

A Collaborative Feature Modeling System

Rafael Bidarra
Eelco van den Berg
Willem F. Bronsvort

Faculty of Information Technology and Systems,
Delft University of Technology,
Mekelweg 4, NL-2628 CD Delft,
The Netherlands
e-mail:
(R.Bidarra/E.vdBerg/W.F.Bronsvort)@its.tudelft.nl

Collaborative systems are distributed multiple-user systems that are both concurrent and synchronized. An interesting research challenge is to develop a collaborative modeling system that offers all facilities of advanced modeling systems to its users, while at the same time providing them with the necessary coordination mechanisms that guarantee effective collaboration. To achieve this, a web-based collaborative feature modeling system, webSpiff, has been developed. It has a client-server architecture, with an advanced feature modeling system as a basis for the server, providing feature validation, multiple views and sophisticated visualization facilities. A careful distribution of the functionality between the server and the clients has resulted in a well-balanced system. On the one hand, the server offers all the functionality of the original feature modeling system. On the other hand, all desirable interactive modeling functionality is offered by the clients, ranging from display of feature model images to interactive model specification facilities. The architecture of webSpiff, the distribution of model data, the functionality of the server and the clients, and the communication mechanisms are described. It is shown that a good compromise between interactivity and network load has been achieved, and that indeed advanced feature modeling with a collaborative system is feasible.

[DOI: 10.1115/1.1521435]

1 Introduction

In the last decade, research efforts in the areas of solid and feature modeling substantially contributed to the improvement of computer-aided design (CAD) systems. A broad range of advanced modeling facilities is now becoming available in high-end commercial systems, amplified by continuous enhancements in interactive and visualization capabilities, and profiting from the availability of ever faster and more powerful hardware. Still, these improvements have their counterpart in the increasing size and complexity of such systems. At the same time, a number of research prototypes are pushing the edge to even more advanced modeling facilities. For example, embodiment of richer semantics in feature models and validity maintenance of such models [1] and physically-based modeling techniques [2] are among the current research issues.

A common characteristic of most current CAD systems is that they run on powerful workstations or personal computers. Interaction with the system is usually only possible if the user is directly working at the CAD station, although remote interaction is sometimes possible through a high-bandwidth local area network. This situation is no longer satisfactory, as nowadays more and more engineers, often at different locations, are getting involved in the development of products. It would be preferable if a user could remotely browse and manipulate a model, via Internet, as if he were working directly at a powerful CAD station. A web-based system would be ideal for this, as it would facilitate access to all sorts of product information in a uniform, simple and familiar framework.

Even more attractive would be the support of collaborative modeling sessions, in which several geographically distributed members of a development team could work together on the design of a product. Typically, in such collaborative sessions, different participants would be provided with their own, application-specific views on the product, e.g. for detailed design, manufacturing planning or assembly planning [3,4]. In addition, each session participant, as in traditional development teams, would be given his own competence and specific privileges by the system.

Some commercial tools that are now emerging provide limited support for collaborative design activities. Meanwhile, prototype systems are being developed which investigate collaborative modeling possibilities.

Considering the state of the art, it is an interesting research challenge to develop a collaborative modeling system that offers all facilities of advanced feature modeling systems to its users, while at the same time providing them with the necessary coordination mechanisms that guarantee effective collaboration. Among these mechanisms, solutions have to be provided to the critical problems of concurrency and synchronization that characterize collaborative design environments. This paper presents a new web-based collaborative feature modeling system that is a major step in the right direction. In Section 2, the main research issues of collaborative modeling systems are surveyed. In Section 3, the architecture of the proposed system is discussed. The structure and functionality of the server and of the clients are described in Sections 4 and 5, respectively. Finally, results and conclusions are presented in Section 6.

2 Survey of Collaborative Modeling Systems

Collaborative systems can be defined as distributed multiple-user systems that are both concurrent and synchronized. *Concurrency* involves management of different processes trying to simultaneously access and manipulate the same data. *Synchronization* involves propagating evolving data among users of a distributed application, in order to keep their data consistent.

These concepts being in general already rather demanding, their difficulty becomes particularly apparent within a collaborative modeling framework, where the complexity and size of model data that has to be synchronized are typically very large, and the concurrent modeling actions taking place may be far-reaching. This section briefly surveys collaborative modeling systems, highlighting the key aspects put forward by recent research, and summarizing the lessons learned from a few tools and prototype systems proposed so far.

2.1 Client-server Architecture. The requirements for concurrency and synchronization in a collaborative modeling context mentioned above lead almost inevitably to the adoption of a *client-server* architecture, in which the server provides the participants in a collaborative modeling session with the indispensable communication, coordination and data consistency tools, in addi-

Contributed by the Computer Aided Product Development (CAPD) Committee for publication in the JOURNAL OF COMPUTING AND INFORMATION SCIENCE IN ENGINEERING. Manuscript received Jun. 2002; revised Sept. 2002. Associate Editor: P. Wright.

tion to the necessary basic modeling facilities. For a survey on client-server architectures, see, for example, Ref. [5].

In client-server systems it is important to balance the complexity of the client application and the network load. In a collaborative modeling context, client complexity is mainly determined by the modeling and interactive facilities implemented at the client, whereas network load is mainly a function of the kind and size of the model data being transferred to/from the clients. A whole range of solutions can be devised between the two extremes, so-called *thin clients* and *fat clients*.

A pure thin-client architecture typically keeps all modeling functionality at the server, which sends an image of its user interface to be displayed at the client. Clicking on the image generates an event, containing the screen coordinates of the interface location the user clicked on. This event is sent to the server, which associates it with an action on a particular widget. Eventually, this action is processed, and an updated image of the resulting user interface is sent back to the client, where it is displayed. This approach requires a continuous information stream between server and clients, and is therefore very expensive in terms of network traffic. The response time would be intolerably high for many model specification actions, thus making it very ineffective to remotely participate in a modeling session.

On the other extreme, a pure fat client offers full local modeling and interaction facilities, maintaining its own local model. Communication with the server is then often required in order to synchronize locally modified model data with the other clients. In a collaborative environment where clients can concurrently modify local model data, preventing data inconsistencies between different clients becomes a crucial problem. In addition, fat clients, to be effective, place on the platform running them the heavy computing power requirements of typical CAD stations. Finally, fat clients are typically platform-dependent applications that require more complex installation and maintenance procedures, and are therefore less practical in a multi-platform environment, in particular across various enterprises.

2.2 Current Tools and Prototype Systems. So far, only a small number of commercial tools have been developed that somehow support collaborative design activities. For example, tools for collaborative model annotation and visualization via Internet are now becoming available, providing concepts such as shared cameras and telepointers [6–8]. However, such tools are primarily focused on inspection, e.g. using simple polygon mesh models, and do not support real modeling activities. In other words, they are valuable assistants for teamwork, but no real CAD systems. To the best of our knowledge, there is currently only one commercial client-server system offering some synchronous collaborative modeling facilities, OneSpace [9], but this system is severely constrained by the model format into which it converts all shared CAD models.

On the other hand, several prototype collaborative modeling systems have been described in literature. Some of these systems will be shortly surveyed here, and their shortcomings identified.

CollIDE [10] is a plug-in for the Alias modeling system, enhancing it with some collaborative functionality. Users of CollIDE have private workspaces, where model data can be adjusted independently from other users. In addition, a shared workspace exists containing a global model, which is synchronized between all users participating in a collaborative modeling session. Users can copy model data between the private and shared workspaces, in order to create and adjust certain model data locally, and add it to the model in the shared workspace. The architecture of CollIDE poses severe restrictions to crucial collaborative modeling issues. In particular, no special measures have been taken to reduce the amount of data sent between the participants of a collaborative modeling session, resulting in delayed synchronization of the shared workspace and of the users' displays. Also, since each user operates on a separate instance of the modeling system, able to

perform modeling operations by itself, concurrency has to be handled by the users themselves in order to keep shared model data consistent.

The ARCADE system [11] defines a *refine-while-discussing* method, where geographically distributed users can work together on a design, interacting with each other in real-time. Every participant uses a separate instance of the ARCADE modeling system, and all ARCADE instances are connected to a session manager via Internet. A message-based approach was chosen, where every change of the product model is converted into a short textual message, which is sent to all other instances of ARCADE through the session manager. ARCADE provides a collaborative environment in which the network load is kept low. This was done by including all modeling functionality in the distributed ARCADE instances, which exchange only textual messages, rather than large sets of polygons. A drawback of this approach, however, is that the user application becomes rather complex, thereby requiring much computational power. In addition, ARCADE provides a primitive concurrency control mechanism, where only one user can edit a particular part at a time.

CSM, the Collaborative Solid Modeling system proposed by Chan et al. [12], is a web-based collaborative modeling system. Within its client-server architecture, the server contains a global model, while every client owns a local copy of this model. When a user has locally modified the model, this is propagated to all other users through the server. Concurrency is managed in two ways: (i) the model can be locked, using token passing, restricting it from being accessed by other users as long as some user is performing a modeling operation; and (ii) functionality can be locked, preventing certain functions from being used by particular users. Clearly, such methods provide a very strict concurrency handling policy. In fact, they turn the clients into several independent modeling systems, just using the same product model alternately. In a real collaborative modeling system, one expects a higher level of coordination support.

NetFEATURE [13] claims to be a web-based collaborative feature modeling system. A server provides basic functions on a central product model, including creation and deletion of features. On the clients, a local model is available, containing a boundary representation of the product, derived from the server-side central model. The local model is used for real-time display, navigation and interaction. For more advanced operations, the server must be accessed. Updating the local model is done incrementally, which required a rather heavy naming scheme, severely reducing the modeling functionality of the system. Furthermore, NetFEATURE uses, just like CSM, very strict concurrency handling methods, thus seriously limiting real collaborative modeling.

2.3 Conclusions. Collaborative modeling systems can support engineering teams in coordinating their modeling activities. Instead of an *iterative* process, sending product data back and forth among several team members, designing becomes an *interactive* process, in which several engineers are simultaneously involved to agree on design issues. Collaborative modeling systems typically have a client-server architecture, differing in the distribution of functionality and data between clients and server.

Concurrency control is still a crucial issue in current collaborative environments. If a user is allowed to change a model entity, while another user is also changing the same entity, problems can easily arise concerning consistency of the model. To avoid this situation, a strict concurrency control mechanism can limit access for other users. It depends on the application, whether all entities of the design should be locked or just some of them. If possible, users should be allowed to simultaneously modify different parts of the design, but this could lead to much more complicated concurrency control mechanisms. Also, one should always bear in mind that designing is a constructive activity. Users can therefore be given some responsibility for establishing a good collaboration.

Current systems also often fall short in adequately handling synchronization of model data among distributed clients. Timely

updating data over a network is difficult, since there is a certain delay between the moment data is sent and the moment it is received at another node of the network; during this time interval, the latter might try to manipulate data that is not up-to-date. Mechanisms to detect such conflicts should be available, and recovery mechanisms provided. Good locking can also help to avoid such situations, but sometimes this may hinder users' flexibility.

For a collaborative CAD system to be successful, it should provide a good level of interactivity. Users will not be able to design properly if they have to wait a long time after every operation. But increasing interactivity by just porting more and more data and functionality to the clients is not a good solution either, as synchronization problems would become critical.

In short, a good solution to the difficulties summarized above can be a web-based client-server approach, where the server coordinates the collaborative session, maintains a shared product model, and provides all modeling functionality, which otherwise could be implemented only in very fat clients. The clients locally specify modeling operations on a graphical model, and only high-level semantic messages, as well as a limited amount of information necessary for updating the client data, is sent over the network. This effectively limits the network load, while guaranteeing good client interactivity at acceptable response times. An important characteristic of such an architecture is that there is only one central product model in the system, which is very advantageous in case of a complex feature model. Clients send their modeling operations to the server, and receive feedback after any modeling operation has been performed on its central model, thus avoiding inconsistency between multiple versions of the same model.

3 webSpiff: A Balanced Architecture

In this section, we discuss the architecture of a system implemented according to the above-mentioned solution: the prototype collaborative feature modeling system webSpiff.

3.1 Overview of webSpiff Architecture. webSpiff has a client-server architecture. As a basis for the server, the Spiff system developed at Delft University of Technology was chosen, which offers several advanced modeling facilities. First, it offers multiple views on a product, each view consisting of a feature model with features specific for the application corresponding to the view. webSpiff provides two such views: one for design and another for manufacturing planning of parts. In the design view, the feature model consists of both additive features (e.g. protrusions) and subtractive features (e.g. slots and holes). In the manufacturing planning view, the feature model consists of only subtractive features. Both views on a part are kept consistent by feature conversion [3]. Second, it offers feature validity maintenance functionality. This can guarantee that only valid feature models, i.e. models that satisfy all specified requirements, are created by a user [1]. Third, it offers sophisticated visualization techniques, which visualize much more specific feature information than most other systems do. For example, feature faces that are not on the boundary of the resulting object, such as closure faces of a through slot, can be visualized too [14]. All these facilities are computationally expensive, and require an advanced product model, including a cellular model with information on all features in all views [15], which uses functionality provided by the 3D ACIS Modeler [16].

In webSpiff, some of the functionality of the original Spiff modeling system, in particular for interaction with feature models, is moved to the clients. However, as soon as real feature model computations are required, such as for executing modeling operations, conversion between feature views, feature validity maintenance and feature model visualization, they are performed at the webSpiff server, on a central product model, and their results are eventually exported back to the clients.

webSpiff consists of several components, as depicted in the global architecture diagram of Fig. 1. On the server side, two main components can be identified: the Spiff modeling system, provid-

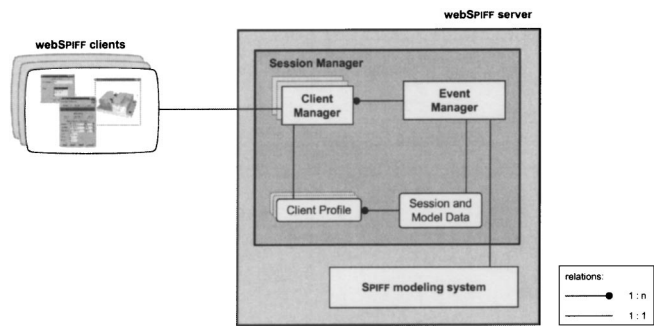


Fig. 1 Architecture of webSpiff.

ing all feature modeling functionality; and the Session Manager, providing functionality to start, join, leave and close a modeling session, and to manage all communication between Spiff and the clients.

The Session Manager stores information about an ongoing session and its participants. It manages all information streams between webSpiff clients and the Spiff process handling the session. Since several session participants can send modeling operations and queries to the webSpiff server at the same time, concurrency must be handled at the Session Manager. Practically, this means that parallel information streams have to be serialized. The Session Manager has been implemented using the Java programming language [17], and therefore its Remote Method Invocation (RMI) facilities were chosen for supporting the communication with webSpiff clients, above alternative approaches such as CORBA and SOAP.

The clients of webSpiff make use of standard web browsers. When a new client registers with webSpiff, a Java applet is loaded, implementing a simple user interface, from which a direct connection with the Session Manager is set up. Different clients can connect from various locations, in order to start or join a modeling session. Using standard web browsers at the clients increases accessibility and platform independence, but limits the complexity of the operations that can be implemented on them. Therefore, careful attention has been paid to make available to the clients, in an interactive way, as much functionality as possible of the original Spiff system.

Once connected to the server, a user can join an ongoing collaborative session, or start a new one, by specifying the product model he wants to work on. Also, the desired view on the part, i.e. design or manufacturing planning, has to be specified. Information on the feature model of that view is retrieved from the server, and used to build the client's graphical user interface (GUI), through which the user can start active participation in the modeling session; see Fig. 2.

Users can specify modeling operations in terms of features and their entities; for example, a feature, to be added to a model, is attachable to entities of features already in the model (e.g. faces and datums), rather than in terms of faces of the evaluated boundary representation of the product. Among other advantages, this approach avoids the well-known problem of persistent naming of model entities [1]. After a feature modeling operation, with all its operands, has been fully specified, the user can confirm the operation. The operation is then sent to the server, where it is checked for validity and scheduled for execution. Notice that this can result in an update of the product model on the server, and thus also of the feature model in the view of each session participant.

In addition to the above functionality, several visualization and interactive facilities have also been ported to the clients. webSpiff clients provide two ways of visualizing the product model, both making use of so-called *camera* windows. A camera window is a separate window in which a graphical representation of the product model is shown. Each client may create as many cameras

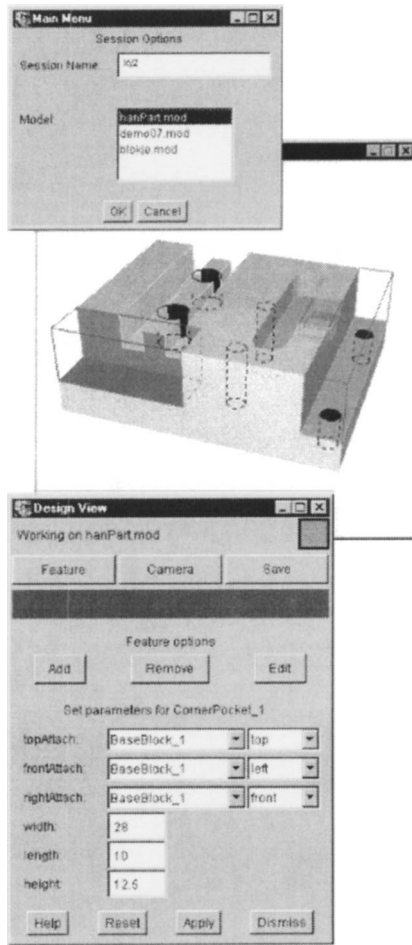


Fig. 2 User interface of a webSpiff client: the user, working on a design view, is modifying certain parameters of the corner pocket at the left-hand side of the displayed model.

as desired. First, a feature model image can be displayed. Second, a model can be rendered that supports interactive modification of camera viewing parameters, e.g. rotation and zoom operations. After the desired viewing parameters have been interactively set, they are sent to the server, where a feature model image corresponding to the new parameter values is rendered and sent back to the client, where it is displayed. Finally, webSpiff cameras also provide facilities for interactive specification of modeling operations, e.g. assisting the user in selecting features or feature entities by having them picked on an image of the model. The visualization and interactive functionality of webSpiff cameras is described in detail by van den Berg et al. [18].

To support all these facilities, the webSpiff clients need to locally dispose of some model data, as described in the next subsection.

3.2 Model Data at the Clients. As explained above, only one central product model is maintained at the server. This model includes all canonical shapes, representing individual features in a specific view, and the cellular model. Some model data, however, is also required at the clients. This data is derived by the server from the central model, but it does not make up a real feature model. webSpiff clients need just enough model information to be able to autonomously interact with the feature model, without continuously requesting feedback from the server.

Model data at the clients can be classified into the following categories.

Textual Data. This data is used for specific sets of model information, mostly in list form. The most important are:

- *List of feature classes:* contains the names of all feature classes available in a given view. It is used to fill a GUI list widget when adding a new feature instance, and is requested from the server at client initialization time. This list does not need to be refreshed during a modeling session.
- *List of feature instances:* contains the names of all feature instances in a given view of the model. It is used to fill a GUI list widget when editing or removing an existing feature instance. This list is set upon initialization of the client, and is refreshed after each modeling operation.
- *List of parameter values:* contains the values of all parameters of a given feature instance, in a pre-defined order. It is used to fill various GUI entry widgets when editing the feature instance, and is always queried before editing is started.

Feature Model Images. Feature model images are rendered at the Spiff server in GIF format, and displayed in camera windows at the clients. These images provide very powerful visualizations of a feature model. Many visualization options can be specified. For example, selected features may be visualized with shaded faces, and the rest of the model as a wire frame or with visible lines only. Also, additional feature information, such as closure faces of holes, can be visualized. A separate image is needed for each camera, and it must be updated every time the model or the camera settings are changed.

Visualization Model. The visualization model represents the global shape of the product model, and is generated by Spiff in VRML format [19]. It is used at the clients for interactively modifying the camera viewing parameters (e.g. rotating and zooming). Real-time rendering of this model is locally feasible, whereas rendering of a smooth sequence of feature model images at the server and transmitting it to the clients would be unfeasible in real time. All cameras on a particular client use the same local visualization model, but each camera displays it with its own viewing parameters.

Selection Model. The selection model is a collection of objects representing the canonical shapes of all features in a given view of the product. Its purpose is to support interactive selection of feature faces on a feature model image, during the specification of a modeling operation. Again, the selection model is identical for all cameras on a client, each applying its own viewing parameters. The selection model is also generated by Spiff in VRML format.

3.3 Data Communication. As obvious from this section, the various components of webSpiff have to exchange information at several stages during a modeling session. Communication among them plays therefore an important role in webSpiff.

webSpiff clients can specify modeling operations, camera operations, and a variety of queries, and send them to the Session Manager. Communication between the Session Manager and the clients uses the RMI functionality provided in Java. Messages sent from clients to the Session Manager are simple textual messages, whereas messages sent in the reverse direction may contain more complex objects, such as model data files.

At the server, a socket connection is used as the communication channel between the Session Manager and the Spiff modeling system. Textual messages are used to pass commands from the Session Manager to the Spiff modeling system. The system replies using also textual messages, but several data structures, such as feature model images and the visualization and selection models, are stored by Spiff into files, so that the Session Manager can easily propagate them to the clients.

4 The Server

As outlined in the previous section, the two main components of the webSpiff server are the Session Manager and the Spiff

modeling system. The functionality of Spiff has been summarized in Subsection 3.1. Therefore, this section is focused on the Session Manager, in particular its synchronization and concurrency management mechanisms.

4.1 The Session Manager. Two important types of processes run on the Session Manager; see Fig. 1. First, it maintains for each client a *Client Manager*, which receives messages from its client, interprets them, and either processes a message itself, or propagates it to the Spiff modeling system. Second, the Session Manager maintains an *Event Manager*, including an event queue that schedules the tasks received from all Client Managers, which have to be passed on to the Spiff modeling system. The Client Managers and the Event Manager run independently from each other, as so-called separate *Threads* in Java.

Each Client Manager maintains a Client Profile, which stores information about an individual client. This includes the user identification, the modeling view he is working on, and a list containing the names of the Cameras the client has opened.

Several types of tasks can be distinguished at the Client Managers. First, session operations have to be handled. These involve starting a session, logging into and out of a session, and closing a session. Second, modeling operations can be received that have to be forwarded to the Spiff modeling system, which, after executing them, returns their result to the Session Manager. Third, camera operations have also to be forwarded to the Spiff modeling system. Fourth, queries about the feature model are directly answered by the Session Manager, which keeps a record of up-to-date model data. In this way, queries can be directly answered, without involving the Spiff modeling system. The results of the last two types of tasks must be sent only to the client that issued the request. In all cases, sending results back to the clients is handled by the Event Manager, through the respective Client Managers, as will be explained in the following subsection.

4.2 Data Synchronization. Model data at the clients is never modified directly by the clients themselves. Instead, it is the task of the Session Manager to synchronize session participants, by sending them updated model data, after a modeling or camera operation has been processed. Several types of feedback are possible here, depending on the type of operation and whether the operation was successful.

When a client modifies any camera settings, the webSpiff server executes the corresponding camera operation, and a new feature model image is generated. Since the feature model remains unaffected by camera operations, the server only needs to send the new feature model image back to the client that requested the camera update.

However, after a modeling operation has been successfully executed, the central model has been changed, so relevant updated model data has to be sent to all other session participants as well. Messages are then sent to each client with the corresponding updates, consisting of, possibly several, new feature model images, a new visualization model, and an incremental update of the selection model (containing only new and/or modified feature canonical shapes). Also, additional information on the feature model is included in the messages, such as an updated list of feature instances. Separate messages have to be created per client, since each client typically has different cameras, and therefore requires different feature model images. Of course, model data will also be different for users of different modeling views. Upon receiving its update message, a client can extract the model data in order to update its data structures.

In order to reduce “collaboration noise,” due to too many distracting updates, clients may analyze updates, and decide to postpone processing them until some convenient synchronization time. In any case, temporary local inconsistencies can still occur at a client. Since transmitting model data from the server to all clients is not instantaneous, for a short period model data on the

clients is not up-to-date. Avoiding conflicts arising from this transitory mismatch will be dealt with in Subsection 5.2.

4.3 Concurrency Handling. Concurrency must be handled at several stages. If this is not done properly, serious problems can arise, such as inconsistency of data structures, or processes indefinitely waiting for each other, i.e. deadlock.

Event Management. The most delicate case of concurrency occurs with event management at the Session Manager. Communication streams from all clients come together, requesting information to be sent back and modeling operations to be carried out. The Session Manager serializes data that arrives in parallel streams.

As seen in Subsection 4.1, each client is represented in the Session Manager by a Client Manager. When two events arrive at the Session Manager at the same moment, only one Client Manager should be able to add an event to the event queue at a time, so the system must determine which Client Manager will be allowed to add the event first. The Java programming language provides useful locking mechanisms for this purpose, the most important being the *synchronized* mechanism. When a class is accessed by a synchronized method, all its methods and data structures are locked, preventing them from being accessed by another process at the same time.

Conflicting Modeling Operations. webSpiff encourages users to coordinate their actions, e.g. using the phone or a chat channel, in order to avoid conflicting operations. To assist in this goal, a *traffic light* icon is displayed on the user interface of every client, informing about the busy state of other clients and of the webSpiff server. This icon switches from green to yellow when another client starts specifying a modeling operation, and to red when the server starts executing a modeling operation. It has been decided not to implement strict token passing policies, since modeling is considered to be a constructive activity. Additional communication between users will always remain necessary, because it does not make much sense to have several users performing modeling operations simultaneously, without any coordination.

Still, it could happen that two clients simultaneously submit a modeling operation. Consider the following situation: the modeling operation that is handled first by the Session Manager removes a certain feature, while the second modeling operation tries to edit this same feature. In a single user modeling system, e.g. Spiff, this situation could not occur, since operations are always performed serially: after a feature has been removed, it is not possible to subsequently edit it, since it cannot be selected anymore in the user interface. In webSpiff, however, where operations are performed concurrently, an operation can be specified on features that no longer exist at the time of its execution. Therefore, the Event Manager always checks the existence of each operation’s arguments, before issuing its execution in Spiff. A user requesting to execute an operation on no longer existing features is notified that the operation is not meaningful anymore.

5 The Clients

The webSpiff clients provide a remote user interface to the users of webSpiff. In order to offer them the same interactive functionality as the Spiff modeling system does, it is not enough to just replicate the user interface of Spiff at the clients. As described in Subsection 3.2, webSpiff clients maintain some data structures with model information, used to provide various interactive facilities. The interactive functionality of webSpiff clients was briefly summarized in Subsection 3.1. This section mainly elaborates some communication issues regarding client model data, in particular its synchronization at the clients.

Several components can be identified within a webSpiff client, as shown in Fig. 3. The user interface is the component used for interactively specifying all operations, by means of panels, menus, buttons and list boxes. Three major components can be distin-

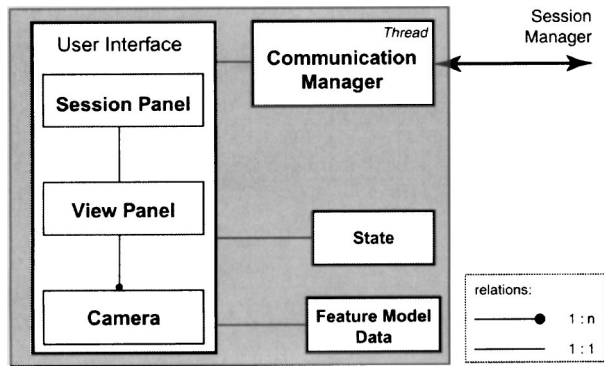


Fig. 3 Architecture of a webSpiff client.

guished, namely the View Panel, the Session Panel and the Cameras. The first two provide plain interfaces with standard widgets, the latter functionality for graphical interaction with the model.

5.1 The Communication Manager. The Communication Manager on the client manages all communication with the Session Manager on the server. Two kinds of messages can be identified here:

- a) messages whose response is not awaited, before the client can be operated again; these messages include all modeling operations, such as adding, removing and editing a feature;
- b) messages whose response is awaited, suspending all activities on the client; these messages include queries and camera update messages.

When a message of type b) is sent, all activity on the client is suspended until a reply to it is received. While in this state, however, it is still possible that other messages arrive at the client earlier than the expected response. These messages are stored in a queue, and processed after the expected response has been received and processed. An exception is made here for messages concerning the update of state information; see the next subsection.

The reason that camera update messages fall into the second category is that, after a camera operation has been specified and sent to the Session Manager, the interactive functionality cannot be used any more to continue operating webSpiff. Specifying camera operations is typically done interactively, using the visualization model, whereas specifying modeling operations is typically done using the feature model image, in combination with the selection model; see Subsection 3.2. However, after the viewing parameters of the visualization model have been modified, the feature model image is inconsistent with both the visualization model and the selection model. As result, no entities for modeling operations can be selected on the feature model image until an updated image arrives. For this reason, the client is suspended until the new image has been received.

Messages sent by a client to the Session Manager take imperceptible transmission times, as they consist of a compact string argument of a single RMI call only. Messages received at a client from the Session Manager are, however, more complex, typically containing multiple objects. The first object in all messages is a command string, which is parsed by the Communication Manager to determine the message type, and how it should be handled. In addition, messages can contain requested information on the feature model, such as a list of feature instances, or model data files. Such files are small, and therefore their transmission times are short. In the example of Fig. 2, a feature model image takes less than 10 Kbytes, the selection model less than 5 Kbytes per feature, and the visualization model less than 50 Kbytes. Taking into account these file sizes, it can be questioned whether compressing

them before transmission to the clients would further improve system throughput, due to the overhead introduced by the compression and decompression algorithms.

5.2 Synchronization. Before data structures can be updated, it must be made sure that the involved clients are in the right state for processing the update. Two types of data can be distinguished here: (i) updated model data, resulting from a modeling or a camera operation, and (ii) updated state information. The order in which these updates are received at the clients is not known in advance, and several scenarios must therefore be handled. Two scenarios are described here.

In webSpiff, the feature model can be modified at any time by one of the users. After such a modification, new model data will have to be sent to all clients. At a client side, however, preparation of a new modeling operation could be underway when the update arrives. It is not convenient to force the user to cancel his operation, since it might well be that the user coordinated his operation with the other users, e.g. using a communication channel outside webSpiff, in a way that the update does not have any influence on the modeling operation being specified. For example, a user could be editing a feature, while the update concerns another feature, having no influence on the edit feature operation that is going on. Canceling the operation in this case, would mean that all the parameters that had been specified so far would have to be specified again. Therefore, the user is allowed to continue specifying the modeling operation, but he is notified of the modeling operation that has been carried out at the server. He can then choose to continue specifying his operation, or to cancel it himself.

Besides updating model data at the client, also state information, such as the traffic light icon mentioned in Subsection 4.3, must be kept up to date. The difference with the other data structures is, however, that state information must always be processed as soon as possible. Whereas other messages that arrive unexpectedly can be put into a queue at the Communication Manager, awaiting their processing, the most recent state information must become available to the client immediately, since its purpose is to inform the user of the current state of the modeling session. Therefore, for every incoming message, it is immediately checked whether it is a state update message.

6 Results and Conclusions

Current trends in product development demand from CAD systems not only advanced modeling facilities, but also that these be concurrently available to distributed multiple users, effectively supporting collaboration sessions among the members of a development team. This paper addresses the new challenges of such requirements. The problems of concurrency and synchronization in a collaborative modeling system can best be handled if a client-server architecture is adopted. Moreover, a web-based approach has additional advantages, although this requires a careful balance between the conflicting requirements of good client interactivity and low network load.

A new web-based collaborative feature modeling system, webSpiff, has been presented that provides a solution for many issues involved in collaborative modeling systems, including concurrency, synchronization and user interaction. The proposed distribution of functionality between the server and the clients has resulted in a well-balanced system. On the one hand, the full functionality of an advanced feature modeling system is offered by the server. On the other hand, all desirable interactive modeling functionality is offered by the clients, ranging from display of feature model images to interactive selection facilities.

All functionality described in this paper has been implemented in the webSpiff prototype system. The webSpiff server runs on a Linux workstation, with a Pentium 4 running at 1.6 GHz. Its performance exceeds that required for supporting synchronous collaborative modeling sessions, which will typically involve no more than 10 users. The Java-based client application is quite simple. So far, webSpiff clients running on Unix, Windows and

Linux platforms have successfully participated in collaborative sessions. The only requirement at the client side is that it needs to have a Java-enabled web browser, with the Java3D API installed. The webSpiff portal has a demo version available on Internet for users to experiment with, at www.webSpiff.org.

Secure transmission and data protection are important issues in any collaborative modeling environment. Although they were not directly considered in the scope of this project, they should receive careful attention, possibly profiting from alternative communication protocols such as SOAP.

As Internet technology rapidly improves, faster and better collaboration becomes possible. It can therefore be expected that, although the development of collaborative modeling systems is still at its early stages, such systems will soon play an important role in the product development process.

Acknowledgments

This article is a revised version of the paper CIE-21286, previously published in the Proceedings of ASME 2001 Design Engineering Technical Conferences. We thank the reviewers for many valuable comments.

References

- [1] Bidarra, R., and Bronsvort, W. F., 2000, "Semantic feature modelling." *Comput.-Aided Des.*, **32**(3), pp. 201–225.
- [2] Kagan, P., Fischer, A., and Bar-Yoseph, P. Z., 1999, "Integrated Mechanically-based CAE System," Proceedings of Solid Modeling '99—Fifth Symposium on Solid Modeling and Applications, Bronsvort, W. F. and Anderson, D.C (Eds.), ACM Press, New York, pp. 23–30. Also in: *Comput.-Aided Des.*, **32**(8/9), pp. 539–552.
- [3] de Kraker, K. J., Dohmen, M., and Bronsvort, W. F., 1997, "Maintaining Multiple Views in Feature Modeling." Proceedings of Solid Modeling '97—Fourth Symposium on Solid Modeling and Applications, Hoffmann, C.M. and Bronsvort, W.F. (Eds.), ACM Press, New York, pp. 123–130.
- [4] Hoffmann, C. M., and Joan-Arinyo, R., 1998, "CAD and the Product Master Model," *Comput.-Aided Des.*, **30**(11), pp. 905–918.
- [5] Lewandowski, S., 1998, "Frameworks for Component-Based Client/Server Computing," *ACM Comput. Surv.*, **30**(1), 3–27.
- [6] Parametric, 2001, Pro/ENGINEER 2001i. Parametric Technologies Corporation, Waltham, MA. <http://www.ptc.com>.
- [7] SDRC, 2001, I-DEAS. SDRC, Milford, OH. <http://www.sdrc.com>.
- [8] Kaon, 2001, HyperSpace-3DForum. Kaon Interactive Inc., Cambridge, MA. <http://www.kaon.com>.
- [9] CoCreate, 2002, One Space Suite Solution, CoCreate Software, Inc., Fort Collins, CO. <http://www.cocreate.com>.
- [10] Nam, T. J., and Wright, D. K., 1998, "CollIDE: A Shared 3D Workspace for CAD," Proceedings of the 1998 Conference on Network Entities, Leeds.
- [11] Stork, A., and Jasnoch, U., 1997, "A Collaborative Engineering Environment," Proceedings of TeamCAD '97 Workshop on Collaborative Design, Atlanta, GA, pp. 25–33.
- [12] Chan, S., Wong, M., and Ng, V., 1999, "Collaborative Solid Modelling on the WWW," Proceedings of the 1999 ACM Symposium on Applied Computing, San Antonio, CA, pp. 598–602.
- [13] Lee J. Y., Kim, H., Han, S. B., and Park, S. B., 1999 "Network-centric Feature-based Modeling." Proceedings of Pacific Graphics '99, Kim, M.-S. and Seidel, H.-P. (Eds.), IEEE Computer Society, Los Alamitos, pp. 280–289.
- [14] Bronsvort, W. F., Bidarra, R., and Noort, A., 2002, "Feature Model Visualization," To be published in: *Computer Graphics Forum*, **21**(4).
- [15] Bidarra, R., de Kraker, K. J., and Bronsvort, W. F., 1998, "Representation and Management of Feature Information in a Cellular Model," *Comput.-Aided Des.*, **30**(4), 301–313.
- [16] Spatial, 2002, 3D ACIS Modeler, Version 7.0. Spatial Technology Inc., Westminster, CO. <http://www.spatial.com>.
- [17] Sun Microsystems, 2002 The Sun Java™ Technology Homepage. <http://java.sun.com>
- [18] van den Berg, E., Bidarra, R., and Bronsvort, W. F., 2000, "Web-based Interaction on Feature Models," In: *From Geometric Modeling to Shape Modeling*, Cugini, U. and Wozny, M. (Eds.), Kluwer Academic Publishers, Dordrecht, pp. 99–112.
- [19] Ames, A., Nadeau, D., and Moreland, J., 1997, *The VRML 2.0 Sourcebook*, Second Edition, John Wiley & Sons, New York.