more flexibility when constructing the reports, it also makes the data processing harder to understand. If queries are built into methods then they can be statically analyzed, but since they are built at runtime, we must reason about the configurations of the objects at runtime. Determining such runtime properties can be difficult if not impossible.

### **Related Patterns:**

- The *Query Objects* pattern helps build objects that query relational databases.
- The *Formula Objects* pattern constructs objects that represent formulas.
- *Metadata and Active Object-Model* patterns [Foote & Yoder 98] can be used to describe the reports and the queries/formulas that compose them. This provides a powerful mechanism for creating and changing the reports at runtime.

### **Known Uses:**

- VisualWorks ReportWriter [PPRW 95] separates the retrieval of data into filter objects and the report view definition into layout objects.
- The Caterpillar/NCSA Financial Model Framework [Yoder] was implemented using this pattern.

---

<sup>1</sup> An artifact that *Works Out of the Box* is one that is immediately able to exhibit useful behavior with minimal arguments or configuration. Enough defaults are provided to get the user up and running without needing to know any details about the system.

---

## Query Objects

---

### **Motivation:**

You are defining reports for your users. Each report queries the database for the information and displays it to the user. While many queries the user will require are known when you are developing the software, there are some that the user will want to add later. For example, you might have a sales query that returns the sales for a specific time interval, but you might want to modify the query to display only Asian sales.

### **Problem:**

You need to create new queries for reports at runtime.

### **Forces:**

If all reports are known in advance, the solution is to create the queries when you build the system. Predefined queries are easy to understand and predict since you just need to analyze the source code. However, they fall short when trying to develop dynamic mappings to the data. Furthermore, changing these queries require program modifications that cannot be performed by the users.

Another option is to create queries at runtime by providing the user with an interface for creating them. This option would require the users to know something about your database in order to build the query. Although users may know what data is stored and what it is used for, they may not know how such information is stored in the database (e.g., they may not grasp the concepts of joining two tables).

### **Solution:**

Create objects that represent queries. Define the operations that will be used to manipulate the queries. For a relational database, you might have operations for selecting which rows should be returned, projecting the fields to be returned, etc. Next, define a method that will return the results of the query. For queries that need to dynamically add new conditions, you can use the *Composable Query Objects* pattern.

For query operations that take expressions such as selections and projections, you can use the *Formula Objects* pattern to create these expressions. Whenever the expression's value changes, you can update the query which will update the report.

This basic intent of this pattern is to convert methods that contain your queries into objects that are easier to manipulate. This is similar to the *Strategy* pattern [GHJV 95], but the main purpose of this pattern is to make objects that are easily manipulated whereas the *Strategy* pattern is more concerned with making them interchangeable.

Although the motivation and forces of this pattern focus on creating a runtime configurable system, *Query Objects* also make it easier for the developer to program. Queries can be described with *Query Objects* by writing object-oriented code rather than SQL. *Query Objects* are similar to Smalltalk collections where common protocol such as *project:*, *select:*, and *join:* create new *Query Objects*.

### **Example:**

For the Asian sales example, you would define an object that represented the sales query, which is just a mapping to retrieve the sales records. You would then add basic manipulation operations such as *project:* that picks which fields are returned; and *select:* that selects which records are returned. In addition to these operations, you need an operation to return the values from the query such as a *value* method. This method converts the query into SQL code that is executed by the database and finally returns the resulting set of values. Using such a protocol, you could modify the sales query object to extract Asian sales simply by evaluating a statement like.

salesQuery select: salesQuery continent = 'Asia'

This statement modifies the salesQuery query object by adding the condition “salesQuery continent = 'Asia'”. This condition is an object that can be modified to support different conditions. For example, the string 'Asia' might be a variable, which can change at runtime. The = method in the “salesQuery continent” object has been overridden so that it is evaluated when the query is performed.

### **Consequences:**

The *Query Object* pattern makes it possible to manipulate queries at runtime. Since the queries are objects, they can be changed to add or remove conditions. This provides the needed flexibility to dynamically create or modify queries for new reports.

One disadvantage with using this pattern is that dynamic queries can be slower than pre-computed queries. The pre-computed queries can be optimized for the database, whereas the dynamic queries might be translated to a slower query. Furthermore, static queries can be pre-computed in database views. If this is a major problem, then you might need to write an optimizing translator that will do a better job translating the query objects into database statements.

### **Related Patterns:**

- *Formula Objects* can create expressions for the select conditions.
- *Composable Query Object* pattern can build new queries out of existing queries.
- The *Strategy* pattern [GHJV 95] is similar since methods are being encapsulated into objects.
- The *Crossing Chasms* pattern language [Brown & Whitenack 95] can be used to map between objects and relational databases.

### **Known Uses:**

- VisualWorks ReportWriter [PPRW 95] created special filters for retrieving data from SQL databases. Although these are objects, they are more like strategies since they do not provide an interface for manipulation.
- Borland Database Engine [Rudraraju 95] provides support for query objects that can be manipulated.
- The Refactoring Browser's BrowserEnvironments [Roberts 99] allow programmers to dynamically query properties of a Smalltalk program.

---

## **Composable Query Objects**

---

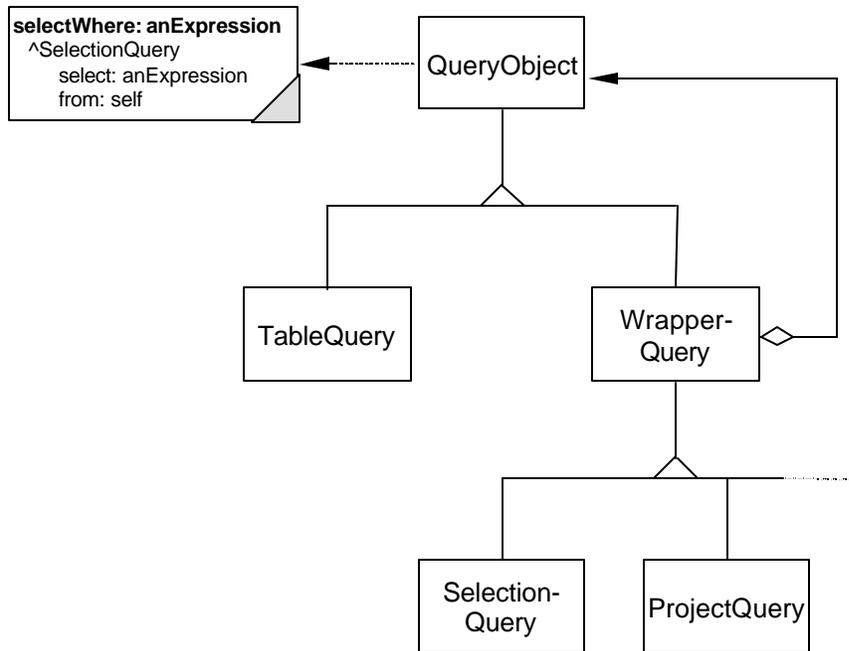
### **Motivation:**

You have decided to use the *Query Objects* pattern to build dynamic queries, but you have many reports with similar queries. Rather than having to define a query for each report, you would like to define one query for the common part and create new queries based on the common query. This would allow you to easily create new queries in some minimalist fashion. Not only would this help during the creation of queries, it would also help during debugging since you only need to fix the common part once and all of the dependent queries are automatically fixed. Furthermore, it allows for the dependent queries to dynamically update whenever the common query modifies its selection criteria.

For example, given the sales query from the previous section, you might want to look at both Asian sales and North American sales during a time interval. Whenever you update the time interval, both reports should also be updated. The only differences between these two reports would be the sales records that are selected.

### **Problem:**

You want to create different queries with simple modifications to common base queries.



**Figure 2: Query Objects Hierarchy**

### **Forces:**

You could use the previous patterns to build a query object for each sales query, but this would be hard to keep up to date with other selections. From the example above, imagine that the user decides to continue to look and compare the North American sales with the Asian sales, but decided to view the fourth quarter instead of the third quarter. This would require triggers or callbacks to make sure that all open views were updated appropriately.

Another solution is to “wrap” new constraints on existing queries. Although wrappers<sup>2</sup> are harder to implement, they do ensure consistency with changing constraints. They also provide a nice way to build a new query from existing queries. If any query needs to be modified to fix a bug, then its wrappers will also be updated without needing to directly modify them.

### **Solution:**

Build queries out of composable parts. Define one class to represent tables in the database, and for each query operation define a class that performs the operation. Additionally, define methods for each operation that will create the appropriate object for that operation on the query. For a relational database, we would define a `TableQuery` class to represent tables, and define classes such as `SelectionQuery` to represent the selection operations, `JoinQuery` for join operations, etc. Since all operations are performed on some query, we can create the abstract superclass, `WrapperQuery`, to handle wrapping these queries with the operations. Finally, you might define methods that create these operations. For example we can define a `selectWhere:` method that will create a `SelectionQuery` on the query (see Figure 2) as opposed to the `select:` method in the previous section that modifies the existing query with the new selection.

To convert a query object into database statements, we use the *Interpreter* pattern [GHJV 95]. Each type of query object will be a node for the interpreter. The table query objects are the terminal nodes for the interpreter, and the operation nodes are the non-terminals since they wrap other query objects. Instead of making an interpreter to create the complete database statements, you can make separate interpreters

<sup>2</sup> We use the terms wrappers and decorators synonymously as done in VisualWorks Smalltalk.

that will create the statement parts (e.g., you would have an interpreter for the *where* clause, another for the projected fields). The primary difference between this use and the *Interpreter* pattern is that the query objects can generate their output without any additional input.

Finally, each of the *WrapperQueries* should be updated whenever its wrapped query changes. This can be accomplished by using the *Observer* pattern [GHJV 95]. The *WrapperQueries* will be observers and their wrapped queries will be subjects. Whenever a wrapped query changes the *WrapperQuery* will also change, allowing the reports to update.

### Example:

For the sales by continent example, we would have a base query object that retrieves all the sales records from the database for a particular time interval. To create a report for a continent, we would just need to create a new query to select those records for that continent. For example, we could create the Asian's sales query with

```
salesQuery select: salesQuery continent = 'Asia'
```

Whenever the *salesQuery*'s date interval changes it will notify its dependents, which are the continent sales queries. Both of the continent sales query objects will change and cause their reports to update. Figure 3 shows the base *salesQuery* object. Both the Asian sales query and the North American sales queries are observers of the *salesQuery*.

### Consequences:

This pattern makes “exploratory programming” easier since you do not need to understand all facets of the database when creating a query. For example, you can extend the sales query without understanding which tables were used, what are the date fields, etc. You would only need to know that there is a continent field. Furthermore, a visual language could be constructed to more easily manipulate these queries.

While this pattern allows you to dynamically create similar queries and automatically maintain them, it also results in objects that are harder to implement. Since the knowledge of a query is spread among several objects, it is harder to create the database statements for the query. We must traverse all of the objects to create the database statements, whereas if all of the details of the query were in one object,

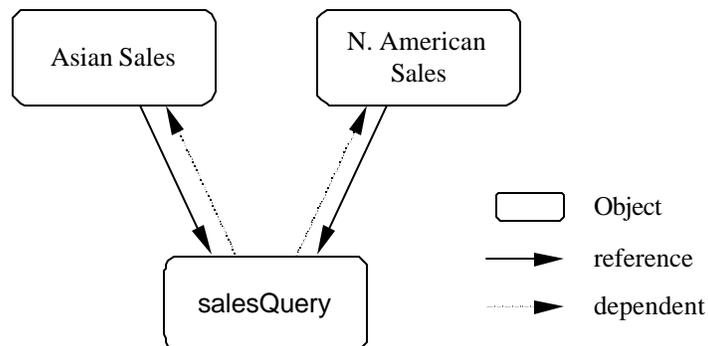


Figure 3: Sales Query Objects

then we could directly create the database statements.

Another disadvantage of this pattern is that it results in many small objects wrapping other query objects. These wrappers make it harder to debug the system, since it is more difficult to traverse the wrappers to see what the query is performing.

### Related Patterns:

- This pattern uses the *Decorator* and *Composite* patterns [GHJV 95] as it wraps new responsibilities and composes new objects.

- Another view is to consider each operation as a parse node using the *Interpreter* pattern [GHJV 95]. Although we could have just referenced these patterns, we felt it was important to include this pattern to show the concrete usage of the design patterns to form *Composable Query Objects*.
- *Observers* [GHJV 95] are used by *Composable Query Objects* to maintain consistency with the *WrapperQueries*.

---

## Formula Objects

---

### **Motivation:**

You are creating a reporting application. Some values are directly available, but many are calculated. For example, PROFIT is a calculated value based on INCOME and EXPENSES. While some of these formulas might be known when the program is created, we want to allow the user to create or modify formulas at runtime.

### **Problem:**

You need to modify formulas and have them automatically maintained.

### **Forces:**

One option would be to define methods for all formulas, but this would require programmer intervention whenever adding or changing formulas. Not only might this introduce new bugs, it also will take some time before it's completed since the change must be explained to the developers. This might require filling out paperwork, attending meetings, etc.

Another option would be to have the user create the formulas at runtime, but many formulas do not change. Users shouldn't have to enter these formulas; they should already be defined. Furthermore, users are more likely to have precedence errors in entering their formulas and are less likely to fully test them.

### **Solution:**

Define an object that represents the results of a computed formula. The result object that is to be used in other formulas should be in the same hierarchy as the objects used to create it. Next, define the basic formula operations to more easily create the *Formula Objects*. For example, you might define methods such as +, -, \*, and / for performing basic arithmetic operations. To make sure that the result object is consistent with its inputs we can use the *Constraint Observer* pattern.

### **Example:**

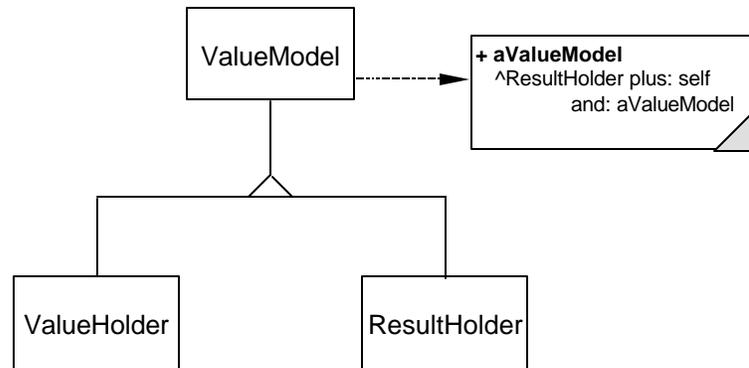
As an example, consider VisualWorks, which stores many of its values in *ValueModels* so that they can easily be plugged into its UI framework. To create a result object we could subclass *ValueModel* to create a *ResultHolder*<sup>3</sup> that would hold the result of the formula (see Figure 4). We could then define basic operations of +, -, \*, and / on *ValueModel* to return a new *ResultHolder* whenever they were executed. Now if both INCOME and EXPENSES are *ValueModels*, then we can create a *ValueModel* that represents PROFIT simply by executing "INCOME - EXPENSES".

### **Consequences:**

This pattern allows both users and programmers to use formulas without knowing how they are maintained. Only the person implementing the pattern needs to understand that. This reduces much of

---

<sup>3</sup> VisualWorks already has a *BlockValue* class that incorporates both the *Formula Object* pattern and the *Constraint Observer* pattern.



**Figure 4: Result Holder Example**

the complexity that would have been required when creating new formulas. Instead, the people creating formulas can concentrate on making them correct, and not on how to maintain them.

*Formula Objects* are not only useful for generating dynamic reports, they can also be used to assist the developer with writing the code by treating these values as objects. For example, rather than getting the value from queries, calculating the resulting value based upon some mathematical function, and then setting the variable to equal the value; you can just assign  $C := A + B$  where *A* and *B* can be *Formula Objects* or *Query Objects*.

One disadvantage of using *Formula Objects* is that it is easy to create too many objects since an object is created for each sub-expression in the formula. For example, a simple formula for computing margin percentage such as “(INCOME - EXPENSES) / INCOME” would create two objects; one that represents “INCOME - EXPENSES” and another for the margin percentage. The extra objects aren’t generally a problem in practice, but could cause problems in applications with many complex formulas. In these applications it may be necessary to “optimize” the formula by making the computed result responsible for the complete formula instead of just expressions in the formula.

### **Related Patterns:**

- *Constraint Observer* pattern can be used to keep the result object consistent with its inputs.
- The *Interpreter* pattern [GHJV 95] can be used to create result objects from user entered formulas.
- *Composable Query Object* pattern is similar. The primary difference is that *Composable Query Objects* operate on queries while *Formula Objects* operate on the results of queries.

### **Known Uses:**

- The WyCash Report Writer [Cunningham 91] has a few formulas that can add additional columns to a report. These are similar to instances of *Formula Objects* in that they define objects for “formula columns” that are based on other columns and/or rows. Using such a system, they were able to build their reports without having to define new types of formulas.
- VisualWorks ReportWriter [PPRW 95] has formula objects that can take either Smalltalk blocks or formulas defined in their own macro language.
- VisualWorks [PPOR 94] has *BlockValues* that can be used to create formulas for *ValueModels*.
- The Accounts framework [Keefer 94] defines functions of time that are used for the attributes of an account.

---

## **Constraint Observer**

---

### **Motivation:**

You have applied the *Formula Objects* pattern to create formulas dynamically, but now you need to update these values whenever one of their input values change. For example, you may have created a

formula "INCOME - EXPENSES" to calculate PROFIT. Whenever INCOME or EXPENSES change, PROFIT should change to maintain the consistency of the formula.

**Problem:**

You need to automatically update calculated values of formulas whenever an input value changes.

**Forces:**

A simple solution would be to create a method for each formula that updates the new values whenever an input changes. This would require a new method for every possible formula, which may not be known in advance.

Constraint systems [Benson & Borning 92 ] are another solution that provides the needed flexibility to automatically maintain the formulas. However, they can limit how programs are written and can be very difficult to understand, especially when trying to follow the control flow while maintaining the formulas.

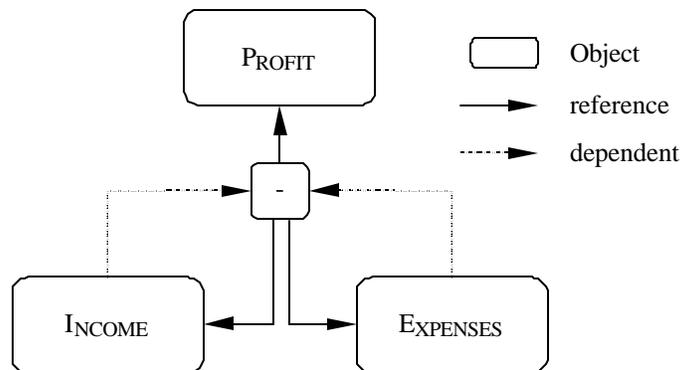
**Solution:**

Create a constraint object that is responsible for the calculation. Use the *Observer* pattern [GHJV 95] to update the constraint whenever an input to the formula changes (i.e., the inputs for the formula will be the subjects and the constraint will be the observer). Whenever the constraint is updated, it evaluates the formula and assigns the result.

The constraint object can either be a part of the result or a separate object. Representing it as separate objects is more flexible since the constraint and the result can be varied independently of each other. Furthermore, each result object could have multiple constraints associated with it. But, if the constraints were created by formulas, then the extra flexibility is not needed since every result object would have only one constraint object.

**Example:**

For the PROFIT example, the constraint depends on both the INCOME and the EXPENSES objects (see Figure 5). Whenever the "-" constraint receives the update message from INCOME or EXPENSES, it evaluates and assigns the result to PROFIT.



**Figure 5: Profit Example**

**Consequences:**

The main advantage of this pattern is simplicity. If you were to use another constraint system, then you would be required to maintain complicated data structures so the constraints would be optimally solved.

Additionally, if you were to use methods, then you would also need other data structures that listed the methods to be executed when an event occurred. This is a maintenance nightmare primarily because you have to maintain many methods that recalculate the values.

The main disadvantage of this pattern is that it can lead to sub-optimal performance. For example, the formula above will change PROFIT whenever INCOME or EXPENSES change. But if every time one value changes the other also changes, then we will be updating PROFIT twice, when one update is all that is required. Although for simple arithmetic formulas, these extra calculations don't have a big impact, for formulas that involve queries these extra calculations might be unacceptable. In these cases a simple flag might be used to delay updating until all objects have changed. If the flag still does not work, you probably need a constraint system that can optimally solve the constraints.

*Constraint Observer* is a simple pattern for implementing a constraint system, but it will not work for many general constraint problems. For example, constraints that are circular would cause this system to get into an infinite cycle of update messages. This is not a problem though, since the constraints were created with formulas. Thus, they guarantee no circularities.

### **Related Patterns:**

- *Observer* pattern [GHJV 95] can be used to update the constraints whenever an input value changes.

### **Known Uses:**

- The *Constraints* pattern described in [Johnson 92] is similar. It is used by HotDraw to update figure locations whenever a dependent figure changes its location. HotDraw only defined simple equality constraints, but it could simulate more complex formulas by using *Locator* objects.
- VisualWorks has a special class that implements *Constraint Observer* for *ValueModels*. The *BlockValue* object [PPOR 94] has both the result and the constraint. The result is stored in the *value* instance variable, and the action performed by the constraint is stored as a *BlockClosure* in the *block* instance variable.
- IBM VisualAge Smalltalk [IBM 97] implements a *Constraint Observer* for maintaining consistency between observable parts (see the connection classes). Visual parts are wired to non-visual parts and whenever the value of one is changed, the value of the dependent part is updated.

## **Summary**

With changing business requirements, users need flexible programs that quickly adapt to their business needs. Users can't wait for programming changes that take weeks. Therefore, it is important to provide the flexibility necessary for the program to adapt to the user's needs. The patterns discussed in this paper help provide such support. They accomplish this by transforming queries and formulas into objects where they can be more easily manipulated.

We are grateful to the members of Professor Johnson's Patterns seminar: Jeff Barcalow, Ian Chai, Brian Foote, Charles Herring, Ralph Johnson, Mark Kendrat, Don Roberts, and Susanne Schacht; our shepherd Kyle Brown; and the members of our writer's workshop who reviewed earlier drafts and provided valuable feedback.

## References

- [Benson & Borning 92] Bjorn Freeman-Benson and Alan Borning, "Integrating Constraints with an Object-Oriented Language", *Proceedings of the 1992 European Conference on Object-Oriented Programming*, June 1992, pages 268-286.
- [Brown & Whitenack 95] Kyle Brown and Bruce G. Whitenack. "Crossing Chasms - A Pattern Language for Object-RDBMS Integration," First Conference on Patterns Languages of Programs (PLoP '94) Monticello, Illinois, August; *Pattern Languages of Program Design* edited by James O. Coplien and Douglas C. Schmidt, Addison-Wesley, 1995.
- [Cunningham 91] Ward Cunningham. "The WyCash Report Writer," OOPSLA '92 Workshop, "Towards an Architecture Handbook." <http://c2.com/doc/ooplsa91.html>.
- [Foote & Yoder 95] Brian Foote and Joseph Yoder. "Evolution, Architecture, and Metamorphosis," Second Conference on Patterns Languages of Programs (PLoP '95) Monticello, Illinois, September 1995; *Pattern Languages of Program Design 2* edited by John M. Vlissides, James O. Coplien, and Norman L. Kerth, Addison-Wesley, 1996.
- [Foote & Yoder 96] Brian Foote and Joseph Yoder. "Selfish Class," Third Conference on Patterns Languages of Programs (PLoP '96) Monticello, Illinois, September 1996; *Pattern Languages of Program Design 3* edited by Robert Martin, Dirk Riehle, and Frank Buschman, Addison-Wesley, 1997.
- [Foote & Yoder 98] Brian Foote and Joseph Yoder. "Metadata and Active Object-Models *Collected papers from the PLoP '98 and EuroPLoP '98 Conference*, Technical Report #wucs-98-25, Dept. of Computer Science, Washington University Department of Computer Science, September 1998. URL: [http://jerry.cs.uiuc.edu/~plop/plop98/final\\_submissions/](http://jerry.cs.uiuc.edu/~plop/plop98/final_submissions/)
- [GHJV 95] Eric Gamma, Richard Helm, Ralph Johnson, John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*, Addison-Wesley, 1995.
- [IBM 97] IBM *VisualAge for Smalltalk User's Reference*. April 1997.
- [Johnson 92] Ralph E. Johnson. "Documenting Frameworks with Patterns," *OOPSLA '92 Proceedings*, SIGPLAN Notices, 27(10): 63-76, Vancouver BC, October 1992.
- [Keefer 94] Paul Dustin Keefer. *An Object Oriented Framework for Accounting Systems*. M.S. Thesis, University of Illinois at Urbana-Champaign, Department of Computer Science, 1994.
- [Roberts 99] Donald Bradley Roberts. *Practical Analysis for Refactoring*. PhD Thesis, University of Illinois at Urbana-Champaign, Department of Computer Science, 1999.
- [Rudraraju 95] Pandu Rudraragu White paper for *Borland Database Engine and IDAPI*. February 1995. <http://www.tietovayla.fi/BORLAND/bbebde/id1ovrvw.htm>.
- [PPOR 94] ParcPlace Systems, Inc. *VisualWorks Object Reference*. June 1994.
- [PPRW 95] ParcPlace Systems, Inc. *VisualWorks ReportWriter: Tutorial and User's Guide*. February 1995.
- [Yoder] Joseph Yoder. *A Framework to Build Financial Models*. URL: [http://www.uiuc.edu/ph/www/j-yoder/financial\\_framework](http://www.uiuc.edu/ph/www/j-yoder/financial_framework).