WYSIWYG Development of Data Driven Web Applications

Fan Yang[*], Nitin Gupta[*], Chavdar Botev[*]
Cornell University, Ithaca, NY
{yangf, niting, cbotev}@cs.cornell.edu

Elizabeth F Churchill, George Levchenko, Jayavel Shanmugasundaram
Yahoo! Research, Santa Clara, CA
{echu, georgel, jaishan}@yahoo-inc.com

ABSTRACT

An emerging trend in Social Networking sites and Web portals is the opening up of their APIs to external application developers. For example, the Facebook Platform, Google Gadgets and Yahoo! Widgets allow developers to design their own applications, which can then can be integrated with the platform and shared with other users. However, current APIs are targeted towards developers with programming expertise and database knowledge; they are not accessible to a large class of users who do not have a programming/database background, but would nevertheless like to create new applications. To address this need, we have developed the AppForge system, which provides a WYSIWYG application development platform. Users can graphically specify the components of webpages inside a Web browser, and the corresponding database schema and application logic will be automatically generated on the fly by the system. The WYSIWYG interface gives instantaneous feedback on what users have created and allows them to run, test and continuously refine their applications. AppForge has been used to create prototype versions of a variety of applications such as an event planning system, a recruiting system, an item trading system and an online course management system. We have also conducted a small and preliminary user study to identify and fix some of the usability aspects of AppForge.

1. INTRODUCTION

As the world moves towards Web 2.0, there is an increasing need to leverage webpages as computing platforms that can enable users to build their own applications. For example, in Facebook and Yahoo! Groups, different groups of users have different needs, and it is difficult for these websites to build applications that satisfy all of these needs. Thus websites are starting to open up their APIs to their advanced users so that they can build new applications that can be deeply integrated with the websites, e.g., the Facebook Platform [30], Yahoo! Widgets [34] and Google Gadgets [20].

However, current APIs and tools are primarily targeted towards developers who have programming and database knowledge. Con-

sequently, they are beyond the reach of the majority of users who lack this knowledge, but would nevertheless like to create and share their own custom applications. For instance, members of a book club in Yahoo! Groups may wish to create a custom application for managing their club events (since no third party application is available to satisfy their specific needs), but the group members may not have the necessary programming expertise to develop this application. Put another way, even though there has been a lot of work on designing languages and tools to simplify application development, ranging from high level programming languages such as Ruby on Rails [28] and Hilda [37] to visual programming tools such as Visual Basic [4] and Oracle Forms [18] to various CASE tools such as UML [7] and WebML [8], the abstractions that these tools provide is still too complex for users with limited programming and database knowledge.

Recently, there has been a flurry of activity on providing online Web application creation services for advanced users[1]. Examples of such websites are Yahoo! Pipes [29], Microsoft Popfly [31], App2You [2], CogHead [11], Zoho Creator [13], Ning [27], Dabble DB [14], WyaWorks [36], JotSpot [23] and SalesForce [33]. These websites allow developers to graphically build web pages and the associated application logic in browsers, thereby greatly lowering the bar for building Web applications. However, these systems suffer from at least one of the following three drawbacks, which limit their applicability and generality.

1. **Non-WYSIWYG development environment.** Most systems (e.g., [2], [11], [13], [27], [29], [33], [36]) have at least two modes: (1) *development mode*, where developers can edit the page structure, application logic and/or database schema, and (2) *execution mode*, where developers and users can actually run and test the application. Consequently, developers have to visualize what they want in the execution mode (i.e., what the end-users will see) and mentally map these into corresponding constructs in the development mode, which results in a significant impedance mismatch. As a loose analogy, consider two popular typesetting tools: LaTeX and Microsoft Word. In LaTeX, users have to mentally map what they want in the final document to the corresponding LaTeX commands, while in Microsoft Word, they directly edit the final document using a WYSIWYG interface. While both approaches have their advantages, the WYSIWYG environment is more accessible to a larger class of users, as also pointed out in [22].

2. **Limited support for creating stateful applications with complex structures such as relationships.** Some systems (e.g.,

---

[*]Work done while the author was at Yahoo! Research.

---

[1]Henceforth, to avoid confusion with end users, we shall refer to advanced users as *developers*; these are not to be confused with professional developers with programming/database knowledge.

[29], [31]) only support stateless web applications with read-only operations, while some other systems (e.g., [23]) only support stateful applications with predefined and restricted structure. A few systems (e.g., [11], [13], [14], [36]) do support sophisticated stateful applications, including advanced features such as relationships between entities, but they require developers to be familiar with relational database schema design. As an illustration, consider a book club event planning application, which includes information about events, speakers and attendees, and also the rating provided by each attendee for every speaker in the event. In effect, the application state contains three entities — event, speaker and attendee — and a 3-way relationship with rating information that connects the three entities. In order to capture this state using the aforementioned systems, developers have to explicitly create a table that connects event, speaker and audience, e.g. using foreign key columns, and also create a column in that table for storing the rating information. In general, such a process is equivalent to creating an Entity-Relationship (E-R) graph [9] and translating it to relational tables, which is challenging for developers.

3. **Limited support for publishing views over multiple related entities.** Many Web applications need to publish pages with complex views of the application state, which could include multiple related entities. For example, in our book club application, we may wish to display the audience for each event, which requires "joining" events with their corresponding participants through a relationship. As pointed out in [22], such join queries in their traditional form are unnatural for average users. Current systems either do not support the creation of such views (e.g., [23]) or assume that developers have database and programming knowledge (e.g. [11], [13], [14], [27], [36]).

To address the above issues, we have developed the AppForge system, which enables developers to graphically build sophisticated applications inside Web browsers. AppForge offers the following advantages over existing systems.

1. **AppForge provides a WYSIWYG environment.** AppForge seamlessly integrates the process of page design, application logic design, schema design, deployment and testing of applications, which greatly lower the bar for building applications. As developers interact with the system by changing the presentation model on web pages, AppForge automatically generates the underlying database schema and application logic on the fly. Developers also get instantaneous feedback on pages when modifications are made to the application logic, database schema and database queries. This allows developers to test, run and continuously refine the application as they are constructing it. The WYSIWYG interface is especially important in our setting, since the developers we are targeting at are expected to constantly make mistakes before producing the desired output.

2. **AppForge enables non-programmers to create sophisticated stateful applications.** Developers just need to focus on building what they want to show in each application page, and AppForge automatically infers the entities and relationships in the underlying database schema. In our book club example, developers can graphically build forms for entering speakers and events, and a view for displaying and editing the speaker for each event. As the pages are being built, AppForge will automatically generate two entities, *Speaker* and *Event*, and a *Presentation* relationship between the two entities. The key technical contribution here is an algorithm that translates a sequence of developers actions based on only *two* simple context-dependent graphical primitives into *arbitrarily complex* schemas in the E-R model. We also prove that this algorithm is capable of generating a large class of E-R models, including those with entities, n-way relationships and aggregations.

3. **AppForge allows non-programmers to create complex views over multiple related entities.** AppForge provides a new navigation paradigm over (automatically generated) E-R models called a Schema Navigation Menu. This menu enables developers to visualize a complex E-R graph as a hierarchical menu and create views with sophisticated operators such as joins, aggregations and selections, without having to understand the details of these operators. The key technical contribution here is an algorithm that generates a Schema Navigation Menu from an arbitrary E-R graph, and then translates developers actions on this menu to sophisticated view definitions. We also prove that this algorithm can generate a large class of nested relational algebra [1] views, including those with primary-foreign key joins, nesting and unnesting.

We have used AppForge to prototype a wide range of applications such as a book club event planning system, a recruiting management system, an online course management system and an item trading system. We have also conducted a small user study to test the usability of AppForge. Based on results of this study, we have identified and fixed some of the main issues that had confused developers, and also identified directions for future exploration.

In summary, the main contributions of this paper are:

- A WYSIWYG interface for developing data driven web applications and an application model for capturing such applications (Sections 2 and 3).

- Schema Navigation Menu, a hierarchical menu-driven way of navigating complex E-R graphs, and algorithms to construct the menu from auto-generated database schema and to translate developers actions on the menu to complex nested relational algebra views (Section 4).

- A core set of two context-dependent graphical primitives for schema design, and an algorithm to translate these primitives into complex E-R graphs (Section 5).

- A preliminary user study to evaluate the usability of AppForge and identify directions for future explorations (Section 6).

2. APPFORGE SYSTEM OVERVIEW

We first provide an overview of the AppForge system architecture before describing the AppForge GUI using a running example.

2.1 System Architecture

Figure 1 shows the architecture of the AppForge System. In the front-end, AppForge provides a graphical interface for building and running applications. As developers build an application, the system automatically generates the schema and application logic, and stores this information in the back-end. Developers can run the application at the same time as they are editing it.

The back-end system consists of two sub-systems: Application Creation System and Application Runtime System. The Application Creation System creates and updates the *application model* based on developers' actions. The application model includes the specification of page views, application logic and database schema. The Page View Creation module provides an interface for creating and updating webpages. Developers' actions at the front-end for creating/editing page views are translated into commands in the Page View Creation module. The Automatic Schema Generation
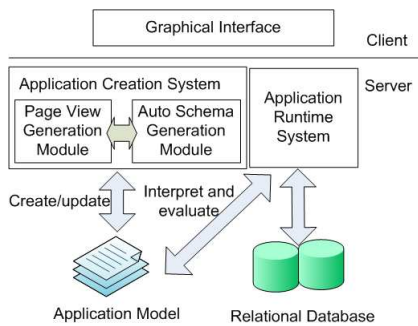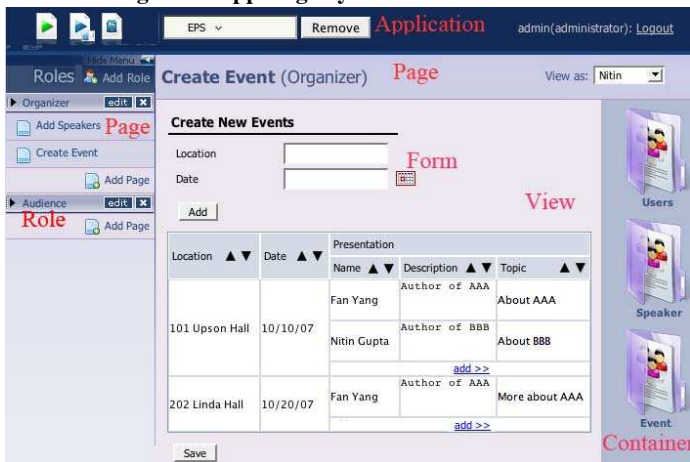
Figure 1: AppForge System Architecture

Figure 2: AppForge GUI

module automatically generates the appropriate relational database schema from page views. Note that in AppForge, building page views and generating the schema is an iterative process: new views are built by navigating the existing schema, and the schema is implicitly updated when page views are updated.

The application model created by the Application Creation System is stored in the file system, while the application state is stored in a relational database system. At start-up time, the Application Runtime System loads the application model into memory, and then serves end users' requests by interpreting the model and issuing SQL queries over the relational database system.

2.2 AppForge GUI

Figure 2 shows a screen shot of the AppForge GUI. As shown, the GUI exposes the following abstractions to developers:

- **Application.** Developers can create and manage multiple applications. Each application can be pre-populated with a list of users. For example, in a Yahoo! Group application, the users can be initialized to be all the members of the group.
- **Role.** Users of applications can be divided into multiple roles. Users in different roles can view pages with different content and allowable actions.
- **Page.** Users in each role can access a set of pages. Each page can contain one or more Forms and Views.
- **Form.** Users can use forms to enter new data. Forms are associated with the logic needed to update the relevant database tables. In AppForge, we support many types of form components such as input fields and drop down boxes.
- **View.** Users can view and update the application state using

views. By default, views are presented as nested tables, but other formats such as unnested lists and charts can also be supported.

- **Container.** Containers corresponds to entities in an application. Containers are automatically created when developers add new forms and views to pages. Containers are used as a visual aid and only developers can see them.

2.3 Running Example

We now illustrate the AppForge GUI using a running example. Consider a book club in Yahoo! Groups that organizes regular events with invited speakers to give presentations on different books. While there are many event planning sites such as Evite [17], none of them support the specific features required by the book club. Consequently, the book club members decide to build their own customized Event Planning System (EPS).

There are two roles in EPS: organizers and attendees. Organizers can add candidate speakers, create events, and view registered attendees and their feedback after each event. Attendees can register for an event and provide feedback on each speaker. They can also volunteer to help speakers in each event, e.g. by providing transportation.

Using AppForge, members of the group can create such an EPS easily. We now illustrate this process by building several key pages for organizers, including the Create Event page (Figure 2), the View Volunteers page (Figure 9), and the View Comments page (Figure 13). Note that the following screenshots show **all** the steps needed to create these pages, and are thus indicative of the easy-to-use aspect of AppForge. In the following discussion, assume that we have already created an application named EPS with a pre-populated container Users, which contains all book club members, and a container Speaker, which contains information about speakers.

Create Events Page (Figure 2). Organizers can create new events, and add/edit speakers and presentation topics for each event (note that adding speakers and presentation topics associated with an event updates not just the relevant entities, but also the relationship between these entities).

1. Create a form named *Event* with fields *Location* and *Date* as in Figure 3. The resulting form is shown in Figure 4.

   Automatic schema updates: a new *Event* entity (container) with attributes *Location* and *Date* is created.

2. Add a view over the *Event* container (Figure 4.1)[2] and select the columns to show in the view (Figure 4.2).

3. Click beside the view (Figure 5.1), add a new column named *Presentation* of type *Speaker*, and select columns in *Speaker* to show in the view (Figure 5.2).

   Automatic schema updates: a 2-way relationship named *Presentation* between the *Event* and *Speaker* entities is created.

4. Click on nested table for *Presentation* (Figure 6.1) and add a new column named *Topic* (Figure 6.2). The resulting page is shown in Figure 2. End users can click on the link $add >>$ under the *Presentation* column to add speakers to each event.

   Automatic schema updates: a *topic* attribute is added to the *Presentation* relationship.

---

[2]In the menu, "insert existing view" means inserting a view over an existing container.
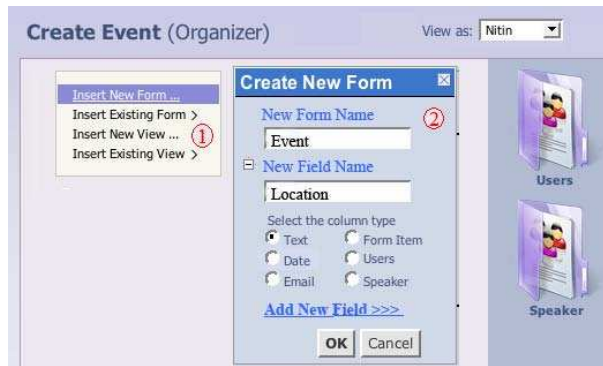
Figure 3: Adding a form for creating new events. The resulting form is shown in Figure 4.

Figure 4: Adding a table to show existing events. The resulting table is shown in Figure 5.

Figure 5: Adding the presentation column to the table. Organizers can add speakers for presentations. The resulting table is shown in Figure 6.
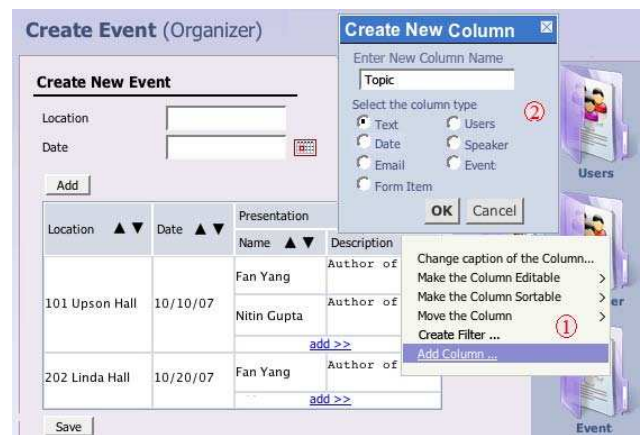

Figure 6: Adding the topic column to the presentation nested table. The resulting table is shown in Figure 2.

Figure 7: Start creating the View Volunteers page by creating a view over *Speaker* and *Event*
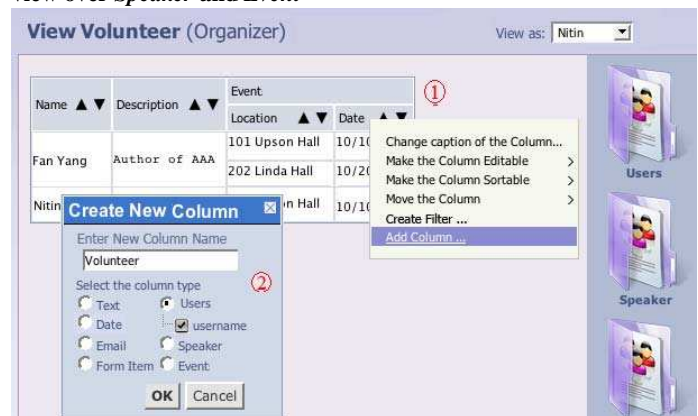

Figure 8: Adding the volunteer column to the *Event* nested table


Figure 9: The View Volunteers Page

View Volunteers Page (Figure 9). Organizers can view the members who volunteered to assist speakers. Volunteers are associated with each speaker in each event.

1. Add a view over *Speaker* and *Event* by navigating from *Speaker* to *Event* in the schema navigation menu (Figure 7.1, 7.2). The resulting view is shown in Figure 8.1.

2. Click on the nested table for *Event* (Figure 8.1) and add a column *volunteer* of type *Users* (Figure 8.2).

   Automatic schema updates: an aggregation of *Speaker* and *Event* is created, and then a 2-way relationship between the aggregation and the *Users* entity is created.

View Comments Page (Figure 13). Organizers can view comments by attendees on event speakers. The view should only show events that have occurred in the past.

1. Add a view over *User*, *Event* and *Speaker* by navigating through the schema navigation menu (Figure 10.1) The resulting view is shown in Figure 10.2.

2. Click on the *Date* column (Figure 11.1) and create a filter that specifies that the event date is earlier than the current date (Figure 11.2).

3. Click on the *Attend* nested table (Figure 12.1) and create a new column named *Rating* 12.2).

   Automatic schema updates: a 3-way relationship with an attribute *Rating* that connects *Speaker*, *Event* and *Users* is created.

The above examples illustrate how AppForge provides a WYSIWYG environment. Developers always view the application the same way as the end users, and they focus on what they want to present in webpages while the underlying schema is created/updated automatically (Figure 14 shows the final schema automatically generated in our running example). Also using the Schema Navigation Menu (Figures 4.2, 7.2 and 10.2), developers can easily navigate through the automatically generated schema and graphically construct complex views (Figures 2, 9 and 13).

3. APPFORGE APPLICATION MODEL

As mentioned in the introduction, two of the key technical contributions of this paper are (a) an algorithm for generating views based on developers' actions and (b) an algorithm for generating the database schema based on page views. In AppForge, the application views and schema is captured formally using an underlying Application Model, which fully characterizes an application. We now introduce the Application Model and describe the algorithms in the subsequent sections.

3.1 Background

The Application Model is an extension of the well-known E-R model [9], which is commonly used to model database entities and relationships, and the Nested Relational Algebra (NRA) [1], which is commonly used to represent nested views. We now briefly review the E-R and the NRA model.

The E-R model models the world in terms of entities and relationships between entities. Figure 14 shows the database schema automatically generated for our running example in the E-R model. Entities are represented as rectangular boxes, e.g., *Speaker*, *Event* and *Users*, and attributes of entities are represented as ellipses.
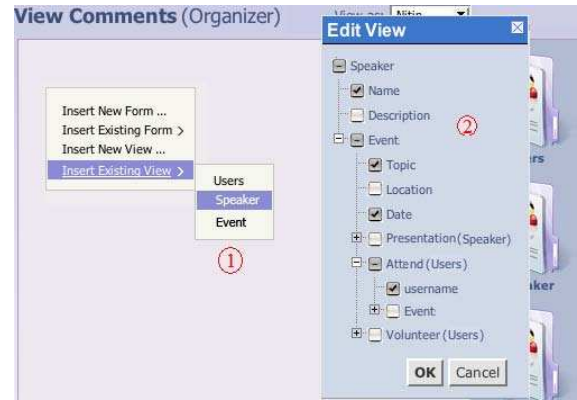
Figure 10: Start creating the View Comments page by creating a view over speaker, event and attendee
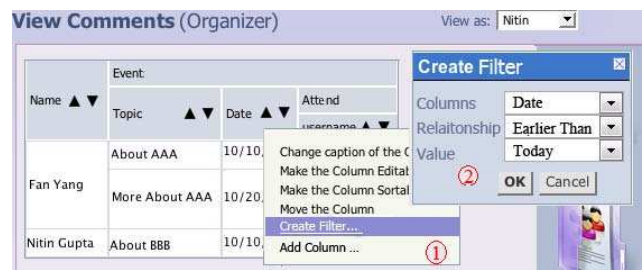
Figure 11: Creating a filter to show only past events

Figure 12: Adding a ratings column for each attendee

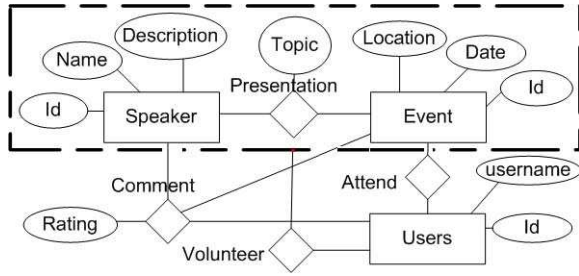**Figure 13: The View Comments page**

Figure 14: Automatically generated database schema

Figure 15: Flat and Nested Tables

Relationships are represented as diamond boxes, e.g., *Presentation* and *Comment*. Presentation is a 2-way relationship that connects Speaker and Event, which captures the meaning that speakers present in events. Comment is a 3-way relationships connecting Speaker, Event and Users, which captures the meaning that an attendee gives ratings for each speaker in each event. In the E-R model, a relationship and all its participating entities can be treated as an *aggregation* for the purpose of taking part in another relationship. For example, the dashed rectangular boxes in Figure 14 is an aggregation that aggregates Speaker and Event pairs. The aggregation participates as an entity in the Volunteer relationship, which captures the meaning that a club member can volunteer to help a speaker who presents in a event.

In AppForge, views are tables in the nested relational model. The nested relational model extends the relational model by relaxing the first normal form assumption i.e. a column can contain a nested table. It is more flexible than the relational model because it can model hierarchical data, which are commonly used in Web applications. The nested relational algebra has two extra operators compared to the relational algebra: nest (ν) and unnest (μ). ν_C groups all the columns other than C based on the value of C. μ is the reverse operation of ν. As an illustration, in Figure 15, the left table is a flat table that shows a list of speakers, and the date and location of the corresponding events. Nesting the table on the name column (ν_{name}) would produce the right table. Columns other than name are aggregated based on name and form a nested table. The effect of unnest is the reverse of nest. Unnesting the right table on the location and the date columns ($\mu_{location,date}$) would produce the left table. The schema of nested tables can be expressed as a nested set of columns. For example, the schema for the right table in Figure 15 is {name, {location, date}}.

3.2 Application Model

The AppForge Application Model contains the following components.

Database Model: It specifies the application state.

- **Schema.** The database schema is represented as an E-R graph. Figure 14 represents the automatically generated schema for our running example.

- **Constraints.** Besides a database schema, an application can have additional constraints on valid application states. In our

running examples, users can provide a rating for a speaker in an event (the 3-way relationship in Figure 14). However, this relationship only makes sense if the speaker presented in the event (a 2-way relationship) and the user attended that event (another 2-way relationship). Such constraints between n-way and $n - 1$ (and lower) way relationship are captured in the application model and enforced by the Application Runtime System. This aspect is illustrated with an example at the end of Section 5.2.

Page Model: The page model specifies the content, structure and presentation of webpages.

- **Content and Structure.** The content and structure of a view (and similarly, a form) is specified as a nested relational algebra expression over the E-R graph. For example, the view in Figure 2 can be defined by the following algebra expression:

$$\nu_{Location,Date}$$

$$(\Pi_{Location,Date,Name,Description,Topic}$$

$$(Event \bowtie_{Left_{Event.id=Presentation.eventid}}$$

$$Presentation \bowtie_{Left_{Speaker.id=Presentation.speakerid}} Speaker)).$$

It joins Event and Speaker through Presentation, projects on necessary columns and nests on columns for Event. The schema of the view is {Location, Date, {Name, Description, Topic}}.

- **Presentation.** These capture presentation aspects of views and forms such as background color, column captions and which columns are updatable.

As mentioned earlier, the application model is automatically generated based on developers' actions such as those illustrated in section 2. Specifically, the Page Model is generated by the Page View Creation module and the Database Model is generated by the Automatic Schema Generation module (Figure 1). Further, the entities and relationships are mapped to relational tables, and nested relational algebra queries are converted into SQL queries at run-time. We now discuss the core abstractions and algorithms used in the Page View Creation and Automatic Schema Generation modules.

4. CONSTRUCTING VIEWS

The Page View Creation module (Figure 1) constructs views based on the database schema and developers' actions. The main challenge is making this functionality accessible to developers without database and programming knowledge. Specifically, we would like to enable developers to (a) navigate through a database schema without exposing the complexity of an E-R graph, and (b) create complex NRA expressions without exposing the details of NRA operators such as join and nest.

We address the above two challenges as follows. First, we introduce the Schema Navigation Menu as a visual utility to transform the E-R graph into a navigational tree menu. Using this menu, developers can easily navigate an E-R graph. Second, we describe a set of three graphical primitives for creating and editing NRA expressions over the schema. We then prove that using only these three primitives, developers can construct views that correspond to the large set of NRA expressions with joins on primary/foreign key.

4.1 Schema Navigation Menu

A Schema Navigation Menu is a tree structured menu whose root is an entity in the E-R graph. The construction of a menu is initiated when a developer selects the root entity (Figures 4.1, 7.1 and 10.1). The options and structure of the menu are determined by the attributes and relationships among entities in the E-R graph (Figures 4.2, 7.2 and 10.2). At each level of the menu tree, the list of checkable options are produced using Algorithm 1. Note that

this algorithm is recursively invoked on demand for each level of the Schema Navigation Menu to produce the hierarchical structure displayed to the developer.

In Algorithm 1, we use term *currentStep* to denote the entity that we are currently expanding. It is initialized to be the root entity. We use term *navigationPath* of *currentStep* to represent the list of entities and relationships through which we have navigated from the root of the menu tree to *currentStep*. *link* represents the relationship through which we just reached *currentStep* from its parent in the menu tree. If the current step is the root entity, *link* is null. At each level of the tree, the following list of checkable options are presented.

- *Entity Attributes.* The list of attributes in *currentStep* (line 2).

- *Relationship Attributes.* The attributes of *link* are shown as if they were attributes of *currentStep* to avoid explicitly exposing the relationship to developers (line 3). For example, in Figure 10.2, topic, which is an attribute of the Presentation relationship, is shown along with other attributes of Speaker. For each n-way ($n > 2$) relationships that *currentStep* participates in, we check if navigationPath of the *currentStep* contains all the entities that the n-way relationship connects. If so, we show the attributes of the n-way relationship as well (lines 12-13).

- *Navigational Link.* If *currentStep* is connected with other entities by 2-way relationships, those entities will be shown in the menu (lines 4-8). For example, in Figure 7.2, Event is shown under Speaker since they are connected by the Presentation relationship. If *currentStep* is connected with an aggregation through a 2-way relationship, all the entities in the aggregation will also be included in the menu (lines 4-10). A navigational link is shown as an expandable item. Selecting this item will expand the menu to show the options for that entity.

  If *link* participates in a relationship as an aggregation, it is treated in the same way as currentStep (lines 15-25). For example, the Presentation relationship forms an aggregation and connects with Users through the Volunteer relationship (Figure 14). When navigating from Speaker to Event (*link* is Presentation), the option volunteer (Users) is shown under Event in the menu (Figure 10.2).

When displaying an entity name, we sometimes also include the relationship name if we can navigate to the same entity through more than one relationship. For example, in Figure 10.2, when starting from Event, we can reach Users as attender or volunteer; the relationship names are used to distinguish these cases.

4.2 Graphical Primitives for Editing Views

AppForge provides the following graphical primitives for developers to edit views. These primitives are automatically translated into NRA expressions.

Select Menu Item. From the Schema Navigation Menu, we can select the following options, each of which updates the underlying view specification.

- *Entity Attributes and Relationship Attributes.* Developers can select what attributes to show in the view. This action corresponds to the projection operator in NRA. For example, in Figure 4, developers can select which attributes of Event are to be shown in the table.

- *Navigational Link.* By navigating to a new entity, the underlying view will be updated by joining the new entity through the navigation relationship. By default, a nested column is created to show the attributes selected after each navigation. For example, if we navigate from Speaker to Event and then to Users

as in Figure 10.1, and select attributes to show along the way, the view will be created by joining the three entities through the Presentation and Attend relationships. Nested columns will be created for columns of Event and Users, producing the view shown in Figure 11.

Move up/down columns. Developers can change the nesting structure of the view by moving columns up and down the view. If they move a column down, they will be asked which nested column it should be moved into, or if the system should create a new nested column. For example, moving down both the Location and Date columns in Figure 15(Left) into a newly created column called Event will produce the nested table in Figure 15(Right). Similarly, moving up the Location and Date columns in Figure 15(Right) will produce Figure 15(Left).

Create filter. We can limit the data shown in a view by specifying a filter predicate of the form (*column operator value*). Operator can be any comparison operators supported by the underlying database system. Developers can input a constant value or select from a list of context variables supported by the system e.g., the current date. Figure 11 shows an example filter that selects past events.

Input : *currentStep* : The current entity being expanded
       link : The relationship through which *currentStep* was reached
Output: *Items* : List of options that can be selected by developers for currentStep

1 **AttrForNextStep** *link, currentStep*
2   Items = currentStep.attributes
3   Items += link.attributes
4   **foreach** *relationship r that currentStep is involved in* **do**
5   **if** *r is 2-way relationship* **then**
6   nextStep = r.otherSide(currentStep)
7   **if** *nextStep is not an aggregation* **then**
8   Items += nextStep
9   **else**
10   Items += all entities in the nextStep aggregation
11
12   **else if** *navigationPath of currentStep contains all entities participating in r* **then**
13   Items += r.attributes
14
15   **if** *link forms an aggregation Agg* **then**
16   **foreach** *relationship r that Agg is involved in* **do**
17   **if** *r is 2-way relationship* **then**
18   nextStep = r.otherSide(Agg)
19   **if** *nextStep is not an aggregation* **then**
20   Items += nextStep
21   **else**
22   Items += nextStep.allEntities
23
24   **else if** *navigationPath of Agg contains all entities participating in r* **then**
25   Items += r.attributes
26
27
28

Algorithm 1: Algorithm for transforming a database schema into a Schema Navigation Menu. The algorithm specifies how to generate options for each step in the menu tree.

4.3 Expressiveness Theorem

We now formally characterize the set of NRA views that can be constructed using AppForge. For ease of exposition, we assume a simple translation from the E-R model to the relational model that maps each entity and each relationship into a separate table.

Definition 1 Let R be a n-way ($n \geq 2$) relationship that relates entities $A_1..A_n$, and let e_1 and e_2 be nested relational algebra expressions whose output schema contains the ids of $A_1...A_m$ and $A_{m+1}...A_n$ ($1 \leq m \leq n$), respectively. We define operators:

- $e_1 \bowtie_{R(A_1...A_m;A_{m+1}...A_n)} e_2$

  $= e_1 \bowtie_{(R.A_1id=A_1.id...\wedge R.A_mid=A_m.id)} R$

  $\bowtie_{(R.A_{m+1}id=A_{m+1}.id...\wedge R.A_nid=A_n.id)} e_2$.

- $e_1 \bowtie_{left_{R(A_1...A_m;A_{m+1}...A_n)}} e_2$

  $= e_1 \bowtie_{left_{(R.A_1id=A_1.id...\wedge R.A_mid=A_m.id)}} R$

  $\bowtie_{(R.A_{m+1}id=A_{m+1}.id...\wedge R.A_nid=A_n.id)} e_2$.

where $\bowtie_b$ and $\bowtie_{left_b}$ are the join and left join operators, respectively, and b is the joining condition.

Intuitively, the two operators represent the join and left join based on foreign key and primary key between two entities that are connected by a relationship. For the rest of the paper, we interpret the left join operator as being right associative, i.e., $A \bowtie_{left} B \bowtie_{left} C = A \bowtie_{left} (B \bowtie_{left} C)$.

The following definition defines the set of NRA expressions that can be constructed using AppForge.

Definition 2 E is recursively defined as follows:

- For every entity en, $en \in E$

- If $e \in E$, then $\Pi_c e \in E$, $\sigma_p e \in E$, $\mu_c e \in E$ and $\nu_c e \in E$, where p is a logical expression on columns in schema of E. c is columns in schema of E.

- If $e_1, e_2 \in E$, then $e_1 \bowtie_{R(A;B)} e_2 \in E$ and $e_1 \bowtie_{left_{R(A;B)}} e_2 \in E$, where A, B are sets of entities.

Theorem 1 Algorithm 1 in conjunction with the graphical primitives in Section 4.2 can construct all and only expressions in E

Proof Sketch: Without loss of generality, we assume that all the attribute names are unique. We first inductively prove that all the expressions that can be constructed using the AppForge graphical primitives are in E. Assume that expressions $e, \tilde{e} \in E$ are constructed using a sequence of AppForge graphical primitives, and after applying another primitive, we get a new expression e'. We need to show that $e' \in$ E. If the operation applied is:

- **Select Menu Item.**

  (a) **Entity or Relationship Attributes** If we select a set of attributes m shown in the menu, then $e' = \Pi_m e$

  (b) **Navigational Link.** If we navigate through a link from entity n (e is an expression over n), and reach entity m ($\tilde{e}$ is an expression over m) by following the link r, and then $e' = e \bowtie_{left_{r(m,n)}} \tilde{e}$

- **Move up/down columns.** Let $NODE(t)$ represent the nested table that contains t as an attribute, $ATTR(T)$ represent all the attributes of table T, $NS(T)$ represent the schema for nested table T, and $PARENT(T)$ represent the table that contains the nested table T as a column. Assume that we want to move column t, and $T_1 = NODE(t)$. So, we have $t \in NS(T_1)$.

(a) Move up columns: We can move t out of the nested column to the upper level in the table. Let $T_2 = PARENT(NODE(t))$, where T_2 is the destination we want to move t to. The resulting expression would be $e' = e[\nu_{NS(T_2)-\{NS(T_1)\}} \rightarrow \nu_{NS(T_1)\cup\{t\}-\{NS(T_2)-\{t\}\}}]$

(b) Move down: We can move t down to an existing nested column or create a new nested column. In the former case, assuming T_2 is the schema tree for the nested column we want to move t into, e' is $e[\nu_{NS(T_1)-\{NS(T_2)\}} \rightarrow \nu_{NS(T_1)-\{t\}-\{NS(T_2)\cup\{t\}\}}]$. In the latter case, e' is $e[\nu_{NS(T_1)-\{NS(T_2)\}} \rightarrow \nu_{NS(T_1)-\{NS(T_2)\}}\nu_{NS(T_1)-\{t\}}]$

- **Create filter.** We can create a filter as a boolean predicate p then $e' = \sigma_p e$

So after each graphical command, the resulting expression is still a valid expression in E.

Next, we inductively prove that for every expression in our algebra E, we can construct it using the set of graphical primitives. Assume we can construct e_1, e_2 using graphical commands, let e be an expression built from e_1, e_2 by following the inductive steps in Definition 2.

- If $e = en$, where en is a relation, we can construct e by selecting en as the root entity or by navigating to en.

- If $e = \Pi_c e_1$, we can construct e by selecting the set of attributes c from the menus for the table corresponding to e_1.

- If $e = \sigma_p e_1$, we can construct e by creating a filter p on the table corresponding to e_1.

- If $e = \mu_c e_1$, we can construct e by moving the columns c down in the table corresponding to e_1.

- If $e = \nu_c e_1$, we can construct e by moving the columns c up in the table corresponding to e_1.

- If $e = e_1 \bowtie_{left_{R(A,B)}} e_2$, we can find an attribute of A in columns/nested columns of e_1. We start navigation from that column and reach B through link R. The we can construct e_2 using graphical commands based on the inductive assumption.

- If $e = e_1 \bowtie_{R(A,B)} e_2$. Since we can use the left join operator and the *not NULL* predicate to represent the join operator, we can use the previous procedure with an extra predicate *B.id is not NULL* to create $e = e_1 \bowtie_{R(A,B)} e_2$. $\square$

Besides proving the expressiveness of the UI operators, Theorem 1 also illustrates how views (NRA expressions) can be constructed through UI operations.

5. AUTOMATIC SCHEMA GENERATION

In the previous section, we described how developers can graphically construct arbitrarily complex views over a given database schema. However, constructing a database schema itself is not an easy tasks for developers. To address this issue, the Automatic Schema Generation module automatically generates complex schemas based on just two simple developers actions: (a) creating forms/views, and (b) adding columns to forms/views. The graphical context (position in form/view) of these two actions is powerful enough to construct arbitrarily complex schemas, including those with n-way relationships and aggregations.

The schema generation algorithm is given in Algorithm 2 . We now walk through this algorithm for the different cases.

5.1 Editing Entities

Entities are created when developers add *new* forms or views (lines 1-3). The columns of tables and fields of forms map to the attributes of entities. The attributes types information are inferred

```
/* Triggered when developers add new
   forms/views to a page.              */
Input : name : Name of the new form/view
        attrs : Columns in the new form/view
1  onNewFormViewEvent name, attrs
2  │  AddEntity (name, attrs)
3
   /* Triggered when developers add columns
      to views.                        */
Input : target : The position in the view where the developer
        clicks
        newAttrName : The name of the column to be added
        type : The type of the column to be added
4  onAddAttributeEvent target, newAttrName, type
5  │  if target is a non-nested column of the view or beside the
   │  view then
6  │  │  targetEntity = root entity of the view
7  │  else
8  │  │  targetEntity = the entity that the target column belongs
   │  │  to
9  │  if NOT isEntity(type) then
10 │  │  if targetEntity is root entity then
11 │  │  │  AddAttribute (targetEntity, newAttrName, type)
12 │  │  else
13 │  │  │  navigationPath = getNavigationPath(targetEntity)
14 │  │  │  if navigationPath contains two entities then
15 │  │  │  │  r = the relationship that connects the two
   │  │  │  │  entities in navigationPath
16 │  │  │  else if exists relationship r that connects all
   │  │  │  entities in the navigationPath AND exists a
   │  │  │  constraints that r depends on all 2-way
   │  │  │  relationships in the navigationPath then
17 │  │  │  │  r = getTheRelationship(navigationPath)
18 │  │  │  else
19 │  │  │  │  r = createRelationship(navigationPath)
20 │  │  │  │  create a constraint that r depends on all 2-way
   │  │  │  │  relationships in navigationPath.
21 │  │  │  AddAttribute (r, newAttrName, type)
22
23 │  else
24 │  │  if targetEntity is root entity then
25 │  │  │  createRelationship (targetEntity, getEntity(type),
   │  │  │  newAttrName)
26 │  │  else
27 │  │  │  navigationPath = getNavigationPath(target)
28 │  │  │  if exists an aggregation over the navigationPath)
   │  │  │  then
29 │  │  │  │  aggregation = getAggregation(navigationPath)
30 │  │  │  else
31 │  │  │  │  aggregation =
   │  │  │  │  createAggregation(navigationPath)
32 │  │  │  createRelationship(aggregation, getEntity(type),
   │  │  │  newAttrName)
33
34
35
```

**Algorithm 2**: Algorithms for automatically generating a database schema when editing views

from the types of graphical components used in the page. The type can be a primitive type such as link, text, email and form components, or it can be an entity type (Figure 5.2) . Developers can edit forms and views later by adding fields and columns. If developers add new columns of a primitive type by clicking besides the view or on the top level columns of the view, new attributes will be added to the root entity of the view (lines 5-6, 10-11).

As an illustration, in Figure 3, adding a form automatically creates an Event entity (Figure 4) and the fields in the form are mapped to the attributes of the entity. An id attribute is also automatically created, which is the key for the entity.

5.2 Editing Relationships

Relationships are created when developers create and edit views that show information about multiple entities.

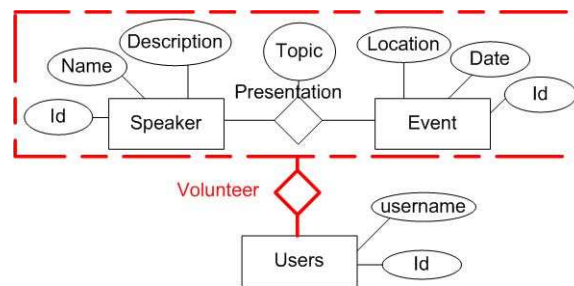### 5.2.1 2-way Relationships without Aggregation

Figure 16: The schema generated for the Create Event page (Figure 2)

When a developer adds a new column of type entity to a table, a new relationship is created to connect the entity associated with the table and the entity associated with the new column (lines 24-25). As an illustration, in the Create Event page (Figure 2), creating a view over Event and then adding a new column to the view (Figure 5.1) of type Speaker (Figure 5.2), creates a 2-way relationship that connects Speaker and Event.

Attributes can be added to 2-way relationships as follows. When developers add a primitive type column to a nested table, the system adds a new corresponding attribute to the relationship between the top level and nested entities (lines 13-14, 21). For example, in Figure 6, adding the topic column to presentation adds a corresponding attribute to the Presentation relationship because this relationship relates Event and Speaker. Note that this is the desired semantics: topic is associated with a speaker-event pair. The schema generated for the Create Events Page is shown in Figure 16.

5.2.2 2-way Relationships with Aggregation

Adding a column of type entity to a nested table establishes a relationship between an entity and the aggregation of the related entities in the view (lines 28-32). As an illustration, in the View Volunteers page (Figure 9), adding a volunteer column of type Users

Figure 17: The schema generated for the View Volunteers page (Figure 9)

to the event nested table creates an aggregation of the Event and Speaker entities, and a 2-way relationship between the aggregation and the Users entity. Note that this is the desired semantics because volunteers are associated with speaker-event pairs. Figure 17 shows the schema generated from the View Volunteers page.

5.2.3 *n-way relationships*

Figure 18: The schema generated for the View Comments page (Figure 13)

n-way relationships are created by adding primitive type columns to nested tables in views. If an n-way relationship that relates all the entities in the nested view already exists, then an attribute corresponding to the new column is added to the relationship; else the n-way relationship is first created (lines 16-21) before adding the new column. In the View Comments page (Figure 9), adding a new column rating to the nested table for attendees creates a three way relationship between users, events and speakers as in Figure 18, and adds the rating attribute to the relationship. Note that this is the desired semantics because the rating is associated with a group member for a particular speaker in a specific event.

Note, however, that there are some semantic constraints that are not captured here in the E-R graph. The 3-way relationship should only connect users that attend the event and speakers that present in the same event. Put another way, the 3-way relationship should connect users, events and speakers that are connected by the two 2-way relationships through which we construct the underlying view (Figure 7.2). Such constraints cannot be captured by participation constraints in E-R model because they related multiple inter-related relationships. So besides the 3-way relationship, AppForge will also create a data constraint that the 3-way relationship depends on the two 2-way relationships. By saying that a n-way relationship depends on a set of $n-1$ 2-way relationships, we mean the instances of entities that are connected by the n-way relationship also have to be connected by the $n-1$ 2-way relationships.

5.3 Expressiveness Theorem

We now formally characterize the set of E-R graphs that can be constructed using AppForge. An E-R graph can be formally defined as a graph $G = (EN, RE, E)$ where EN represents the set of entities and RE represents the set of relationships. E represents the set of edges that connects entities with relationships and edges that connects relationships with relationships as in the case of aggregations, i.e., $E \subseteq \{(u,v)|u \in EN$ and $v \in RE$ or $u \in RE$ and $v \in RE\}$.

Definition 3 For $e_1, e_2 \in EN \cup RE$, we define $R(e_1, e_2) = \{r|(e_1, r) \in E$ and $(e_2, r) \in E\}$. $R(e_1, e_2)$ is the set of 2-way relationships that exist between entities/aggregations e_1 and e_2.

Definition 4 For $EA \subseteq EN \cup RE$ and $|EA| > 2$, we define $M(EA) = \{r|\forall e \in EA \exists (e, r) \in E\}$. $M(EA)$ is the set of n-way $(n = |EA|)$ relationships that connects all the entities/aggregates in EA.

The following theorem fully characterizes the set of E-R graphs that can be constructed using AppForge.

Theorem 2 Algorithm 2 generates all and only E-R diagrams that satisfy the following constraints:

$$\forall EA \subseteq EN \cup RE \text{ where } |EA| > 2, |M(EA)| \leq \Pi_{e_1, e_2 \in EA}|R(e_1, e_2)|$$

The intuition is that n-way relationship created will depend on $n-1$ 2-way relationships, so the number of n-way relationships that could be created on top of a set of n entities cannot be more than the product of the number of 2-way relationships between any 2 entities in the set.

6. PRELIMINARY USER STUDY OF THE APPFORGE INTERFACE

Given that our primary aim was to support developers who are not experts in databases, we carried out a preliminary user study to test our first interface iteration. The user study consisted of three groups of two people, pairs, who were given three tasks to complete. The tasks were described as follows:

Members of a Yahoo! Group would like to give away unwanted stuff for free. Please create an application that provides the following functionality to members:

1. Post items that they want to give way. Each item includes a name, a description and the original owner (who posted the item).

2. List all the items posted by everyone up to now. Each listing should include the name, description and the owner of the item, and the list of members who have placed a request for the item. The current member can add herself to the requesters list.

3. List the items given away by the current member. Each listing should show the name and description of an item, and the persons requesting the item.

Group 1: Our first pair were two researchers who have advanced degrees in computer science. Both are actively involved in designing, programming and using databases.

Group 2: Our second pair were both researchers trained with advanced degrees in computer science, but neither is a database expert.

Group 3: Our third pair were both experienced computer users. One trained in computer science, but currently in a managerial position with no programming responsibilities; and the other a recruiter familiar with using complex database-backed web applications, but with little formal training in computer science.

Groups were given up to an hour to complete the tasks. Each group was videotaped interacting with the AppForge interface, and all conversation and questions were recorded. The system developers were present to listen to the user interactions, with one of our developers providing advice when needed. Following the trials the development team watched the videos together, made notes and excerpted issues from the sessions, and various redesigns of the interface and considerations for the application were generated and prioritized.

The main finding of our user study was that people who had extensive database experience found mapping the visual presentation we offered to the underlying system structure and logic very easy, completing all three tasks within 20 minutes. Those who were less experienced with database programming did not find the visual presentation quite so intuitive. Group 2 had minor issues with terminology and interface presentation, taking slightly longer to complete the tasks, and asking more questions of us.

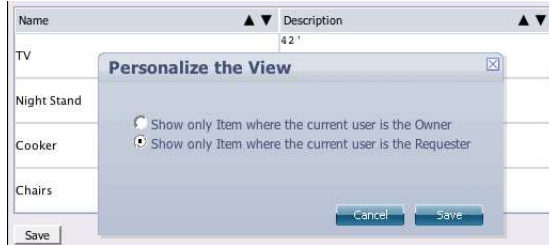**Figure 19: Multiple levels of abstractions for developers**

Figure 20: Personalization

Group 3 were the most challenged, and for us the most interesting of the groups, as they most closely represent our target audience. Therefore we paid special attention to the issues they encountered and have addressed these in our interface redesign.

In particular, Group 3 were confused by the different levels of abstraction that they were required to switch between while developing the application. These levels are illustrated in Figure 19. When creating/adding pages, AppForge developers are the creators of the application, while when viewing and interacting with the pages they have created, they are viewing the pages as members of the Yahoo! Group. In addition, some pages are non-parameterized as for all members (Task 1), and others are parameterized pages where the parameter is the current logged in member (Tasks 2 and 3). Creating parameterized pages proved confusing, with our Group 3 participants struggling to understand the difference between an instance and a variable in place of an instance. We note that these issues are commonly noted in research with novice programmers, and often require careful interface and instruction based scaffolding.

To address this problem, we first redesigned our initial interface to distinguish between the operations AppForge developers can perform as creators and as the intended end users of the applications that they create. We put all the operations for creators in pop-up windows and accessible by right mouse clicks, while all operations for end users were interactive components in the page, e.g. input fields, buttons. To help AppForge users create parameterized pages, we developed the personalization pop-up window, Figure 20, to give suggestions for how to personalize the views (Owner and Requestor are relationships that relate members to items). Our AppForge developers can now choose to view the parameterized pages as a specific user (Figure 21, an instantiated page) which effectively fills in the parameter with a specific user. Using this method, the AppForge developer can see what their intended user

Figure 21: Viewing Pages as a Specific User

would see for the page. Essentially, we developed a WYSIWYG and also a WYSIWTS (what you see is what they see) interface.

In addition to the problem discussed above, we developed a clearer model of containers and views. In our original implementation, we tried to hide the concept of the containers from AppForge users, but our Group 3 participants got confused when multiple forms or views were mapped from the same container. We therefore exposed the notion of containers as collections of data in the visual interface.

Other minor issues exposed during the user study include confusion with the database terminology and poorly delineated interactive areas in the application window (right-clicking on different areas of the application interface revealed different menus). We have addressed these by creating an introductory help panel and a wizard where terms are explained in a Tool-Tips fashion. We also created visual indications of interactive/non-interactive areas, and created consistent menu pop-up and selection.

Having implemented these changes in response to our study results, we are planning a further user study to assess the effect of our modifications on AppForge usability.

7. RELATED WORK

Many CASE tools such as UML [7] and WebML [8] have been developed over the past few decades to help developers build applications. WebML extends UML with links and operations — abstractions tailored specially for web applications — and provides a graphical way of specifying the database schema, application logic and navigational structure of web applications. The main difference between WebML and AppForge is that WebML separates the phases for designing the database schema and designing Web page content. Further, WebML separates the query specification from the output and hence does not provide WYSIWYG interface for creating web pages. In contrast, in AppForge, the database schema is generated implicitly, and changing the queries that populate the page contents will result in instantaneous changes in web pages, which allows users to continuously refine the query as they are constructing it.

There has been a lot of work on graphically creating SQL queries such as Query-By-Example [38] ,Visual Query Builder [6], Visual Query Language [5, 25]. While these approaches hide the SQL syntax from users, they still expose the full schema in terms of relational tables. This is especially confusing when relationships are normalized into tables and users are required to use joins to "stitch" information back together [22]. In contrast, AppForge hide the complexity of the E-R and the relational models, and instead exposes a simple hierarchal Schema Navigation Menu. Another major difference is that AppForge provides a WYSIWYG experience that is tightly integrated with schema generation.

Forms-based approaches [10, 16, 24] for query interface design have been proposed to provide users with visual tools to frame queries and to perform tasks such as database design and view definition. However, like Query-By-Example based methods, they require the users to deal with joins across multiple normalized tables, and they are not truly WYSIWYG, which reduces their usability for our target audience.

In [26], an instantaneous-response interface is proposed to allow users to continuously refine the query as they are typing the initial query. By the time the user has typed out the entire query, the query has been correctly formulated and the results have returned. We share the same philosophy in making the database more usable. Our system extends the same WYSIWYG methodology for query formulation (view creation) to other aspects of creating web applications such as schema creation and form creation.

Many other commercial website creation tools such as Dreamweaver [15]

11

and Frontpage [19] provide a WYSIWYG interface for creating Web pages. However, they are mainly used for creating static Web pages, and the backend application server and database have to be developed separately. Our system take these systems a step further by providing WYSIWYG development not only for Web pages, but also for application logic and backend database development.

Zoho Creator [13], CogHead [11], App2You [2] DabbleDB [14], and Wyaworks [36] provide developers with a form-oriented, drag-and-drop interface to build data driven Web applications. Salesforce [33], QuickBase [32] and Instant Application Platform (IAP) [21] provides extensive solution libraries for developers to customize applications to fit their business requirements. While a few of these systems provide a WYSIWYG environment, and most of them do not require developers to edit the database schema directly, they do not provide an abstraction for complex schemas, including n-way relationships and aggregation, and complex views including joins, aggregations and nesting.

Ning [27] is a website that allows developers to create and customize their own social network portal. While simple customization can be performed using templates, more sophisticated customization involving new entities and relationships requires explicit programming. JotSpot [23] is a related website that extends Wiki [35] with rich structured content, forms [3], and a WYSIWYG interface. However, it is not designed for general Web applications with multiple entities and complex relationships. There are also many other enterprise tools designed to improve developer productivity, e.g., SAP Visual Composer [12] and Oracle Forms [18]. While these tools are more powerful, they are mostly targeted towards professional developers.

8. CONCLUSION AND FUTURE WORK

A growing breed of advanced users are increasingly facing the following dilemma: use a simple graphical tool to build a stripped down version of an application, or go through a steep learning curve and build the more sophisticated application they really want. AppForge tries to provide a solution to this dilemma by expanding the boundary of applications that can be built using a graphical WYSISYG framework. As we have illustrated, AppForge can be used to build fairly sophisticated applications, involving complex schemas and sophisticated page views, without programming or database knowledge.

We have also conducted a small and preliminary user study to evaluate the effectness of AppForge. Based on this study, we have identified some concepts that can be confusing to developers, such as multiple levels of user abstraction. While we have made some changes based on this feedback, fully addressing and evaluating these aspect is an interesting topic for future work. We are also exploring graphical primitives for capturing more sophisticated application logic such as notifications, workflows, and other forms of information passing between pages (e.g., allowing a user to select an event from a list and navigate to a new page that shows all the events that occur on the same day as the selected event).

9. REFERENCES

[1] Serge Abiteboul and Nicole Bidoit. Non first normal form relations: An algebra allowing data restructuring. *J. Comput. Syst. Sci.*, 33(3):361–393, 1986.

[2] App2You. http://app2you.com/site/.

[3] Form Assembly. http://www.formassembly.com/.

[4] Visual Basic. http://msdn2.microsoft.com/en-us/vbasic/default.aspx/.

[5] Francesca Benzi, Dario Maio, and Stefano Rizzi. Visionary: a viewpoint-based visual language for querying relational databases. *Journal of Visual Languages and Computing*, 1999.

[6] Active Query Builder. http://www.activequerybuilder.com/.

[7] Rainer Burkhardt. *UML: Unified Modeling Language*. Addison-Wesley, 1997.

[8] Stefano Ceri, Piero Fraternali, and Aldo Bongio. Web modeling language (webml): a modeling language for designing web sites. In *WWW'00*, 2000.

[9] Peter P. Chen, editor. *Entity-Relationship Approach to Information Modeling and Analysis, Proceedings of the Second International Conference on the Entity-Relationship Approach (ER'81)*. North-Holland, 1983.

[10] J. Choobineh, M. V. Mannino, and V. P. Tseng. A form-based approach for database analysis and design. In *CACM, 35(2)*, 1992.

[11] CogHead. http://www.coghead.com.

[12] SAP NetWeaver Visual Composer. https://www.sdn.sap.com/irj/sdn/visualcomposer.

[13] Zoho Creator. http://creator.zoho.com/index.jsp?serviceurl=

[14] DabbleDB. http://dabbledb.com/.

[15] Adobe Dreamweaver. http://www.adobe.com/products/dreamweaver/.

[16] D. W. Embley. Nfql: The natural forms query language. In *ACM Trans. Database Syst.*, 1989.

[17] Evite. http://www.evite.com/.

[18] Oracle Forms. http://www.oracle.com/.

[19] Microsoft Office FrontPage. http://msdn2.microsoft.com/en-us/office/aa905421.aspx.

[20] Google Gadgets. http://www.google.com/apis/gadgets/.

[21] Interneer. http://www.interneer.com/.

[22] H. V. Jagadish, Adriane Chapman, Aaron Elkiss, Magesh Jayapandian, Yunyao Li, Arnab Nandi, and Cong Yu. Making database systems usable. In *SIGMOD '07*, 2007.

[23] JotSpot/Google. http://www.jot.com.

[24] K. Mitchell and J. Kennedy. Drive: An environment for the organized construction of user-interfaces to databases. In *Interfaces to Databases (IDS-3)*, 1996.

[25] N. Murray, N. Paton, and C. Goble. Kaleidoquery. A visual query language for object databases. In *Advanced Visual Interfaces*, 1998.

[26] A. Nandi and H. V. Jagadish. Assisted querying using instant-response interfaces. In *SIGMOD*, 2007.

[27] Ning. http://www.ning.com/.

[28] Ruby on Rails. http://www.rubyonrails.org/.

[29] Yahoo! Pipe. http://pipes.yahoo.com/pipes/.

[30] Facebook Platform. http://developers.facebook.com/.

[31] Microsoft Popfly. http://www.popfly.ms/.

[32] QuickBase. http://www.quickbase.com/p/home.asp.

[33] Salesforce. http://www.salesforce.com/.

[34] Yahoo! Widgets. http://widgets.yahoo.com/workshop/.

[35] Wiki. http://www.wiki.org/.

[36] WyaWorks. http://www.wyaworks.com/.

[37] Fan Yang, Jayavel Shanmugasundaram, Mirek Riedewald, and Johannes Gehrke. Hilda: A high-level language for data-drivenweb applications. In *ICDE '06*.

[38] M. M. Zloof. Query-by-example: the invocation and definition of tables and forms. In *VLDB*, 1975.

Appendix A:

In this appendix, we will prove Theorem 2. An E-R graph G can be represented as (EN, RE, E), where EN represents the set of entities and RE represents the set of relationships. E represents the set of edges that connects entities with relationships and edges that connects relationships with relationships as in the case of aggregations, i.e., $E \subseteq \{(u,v)|u \in EN \text{ and } v \in RE \text{ or } u \in RE \text{ and } v \in RE\}$. Theorem 2 fully characterizes the set of E-R graphs that can be constructed using AppForge.

Theorem 2 Algorithm 2 generates all and only E-R diagrams that satisfy the following constraints:

$$\forall EA \subseteq EN \cup RE \text{ where } |EA| > 2, |M(EA)| \leq \Pi_{e_1,e_2 \in EA}|R(e_1,e_2)|$$

Proof: We first prove that all schemas constructed by graphical primitives in Algorithm 2 are E-R graphs that satisfy the above constraints inductively. For the rest of this proof, all the line numbers we refer to are for Algorithm 2.

Assume the E-R graph G are constructed using a sequence of graphical primitives in Algorithm 2 and G satisfies the above constraint. After applying another operation as follows, we get a new E-R graph G' and we need to show G' is also a E-R graph satisfying the above constraint. If the operation applied is:

- Create a form/view N. G' is formed by adding entity N to G. For $\forall EA' \subseteq EN' \cup RE', |EA'| > 2$, if $N \notin EA'$, the constraints hold based on the induction hypothesis. Otherwise, $|M(EA')| = |M(EA' - N)|$ and $\Pi_{e_1,e_2 \in EA'}|R(e_1,e_2)| = \Pi_{e_1,e_2 \in EA'-N}|R(e_1,e_2)|$, so the inequation still hold. Therefore G' is still a valid E-R graph with the constraint satisfied.

- Adding a new column to a view. According to Algorithm 2, the operations can result in following updates to the schema based on where the column is added to.

  - A new attribute is added to an entity (Line 11) or to a relationship (Line 21). Both sides of inequation are left unchanged. So it still holds.

  - A new n-way relationship r is added (Lines 19, 20). Let $S(r)$ be the set of entities/aggregations that r connects. Based on Lines 16-18, the new n-way relationship r can be created only when there are no other n-way relationship that connects all entities in the navigationPath and also depends on all the 2-way relationships in the navigationPath. This means $|M(S(r))| < \Pi_{e_1,e_2 \in S(r)}|R(e_1,e_2)|$ holds in G. In G', $M(S(r))$ is increased by one. The inequation in the constraint would still hold. For other $EA' \subseteq EN' \cup RE'$, both sides of the inequation are unchanged thus the constraint holds trivially.

  - A new 2-way relationship is added (Line 25). $\forall EA \in EN \cup RE$, $M(EA)$ is unchanged while the right side can only be increased. Thus the inequation still holds.

  - A new aggregation (Line 31) and a 2-way relationship (Line 32) is added. $\forall EA' \in EN' \cup RE'$, $M(EA')$ is unchanged while the right side can only be increased. Thus the inequation still holds.

Next we prove for every E-R graph satisfying the given constraints can be constructed by a sequence of graphical primitives according to Algorithm 2.

We can use the following sequence of operations to construct a given E-R graph $G = (EN, RE, E)$. First, for every EN, we construct a corresponding view (Lines 1-3). Next, for all $r \in RE$ and r is a 2-way relationship connecting entity A and entity B, we

add a column of type B to the view for A (Line 25). Then for all $r \in RE$ and r is a 2-way relationship connecting entity A and aggregation B, we first create a view for entities and relationships in B and then add a column of A to the view (Lines 31-32). Last, for all $r \in RE$, r is a n-way relationship and r depends on n-1 2-way relationships R, we first create a view for entities and relationships in R and then add an attribute to the view (Lines 19-21). $\quad\square$