

Fine-Grained Specification and Control of Data Flows in Web-Based User Interfaces

Matthias Book, Volker Gruhn, and Jan Richter

Chair of Applied Telematics/e-Business, University of Leipzig
Klostergasse 3, 04109 Leipzig, Germany

{book,gruhn}@ebus.informatik.uni-leipzig.de, jan.richter@saxess.ag

Abstract. When building process-intensive web applications, developers typically spend considerable effort on the exchange of specific data entities between specific web pages and operations under specific conditions, as called for by business requirements. Since the WWW infrastructure provides only very coarse data exchange mechanisms, we introduce a notation for the design of fine-grained conditional data flows between user interface components. These specifications can be interpreted by a data flow controller that automatically provides the data entities to the specified receivers at run-time, relieving developers of the need to implement user interface data flows manually.

1 Introduction

Web-based user interfaces have become popular front-ends for information systems that require convenient access at any time from anywhere [1]. Especially in business-to-business and intranet applications that are designed to support elaborate business processes, these user interfaces can turn out to be quite complex. Their complexity is typically twofold: Most obviously to the user, they have intricate dialog structures that include nested dialog sequences, wizards, context-sensitive links and other navigation patterns. More transparent for the user, but all the more palpable for the developer, are the complex data flows between the interface and the business logic. No matter if these data flows mirror major business process features or serve minor technical purposes, the developer must ensure that the right data is available for the right component at the right time, all the while keeping an eye out for security issues, validity concerns, performance considerations and persistence strategies.

Despite some progress over the past years, the WWW infrastructure itself still provides only rudimentary data flow support: Originally, the only available data flow mechanism was the transmission of parameter strings from clients to servers in **HTTP requests**. Soon, web browsers became capable of receiving short **cookie** strings from a server and sending them back to the same server with every subsequent request. This in turn enabled servers to unambiguously associate multiple requests with the same client, and keep individual state information for all clients in **sessions**. Since sessions usually expire after some

time without user activity, **persistence** of selected data is typically ensured by integration of databases or other storage technologies in the back-end.

These basic data flow mechanisms are technically sufficient to build any web-based application: Requests provide a channel for data flows from client to server, sessions provide a scope for data flows among server components, and if necessary, cookies can provide an additional channel for client-server data exchange, client-side state-keeping, and even simple client-side persistence. However, these mechanisms are only convenient for a small subset of conceivable data flows: Requests are ideal for sending data from a web page to those business operations responsible for building the server's response; and sessions are ideal for making data accessible throughout the application. These alternatives mark two ends of a spectrum – in our experience, however, most data flows described by business processes lie somewhere in the middle: Often, data generated in a certain process step is intended only for a clearly defined, but not necessarily immediately following set of pages and/or operations. Thus, the respective data flows must reach beyond a single request-response cycle, but do not require session-wide data publication.

When mapping such process requirements to the technical level, developers currently need to choose the lesser of two evils: Passing data along a chain of requests toward its ultimate destination violates the principle of encapsulation, as the intermediate pages and operations need to handle data that they are not actually responsible for – an unclean and error-prone solution that requires high implementation and maintenance effort. Alternatively, storing data in the session for use at a later time relieves the intermediate steps from a lot of hassle, but bears the danger of memory leaks if the data is not removed from the session when it is no longer required. In addition, both approaches pose inherent security risks as data is exposed to pages or operations that do not need to know about it (and in the first case, even repeatedly sent over the network). Even if other application components are considered “friendly”, this unnecessary exposure multiplies the possible points of failure or attack.

Since neither the request nor the session scope provides satisfactory dialog flow support, we propose a supplemental method for realizing data flows that match the process requirements of application domains more closely, while at the same time reducing the required implementation effort. In this paper, we will first show how arbitrary point-to-point data flows can be specified in a graphical notation (Sect. 2). We then show how this specification can be interpreted at run-time by a data flow controller that provides just the specified data to each page and operation (Sect. 3). We conclude with a discussion of related work (Sect. 4) and an overview of further research opportunities (Sect. 5).

2 Data Flow Specification

In the introduction, we discussed data flows related to the user interface layer. Obviously, data also flows along the control structures within the application logic, may be exchanged with third-party, legacy or back-end systems, and can

be stored in and retrieved from persistent memory. However, since these data flows are the responsibility of the application and persistence logic, established specification methods (e.g. UML data flow diagrams, sequence diagrams etc.) and implementation techniques (e.g. SOAP, EJB etc.) can be used for these layers. Our work instead focuses on data flows between the presentation and application logic, where suitable implementation techniques (apart from the coarse request and session scopes) and specification methods have not yet been widely established.

To specify user interface data flows on a more fine-grained level than the one provided by the request or session scopes, we first need to define the concept of a data flow more clearly. Given a data source A , a data sink B and a data entity d , we define that the data flow of d from A to B is the provision to B of the d available to A (i.e. B gets to know the d that is known to A). As we will soon see, it is helpful to make data flows conditional, i.e. to execute them only if a certain constraint is fulfilled.

In designing a notation that maps this abstract definition onto a technical level, we need to answer a number of questions:

- What are concrete data sources and data sinks?
- Which constraints determine if a data flow is executed?
- How are sources and sinks related?
- What are concrete data entities?
- How is the “provision” of data entities realized?

The first three questions relate to the specification of data flows, and will be answered in the following subsections. The last two questions relate to the implementation of data flows, which will be discussed in Sect. 3.

2.1 Data Sources and Sinks

The data sources and data sinks used to model data flows should be entities that represent the structure of web-based user interfaces. In our past work, we have found it natural to distinguish between web pages and application logic operations in order to model the navigation structure of web applications: In the Dialog Flow Notation (DFN) [2], web pages are symbolized as dog-eared sheets (the so-called **masks**), while application logic operations (**actions**) are symbolized by circles. To specify all possible navigation paths, these elements are connected by arrows symbolizing **dialog events** generated by masks or actions when the user submits a request, or a business operation produces a result. In the dialog graphs specified this way, masks and actions do not have to alternate since it is conceivable for a mask to trigger a succession of several server-side operations, and also conceivable for a mask to link directly to another mask without the need for any intermediary business operations. The dialog graphs are encapsulated in **dialog modules** that can call each other at run-time to form hierarchically nested dialog structures.

We found that the DFN provides an ideal basis for specifying the data flows within a web-based user interface: The masks, actions and modules can serve as

data sources and sinks, while the dialog events can be interpreted as constraints that determine when a data flow is executed. In the following subsections, we will show how the original DFN was extended to specify different kinds of data flows. Our aim was to extend the notation in such a way that the data flows are intuitively readable and conceptually integrated with the DFN semantics, but do not add unnecessary optical clutter to the diagrams.

2.2 Data Flow Types

Depending on the types of the data sources and sinks, and their proximity within the topology of the dialog graph, we distinguish several kinds of data flows:

Parallel Data Flows. The simplest type of data flow is running in parallel with the dialog flow, i.e. a data flow from the generator of a dialog event to the receiver of the same event. It is specified in the DFN by adding the label of the data entity in square brackets to the event’s name (if the data flow shall comprise several data entities, we separate them with commas). For example, in the *login* module in Fig. 1, the *name* and *passwd* data entities entered by the user on the *login* mask shall only be provided to the *check name, passwd* action upon generation of the *submit* event. At first sight, this data flow may look equivalent to the request scope, but it is actually more finely grained: While the request scope always encompasses all intermediate actions up to the next mask, the parallel data flow is restricted to the current dialog event’s receiver only. Therefore, the *passwd* data in the example is only provided to the succeeding action, but not accessible beyond it (as it would be in the request scope).

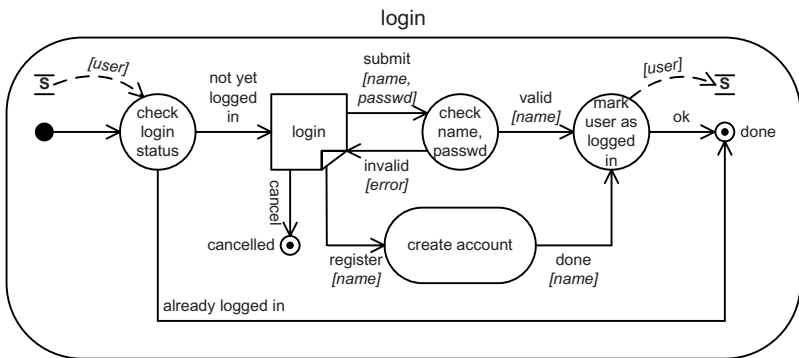


Fig. 1. Login module performing user authentication

Request scopes and parallel data flows also differ in their starting points: While a request can only be generated by a mask, a dialog event with associated parallel dialog flow can also be generated by actions or modules. For example, if the login credential check fails, the *error* data is provided by the *check name, passwd* action to the *login* mask with the *invalid* event.

Finally, parallel data flows can be used more flexibly than request scopes since developers can specify which data is provided with which events: Using only the request scope, the developer may not be able to prevent that the credentials entered by the user on the *login* mask are provided to subsequent elements both for the *submit* and (unnecessarily) the *cancel* event. Using parallel data flows, however, the developer can specify that the *name* and *passwd* data is provided to the subsequent action with the *submit* event, just the *name* data is provided to the *create account* module with the *register* event, but no data is provided with the *cancel* event.¹

Divergent Data Flows. As we have just seen, parallel data flows enable developers to specify which data entities shall be provided to the receivers of which events. Often, however, the data generated by one element should not be provided to any of its immediate successors, but rather to some more distant element in the dialog graph that is responsible for the actual data processing.

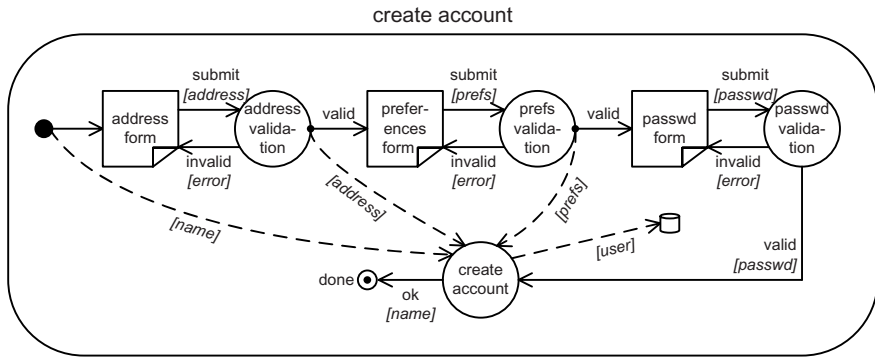


Fig. 2. *Create account* module employing a “wizard” navigation pattern for gathering user input

As an example, consider the “wizard” navigation pattern of the *create account* module in Fig. 2: Here, we prompt the user for his address, preferences and password, validate each of the inputs, and finally create an account from all the collected data. In order not to burden the various forms and validation actions with forwarding the data of their predecessors, we would like to provide the data gathered in each step directly to the final processing action. For the validation actions, the outgoing dialog and data flows therefore diverge: For example, if the *address validation* action produces a *valid* event, the dialog flow shall continue to the next wizard step, i.e. the *preferences form*, but the *address* data shall be provided directly to the *create account* action, bypassing the other steps of the

¹ Note that these restrictions cannot prevent the respective data from being transmitted across the network in the request, as this would require client-side logic. However, the data flow controller will ensure that any submitted data is only available to the specified receivers on the server side.

wizard. To specify that the data flow diverges from the dialog flow in this case, we draw a dashed line to the element that the data shall be provided to, and annotate it with the data entity's label in square brackets.

Just like a parallel data flow is always tied to a dialog event, a divergent data flow is also always associated with a particular dialog event and only executed if that event is traversed. To specify this constraint in the DFN, the data flow arrow must always begin in the same spot (marked by a black dot) as the event arrow that it shall be associated with. To maintain consistency with parallel data flows, which are always implicitly associated with an event, the DFN does not allow the specification of divergent data flows without an associated event.

Inter-module Data Flows. Of course, data does not only flow among the dialog elements within a module, but also between modules. Therefore, when a module is called, it should be able to accept data provided to it, and when a module terminates, it should be able to provide data to subsequent dialog elements. To express this behavior in the DFN, compatible data flows must be specified both in the exterior dialog graph that a module is embedded in, and in the interior dialog graph of the module itself: In the exterior dialog graph, modules can simply serve as sources and sinks of parallel and divergent dialog flows just like masks and actions. In Fig. 1, for example, the *create account* submodule is provided *name* data with the incoming *register* event and provides *name* data with its outgoing *done* event.

To specify which data entities a module can accept and provide, the initial and terminal anchors of its interior dialog graph serve as data flow interfaces: Parallel and divergent data flows can originate from the **initial anchor** (the black disk marking the starting point of the dialog graph traversal) to specify to which elements the incoming data shall be provided. Analogously, parallel and divergent data flows can lead to the module's **terminal anchors** (the circled small black dots marking the end of the dialog graph's traversal). They specify which data entities the module will provide to its exterior dialog graph, and who provides those data entities internally. In the definition of the *create account* module in Fig. 2, for example, the incoming *name* data flows directly from the initial anchor to the *create account* action, which will process it later together with the other data collected by the wizard. The *name* data will then flow from the *create account* action to the *done* terminal anchor, so it can be provided to the module's successor in the exterior dialog graph upon its termination.

Through these incoming and outgoing interfaces, data can be flexibly passed from several sources outside a module to several sinks inside a module and vice versa. Developers need to ensure that the data flows specified in the interior and exterior dialog graphs match, as unmatched entities will otherwise be unavailable to their sinks.

2.3 Shared Scope Access

Parallel and divergent data flows enable developers to define the propagation of data entities on a very fine-grained level. However, their specification may

become cumbersome if the same data entities shall be provided to many elements, as a profusion of explicit data flow arrows would be required. For this reason, the DFN provides constructs to symbolize data exchange through various shared scopes.

Module Scope. Often, the same data entities should be provided to virtually all elements within the same module. As we mentioned earlier, storing such data entities in the session scope is not an optimum solution for this problem since the developer would be responsible for removing them from the session before the module is terminated. Instead, it would be preferable if all elements within a module had a shared scope that is cleared out automatically once the module terminates. Our Dialog Control Framework (DCF) described in Sect. 3 provides the **module scope** for this purpose: It is comparable to the session scope in that it is associated with the current user, but exists only during the traversal of the current module. An empty module scope is instantiated whenever a module is entered, and discarded again once a terminal anchor has been reached.

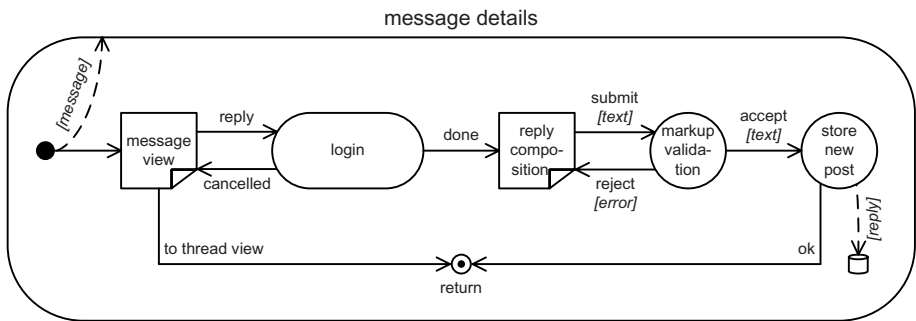


Fig. 3. *Message details* module for viewing and replying to messages in a discussion forum

To specify in the DFN that a data entity should be provided through the module scope, developers can simply draw a divergent data flow with the respective entity label in square brackets from any dialog event to the module’s contour. In the *message details* module in Fig. 3, for example, the incoming *message* data is provided to all other elements through the module scope right away.

Of course, data flows to the contour can also originate from any other source, indicating that the respective data will be available throughout the module as soon as the associated event is traversed. Since all data in the module scope is implicitly available to all dialog elements, no explicit notation construct is required to indicate that an element accesses data from the module scope. In the *message details* module, for example, all masks and actions can access the original *message* data in the module scope to display it, quote it in the text of the reply, and reference it with the new post.

In the DCF implementation, module scopes are stacked to reflect the nested module call structure: When a module calls a sub-module, a fresh module scope for the sub-module is pushed onto the stack, rendering the calling module's scope temporarily inaccessible. When the sub-module terminates, its module scope is removed from the stack, and the calling module's scope becomes available again.

Session, Application and Cookie Scope; Storage Access. The mechanisms described so far are helpful in many situations where the data scopes provided by the application server are too coarse for the data flow requirements at hand. However, there are obviously also a number of situations where those more generous scopes are the perfect choice – information about the user's login status, for example, should be available throughout the application and is therefore ideally stored in the session scope. Other data may best be globally provided through the application or cookie scope, and business objects often need to be stored in or retrieved from persistence storage by the application logic.

The data flows for these scenarios usually concern only the application (and possibly persistence) layer, so the dialog control logic that couples the presentation and application layer should not be involved in them. Consequently, the DCF does not handle data flows through these larger scopes, but lets the business logic interact directly with the application server or persistence layer API. This restriction of responsibilities allows for a clean separation of concerns on an architectural level, and allows developers to choose whichever scope and persistence framework is suited best to their needs, instead of being bound to the DCF's data flow mechanisms throughout the application.

Due to the framework's limitation to user interface management, the DFN technically would not need to provide additional constructs related to the larger scopes. Intuitively, however, developers will expect to be able to specify not only fine-grained data flows through the user interface, but also data shared through the more coarse scopes.

The DFN solves this dilemma by providing a compromise – developers cannot specify the complete application logic's internal data flows within the dialog flow, but the DFN enables them to specify where the dialog flow interfaces with the larger scopes: Data flow arrows leading to or from the letters **S**, **A** or **C** enclosed by horizontal bars indicate that the respective data entities are stored in or retrieved from the session, cookie or application scope, respectively. For example, in Fig. 1, the *check login status* action retrieves *user* data from the session scope, and the *mark user as logged in* action stores *user* data in the session scope. In addition, data flow arrows leading to or from a “can” symbol indicate that the respective data entity is stored in or retrieved from persistent storage – for example, the *user* data constructed by the *create account* action in Fig. 2 is sent to the back-end for storage. Note that in contrast to divergent data flows, the data flows leading to those larger scopes cannot share their starting point with an event, since the respective data handling operations are beyond the control of the dialog controller that is only aware of dialog events.

Since the DCF does not provide mechanisms for handling these scopes, the corresponding notation constructs have only illustrative character – the actual

data handling will have to be implemented manually by the developer. In contrast, the module scope and parallel, divergent and inter-module data flows are executable specifications: The developer only needs to specify the desired behavior, but does not need to implement it since the framework handles the respective data flows automatically.

2.4 Summary

With the traditional scopes, the module scope and the new finely-grained data flows, developers now have a full spectrum of data flow mechanisms at their disposal, as Fig. 4 illustrates: While the application and session scopes are suitable for publishing application-global and user-global data, parallel and divergent data flows are ideally suited for data propagation between arbitrary dialog elements. Since multiple parallel and divergent data flows can be associated with any dialog event, developers can control which data is propagated where under which conditions very flexibly.

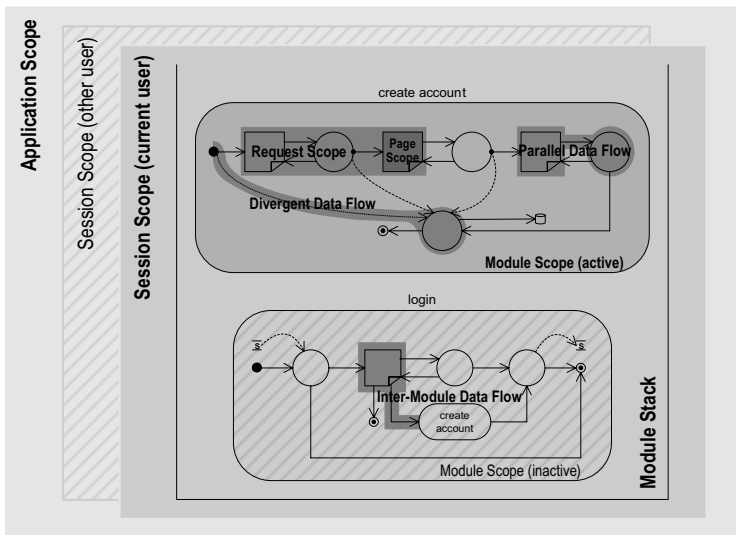


Fig. 4. Scopes and data flows supported by the Dialog Flow Notation and Dialog Control Framework

The page scope remains an important tool for data exchange between the components involved in constructing a response (e.g. a JavaServer Page (JSP) and its tag libraries). The request scope, however, has been rendered virtually redundant since its visibility is determined by purely technical criteria that only seldomly map exactly to business requirements for data propagation. This allows us to retrofit it for conveniently accessing the new data flow mechanisms, as we will show in the following section.

3 Data Flow Controller

Having introduced the new data flow constructs of the extended Dialog Flow Notation, we will now show how their semantics are supported by the Dialog Control Framework (DCF). Revisiting our initial questions from Sect. 2, we first need to address the technical representation of data entities, and discuss how the provision of those entities from sources to sinks is interpreted. After this, we will give a brief overview of the technical implementation of the data flow control mechanism within the DCF.

3.1 Data Entity Representation and Provision

In the previous sections, we identified data entities only by their label (e.g. “*user data*”). In concrete applications, a data entity can be any object – in Java-based applications, it will typically be an instance of a `JavaBean` holding various attributes, which is stored in or retrieved from a data scope using its label as a look-up key. In Java Enterprise Edition-compatible application servers, for example, `HttpSession` instances provide the `Object getAttribute(String name)` and `void setAttribute(String name, Object value)` methods for this purpose. Just like the DFN relies on existing notations to specify the layout of dialog masks (e.g. visual design sketches or XForms) or the control flow within actions (e.g. UML activity or state diagrams), it does not provide its own constructs for the specification of data entities’ internal structure. Instead, we recommend using existing notations such as entity-relationship or UML class diagrams for this purpose.

Since we assume that the data entities in our data flows are objects (i.e. reference types), we have two alternatives for interpreting the flow of data from a source to a sink: “providing a data entity” could mean

- forwarding the data entity itself from the source to the sink, or
- extending the data entity’s scope so it is not only available to the source, but also to the sink.

While these alternatives may at first sight look like equivalent implementation variants, they exhibit different behavior if a source provides the same data entity to several sinks, as illustrated in Fig. 5. In a situation like this, the first approach (forwarding the data entity) intuitively implies that both *B* and *C* receive identical copies of *d* from *A*, and that any changes *B* makes to *d* will not affect *C*’s copy of *d*. In contrast, the second approach (extending the entity’s scope) implies that *B* and *C* can now both access the same instance of *d* that is already known to *A*, so any changes that *B* makes to *d* will also affect the *d* available to *C*. In short, the first approach requires a copy-by-value mechanism, while the second can be implemented as copy-by-reference.

The copy-by-value implementation has the advantage that side effects are avoided. However, since we are dealing with objects that may have arbitrary complexity, copying data entities is not trivial. We could simply require that all data entities are cloneable and thereby put the responsibility on application




Fig. 5. Data flow with different copy-by-value and copy-by-reference behavior

developers, however, this is hardly a satisfactory solution for a framework that strives to make developers' jobs easier. Furthermore, cloning comes at a high performance and memory cost (especially since some clones may turn out to be unnecessary if the respective data sinks happen never to be visited in the user's subsequent traversal of the dialog graph), and some data structures may not be cleanly cloneable with justifiable effort. A restriction to flat copies or primitive types is not a realistic option either, since it would limit application developers' design freedom severely.

For these reasons, we prefer a copy-by-reference implementation of data flows. While it may involve a bit more subtle semantics, its basic concept and the situations in which side effects may occur are well known to experienced developers. This approach is much easier to implement for framework and application developers, and does not cause the performance and memory overhead of cloning.

The above considerations were confirmed by our initial prototype of a data flow control extension for the DCF: Here, clones of JavaBeans were only flat copies, so any nested references were not cloned, and the isolation between the data entities available to the dialog elements was not perfect. We are therefore currently switching to a data flow controller implementation that copies only object references instead of whole instances, and thus realizes the scope extension.

3.2 Data Flow Controller Design

Having determined the operational semantics of data flows, we still need a way to actually provide the specified data to the respective dialog elements in a web application. At first sight, an obvious solution would be to equip every mask and action instance with a look-up table that holds references to all data entities available to that element. This individual "element scope" could then be accessed by the application and presentation logic just like the session or application scope, using `setAttribute` and `getAttribute` methods. A data flow controller would ensure that object references are copied from one element scope to another according to the data flow specifications.

While this approach seems straightforward, it proves cumbersome in practice since the element scope itself cannot be easily made available to the presentation logic: Servlets, JSPs and other web resources provide convenient mechanisms for accessing the standard scopes (e.g. through the implicit `application`, `session` and `request` objects of the Java EE Expression Language), but do not provide as convenient means for accessing custom-built scopes.

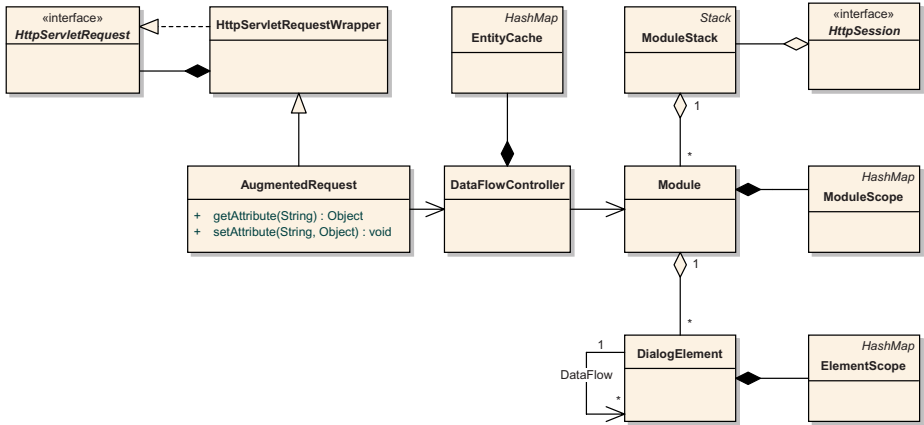


Fig. 6. Data flow control logic of the Dialog Control Framework

In the data flow extension to the DCF, we therefore took a slightly different approach: The data flow controller still manages look-up tables that contain the object references available to each dialog element. However, these tables are not directly accessible to the dialog elements. Rather, we project their contents into the existing request scope that is already conveniently available to all masks and actions.

As Fig. 6 shows, this can be achieved by wrapping the original HTTP request into a wrapper object (here, `AugmentedRequest`) that looks and behaves just like an HTTP request, since it passes most method calls directly to their counterparts in the original `HttpServletRequest`. The only exceptions are the `setAttribute` and `getAttribute` methods that are normally used to access the request scope, but now rerouted to the `DataFlowController`. This central controller has (via the current `Module` on top of the `ModuleStack`) access to the data flow model (a graph of `DialogElements` linked with `DataFlow` edges), the current `ModuleScope` and the current element’s `ElementScope` (two `HashMap`s containing the object references associated with the data entity labels).

Whenever a mask or action calls the request’s `getAttribute` method with some data entity label, the `DataFlowController` looks up the respective object reference in the current `ModuleScope` and `ElementScope` (the element scope takes precedence over the module scope in case both contain a reference with the same label), and returns the found reference.

Populating the look-up tables requires slightly more effort: Whenever a dialog element calls the request’s `setAttribute` method, the `DataFlowController` stores the respective data entity in the current element’s scope. It also looks up any departing data flow edges in the module’s data flow model and caches the entity’s object reference if it is associated with any outgoing event. Once an event is generated by the current dialog element, the associated data references

from the `EntityCache` are stored in the element scopes of the respective data flow receivers, and the cache is cleared.² This way, the data flow controller ensures that every element can access through the request context all data entities available to it according to the specification.

Besides giving the presentation and application logic convenient access the element and module scopes, projecting the element scope into the request scope has the additional benefit that the mechanism is transparent to other web application frameworks: In our prototype implementation, we integrated the DCF with JavaServer Faces (JSF) to make use of its UI component model and validation logic. Since JSF pages exchange data through the request scope, they can work with the new module scope and data flows seamlessly.

4 Related Work

Many modeling languages for web-based applications have traditionally had a strong focus on data-intensive web applications [3], allowing developers to specify how users navigate through complex data schemas. More recently, a number of languages such as OO-HDM [4], OO-H and UWE [5], and WebML [6] have also incorporated aspects of business process modeling, thereby narrowing the gap between process and navigation specifications that developers need to bridge. The data focus of these modeling languages is reflected in the variety of constructs they provide for specifying relations between and manipulations of data entities, most of which reside in the back-end of a web application.

However, apart from the transport links in WebML, which have similar semantics as our divergent data flows, the above modeling languages do not seem to provide explicit notation constructs for the fine-grained specification of inter-element data scopes and provision: The navigational model and process flow model in UWE, for example, specify how to navigate across and manipulate a data space provided by the back-end, but does not show which data instances are provided from one interface component to another. In OO-H, activity and navigation access diagrams enable developers to specify which navigation nodes will invoke which data-manipulation methods on data objects, but any scoping of these instances is not explicitly modeled.

The extension to the Dialog Flow Notation we introduced here focuses on the fine-grained specification of data scopes and data provisioning (i.e. which data instances are made available to which dialog elements). These data flow specifications do not have to be implemented manually, but are automatically enforced by the Dialog Control Framework at run-time.

In contrast to the above notations, the DFN does not provide constructs for specifying how the data that is provided to the various dialog elements is

² The actual data flow controller is a bit more complex than the diagram and this brief description suggest – among other things, it also moves incoming request parameters into request attributes, and ensures that data references are also copied to other scopes if they have not been set by calls to `setAttribute`, but received through data flows from other elements.

manipulated. This is in keeping with our philosophy that the dialog flow is what distinguishes web applications most from traditional applications – we therefore focus the DFN on the typical and unique challenges of navigation and data flow in web applications, and encourage developers to use other established notations for modeling those aspects that go beyond this layer (e.g. by using activity or state diagrams to specify how data entities are manipulated in actions).

Regarding run-time support for complex dialog flows, popular web application frameworks have recently also adopted the notion of encapsulating dialog sequences in so-called “flows” (in Spring Web Flow [7]) or “conversations” (in the Shale Framework [8]), which have associated data scopes. However, these frameworks do not provide mechanisms for the realization of parallel or divergent data flows, and are lacking a corresponding notation that would provide executable specifications for such data flows.

5 Conclusion

In this paper, we presented a data flow extension to the Dialog Flow Notation (DFN) that enables web application developers to specify data flows between arbitrary elements in a dialog graph, as well as between nested dialog modules. Since all data flows are associated with events whose traversal is the condition for data propagation, and the receiver of a data flow may be different from the receiver of the associated dialog event, developers can build on the existing DFN semantics to specify fine-grained conditional data propagation within a web application’s user interface. A new module-wide data scope complements the existing coarse scopes provided by common application servers.

To relieve developers of the tedious and error-prone effort of manually implementing secure and correct data flows throughout a web-based user interface, the data flow specifications created with this notation are executable: Using a graphical editor, they can be transformed into machine-readable specifications that are interpreted by our Dialog Control Framework (DCF), which manages the dialog and data flow automatically. At run-time, the framework ensures that data entities are always provided to their specified receivers, and that dialog elements can only access those data entities that they are entitled to. Since the new data flow mechanisms are transparently added to the existing request scope, other web application frameworks relying on this scope can also benefit from the specified data flows.

Our work aims to support the development process for web-based applications in both the design and the development phase: We expect it to be easier for designers to communicate with clients and other non-technical stakeholders in the software development process, since data flow information that was previously contained only in business process models can now be mapped to dialog flow diagrams, thereby providing a smoother transition from requirements to implementation. Given the specifications created in this way, we are confident that developers can reduce implementation, testing and maintenance efforts since they do not need to be concerned with the actual data flow implementation.

A prototype version of the enhanced Dialog Control Framework has already demonstrated that the described concepts are technically feasible. In our ongoing work, a first application for the data flow mechanism will be the DiaGen extension to the DCF, which automatically breaks dialog masks into wizard-style interaction sequences at run-time to cater to different device capabilities [9] – the use of divergent data flows will greatly simplify the auto-generated data propagation code in this context. Apart from this, we are striving to test our hypotheses regarding the positive impact on the software development process in real-life projects. We are also considering further refinements of the data flow notation to provide interfaces to other modeling languages for the specification of the application logic's internal data handling. Another major upcoming research effort based on the data flow engine will be the development of algorithms for recognizing and handling backtracking, cloned browser windows and other unexpected user activities in web applications.

References

1. Gaedke, M., Beigl, M., Gellersen, H., Segor, C.: Web content delivery to heterogeneous mobile platforms. In: Kambayashi, Y., Lee, D.-L., Lim, E.-p., Mohania, M.K., Masunaga, Y. (eds.) *Advances in Database Technologies*. LNCS, vol. 1552, Springer, Heidelberg (1999)
2. Book, M., Gruhn, V.: Modeling web-based dialog flows for automatic dialog control. In: *ASE 2004. 19th IEEE International Conference on Automated Software Engineering*, pp. 100–109. IEEE Computer Society Press, Los Alamitos (2004)
3. Fraternali, P.: Tools and approaches for developing data-intensive web applications: A survey. *ACM Computing Surveys* 31(3), 227–263 (1999)
4. Rossi, G., Schmid, H., Lyardet, F.: Engineering business processes in web applications: Modeling and navigation issues. In: *Third International Workshop on Web-Oriented Software Technology*, pp. 81–89 (2006)
5. Koch, N., Kraus, A., Cachero, C., Meliá, S.: Integration of business processes in web application models. *Journal of Web Engineering* 3(1), 22–49 (2004)
6. Brambilla, M., Ceri, S., Fraternali, P., Manolescu, I.: Process modeling in web applications. *ACM Transactions on Software Engineering and Methodology* 15(4), 360–409 (2006)
7. Vervaet, E., et al.: Spring web flow. <http://www.springframework.org/go-webflow>
8. Apache Software Foundation: Shale framework. <http://shale.apache.org>
9. Book, M., Gruhn, V., Lehmann, M.: Automatic dialog mask generation for device-independent web applications. In: Book, M., Gruhn, V., Lehmann, M. (eds.) *ICWE 2006. Proceedings of the 6th International Conference on Web Engineering*, pp. 209–216. ACM Press, New York (2006)