

A FRAMEWORK FOR INTERACTIVELY HELPFUL WEB FORMS

MORTEN BOHØJ
Alexandra Institute
Åbogade 34, 8200 Aarhus N, Denmark
morten.bohoej@alexandra.dk

NIELS OLOF BOUVIN
Department of Computer Science, Aarhus University
Åbogade 34, 8200 Aarhus N, Denmark
bouvin@cs.au.dk

HENRIK GAMMELMARK
Dat1
Åbogade 15, 8200 Aarhus N, Denmark
gammelmark@dat1.dk

Received November 3, 2011

Revised January 9, 2012

AdapForms is a framework for adaptive forms, consisting of a form definition language designating structure and constraints upon acceptable input, and a software architecture that continuously validates and adapts the form presented to the user. The validation is performed server-side, which enables the use of complex business logic without duplicate code. Thus, the state of the form is kept persistently at the server, and the system ensures that all submitted forms are valid and type safe.

Keywords: Adaptivity, forms, framework

Communicated by: O. Diaz & S. Auer

1 Introduction

The World Wide Web has gained its remarkable success not only because it is a (relatively) simple and scalable publishing system, but also because it offers easy access for users to add new content or interact in other ways by filling on-line forms. This is done many times daily in search fields, e-shops, forum postings, application forms, and so on.

While many of these text fields and forms are straightforward, some are more challenging for the user to correctly interpret and fill. Likewise, an essential part of form handling at the developer's end is the checking and validation of the input, before it can be added to a database, or used as basis for other processing. These validation steps range from trivial tests that can be evaluated using e.g., string length over more complex rules checked with regular expressions, and finally input that can only be checked against the “business logic” of the particular application. In practical terms, this can range from checking whether a family name has been entered, whether an amount entered consists only of digits, or whether a desired vacation period is actually available.

As many administrative processes move on-line, the associated forms follow. Forms can be simple or complex with opaque and non-trivial interrelationships that can be difficult to communicate effectively to the user. When is a form filled out correctly or sufficiently; and how do the choices made in one field affect the rest of the form? Which parts of the form are relevant to the user's situation, and what can be safely ignored? Online forms will usually, at least initially, be derived from their paper-based counterparts, as this will most closely match the established processing of the data. While neatly fitting in with existing bureaucratic practices, this ignores the opportunities for interactively helpful forms.

The notion of an interactive form can be broken down into the following dimensions:

Validation Ensure that the entered value is valid in the given context

Adaptation Change the form depending on the user's input in order to minimise cognitive overhead

Persistence Ensure that data entered is not lost

We describe in this paper AdapForms, a general solution to handling forms in a fashion that enables developers to clearly designate accepted types of input, reuse existing templates (e.g., sub-forms for postal addresses), use arbitrarily complex validation rules, allow users to resume filling out complex forms at a later point, and communicate the state of the form to the user unobtrusively. An open source demonstrator of the framework is available online^a and an exhaustive description can be found in [1].

The paper is structured as follows: Our system is described from a functional point of view in Section 2 and an architectural one in Section 3, the system has been evaluated as described in Section 5, we describe related work in Section 6, and the paper is concluded in Section 7.

2 The AdapForms Framework

Validation of user submitted data on the Web is usually a two-step process—input is validated in the browser (usually through custom JavaScript code), and validated again when submitted to the server. Client-side validation has certain advantages, especially responsiveness as no network latency is involved. However, client-side validation cannot replace server-side validation, as the input sent back to the server may be subverted or corrupted accidentally or maliciously. Thus, server-side validation is still necessary—at the very least to thwart SQL injections and similar attacks.

AdapForms is a server-centric framework, which allows direct code access to the form and data structures, described further in Section 3. This also means that all validation is done on the server instead of the client, which provides a higher degree of safety regarding validity of the input. An advantage of having the state constantly up-to-date server-side, is that it allows direct transparent integration with the host application instead of relying on additional server calls for adaptation and validation, as described below. Furthermore, business logic is not exposed in JavaScript accessible in the browser, but is kept on the server. To help illustrate the features of the framework an example web UI is used throughout this paper.

^a<http://adapforms.gammelmark.eu/>

Job application

Applicant information

Social Security Number Social Security Number: Required

Full name

First name

Middle names

Last name

Email address

Phone number Phone number: Incorrect format

Monthly income Monthly income: Required

Date of birth (dd/mm/yyyy)

You are under the age of 18. Please name your parent or legal guardian

Parent or legal guardian

First name First name: Required

Middle names

Last name Last name: Required

Country Country: Required

Form contains invalid values

Fig. 1. An AdapForms form. Yellow markers signify missing data, red markers invalid data. The Parental form elements has appeared because the user is a legal minor.

Job application

Applicant information

Social Security Number Social Security Number: Required

Full name

First name

Middle names

Last name

Email address

Phone number

Monthly income Monthly income: Required

Date of birth (dd/mm/yyyy)

Country Country: Required

Civil status

Marital status Marital status: Required

Form is not yet completed

Fig. 2. The applicant is of legal age, and thus Civil status is relevant.

```

1 <adapforms>
2 <include file="templates.xml" />
3 <form title="Example Form">
4 <defaults readonly="false" write-roles="applicant" />
5 <template name="anAddress">
6 <group>
7 <text id="street" label="Street">
8 <validator criteria="lower-case(.) = ." message="Must be all lower-case" />
9 </text>
10 <text id="city" label="City" pattern="[a-z]*" />
11 <integer id="zip" label="ZIP/postal code" />
12 </group>
13 </template>
14 <group id="info" label="Personal information">
15 <use template="anAddress" id="address" label="Current address" />
16 <integer id="age" label="Age" minValue="0" />
17 <toggle id="parental" label="Parental accept" relevant="../age &lt;= 18"
18 write-roles="parents" />
19 <text id="parentName" label="Name of parent" relevant="../parental/@relevant" />
20 </group>
21 <bean id="payment" label="Payment details" type="myapp.model.PaymentInfo" />
22 <repeat id="kids" label="Children" entryLabel="Child" minRepeats="0">
23 <text id="name" label="Childs name" />
24 </repeat>
25 </form>
26 </adapforms>

```

Fig. 3. An AdapForms Sample Form

```

1 <adapforms-module>
2 <template name="coAddress">
3 <group>
4 <text id="co" label="Co name" />
5 <text id="street" label="Street">
6 <validator criteria="lower-case(.) = ." message="Must be all lower-case" />
7 </text>
8 <text id="city" label="City" pattern="[a-z]*" />
9 <integer id="zip" label="ZIP/postal code" />
10 </group>
11 </template>
12 </adapforms>

```

Fig. 4. An AdapForms Module Sample

Two screenshots of the UI are shown in Figure 1 and 2 and are used as reference in Section 2.2 and Section 2.3. The web UI implementation is described in further detail in Section 4.

2.1 Defining Forms

The form structure in AdapForms is defined in a XML document, which is parsed by the host application and forms the basis for the internal representation of the form. The form structure is defined by a XML Schema. The schema is too large to reproduce here, but it is available online^b. The schema defines the types of elements allowed in the framework and structures them in two types of allowed overall documents. There is the `adapforms` and the `adapforms-module` documents. The Figure 3 is an example of the `adapforms` document, and the `adapforms-module` defines the format of files included in `adapforms`. A sample `adapforms-module` is shown in Figure 4 and it is included in line 2 of Figure 3.

The sample in Figure 3 is meant to demonstrate features rather than show a typical form. The XML form can contain the following elements representing different types of data:

Label Read-only text label.

^b<http://adapforms.gammelmark.eu/ns/adapforms.xsd>

Text Regular text element, single- or multi-line.

Secret Text field for masked input.

Integer & Decimal Ranged or unranged integer or decimal respectively, equivalent to Java `long` and Java `double`.

Toggle True or false.

Choice Single selection among fixed, mutually exclusive, choices.

Multi-choice Multiple selections among fixed choices.

Date Single date selection.

HelpText Text to help the user fill out the form.

Group Logical grouping of elements in the form. **Group** can contain the elements described above and group them together. This grouping allows for an easier overview of the form and allows for making an entire section e.g., **relevant** or **required**. An example of the **group** element is shown in line 14 of Figure 3.

Repeat Repeats are a collection of form elements that can be repeated in a list-like manner. This can e.g., be relevant for allowing users to add a variable number of children, without knowing how many elements to create beforehand. The number of repeats allow can be limited or unlimited and both as a required minimum or limited maximum. A small example of a **repeat** element is shown in line 22 of Figure 3.

Bean A bean is a **group** generated from a JavaBean—this is elaborated upon below.

Each of these form elements has a number of possible attributes. Some of these attributes are specific to a particular form element, while others are shared among all elements. Most are optional and only **id** and **label** are required. The common attribute set is as follows:

id A locally unique ID. The ID must be unique among siblings in the XML tree. The ID is required for all form elements.

label A textual label explaining the element’s purpose. The label attribute is optional for the *group* and *bean* elements and required for the rest.

relevant Boolean value indicating the relevance for the user. Relevance can e.g., be used to hide irrelevant elements from the user.

required Indicates whether or not the element is required.

read-roles & write-roles List of user roles allowed to read or write, respectively, the specific element. The user roles are handled by the host application when loading the form. The use of roles is not required. User roles may be used to allow more rights to some users, such as administrators.

readonly Indicates that no user is allowed to change the element value.

uiflags Text string which may hold UI specific information.

If an optional attribute is not present in an element, the value is inherited from the parent node. If no parent node is present, the value is taken from the nearest **defaults** element (see line 4 in Figure 3). The sole purpose of the **defaults** element is to provide default values, and if no default value is found, the framework defaults apply. When designing forms, one might experience defining the same element structure more than once. This could e.g., be elements of an address section. To allow for reuse of such reoccurring structures, the framework has a template option. Templates can be defined within the same XML document as the rest of the form, as in line 5–13 of Figure 3, with the template used in line 15, or they can be placed in a separate XML document and included in the form using the **include** element, as in line 2 in Figure 3. This supports the reuse of common structures, such as addresses, across forms. Similar to the use of templates, the framework allows for the use of JavaBeans. A bean is included in the form using the **bean** element as shown in line 21 of Figure 3. The use of JavaBeans results in a group node in the form structure, which contains the public available properties. The framework supports properties with the following Java types: **String** (Represented as **Text** element), **int** or **long** (**Integer** element), **float** or **double** (**Decimal** element), **boolean** (**Toggle** element), **java.util.Date** (**Date** element) and **enum** types (**Choice** element). The framework does not currently support beans within beans. The individual properties can be annotated in the Java class using **@AdapFormsProperty** to support e.g., renaming the label or designating required attributes. Annotating the bean class itself defines defaults for all properties.

2.2 *Adaptation*

One of the core features of AdapForms is the ability to adapt the form according to user input. These adaptations could be changing/setting element values based on input in other fields, or hiding/showing parts of the form that may be irrelevant/relevant based on user input. An example of the latter could be to hide the workplace address field, if the user indicates she is currently unemployed. Figure 1 and Figure 2 shows an example of this adaptation, where in Figure 1 the date of birth is set to 2009, thus giving an age under 18 and therefore requiring a parental section. In Figure 2, the age is above 18 and the parental section is replaced by the question of marital status. By hiding fields not relevant to the user, we eliminate some potential confusion for the user, so she can concentrate on the fields that are relevant to her. Likewise, by filling out some fields based on either entered information or previous knowledge, we minimise the typing needed by the user, as she can merely verify the information.

2.2.1 *XPath*

One way to perform form adaptation is to do it directly in the XML document using AdapForms' integration of XPath 2.0[2] expressions. An example of the use of XPath expressions is shown in line 17 of Figure 3. Here the relevance of the toggle element **parental** is dependent on the value of the sibling element **age** being less than or equal to 18. XPath is very expressive and has the advantage of placing the validation rules close to the adaptation targets. As XPath is often used in connection with XML, the developer is also more likely to be familiar with it than having to learn a completely new language. While XPath is powerful, it has some limitations in this setting, e.g., it can only affect a single state parameter as it only has

one output value. It can also only look at the information typed in the form and not access previously known information.

2.2.2 Form Hooks

The adaptations mentioned above are pre-coded adaptations meaning that they are based on pre-coded rules and anything learnt from other users filling out the same form can not be taken into account. This is where the AdapForms goes beyond using XForms and XPath. AdapForms supports more complex adaptations through the use of form hooks. A hook is defined by the Java interface `FormHook` which defines three methods `onValueChange`, `onRepeatEntryAdd` and `onRepeatEntryRemove`. Each hook is registered with a specific form path, described in Section 3.2, before the form instance is initialised. The same hook may be registered with more than one form path or no path, if, for some reason, the developer wants a single hook to receive all change events generated. The hook is notified through the aforementioned methods, when a change occurs at the registered path. `onValueChange` is called on value changes, whereas `onRepeatEntryAdd` and `onRepeatEntryRemove` is called for changes on repeat entries. Changes on form-level, such as submission of the form, is handled through the `InstanceCallback` interface, and here only a single callback can be registered to a form instance. Because hooks are created using Java, it allows for more options with regards to adaptation, as a hook can manipulate several values or look up previously entered information in a database. Hooks also make it possible to base the adaptations on arbitrary business logic, or more complex computations using machine learning, if these solutions are Java based. Machine learning would e.g., allow for the form to change a default choice in a multiple choice to reflect the choice of the majority of users.

2.3 Validation

When receiving data through web forms, it is essential to validate the input prior to processing. This validation can be anything from checking whether a field is a valid date to checking that a supplied user name is known. Often frameworks, such as PowerForms[3], perform this validation client-side in order to provide quick user feedback, but this means that the same validation must be done at the server as well to make sure no invalid data slips through. AdapForms performs the validation server-side, while continuously providing feedback to the client. This adds some time for validation, but ensures that the form instance on the server holds the correct information. The validation is done asynchronously, so the latency is usually not noticed. An overview of the available validation in the framework is shown in Table 1 and explained in more detail below.

The framework operates with three different kinds of problem definitions: **Error**, **Warning** or **Required**. **Error** indicates that the input is invalid. This could be entering a string where a number is expected. **Warning** indicates that while the input is not invalid it may be problematic. This could e.g., be if the user have put the year 1010, where it is more likely that he meant 2010. **Required** indicates that an element that is required is missing some input. These problem classifications lead to a form instance being able to be in the following states:

Uninitialised The form instance has been created, but is not yet initialised and can thus not handle input

Validation	Explanation
Required	By setting the <code>required</code> attribute to true, the existence of a value can be ensured
Type check	The framework will check e.g., that <code>Integer</code> elements only holds integers
Range check	Specifying a range for integer elements helps validate the input
Format check	<code>Text</code> elements allows for regular expressions to be defined
Validators	Validators offers a more complex validation through the evaluation of XPath expressions
Form hooks	Form hooks are the most flexible way of validation and can incorporate complex business logic if needed

Table 1. Overview of framework validations

Initialising The initialising process is ongoing

Invalid The form contains one or more problems. In Figure 1 this is highlighted by displaying an error message next to the submission button.

Incomplete The form is valid except for errors of the `Required` kind. Figure 2 shows how this could be shown to the user by a warning message next to the submission button.

IncompleteWithWarnings The form may hold `Required` or `Warning` errors

ValidWithWarnings The form only holds errors of the `Warning` type.

Valid The form holds no errors

On submission of the form, the host application may choose to accept or reject the form based on the form status. The host application can use the form status to perform different actions based on the status, e.g., saving incomplete forms for later completion instead of submitting it.

The validation in `AdapForms` takes place on three different levels: basic framework validation, validation rules, and host application. The final form status is based on the union of the errors on these three levels.

The basic framework validation includes:

Required field validation Checking that required fields contain a value. This could be shown as in both Figure 1 and Figure 2 with yellow warnings.

Type check Checks that e.g., integer elements only holds integers or date elements contains only dates.

Range check Integer elements may be upper- or lower-bounded, as shown in line 16 of Figure 3, as can the length of text elements

Format check Text elements may define regular expressions the input must uphold as in line 10 of Figure 3

This built-in validation will trigger when the form is loaded and when values are changed. The host application may not override these mechanisms to avoid inconsistencies. The XML form itself may also contain more complex validation rules by using XPath expressions. These can be placed in `validator` elements, which is then placed as children of the element it validates. Each `validator` element has a criteria, which contains an XPath expression, and a message to give the user in case of an error. An example of a validator is shown in line 8 of Figure 3, where it checks that the street is written only in lower-case letters. Validators are mostly used to validate values, but are not limited to this. Validators are evaluated on every change of values in the form, and not just on change of the element the validator is attached to. As a third option of validation, the host application may set or remove validation rules for individual form elements. This is most easily done using `FormHooks`, attached to the element that needs checking. Hooks would typically be used to perform more elaborate validation that is not possible with the built-in validation, such as comparing input to values in a database. The host application may decide to perform validation any time and is not limited to the use of hooks.

3 Architecture

The framework is built using the Model-View-Controller pattern, where the core framework essentially comprises the Model and Controller by providing the data structures as well as the logic making the model dynamic. The View part is UI-specific and thus highly coupled to the display- and interaction-capabilities of the technology chosen. How to specifically display data to, and interact with the user, is not the primary concern of the framework. Instead the focus is on determining what to display and what to do with the input gathered. An important goal of the architecture is to have the host application manipulate the form and interact with the user exclusively via local proxy interfaces, regardless of how the UI exposes the form and where it is physically being executed.

3.1 Form Life Cycle

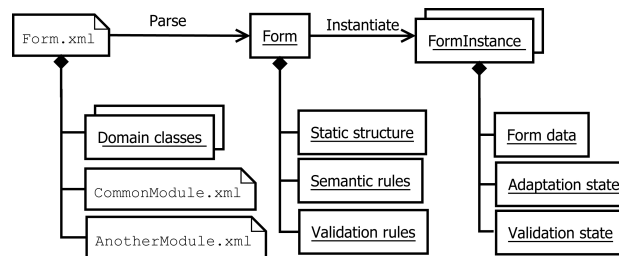


Fig. 5. Parsing and instantiation of a form

The software architecture of AdapForms can best be understood by considering the life cycle of an adaptive form, so that the various phases can be related to concrete software elements. The Instantiation and Initialisation phases can be thought of as a single phase in the abstract sense.

Definition The form syntax and semantics are defined in a XML file external to the frame-

work.

Parsing and loading The form is parsed and its dependent definitions and objects are parsed, ending up in a single **Form** entity representing the static form, with all its form elements (explicitly specified, or inferred from templates, beans etc.). The form structure may be viewed as an abstract structured tree, where each node is a form element. See Figure 5 for a visual illustration of the process.

Instantiation A concrete instance of the form is created. Using the static form as a basis, the instance will hold current values, validation status etc. of all nodes in the static form structure. The **FormInstance** is essentially holding the collective state of the instantiated form. Multiple users can share the same **Form**, but they will each have a unique **FormInstance**.

Initialisation The form instance is initialised with default values and, if available, stored form data from a previous session. The role of the current user (if any) is also selected at this point. All relevant semantic rules (i.e., XPath and hooks) are triggered, causing the form to be pre-adapted as much as possible before the user encounters the form. The initialisation is explained in more detail below.

User interaction Input from the user (or the host application) is passed to the instantiated form, and the instance may respond with a list of validation errors and form adaptations, triggered by the loaded semantics. Each user interaction triggers a new adaptation cycle.

Adaptations may be scripted with XPath expressions and queries. To facilitate complex domain-specific validation and reasoning, the host application may register hooks with the instance, which are triggered whenever a given form element changes its value or state.

Submission At submission, the instance is checked for any consistency problems and validation errors, and the event is then passed (along with the instance) to a callback handler, specified by the host application.

3.1.1 Initialisation Sequence

During the initialisation process the internal state tree of the form instance is constructed and the default values are set/calculated along with the state parameters. In detail the sequence is as follows:

1. **Create all state tree nodes.** At creation, each node is initialized with the framework defaults, to make sure they have a value for each state parameter (such as `ReadOnly`, `Relevant`, `Required`, etc). If the form definition supplies a static parameter value (which is often the case), this value will be used instead. The rationale is that this presumably closely resembles the truth, even if they may be changed later by form hooks.
2. **Input values.** The default element values (or values previously entered and now re-loaded) are inserted into the form sequentially. In the process, only the basic validation code is triggered for each element, updating the validation status for the element.

3. **Evaluate dynamic state parameters.** All dynamic state parameters are evaluated and any complex validation rules are triggered. Validation is described previously in Section 2.3. The rationale for waiting with these evaluations is, that the expression will often depend on the form values entered, and thus should be re-evaluated later anyway.
4. **Execute relevant hooks.** For each element that has been assigned a value in the previous step, any associated hooks are triggered in registration order. At this point in time, we have initialized as many values as possible, providing the best possible premise for the hook to execute. Notice that hooks yet to be called, may still change some of the state the current hook is reacting upon, and the problem has thus not been completely eliminated. Whenever a hook has performed changes to the instance, the dynamic state parameters are re-evaluated to keep them updated. Hooks may choose to respond differently whether the form is being initialised or is performing a normal adaptation cycle.
5. **Repeat entries are initialized.** If the form contents were loaded from previous stored data, any repeat entries that were entered, are now initialized by recursively running this procedure on each entry.

3.2 *Form Paths*

A recurring scheme is the use of *form paths*. A path uniquely identifies a form element, and is used throughout the framework for multiple purposes. Examples include look-up of form elements (form structure tree), addressing nodes in the internal state tree, defining semantic rules etc. A path resembles the hierarchical file names, and is heavily inspired by XPath and by h-maps [4]. An example form path looks like this: `/applicant/spouse/name/first`. Paths may be fully qualified (indicated by a leading slash) or relative to the context they are used in. For the most part it is possible to create a 1-to-1 mapping between a form element in the static form with the representation of that field in the instantiated form, simplifying things. However, this is not possible in the case of repeated sub-forms, where the user may specify any number of occurrences of a given data type. In these cases, an extra annotation level to the instance paths is introduced, so the paths become `/applicant/jobs[1]/company` and `/applicant/jobs[2]/company` for the first and second repeat, respectively. This syntax is also similar to XPath. When making references to the static form definition, the numbering scheme is omitted. When defining semantic rules, both versions may be used, depending on whether it is a general rule, or a rule specific to a single repeated entry.

3.3 *Element State Tree*

The form instance holds state for each form element present in the instance. The actual state is held in a hierarchical data structure with form paths as keys.

This is also inspired by how FormGen [5] stores state, but with the major difference that the host application need not be implemented against specifically generated tree classes. Thus, any host application can load any adaptive form and process it at run-time, although it may of course not always be able to do anything sensible with the entered data besides saving or forwarding it. The trade-off of this approach is that the API exposed to the host application must be more general and therefore may be less natural to use, than if the actual tree was generated specifically for the domain.

At any given time, the following state parameters are available for all elements: The unique form path and a **Relevant** state parameter, specifying whether the element is currently visible to the user. In addition, any form element capable of holding a value, also holds the following state: Current value, **ReadOnly** state parameter, **Required** state parameter (for validation), list of associated hooks, and the validation state. In Figure 5, the state tree is represented by the Form data, Adaptation state, and Validation state collectively.

All of the mentioned state entries can be read and manipulated by the host application, with the exceptions stated above. Changing the state will in most cases trigger an adaptation of the form instance. Access to reading and manipulating the state of a given element is exposed by the **ElementState** interface, allowing direct API manipulation.

As mentioned above, the tree is partially inspired by hierarchical maps [4]; a data structure intended for multi-purpose storage of data and state of varying complexity using an ad-hoc tree-like structure, that changes over time. Each piece of information is stored in its own leaf node, and is identified by a unique path, where each part of the path represents a level in the tree, starting from the root. In addition to holding values, a leaf may have registered one or more message handlers, which are invoked in order, when a message is destined for the corresponding tree path. Due to the hierarchical nature of the data structure, serialization to and from XML is straightforward, thus allowing the host application to easily save an entire instance in textual form to disk, and reload it at a later time, whether the form has been completed or not.

Furthermore, it allows XML-based applications easier integration with the framework, as the XML nodes can be moved directly between application and framework (and thereby indirectly the user interface). Because the developer is free to design the form structure (and thus the data model) as needed, within the boundaries of the form language, in many cases no XML transformation is needed, simplifying the integration. In practice this access is provided by superimposing the Java DOM API interfaces upon the state tree, essentially making the state tree behave like an XML document. This also means, that it is possible for a **FormHook** to use XPath expressions of its own, thus combining the strength of the two otherwise different mechanisms.

In contrast to the original h-map, the set of possible nodes is restricted to match the paths dictated by the static form structure, with the exception of **Repeats**, as indicated above. Thus, every update to the structure is checked against the form structure. This serves several purposes:

- Type safety is guaranteed, as only data of the type dictated by the corresponding form element can be written. When retrieving data from the structure, there is thus no need for type checking, and any type errors are caught early. Automatic data type conversion is applied, where applicable.
- Upon querying or updating a value in the structure, the action is checked against the role permissions of the element. This guarantees that no data is accidentally overwritten by a user not having the correct role.
- It is guaranteed that no “phantom” values are inserted, that are not logically linked to the form due to path typos, logic errors, or similar.

Although most of this should already be implicitly guaranteed by the UI layer implementation, it serves as a final centralised sanity check and guards against both programming errors and attempts to craft malicious input. Furthermore, the host application may manipulate parts of the data structure directly, and once again this therefore guarantees a sound state of the structure at all times.

While each state tree node has its own set of parameters, these may be overruled at runtime due to the semantic inheritance. Consider the case of two nested form elements; the logical group `/foo` and its contained text element `/foo/bar`. If the `relevant` flag of the former element is changed to `false` this value will propagate to the latter as well.

Mechanisms are provided for the form semantics logic to query and modify both the stored (possibly overruled) and the active value, but regardless of how the parameters are manipulated, the inheritance overrule remains in effect. Inheritance rules are out of scope for this article, but are described in [1].

3.4 Adaptation Cycles

All adaptations within a form instance are initiated by a user action. This means that the form cannot suddenly change in front of the user, unless he or she has triggered the adaptation by changing one or more values in the form, submitting it, or have otherwise actively performed an action.

The term adaptation cycle denotes a complete cycle from the reception of user input (value change, or submission request), and until the user experiences the adaptations. Although the adaptation cycle is always triggered by a single user action, the triggering of form hooks or semantic rules may cause a cascade of adaptations to take place before the cycle ends. The host application can perform changes to the form instance at any given time. However, these adaptations will be accumulated in a buffer, and will only be presented to the user at the next adaptation cycle to avoid confusion.

An adaptation is conceptually anything from a value change, validation status change, relevance change, read-only status change to more volatile events such as displaying a message to the user and more UI-specific adaptations, such as redirecting a browser user to another website, closing a window, or playing audio.

Figure 6 illustrates the interaction sequences and the role of the iterative adaptation cycles within a control flow context. Each iteration of the outer loop represents a complete adaptation cycle. A single user-initiated adaptation may trigger further adaptations and form hooks.

3.5 UI-specific View Component

The UI layer responsible for interfacing between the user and the core framework is only defined abstractly in the general architecture, as a component with the following primary responsibilities:

- Presenting a pre-adapted form instance to the user.
- Populate the form elements with data and sending it to the core component.
- Adapting the presented form, as indicated by the core component.

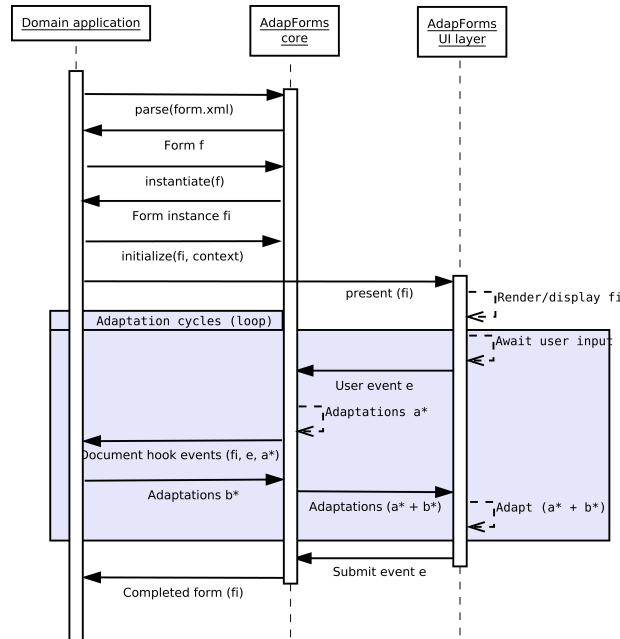


Fig. 6. Interaction of the major components

- Highlighting validation errors within the presented form, based on feedback from the core component.
- Submitting the form, when the user deems it to be completed.

Note, that we are not restricted to graphical interfaces. Any interface, e.g., audible or tactile, could theoretically be implemented, although the mapping is not necessarily straightforward. A proof-of-concept implementation for a XHTML/ Ajax web platform has been implemented, as mentioned previously and described in detail below.

4 Web UI Implementation

The AdapForms framework is focused on the underlying handling of form structure and data and not on the user interface the user is presented with. The framework is independent of the UI the user is interacting with and the interaction between the UI and the core framework is defined abstractly as described later. Seeing how more and more applications migrate from desktop applications to Internet applications, e.g., Google Apps and Office 365, the primary use of AdapForms is likely to be web-based. The ideas behind AdapForms are also inspired by different web technologies, e.g., the use XPath [2] for validation and adaption, and the adaption and feedback presented by PowerForms [3]. This made the choice easy for when choosing a platform for an example UI implementation.

The example implementation was used throughout development for testing features and for the user evaluation presented in Section 5. The adaptive form is presented as a XHTML 1.1 input form in a table layout. The implementation is based on Java servlets, JSP pages, and

JavaScript and thus requires a Java enabled web server, e.g., Apache Tomcat. The JavaScript processes the form state it is given and thus no validation rules need to be written using JavaScript. This also means that the no changes needs to be made to the JavaScript for each different form.

At the time of development HTML5 wasn't available, but had it been available it could have substitute some of the client UI implementation, but not the core framework. HTML5 could have provided a richer user experience by using some of the form elements available, but a HTML5 form does not have native adaptation and would still need server-side validation of the data submitted. Screenshots of the sample web UI are shown in Figure 1 and 2. The UI uses colours to inform the user of errors and warnings with yellow for warnings and red for errors. Examples of which information these errors and warnings can inform the user about is given in Section 2.3. If the web implementation is used the styling can be change by providing a customised CSS stylesheet.

Figure 7 illustrates the flow of how users interact with the form and the communication with the core framework. Initially the requests the web page containing the form from the host application. The host application instantiates the framework, which parses the requested form. The XHTML page containing the form and JavaScript is returned to the browser and presented to the user. The user enters value in a field. The change is sent to the server using AJAX. This change triggers an adaptation cycle (described in Section 3.4), and the adaptations and validations are returned to the browser. Here JavaScript adapt the UI based on the changes and the user sees the form change to reflect the current state. This process repeats as the user fills out the complete form. When the user submits the form a submission event is sent to the server using AJAX, the submit is processed and the user is redirected to the submission complete page.

The web UI is not part of the core framework and is only meant as a demonstration. The benefit of this separation from the UI allows for greater flexibility. The same form definition can be used across several platforms. E.g., a smartphone application could be created to access the same form as the web interface. The core framework would run server side and simply react to change events received and return the XML representation of the form with adaptations and validations. Implementing different UIs also allows for changing the rendering of how feedback is given to the users.

5 Evaluation

A framework such as AdapForms must be viewed from at least three perspectives, when considering its fitness as a tool, namely the perspectives of the developer (creating the forms), the form filler (the user filling in information), and the form handler (responsible for handling the input). We have conducted some initial user feedback studies on each of the three perspectives.

5.1 Form Filler Feedback

As mentioned above, a prototype user interface for AdapForms has been implemented The form filler feedback was conducted on a version prior to the demonstrator web UI included in the current framework version. Changes were only made to the demonstrator UI and not to the handling framework. The screen shots in Figure 1 and 2 are from the new version. In

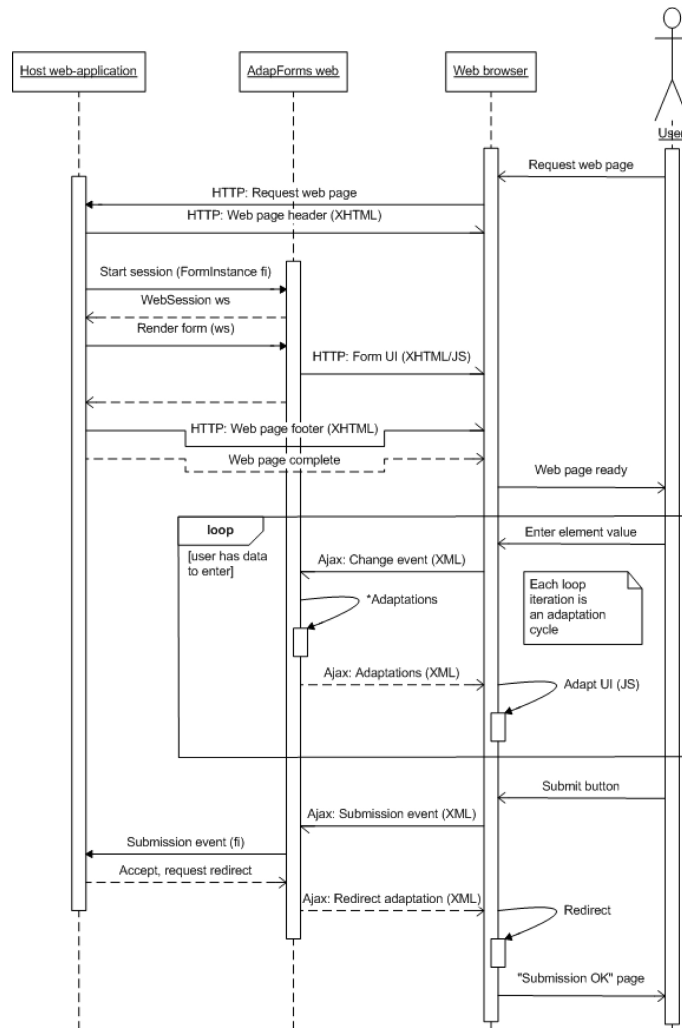


Fig. 7. Interaction model for the web-based UI implementation

the version used for the test, the framework indicated correctly filled fields with green icons, and indicated errors and warnings with red and yellow icons respectively. Error and warning messages were only available when hovering over the icons next to the field.

In order to get feedback from form fillers, the web UI prototype was used for experimentation. Three users were presented with a form produced with the AdapForms framework, using the example UI, and a similar form implemented in standard static HTML. In order to do a reasonable comparison, the static form performed the same validation on submission as the adaptive form did during completion and gave the same messages at the top of the document.

One of the major issues discovered with the adaptive form was with the validation feedback given to the user. The coloured icons indicating validation status next to the fields were largely ignored by the users, and the users were puzzled as to why they could not submit the form. When the users noticed the icons, it was unclear to them what the different icons signified. The fact that the error messages were available by hovering over the icons was not immediately clear to the users and one user tried to click the icon. On the other hand, the users had no problems with the validation error messages shown at the top of the traditional web form after submission. This is probably because this seems to be the norm for displaying error messages, but it would also indicate that error messages should be displayed clearly in order for the user to notice and react to them. This feedback has resulted in changes in the web implementation to display the error messages directly beneath the field the error concerns, as can be seen in Figure 1. The green icons indicating correct fields have also been removed to avoid too much visual overload.

Another issue identified during the user testing was how to indicate the expected format. The tested version did not indicate the expected format, which especially caused problems with date fields, as dates can be formulated in a number of different ways. This observation resulted in showing the expected format next to the date field in the current version.

An important observation during the user testing was how the form was filled out top to bottom, under the assumption that the fields already filled out would not change. This is as expected, but is important to remember when defining forms and avoid creating rules that will change fields already filled out “on the way down”. The lessons from this evaluation are mainly about the construction of the UI, which is not directly the focus of the framework. The evaluation does however show the usefulness of the idea behind the framework and the use of adaptive forms. A majority of users in the evaluation liked the adaptive forms after having adjusted themselves to the different mindset of getting instant feedback and not having to wait for submission validation.

5.2 Developer Feedback

A web developer was given the framework including the UI web implementation and was asked to develop a web form.

The developer found the learning curve of using the web implementation quite steep due to the many Ajax HTTP calls exchanged between the client and the server. It needed a different way of thinking about control in the servlet environment and a change to letting the hooks do much of the work. This would indicate that the control flow should be made more clear in the future along with some more detailed tutorials.

The work with defining the forms using XML worked well and the structure was quickly grasped. The main complaint was the lack of a more convenient way of verifying the correctness of the form, prior to instantiation. This could be solved by including a tool in the framework that would control the form against the XML Schema available and at the same time ensuring that referenced JavaBeans actually exist.

As for the use of the validation rules, the simple rules were easily understood and after looking into XPath, which was unknown to the developer beforehand, this was also easily understood. One complaint was however how this would fit into already defined business logic without having to write the same thing twice. One way of avoiding defining business rules more than once could be to rely mainly on JavaBeans and then reuse these beans throughout.

The general response to the use of the framework, was that after passing the initial hurdle of learning the work flow of the framework, it can save a developer a lot of time, compared to having to writing the validation of client and server side by hand.

One of the authors (not the original developer of AdapForms) have utilised the AdapForms framework in the context of the eGov+ project, as reported in [6, 7]. Our experience with the framework closely resembles those reported above.

5.3 Form Handler Feedback

The framework has also been used in connection with the parental leave case of the eGov+ project and included in the prototype developed in connection with this case and described in [6, 7].

In connection with the development of this prototype, the concept of adaptive documents was introduced to the caseworkers set to handle the incoming forms. They were generally positive about adaptive documents and some effort was made trying to annotate existing physical documents, in order to establish relationships between fields.

The caseworkers found it challenging to analyse the existing documents and identify fields which were e.g., dependent on former input or pairs of fields depended on each other. This shows that the change from creating one static form to handle all potential users and their input and then designing a form which is capable of adapting based on input such as gender or user role is non-trivial. This challenge is something that has to be worked out between the developers in charge of developing the form and the people with the domain knowledge specific to the form in question.

6 Related Work

Form handling is an inherent part of information processing, and predates computers considerably. Automatic, interactive form validation is much younger, but can certainly be traced back to systems such as the 1983 COUSIN interface building system [8], wherein the syntactic and semantic values of form fields could be specified to ease data entry by providing e.g., error correction, or appropriate help.

FormGen [5] is a Java GUI tool, which can generate a dynamic form based on data structured as a context-free grammar (CFG). FormGen uses this CFG to generate Java classes for both data model and GUI components. A disadvantage here is the need to regenerate the Java classes every time a change is made to the CFG unlike our AdapForms which, with




Fig. 8. Annotating paper documents for translation to adaptive documents

its XML base, is evaluated to create the output each time the form is requested. The forms are in AdapForms XML based and evaluated to create the output form each time the form is requested. This gives a more dynamic options for changing forms with the same underlying backend. As in AdapForms, the data model in FormGen is a tree, and as the user adds data to the form, new nodes are added to the tree.

Dynamic Forms [9] is an example of a validating form generator using the Form Descriptive Language (FDL). FDL can be edited with an interactive editor written in FDL. When users are interacting with the finished form, Dynamic Forms provides feedback of completion and validity to the user by colour coding. The idea of colour coding is also known from other frameworks and is also used in the example implementation of the web user interface for AdapForms. FormGen was designed with desktop applications in mind, whereas AdapForms is designed to be cross platform. The form controller of AdapForms is separated from the user interface, so that any implementation can be put on top. An example user interface implementation for web is included in the framework.

XForms [10] is a W3C alternative to the common HTML form. XForms is, as AdapForms, XML based, and differs from regular HTML forms by separating data and markup. Thus, one form can be specified and used several different places with different markup. Support for repeated and nested elements exists, but the recursion level is limited, thus arbitrary tree structures are not possible [11, chapter 3] XForms extends the standard HTML form with validation and adaptation. The validation can be performed using XML Schema. Types are defined using XML Schema, and input can later be validated against these types. XForms can also mark input as **read-only**, **required**, or **relevant**. Adaptation of the document can be accomplished by changing the value of a field's relevance at run-time. The GUI may hide all irrelevant fields as not to confuse the user. AdapForms have the same capabilities for in-document validation and adaptation, but extends these capabilities by allowing more complex validation and adaptation within the framework, such as using database look-ups and other server-side techniques, as described in Section 2.2 and 2.3. XForms requires native browser support in order to work, and at the time of writing no widely used browser supports it. Support is however available in modern browsers through plug-ins or through client-side simulations such as FormFaces [12], where XForms' behaviour is simulated using JavaScript.

JavaScript is widely used to make regular static HTML forms more responsive and adaptive. Such JavaScript can however be tedious and error prone to develop, which is addressed by PowerForms [3] by supporting client side validation for user input in HTML forms. It generates an interactive form based on a HTML static form and an PowerForms specification. The validation is continuous, showing the validity as the form is filled. The status is again indicated by red, yellow or green, based on the validity of the input: green indicating valid, yellow indicating not complete, and red indicating invalid. The rules are constructed using a simple **if-then-else** structure, or using regular expressions. Regular expressions are used to validate the format of the text being entered. The rules are transformed into a finite-state machine expressed in JavaScript, which validates the form input. The **if-then-else** constructs can be chained to make the validity of fields depend on other fields.

Whereas AdapForms is solely concerned with the adaption of forms and the ease of data entry therein, it shares some similarities with the more general area of multi-target interface building systems, wherein the UI is more or less abstractly defined so that it may used to

generate specific interfaces for different platforms. Such systems, e.g., [13] and UsiXML,^c while more far reaching in scope, still need to specify forms, preferably with enough syntactic/semantic information to generate appropriate interfaces. An example of such verification is found in GrafiXML [14], where entered data may be verified using masking, or Perl regular expressions. These systems have done extensive work on the ability to retarget a user interface to devices with different capabilities, which AdapForms has yet to do. In the terms found in [15], we have with the AdapForms form specification (as the one found in Figure 3) created an Abstract UI, which went introduced to our framework can generate a Final UI, accessible to the user.

AdapForms is similar to adaptive hypermedia systems [16, 17] in that both maintain a model of the user’s actions and both can affect changes in what is presented to the user based on a combination of the user model and predefined rules. However, AdapForms is specialised for forms and their contents and rules, whereas adaptive hypermedia systems (as the name suggests) are general hypermedia systems. It is however entirely possible to consider a combination of an adaptive hypermedia system and AdapForms, each informing the other, though this lies outside the scope of the present paper.

7 Conclusion

Forms are an integral part of modern living, yet the basic form has changed little over the decades—it may these days reside on a web page, but the general form remains the same. Validation of forms is usually either a lengthy process (fill-out, submit, review error messages, revise, resubmit). or relatively simple-minded using specialised JavaScript code. We have in this paper presented a general framework to define forms, designate acceptable field values, the interrelationships between fields, as well as an architecture to transparently validate, report state of, and adapt a form as the user is filling it out. Depending on the needs of the developer, the validation can be simple (type and range checks), XPath based, or advanced such as integrating server-side business logic through JavaBeans. Depending on the user’s input, the form can adapt itself, freeing the user from having to deal with an over-general form covering all possible permutations. Our initial evaluation leaves us hopeful that this is a valid approach, though there certainly still are many aspects of effectively communicating expected values and form state to the user that need to be explored. Likewise, the prototype implementation, while suited for our experiments has room for improvement, notably in terms of scalability.

Acknowledgments

The eGov+ project is financed by a grant from NABIIT, the Danish strategic research programme for nano-, bio-, and IT-sciences.

References

1. Henrik Gammelmark. Adaptive forms. Master’s thesis, Department of Computer Science, Århus, Denmark, April 2009. <http://adapforms.gammelmark.eu/files/adapforms-thesis.pdf>.

^c<http://www.usixml.org/>

2. Anders Berglund, Scott Boag, Don Chamberlin, Mary F. Fernández, Michael Kay, Jonathan Robie, and Jérôme Siméon. XML path language (XPath) 2.0, W3C recommendation. Technical report, The World Wide Web Consortium (W3C), 2007.
3. Claus Brabrand, Anders Møller, M. Ricky, and Michael Schwartzbach. Powerforms: Declarative client-side form field validation. *World Wide Web*, 3(4):205–214, 2000.
4. Peter Ørbæk. Programming with hierarchical maps. Technical report, Aarhus University, PB-575, 2005. <http://www.daimi.au.dk/publications/PB/575/PB-575.pdf>.
5. A. Brandl and G. Klein. FormGen: A Generator for Adaptive Forms Based on EasyGUI. In *Human-Computer Interaction: Ergonomics and User Interfaces. Proceedings of HCI International*, volume 99, pages 22–26, 1999.
6. Morten Bohøj, Nikolaj G. Borchorst, Niels Olof Bouvin, Susanne Bødker, and Pär-Ola Zander. Time collaboration. In *Proceedings of the 28th international conference on Human factors in computing systems*, Atlanta, Georgia, USA, April 2010. ACM.
7. Morten Bohøj and Niels Olof Bouvin. Collaborative time-based case work. In *Proceedings of the Hypertext conference 2009*, pages 141–146, New York, NY, USA, 2009. ACM.
8. Philip J. Hayes and Pedro A. Szekely. Graceful interaction through the COUSIN command interface. *International Journal of Man-Machine Studies*, 19(3):285 – 305, 1983.
9. A. Girgensohn, B. Zimmermann, A. Lee, B. Burns, and M.E. Atwood. Dynamic forms: An enhanced interaction abstraction based on forms. In *Proceedings of Interact'95*, pages 362–367. Chapman & Hall, 1995.
10. W3C. *XForms 1.1*, 2007. <http://www.w3.org/TR/xforms11/>.
11. Mikko Honkala. *Web User Interaction a Declarative Approach Based on XForms*. PhD thesis, Helsinki University of Technology, jan 2007.
12. FormFaces. <http://www.formfaces.com> (*Accessed January 2010*).
13. François Bodart, Isabelle Provot, Jean-Marie Leheureux, François Bodart, Anne-marie Hennebert, Anne-marie Hennebert, Isabelle Provot, Jean Vanderdonckt, Giovanni Zucchinetti, and Giovanni Zucchinetti. *Key Activities for a Development Methodology of Interactive Applications*, chapter 7, pages 109–134. Springer, 1996.
14. B. Michotte and J. Vanderdonckt. GrafiXML, a multi-target user interface builder based on UsiXML. In *Autonomic and Autonomous Systems, 2008. ICAS 2008. Fourth International Conference on*, pages 15 –22, mar 2008.
15. Gaëlle Calvary, Joëlle Coutaz, David Thevenin, Quentin Limbourg, Laurent Bouillon, and Jean Vanderdonckt. A unifying reference framework for multi-target user interfaces. *Interacting with Computers*, 15(3):289 – 308, 2003.
16. Paul De Bra, Ad Aerts, Bart Berden, Barend de Lange, Brendan Rousseau, Tomi Santic, David Smits, and Natalia Stash. AHA! the adaptive hypermedia architecture. In Les Carr and Lynda Hardman, editors, *Proceedings of the 14th ACM Hypertext Conference*, pages 81–84, Nottingham, UK, August 2003. ACM Press.
17. Paul De Bra, Geert-Jan Houben, and Hongjing Wu. AHAM: a Dexter-based reference model for adaptive hypermedia. In Klaus Tochtermann, Jörg Westbomke, Uffe K. Wiil, and John J. Leggett, editors, *Proceedings of the 10th ACM Hypertext Conference*, pages 147–156, Darmstadt, Germany, February 1999. ACM Press.