EdiFlow: data-intensive interactive workflows for visual analytics

Veronique Benzaken, Jean-Daniel Fekete, Pierre-Luc Hémery, Wael Khemiri, Ioana Manolescu

► To cite this version:

Veronique Benzaken, Jean-Daniel Fekete, Pierre-Luc Hémery, Wael Khemiri, Ioana Manolescu. Ed-iFlow: data-intensive interactive workflows for visual analytics. International conference on Data Engineering, IEEE, Apr 2011, Hannover, Germany. inria-00532552

HAL Id: inria-00532552

<https://hal.inria.fr/inria-00532552>

Submitted on 28 Sep 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

EdiFlow: data-intensive interactive workflows for visual analytics

Vronique Benzaken ², Jean-Daniel Fekete ¹, Pierre-Luc Hmery ¹, Wael Khemiri ^{1,2}, Ioana Manolescu ^{1,2}

¹INRIA Saclay-Île-de-France
4, rue Jacques Monod 91893 Orsay, France
firstname.lastname@inria.fr

²LRI, Université de Paris Sud-11
Bat 490, Université Paris-Sud 11 91405 Orsay, France
firstname.lastname@lri.fr

Abstract—Visual analytics aims at combining interactive data visualization with data analysis tasks. Given the explosion in volume and complexity of scientific data, e.g., associated to biological or physical processes or social networks, visual analytics is called to play an important role in scientific data management.

Most visual analytics platforms, however, are memory-based, and are therefore limited in the volume of data handled. Moreover, the integration of each new algorithm (e.g. for clustering) requires integrating it by hand into the platform. Finally, they lack the capability to define and deploy well-structured processes where users with different roles interact in a coordinated way sharing the same data and possibly the same visualizations.

We have designed and implemented *EdiFlow*, a workflow platform for visual analytics applications. *EdiFlow* uses a simple structured process model, and is backed by a persistent database, storing both process information and process instance data. *EdiFlow* processes provide the usual process features (roles, structured control) and may integrate visual analytics tasks as activities. We present its architecture, deployment on a sample application, and main technical challenges involved.

I. INTRODUCTION

The increasing amounts of electronic data of all forms, produced by humans (e.g., Web pages, structured content such as Wikipedia or the blogosphere) and automatic tools (loggers, sensors, Web services, scientific programs or analysis tools) lead to a situation of unprecedented potential for extracting new knowledge, finding new correlations, and interpreting data. Visual analytics is a new branch of the information visualization / human-computer interaction field [1]. Its aim is to enable users to closely interact with vast amounts of data using visual tools. Thanks to these tools, a human may detect phenomena or trigger detailed analysis which may not have been identifiable by automated tools alone. Visual analytics tools routinely include some capacity to mine or analyze the data; however, most applications require specific analysis functions.

Though, most current visual analytics tools have some conceptual drawbacks. Indeed, they rarely rely on persistent databases (with the exception of [2]). Instead, the data is loaded from files or databases and is manipulated directly in

memory because smooth visual interaction requires redisplaying the manipulated data 10-25 times per second. Standard database technologies do not support continuous queries at this rate; at the same time, ad-hoc in-memory handling of classical database tasks (e.g., querying, sorting) has obvious limitations. Based on our long-standing experience developing information visualisation tools [3] [4] [5], we argue connecting a visual analysis tool to a persistent database management system (DBMS) has many benefits:

- Scalability: larger data volumes can be handled based on a persistent DBMS
- Persistence and distribution: several users (possibly on remote sites) can interact with a persistent database, whereas this is not easily achieved with memory-resident data structures. Observe that users may need to share not only raw data, but also visualizations built on top of this data. A visualization can be seen as an assignment of visual attributes (e.g., X and Y coordinates, color, size) to a given set of data items. Computing the value of the visual attributes may be expensive, and/or the choice of the visualized items may encapsulate human expertise. Therefore, visualizations have *high added value* and it must be easy to store and share them, e.g., allowing one user to modify a visualization that another user has produced.
- Data management capabilities provided by the database: complex data processing tasks can be coded in SQL and/or some imperative scripting language. Observe that such data processing tasks can also include user-defined functions (UDFs) for computations implemented outside the database server. These functions are not stored procedures managed by the database (e.g., Java Stored Procedure). These are executable programs external to the database.

The integration of a DBMS in a visualisation platform must take into account the following prevalent aspects in today's visual analytics applications:

- Convergence of visual analytics and workflow: current visual analytics tools are not based on workflow (process)

models. This fits some applications where datasets and tasks are always exploratory and different from one session to the next. Several visual analytics applications however, require a recurring process, well supported by a workflow system. The data processing tasks need to be organized in a sequence or in a loop; users with different roles may need to collaborate in some application before continuing the analysis. It also may be necessary to log and allow inspecting the advancement of each execution of the application. (Scientific) workflow platforms allow such automation of data processing tasks. They typically combine database-style processing (e.g., queries and updates) with the invocation of external functions, implementing complex domain-dependent computations. Well-known scientific workflow platforms include Kepler [6], Taverna [7], or Trident [8]. These systems build on the experience of the data and workflow management communities; they could also benefit from a principled way of integrating powerful visualisation techniques.

- Handling dynamic data and change propagation: an important class of visual analytics applications has to deal with dynamic data, which is continuously updated (e.g., by receiving new additions) while the analysis process is running; conversely, processes (or visualisation) may update the data. The possible interactions between all these updates must be carefully thought out, in order to support efficient and flexible applications.

Our work addresses the questions raised by the integration of a DBMS in a visual analytics platform. Our contributions are the following:

- 1) We present a generic architecture for integrating a visual analytics tool and a DBMS. The integration is based on a core data model, providing support for (i) visualisations, (ii) declaratively-specified, automatically-deployed workflows, and (iii) incremental propagation of data updates through complex processes, based on a high-level specification. This model draws from the existing experience in managing data-intensive workflows [9], [10], [11], [12].
- 2) We present a simple yet efficient protocol for swiftly propagating changes between the DBMS and the visual analytics application. This protocol is crucial for the architecture to be feasible. Indeed, the high latency of a "vanilla" DBMS connection is why today's visual analytics platforms do not already use DBMSs.
- 3) We have fully implemented our approach in a bare-bones prototype called EdiFlow, and de facto ported the InfoVis visual analytics toolkit [4] on top of a standard Oracle server. We validate the interest of our approach by means of three applications.

This article is organized as follows. Section II compares our approach with related works. Section III describes three applications encountered in different contexts, illustrating the problems addressed in this work. Section IV presents our proposed data model, while the process model is described

in Section V. We describe our integration architecture in Section VI, discuss some aspects of its implementation in our EdiFlow platform, we then conclude in Section VIII.

II. RELATED WORKS

Significant research and development efforts have resulted in models and platforms for workflow specification and deployment. Recently, *scientific workflow* platforms have received significant attention. Different from regular (business-oriented) workflows, scientific workflows notably incorporate data analysis programs (or scientific computations more generally) as a native ingredient. Moreover, scientific workflows are meant to be specified by scientists: their end users. This contrasts with business workflows, usually specified by business analysts which do not enact them. Both business and scientific workflows are, by now, commonly deployed relying on a DBMS for data storage and/or process control. Due to the importance of visualisation and interaction, and to the exploratory nature of visual analytics, we position our work with respect to scientific workflows, to which it relates more closely than usual business workflows.

One of the first integration of scientific workflows with DBMSs was supported by [9]. Among the most recent and well-developed scientific workflow projects, Kepler [13] is designed to help scientists, analysts, and computer programmers to create, execute, and share models and analyses across a broad range of scientific and engineering disciplines. Kepler provides a GUI which helps users to select and then connect analytical components and data sources to create a scientific workflow. In this graphical representation, the nodes in the graph represent actors and the vertices are links between the actors.

SciRun [14] is a Problem Solving Environment, for modeling, simulation and visualization of scientific problems. It is designed to allow scientists to interactively control scientific simulations while the computation is running. SCIRun was originally targeted at computational medicine but has, later, been expanded to support other scientific domains. The SCIRun environment provides a visual interface for dataflow network's construction. As the system will allow parameters to be changed at runtime, experimentation is a key concept in SCIRun. As soon as a parameter is updated, at runtime, changes will be propagated through the system and a re-evaluation induced.

GPFlow [15] is a workflow platform providing an intuitive web based environment for scientists. The workflow model is inspired by spreadsheets. The workflow environment ensures interactivity and isolation between the calculation components and the user interface. This enables workflows to be browsed, interacted with, left and returned to, as well as started and stopped.

VisTrails [16] combines features of both workflow systems and visualization fields. Its main feature is to efficiently manage exploratory activities. The user interaction in VisTrails is performed by iteratively refining computational tasks and formulating test hypotheses. VisTrails maintains detailed

provenance of the exploration process. Users are able to return to previous versions of a dataflow and compare their results. However, VisTrails is not meant to manage dynamic data. In VisTrails, dynamicity is performed by allowing users to change some attributes in order to compare visualization results. It does not include any model to handle data changes. Indeed, when the user starts its workflow process, VisTrails does not take into account the updated data in activities that have already started: there is no guarantee that the model for updates is correct.

Trident [8], [17] is a scientific workflow workbench built on top of a commercial workflow system. It is developed by Microsoft corporation to facilitate scientific workflows management. Provenance in Trident is ensured using a publication/subscription mechanism called the Blackboard. This mechanism allows also for reporting and visualizing intermediate data resulting from a running workflow. One of the salient features of Trident is to allow users to dynamically select where to store results (on SQL Server for example) issued by a given workflow. However, it does not support dynamic data sources nor does it integrate mechanisms to handle such data.

Orchestra [18], [19] addresses the challenge of mapping databases which have potentially different schemas and interfaces. Orchestra is specially focusing on bioinformatics applications. In this domain, one find many databases containing overlapping informations with different level of quality, accuracy and confidence. Database owners want to store a relevant ("alive") version of relevant data. Biologists would like to download and maintain local "live snapshots" of data to run their experiments. The Orchestra system focus on reconciliation across schemas. It is a fully peer-to-peer architecture in which each participant site specifies which data it trusts in. The system allows all the sites to be continuously updated, and on demand, it will propagate these updates across sites. User interaction in Orchestra is only defined at the first level using trust conditions. Moreover, the deployed mechanism is not reactive. Indeed, there is no restorative functions called after each insert/update operation.

Several systems were conceived to create scientific workflows using a graphical interface and enabling data mining tasks (e.g., Knime [20] and Weka [21], RapidMiner [22], Orange [23], [24]). However, none of these systems includes a repair mechanism to support the change in data sources during a task or process execution.

To summarize, all these platforms share some important features, which we also base our work on. Workflows are *declaratively specified*, *data-intensive* and (generally) *multi-user*. They include *querying and updating* data residing in some form of a database (or in less structured sources). Crucial for their role is the ability to invoke *external procedures*, viewed as black boxes from the workflow engine perspective. The procedures are implemented in languages such as C, C++, Matlab, Fortran. They perform important domain-dependent tasks; *procedures may take as input and/or produce as output large collections of data*. Finally, current scientific workflow

Fig. 1. US Election screen shot.

platforms do provide, or can be coupled with, some *visualisation* tools, e.g., basic spreadsheet-based graphics, map tools.

With respect to these platforms, our work makes two contributions: (i) we show how a generic data visualisation toolkit can be integrated as a *first-class citizen*; (ii) we present a *principled way of managing updates* to the underlying sources, throughout the enactment of complex processes. This problem is raised by the high data dynamicity intrinsic to visual analytics applications. However, the scope of its potential applications is more general, as long-running scientific processes may have to handle data updates, too. *None of these platforms are currently able to propagate data changes to a running process*. The process model we propose could be integrated, with some modest programming effort, in such platforms, hence offering complementary functionalities to their existing ones.

Most of existing interactive platforms for data visualization [3], [25] focus on the interaction between the human expert and a data set consisting of a completely known set of values. They do not ease the inclusion of data analysis programs on the data. Moreover, as previously explained, most of them do not support the definition of structured processes, nor (by absence of an underlying DBMS) do they support persistence and sharing. An exception is [2] which is a visualization tool combining database technology. However, there is no repair mechanism and the change propagation is not supported.

Unlike current data visualisation platforms, our work provides a useful coupling to DBMSs, providing persistent storage, scalability, and process support. Our goal is to drastically reduce the programming effort actually required by each new visual analytics application, while enabling them to scale up to very large data volumes. In this work, we present an architecture implementing a repair mechanism, to propagate data source changes to an executing process.

III. USE CASES

The following applications illustrate the data processing and analysis tasks which this work seeks to simplify.

a) *US Elections*: This application aims at providing a dynamic visualisation of elections outcome, varying as new election results become available. The database contains, for

Fig. 2. Wikipedia screen shot.

each state, information such as the party which won the State during the last three elections, the number of voters for the two candidates, the total population of the state. On the voting day, the database gradually fills with new data. This very simple example uses a process of two activities: computing some aggregates over the votes, and visualizing the results. Upon starting, a TreeMap visualisation is computed over the database (distinguishing the areas where not enough data is available yet), as shown in Figure 1. The user can choose a party, then the 51 states are shown with varying color shades. The more the states vote for the respective party, the darker the color. When new vote results arrive, the corresponding aggregated values are recomputed, and the visualisation is automatically updated.

b) Wikipedia: The goal of the application is to propose to Wikipedia readers and contributors some measures related to the history of an article. e.g., how many authors contributed to an article? How did a page evolve over time? A more complex metric is how "durable" are the contributions of a given user? This last metric corresponding to the inverse number of characters inserted by the user divided by the characters remaining in the latest version. The challenge is then, to compute and store such metrics for the whole Wikipedia database. It can then be displayed to users close to the pages (s)he consults [26] or explored more thoroughly [27]. One must also update those metrics as the Wikipedia site is updated. The metrics are produced and visualized by the application, whereas the (current) Wikipedia page is displayed directly from the original site, as shown in Figure 2.

This application can be decomposed in four elementary tasks: (i) compute the differences between successive versions of each article; (ii) compute a contribution table, storing at each character index, the identifier of the user who entered it; (iii) for each article, compute the number of distinct effective contributors; and (iv) compute the total contribution (over all contribution tables) of each user. All these computations' results must be continuously updated to reflect the continuous changes in the underlying data. A total recomputation of the

aggregation is out of reach, because change frequency is too high (10 edits per second on average for the French Wikipedia, containing about 1 million pages). Moreover, updates received at a given moment only affect a tiny part of the database. Thus, the Wikipedia application requires: a DBMS for storing huge amounts of data; a well-defined process model including ad-hoc procedures for computing the metrics of interest; incremental re-computations; and appropriate visualisations.

c) INRIA activity reports: We have been involved in the development of an application seeking to compute a global view of INRIA researchers by analysing some statistics. The data are collected from Raweb (INRIA's legacy collection of activity reports available at <http://ralyx.inria.fr>). These data include informations about INRIA teams, scientists, publications and research centres. Currently, the report of each team from each year is a separate XML file; new files are added as teams produce new annual reports. Our goal was to build a self-maintained application which, once deployed, would automatically and incrementally re-compute statistics, as needed. To that end, we first created a database out of all the reports for the years 2005 to 2008. Simple statistics were then be computed by means of SQL queries: age, team, research center distribution of INRIA's employees. Other aggregates were computed relying on external code such as the similarity between two people referenced in the reports in order to determine whether an employee is already present in the database or needs to be added.

All these applications feature data- and computation-centric processes which must react to data changes while they are running and need visual data exploration. The Wikipedia application is the most challenging, by the size of the database, the complexity of its metrics, and the high frequency of updates requiring recomputations.

IV. DATA MODEL

In this Section, we describe our conceptual data model in Section IV-A, and its concrete implementation in a relational database in Section IV-B.

A. Conceptual data model

The conceptual data model of visual analytics application is depicted in Figure 3. For the sake of readability, entities and relationships are organized in several groups.

The first group contains a set of entities capturing *process definitions*. A process consists of some activities. An activity must be performed by a different group of users (one can also see a group as a role to be played within the process). Process control flow is not expressed by the data model, rather, it is described in the process model (see Section V). An activity instance has a start date and an end date, as well as a *status* flag ranging in the following set of values: {*not_started*, *running*, *completed*}. The flag *not_started* states that the activity instance is created by a user who assigns it to another for completion, but the activity's task has not started yet. The *running* flag indicates that the activity instance has started and

Fig. 3. Entity-relationship data model for EdiFlow.

has not yet finished. Finally, the flag *completed* means that the activity instance has terminated. Process instances will also take similar values.

Entities in the second group allow recording *process execution*. Individual users may belong to one or several groups. A user may perform some activity instances, and thus be involved in specific process instances. A *ConnectedUser* records the host and port from which a user connects at a given time. This information is needed to propagate updates, received while the process is running, to a potentially remote visualisation component running on the remote user's desktop. This point will be further discussed in Section VI.

The gray area can be seen as a meta-model, which has to be instantiated for any concrete application with one or several entities and relationships modelling it. For instance, in the Wikipedia application, one would use the entities *Article*, *User*, and *Version*, with relationships stating that each version of an article is produced by one user's article update. Black-box functions, such as Wikipedia user clustering functions, must also be captured by this application-dependent part of the data model. Tracking workflow results requires at a simple level that for each data instance, one may identify which activity instance which created it, updated it etc. To that purpose, specific customized relationships of the form *createdBy*, *validatedBy* may be defined in the conceptual model. They are represented in Figure 3 by the gray background relationship between *ApplicationEntity* and *ActivityInstance*. Of course, many more complex data provenance models can be devised e.g., [16], [28]. This aspect is orthogonal to our work.

The third group of entities is used to model visualization. A *Visualization* consists of one or more *VisualisationComponent*

ments. Each component offers an individual perspective over a set of entity instances. For example, in Figure 2, three visualisation components are shown in the bar at the left of the article, making up a given visualization associated with the article's edit history. Components of a same visualisation correspond to different ways of rendering the same objects. In each visualisation component, a specific set of *VisualAttributes* specifies how each object should be rendered. Common visual attributes include (x, y) coordinates, width, height, color, label (a string), whether the data instance is currently *selected* by a given visualisation component (which typically triggers the recomputation of the other components to reflect the selection).

Finally, the *Notification* entity is used to speedily propagate updates to the application entities in various places within a running process. A notification is associated with one or more instances of a particular application entity. It refers to an update performed at a specific moment indicated by the *seq_no* timestamp, and indicates the kind of the update (insert/delete/modify). Its usage is detailed in Section VI.

B. Concrete data model

We assume a simple *relational* enactment of this conceptual model. We have considered XML but settled for relations since performant visualisation algorithms are already based on a tabular model [4]. Thus, a relation is created for each entity endowed with a primary key. Relationships are captured by means of association tables with the usual foreign key mechanism. By issuing a query to the database, one can determine "which are the completed activity instances in process *P*", or "which is the *R* tuple currently selected by the user from the visualization component *VC*₁".

Process	::=	Configuration	Constant*
Configuration	::=	Variable+ Relation+	
Constant	::=	Function* StructProcess	
Variable	::=	DBdriver DBuri DBuser	
Relation	::=	name value	name $\in N$, value $\in V$
RelationType	::=	name type	name $\in N$, type $\in T$
Function	::=	name primaryKey RelType	
StructuredProcess	::=	(attName attType)*, attName $\in N$, attType $\in T$	
Sequence	::=	Activity	
AndSplitJoin	::=	Activity Sequence AndSplitJoin OrSplitJoin ConditionalProcess	
OrSplitJoin	::=	AND-split (StructuredProcess)+	
ConditionalProcess	::=	AND-join	
Activity	::=	OR-split (StructuredProcess)+	
Expression	::=	OR-join	
	::=	IF Condition StructuredProcess	
	::=	activityName Expression	
	::=	askUser callFunction run-Query	

Fig. 4. XML schema for the process model.

We distinguish two kinds of relations. DBMS-hosted relations are by definition persistent inside a database server and their content is still available after the completion of all processes. Such relations can be used in different instances, possibly of different processes. In contrast, temporary relations are memory-resident, local to a given process instance (their data are not visible and cannot be shared across process instances) and their lifespan is restricted to that of the process instance which uses them. If temporary relation data are to persist, they can be explicitly copied into persistent DBMS tables, as we shortly explain below.

V. PROCESS MODEL

We consider a process model inspired by the basic Workflow Management Coalition model [29]. Figure 4 outlines (in a regular expression notation) the syntax of our processes. We use a set of variables, constants and attribute names N , a set of atomic values V , and a set of atomic data types T ; terminal symbols used in the process structure are shown in boldface. The main innovative ingredient here is the treatment of *data dynamics*, i.e., the possibility to control which changes in the data are propagated to which part(s) of which process instances. We now describe the process model in detail.

Relations and queries A process is built on top of a set of relations implementing the data model. Relations are denoted by capital letters such as R, S, T , possibly with subscripts. A query is a relational algebraic expression over the relations. We consider as operators: selection, projection, and cartesian product. Queries are typically designated by the letter Q possibly with subscripts.

Variables A variable is a pair composed of a name, and of an (atomic) value. Variables come in handy for modelling useful constants, such as, for example, a numerical threshold for a clustering algorithm. Variables will be denoted by lower-case letters such as v, x, y .

Procedures A procedure is a computation unit implemented by some external, black-box software. A typical example is the code computing values of the visual attributes to be used in a visualisation component. Other examples include e.g., clustering algorithms, statistical analysis tools.

A procedure takes as input l relations $R_1, R_2, \dots, R_l$ which are read but not changed and m relations $T_1^w, T_2^w, \dots, T_m^w$ which the procedure may read *and* change, and outputs data in n relations:

$$p : R_1, R_2, \dots, R_l, T_1^w, T_2^w, \dots, T_m^w \rightarrow S_1, S_2, \dots, S_n$$

We consider p as a black box, corresponding to software developed outside the database engine, and outside of EdiFlow by means of some program expressed e.g., in C++, Java, MatLab. Functions are processes with no side effects ($m = 0$).

Delta handlers Associated to a procedure may be *procedure delta handlers*. Given some update (or delta) to a procedure input relation, the delta handler associated to the procedure may be invoked to propagate the update to a process. Two cases can be envisioned:

- 1) Update propagation is needed while the procedure is being executed. Such is the case for instance of procedures which compute point coordinates on a screen, and must update the display to reflect the new data.
- 2) Updates must be propagated after the procedure has finished executing. This is the case for instance when the procedure performs some quantitative analysis of which only the final result matters, and such that it can be adjusted subsequently to take into account the deltas.

The designer can specify one or both of these handlers. Formally, each handler is a procedure in itself, with a table signature identical to the main procedure. The convention is that if there are deltas only for some of p 's inputs, the handler will be invoked providing empty relations for the other inputs. With respect to notations, $p_{h,r}$ is the handler of p to be used while p is running, and $p_{h,f}$ is the handler to be used after p finished. Just like other procedures, the implementation of handlers is opaque to the process execution framework. This framework, however, allows one to recuperate the result of a handler invocation and inject it further into the process, as we shall see.

Distributive procedures An interesting family of procedures are those which distribute over union in all their inputs. More formally, let X be one of the R_i inputs of p , and let ΔX be the set of tuples added to X . If p is distributive then:

$$p(R_1, \dots, X \cup \Delta X, \dots, T_m^w) = \begin{aligned} & p(R_1, \dots, X, \dots, T_m^w) \\ & \cup p(R_1, \dots, \Delta X, \dots, T_m^w) \end{aligned}$$

There is no need to specify delta handlers for procedures which distribute over the union, since the procedure itself can serve as handler.

Expressions We use a simple language for expressions, based on queries and procedures. More formally:

$$e ::= Q \mid p(e_1, e_2, \dots, e_n, T_1^w, T_2^w, \dots, T_p^w).t_j, 1 \leq j \leq m$$

The simplest expressions are queries. More complex expressions can be obtained by calling a procedure p , and retaining only its j -th output table. If p changes some of its input table, evaluating the expression may have side effects. If the side effects are not desired, p can be invoked by giving it some new empty tables, which can be memory-resident, and will be silently discarded at the end of the process. Observe that the first n invocation parameters are expressions themselves. This allows nesting complex expressions.

Activities We are now ready to explain the building blocks of our processes, namely activities.

$$a ::= v \leftarrow \alpha \mid \text{upd}(R) \mid (S_1, S_2, \dots, S_n) \leftarrow p(e_1, e_2, \dots, e_n, T_1^w, T_2^w, \dots, T_n^w)$$

Among the simplest activities are *variable assignments* of the form $v \leftarrow \alpha$. Another simple activity is a declarative update of a table R , denoted $\text{upd}(R)$. Unlike the table modifications that an opaque procedure may apply, these updates are specified by a declarative SQL statement. Finally, an activity may consist of invoking a procedure p by providing appropriate input parameters, and retaining the outputs in a set of tables.

Visualisation activities must be modeled as procedures, given that their code cannot be expressed by queries.

Processes A process description can be modelled by the following grammar:

$$P ::= \epsilon \mid a, P \mid P \parallel P \mid P \vee P \mid e?P$$

In the above, a stands for an activity. A process is either the empty process (ϵ), or a *sequence* of an activity followed by a process ($,$), or a *parallel (and) split-join* of two processes ($\parallel$), or an *or split-join* of two processes (with the semantics that once a branch is triggered, the other is invalidated and can no longer be triggered). Finally, a process can consist of a *conditional block* where an expression e (details below) is evaluated and if this yields true, the corresponding process is executed.

Reactive processes A *reactive process* can now be defined as a 5-tuple consisting of a set of relations, a set of variables, a set of procedures, a process and a set of *update propagations*. More formally:

$$RP ::= R^*, v^*, p^*, P, UP^*$$

An *update propagation* UP specifies what should be done when a set of tuples, denoted ΔR , are added to an application-dependent relation R , say, at $t_{\Delta R}$. Several options are possible. We discuss them in turn, and illustrate with examples.

- 1) Ignore ΔR for the execution of all processes which had started executing before $t_{\Delta R}$. The data will be added to R , but will only be visible for *process instances having started after $t_{\Delta R}$* . This recalls locking at process instance granularity, where each process operates on exactly the data which was available when the process started. We consider this to be the default behavior for

all updates to the relations part of the application data model.

Use case: A social scientist applies a sequence of semi-automated partitioning and clustering steps to a set of Wikipedia pages. Then, the scientist visualises the outcome. In this case, propagating new items to the visualisation would be disruptive to the user, which would have to interrupt her current work to help apply the previous steps to the new data.

- 2) Ignore ΔR for the execution of all *activities* which had started executing (whether they are finished or not) before $t_{\Delta R}$. However, *for a process already started, instances of a specific activity which start after $t_{\Delta R}$ may also use this data.*

Use case: The social scientist working on a Wikipedia fragment first has to confirm personal information, give some search criteria for the pages to be used in this process. Then, she must interact with a visualisation of the chosen pages. For this activity, it is desirable to provide the user with the freshest possible snapshot, therefore additions between the beginning of the process instance, and the moment when the user starts the last activity, should be propagated.

- 3) As a macro over the previous option and the process structure, one could wish for ΔR to be propagated to *instances of all activities that are yet to be started in a running process.*

Use case: Intuitively, data should not "disappear" during the execution of a process instance (unless explicitly deleted). In the previous use case, if we add an extra activity at the end of the process, that activity would typically expect to see the whole result of the previous one.

- 4) Propagate the update ΔR to *all the terminated instances of a given activity*. We can moreover specialize the behavior on whether we consider only activity instances whose *process instances have terminated*, only activity instances whose *process instances are still running*, or both.

Use case: We consider a process whose first activities are automatic processing steps, e.g., computing *diffs* between the old and the new version of a Wikipedia page, updating a user's contribution, the page history etc. The last activity is a visualisation one where the scientist should be shown fresh data. Typically, the visualisation activity will last for a while, and it may refresh itself at intervals, to reflect the new data. In this case, it makes sense to apply the automated processing activities to the new pages received while running the process instance, even after the respective activities have finished.

- 5) Propagate the update ΔR to all the running instances of a given activity, whether they had started before $t_{\Delta R}$ or not.

Use case: This may be used to propagate newly arrived tuples to all running instances of a visualisation activity,

Fig. 5. EdiFlow architecture.

to keep them up-to-date.

Formally then, an update propagation action can be described as:

$$UP ::= R, a, ((\text{'ta'}, (\text{'rp'} | \text{'tp'})) | \text{'ra'} | (\text{'fa'}, \text{'rp'}))$$

where R is a relation and a is an activity. An update propagation action describes a set of instances of activity a , to which the update ΔR must be propagated. The possible combinations of terminal symbols designate:

- ta rp:** terminated activity instances part of running processes;
- ta tp:** terminated activity instances part of terminated processes;
- ra:** running activity instances (obviously, part of running processes);
- fa rp:** future activity instances part of running processes.

It is possible to specify more than one compensation action for a given R and a given activity a . For instance, one may write: $(R, a, \text{'ra'}), (R, a, \text{'fa'}, \text{'rp'})$.

For simplicity, the syntax above does not model the macro possibility numbered 3 in our list of options. One can easily imagine a syntax which will then be compiled into UP 's as above, based on the structure of P .

VI. AN ARCHITECTURE FOR REACTIVE PROCESSES

Our proposed architecture is depicted in Figure 5. This architecture is divided into 3 layers:

- The DBMS: The workflow management logic runs on top of the DBMS. The database ensures the relation between the pther layers. The databsase contains all informations about thr process execution and data tables of several entities.
- The ediflow process: It corresponds to the XML specification of the process. Processes are specified in a high-level syntax following the structure described in Section V. Processes are specified in a high-level syntax following the structure described in Section V.
- The modules: This is a set of procedures and functions invoked by the user through the process file. These modules may correspond to visualization softwares.

Processes are specified in a high-level syntax following the structure described in Section V. The enactment of a process

thus specified consists of adding the necessary tuples to the Process and Activity relations. During process executions, the necessary data manipulation statements are issued to (i) record in the database the advancement of process and activity instances, (ii) evaluate on the database queries and updates, allow external procedures to read and update the application-driven entities and (iii) record the connections between users and application instances, and application data.

In the sequel, Section VI-A shows how to implement various degrees of isolation between concurrent processes operating on top of the same database. Section VI-B outlines update propagation. Section VI-C considers an important performance issue: efficient synchronization between memory-resident tables, that visualisation uses, and disk-resident tables.

A. Isolation

Applications may require different levels of sharing (or, conversely, of isolation) among concurrent activities and processes.

Process- and activity-based isolation Let a_1 be an instance of activity a , such that a_1 is part of a process instance p_1 . By default, queries evaluated during the execution of p carry over the whole relations implementing the application-specific data model. Let R be such a relation.

It may be the case that a_1 should only see the R tuples created as part of executing p_1 . For instance, when uploading an experimental data set, a scientist only sees the data concerned by that upload, not the data previously uploaded by her and others. Such isolation is easily enforced using relationships between the application relations and the *ActivityInstance* table (recall Figure 3 in Section IV). A query fetching data from R for a_1 should select only the R tuples created by p_1 , the process to which a_1 belongs, etc. These mechanisms are fairly standard.

Time-based isolation As discussed in Section V, the data visible to a given activity or process instance may depend on the starting time of that instance. To enable such comparisons, we associate to each application table R a *creation timestamp*, which is the moment when each R tuple entered the database (due to some process or external update). R tuples can then be filtered by their creation date.

Isolating process instances from tuple deletions requires a different solution. If the process instance p_3 erases some tuples from R , one may want to prevent the deleted tuples from suddenly disappearing from the view of another running process instance, say p_4 . To prevent this, tuples are not actually deleted from R until the end of p_3 's execution. We denote that moment by $p_3.end$. Rather, the tuples are added to a *deletion table* R_- . This table holds tuples of the form $(tid, t_{del}, pid, \perp)$, where tid is the deleted R tuple identifier, t_{del} the deletion timestamp, pid the identifier of the process deleting the tuple. The fourth attribute will take the value $p_3.end$ at the end of p_3 . To allow p_3 to run as if the deletion occurs, EdiFlow rewrites queries of the form $\text{select } * \text{ from } R$ implementing activities of p_3 with:

```
select * from R where tid not in
(select tid from R_ where pid=p3)
```

When p_3 terminates, if no other running process instance uses table R^1 , then we delete from R and R_- the tuples $\sigma_{pid=p_3}(R_-)$. Otherwise, R and R_- are left unchanged, waiting for the other R users to finish. However, a process instance started after $t_0 > p_3.end$ should not see tuples in R_- deleted by p_3 , nor by any other process whose end timestamp is smaller than t_0 . In such a recently-started process, a query of the form `select * from R` is rewritten by EdiFlow as:

```
select * from R where tid not in
(select tid from R_ where processend < t0)
```

We still have to ensure that deleted tuples are indeed eventually deleted. After the check performed at the end of p_3 , EdiFlow knows that some deletions are waiting, in R_- , for the end of a process instances started before $p_3.end$. We denote these process instances by $wait_{R,p_3}$. After $p_3.end$, whenever a process in $wait_{R,p_3}$ terminates, we eliminate it from $wait_{R,p_3}$. When the set is empty, the tuples $\sigma_{pid=p_3}(R_-)$ are deleted from R and R_- .

B. Update propagation

We now discuss the implementation of the update propagation actions described in Section V. EdiFlow compiles the *UP* (update propagation) statements into statement-level triggers which it installs in the underlying DBMS. The trigger calls EdiFlow routines implementing the desired behavior, depending on the type of the activity (Section V), as follows. Variable assignments are unaffected by updates. Propagating an update ΔR_i to relation R_i to a query expression leads to incrementally updating the query, using well-known incremental view maintenance algorithms [30]. Propagating an update to an activity involving a procedure call requires first, updating the input expressions, and then, calling the corresponding delta handler.

C. Synchronizing disk-resident and in-memory tables

The mechanisms described above propagate changes to (queries or expression over) tables residing in the SQL DBMS. However, the visualisation software running within an instance of a visualisation activity needs to maintain portions of a table in memory, to refresh the visualisation fast. A protocol is then needed to efficiently propagate updates made to a disk-resident table, call it R_D , to its possibly partial memory image, call it R_M . Conversely, when the visualisation allows the user to modify R_M , these changes must be propagated back to R_D . Observe that R_M exists on the client side and therefore may be on a different host than R_D .

To that end, we install CREATE, UPDATE and DELETE triggers monitoring changes to the persistent table R_D . Whenever one such change happens, the corresponding trigger adds

to the Notification table stored in the database (recall the data model in Figure 3) one tuple of the form (seq_no, ts, tn, op) , where seq_no is a sequential number, ts is the update timestamp, tn is the table name and op is the operation performed. Then, a notification is sent to R_M that "there is an update". Smooth interaction with a visualization component requires that notifications be processed very fast, therefore we keep them very compact and transmit no more information than the above. A notification is sent via a socket connected to the process instance holding R_M . Information about the host and port where this process runs can be found in the Client table (Figure 3). When the visualisation software decides to process the updates, it reads them from the Notification table, starting from its last read seq_no value.

The synchronization protocol between R_M and R_D can be summarized as:

- 1) A *memory object* is created in the memory of the Java process (R_M).
- 2) It asks the *connection manager* to create a connection with the database.
- 3) The connection manager creates a network port on the local machine and associates locally a quadruplet to R_M : $(db, R_D, ip, port)$.
- 4) The quadruplet is sent to the DBMS to create an entry in the ConnectedUser table.
- 5) The DBMS connects back to the client using at the $ip : port$ address, and expects a HELLO message to check that it is the right protocol.
- 6) The connection manager accepts the connection, sends the HELLO message and expects a REPLY message to check that it is the expected protocol too.
- 7) When the R_D is modified, the DBMS trigger sends a NOTIFY message with the table name as parameter to client at $ip:port$, which holds R_M .
- 8) The visualization software may decide what are the appropriate moments to refresh the display. When it decides to do so, it connects to the DBMS and queries the created/updated/deleted list of rows, and propagates the changes to R_M .
- 9) When R_M is modified, it propagates its changes to the R_D and processes the triggered notifications in a smart way to avoid redundant work.
- 10) When R_M is deleted, it sends a disconnect message to the database that closes the socket connection and removes the entry in the ConnectedUser table.
- 11) The Notification table can be purged of entries having seq_no lower than the lowest value in the Client table.

At first glance, this mechanism may look similar to updates over views (a.k.a. *materialized views*). However, our architecture has two main differences compared to materialized views:

- Propagation process. The propagation process for materialized views is relatively simple. Indeed, when changes occur on relations, the corresponding relevant views are updated. The difficulty is to know "when" and "how" the view should be updated. Moreover, updates are gener-

¹The definition of a process explicitly lists the tables it uses, and from the process, one may retrieve the process instances and check their status (Figure 3).

Fig. 6. EdiFlow architecture for managing several visualization views.

ally limited to insertions aggregations. However, in our architecture, a change that occurs on a relation may invoke many different update operations which generally correspond to external program's invocations. This is what we call *repair mechanism*.

- Two-way propagation. In the framework of materialized views, updates are usually done in one way (relation towards view). However, our architecture allows to manage changes that occur on the database while the analysis process is running. Moreover, it allows to update the database when users perform visual interaction.

EdiFlow can maintain several visualization views for one visualization. As shown in Figure 6, the visual attributes can be shared by several visualization views and by several users that may choose to visualize some or all of the data (e.g. on an iPhone showing 10% of the data, on a laptop showing 30% and on our WILD Wall-Sized display [31]) showing all of the data.

Moreover, in applications such as the INRIA co-publications example outlined in Section III, a user may want to visualize a scatter plot displaying the number of publications per year on one machine and displaying the number of publication by author on another machine. The two are obtained from the same data but using two different views. To this purpose, the visualization component computes and fills the visual attributes only once regardless of the number of generated views. For each view, a display component is activated to show the data on the associated machine using a visualization toolkit such as Prefuse [3] or the InfoVis Toolkit [4]. This architecture offers several advantages:

- It allows sharing visual attributes by different views and maintaining consistency between data and views.
- The computation of visual attributes is done only once. If an update occurs, the VisualAttributes table is updated and all associated views will be automatically updated.
- Such architecture can satisfy the principle of visualization: a visualization may have several views.

In practice, to display the co-publications graph on the WILD, we used a workstation running the visualization module and a cluster of 16 machines to display the graph over the

32 screens of the WILD. Each machine controls two screens and runs an EdiFlow instance to launch visualization view modules. When the data is updated, the DBMS notifies the visualization module to compute new visual attributes and to insert them into the VisualAttributes relation. Then, the database notifies the running visualization view modules that they need to refresh all displays.

D. EdiFlow tool implementation

EdiFlow is implemented in Java, and currently we have deployed it on top of both Oracle 11g and MySQL 5. EdiFlow processes are specified in a simple XML syntax, closely resembling the XML WfMC syntax XPD [32].

Procedures are implemented as Java modules using the Equinox implementation of the OSGi Service Platform [33]. A procedure instance is a concrete class implementing the EdiFlowProcess interface. This interface requires four methods: initialize(), run(ProcessEnv env), update(ProcessEnv env) and String getName(). The class ProcessEnv represents a procedure environment, including all useful information about the environment in which the processes are executed. An instance of ProcessEnv is passed as a parameter to a newly created instance of a procedure. Integrating a new processing algorithm into the platform requires only implementing one procedure class, and serving the calls to the methods. All the dependencies in term of libraries (JAR files) are managed by the OSGi Platform.

The implementation is very robust, well documented, efficient in term of memory footprint and lightweight for programming modules and for deploying them, which is important for our goal of sharing modules. We have implemented and ran the sample applications described in Section III.

VII. EXPERIMENTAL VALIDATION

A. Experimental setup

In this Section, we report on the performance of the EdiFlow platform in real applications.

Hardware Our measures used a client PC with Intel 2.66GHz Dual Core CPUs and 4GB memory running Java heap size was set to 850MB. The Oracle database is mounted

on a workstation with 8 CPUs equipped with 8GB RAM. The PC is connected to the database through the local area network.

Dataset We used a dataset of co-publications between INRIA researchers. We analyse this data set to produce visual results which have interesting insight for the INRIA scientific managers, and has to proceed while new publications are added to the database. This dataset includes about 4500 nodes and 35400 edges. The goal is to compute the attributes of each node and edge, display the graph over one or several screens and update it as the underlying data changes.

B. Layout procedure handlers

Our first goal was to validate the interest of procedure handlers in the context of data visualization. In our INRIA co-publication scenario, the procedure of interest is the one computing the positions of nodes in a network, commonly known as *layout*. We use the Edge LinLog algorithm of Noack [34] which is among the very best for social networks, and provides aesthetically good results. What makes EdgeLinLog even more interesting in our context is that it allows for effective delta handlers (introduced as part of our process model in Section V).

In our implementation, the initial computation assigns a random position to each node and runs the algorithm iteratively until it converges to a minimum energy and stabilizes. This computation can take several minutes to converge but, since the positions are computed continuously, we can store the positions in the database at any rate until the algorithm stops. Saving the positions every second or at every iteration if it takes more than one second allows the system to appear reactive instead of waiting for minutes before showing anything.

If the network database changes, for example when new publications are added to/removed from the database, the handler proceeds in a slightly different manner. First, it updates the in-memory co-publication graph, discards the nodes that have been removed and adds new nodes. To each new node it assigns a position that is close to their neighbors that have already been laid-out. This is to improve the time and quality of the final layout. If disconnected nodes are added, they are assigned a random position. Then, the algorithm is run iteratively like for the initial computation, but it terminates much faster since most of the nodes will only move slightly: the convergence of the iterations will be much faster. Like before, we store in the DBMS the results of some of the iterations to allow the visualization views to show them.

Using this strategy, we have obtained an incremental layout computation, remarkably stable and fast.

C. Robustness evaluation

Our second experimental goal was to study how the EdiFlow event processing chain scales when confronted with changes in the data. For this experiment, the DBMS is connected via a 100 MHz Ethernet connection to two EdiFlow instances running on two machines. The first EdiFlow machine computes visual attributes (runs the layout procedure), while the second

Fig. 7. Part of the graph of INRIA co-publications.

Fig. 8. Time to perform insert operation.

extracts nodes from VisualAttributes table and displays the graph. This second EdiFlow machine is a laptop.

We study the robustness of our architecture when adding increasing numbers of tuples to the database. Inserting tuples requires performing the sequence of steps below, out of which steps 1, 2 are performed on the first EdiFlow machine, while steps 3, 4 and 5 are performed on all machines displaying the graph.

- 1) Parsing the message involved after insertion in nodes table. It refers to step 7 in the protocol described in section VI-C.
- 2) Inserting the resulting tuples in the VisualAttributes table managed by EdiFlow in the DBMS.
- 3) Parsing the message involved after insertion in VisualAttributes table. After inserting tuples, in VisualAttributes, a message is sent to all machines displaying the graph. The message is parsed to extract the new tuple information. It refers to step 9 in the protocol described in section VI-C.
- 4) Extracting the visual attributes of the new nodes, from the VisualAttributes table, in order to know how to

display them at the client.

- 5) Inserting new nodes into the display screen of the second machine.

The times we measured for these five steps are shown in Figure 8 for different numbers of inserted data tuples. The Figure demonstrates that the times are compatible with the requirements of interaction, and grow linearly with the size of the inserted data. The dominating time is required to write in the VisualAttributes table. This is the price to pay for having these attributes stored in a place from where one can share them or distribute them across several displays.

VIII. CONCLUSION

In this article, we have described the design and implementation of EdiFlow, the first workflow platform aimed at capturing changes in data sources and launching a repair mechanism. EdiFlow unifies the data model used by all of its components: application data, process structure, process instance information and visualization data. It relies on a standard DBMS to realize that model in a sound and predictive way. EdiFlow supports standard data manipulations through procedures and introduces the management of changes in the data through *update propagation*. Each workflow process can express its behavior w.r.t data change in one of its input relations. Several options are offered to react to such a change in a flexible way.

EdiFlow reactivity to changes is necessary when a human is in the loop and needs to be informed of changes in the data in a timely fashion. Furthermore, when connected to an interactive application such as a monitoring visualization, the human can interactively perform a command that will change the database and trigger an update propagation in the workflow, thus realizing an interactively driven workflow.

We are currently using EdiFlow to drive our Wikipedia aggregation and analysis database as a testbed to provide real-time high-level monitoring information on Wikipedia, in the form of visualizations or textual data [26]. We are also designing a system for computing and maintaining a map of scientific collaborations and themes available on our institutions.

We still need to experiment with it to find out the limitations of EdiFlow in term of performances, typical and optimal reaction time and ability to scale with very large applications.

We strongly believe that formally specifying the services required for visual analytics in term of user requirements, data management and processing, and providing a robust implementation is the right path to develop the fields of visual analytics and scientific workflows together. For more details, examples, pictures and videos of the usage of EdiFlow, see the EdiFlow website: <http://scidam.gforge.inria.fr/>.

REFERENCES

- [1] J. Thomas and K. Cook, Eds., *Illuminating the Path: Research and Development Agenda for Visual Analytics*. IEEE Press, 2005.
- [2] S.-M. Chan, L. Xiao, J. Gerth, and P. Hanrahan, "Maintaining interactivity while exploring massive time series," in *VAST*, 2008.
- [3] J. Heer, S. Card, and J. Landay, "Prefuse: a toolkit for interactive information visualization," in *SIGCHI*, 2005.
- [4] J.-D. Fekete, "The InfoVis toolkit," in *InfoVis*, 2004.
- [5] "Protovis," <http://vis.stanford.edu/protovis/>.
- [6] B. Ludäscher, I. Altintas, C. Berkley, D. Higgins, E. Jaeger, M. Jones, E. Lee, J. Tao, and Y. Zhao, "Scientific workflow management and the Kepler system," *Concurrency and Computation: Practice and Experience*, 2006.
- [7] D. Hull, K. Wolstencroft, R. Stevens, C. Goble, M. Pocock, P. Li, and T. Oinn, "Taverna: a tool for building and running workflows of services," *Nucleic Acids Research*, 2006.
- [8] "Project Trident: A scientific workflow workbench," <http://research.microsoft.com/en-us/collaboration/tools/trident.aspx>.
- [9] A. Ailamaki, Y. E. Ioannidis, and M. Livny, "Scientific workflow management by database management," in *SSDBM*, 1998.
- [10] M. Brambilla, S. Ceri, P. Fraternali, and I. Manolescu, "Process modeling in web applications," *ACM Trans. Softw. Eng. Methodol.*, 2006.
- [11] S. Ceri, P. Fraternali, A. Bongio, M. Brambilla, S. Comai, and M. Mätere, *Designing Data-Intensive Web Applications*. Morgan Kaufmann, 2003.
- [12] S. Shankar, A. Kini, D. DeWitt, and J. Naughton, "Integrating databases and workflow systems," *SIGMOD Record*, 2005.
- [13] I. Altintas, C. Berkley, E. Jaeger, M. Jones, B. Ludäscher, and S. Mock, "Kepler: An extensible system for design and execution of scientific workflows," in *SSDBM*, 2004.
- [14] S. G. Parker and C. R. Johnson, "SCIRun: a scientific programming environment for computational steering," in *ACM SC*, 1995.
- [15] A. Rygg, P. Roe, and O. Wong, "GPFlow: An intuitive environment for web based scientific workflow," in *Proceedings of the 5th International Conference on Grid and Cooperative Computing Workshops*, 2006.
- [16] E. Anderson, S. Callahan, D. Koop, E. Santos, C. Scheidegger, H. Vo, J. Freire, and C. Silva, "VisTrails: Using provenance to streamline data exploration," in *DILS*, 2007.
- [17] R. S. Barga, J. Jackson, N. Araujo, D. Guo, N. Gautam, K. Grochow, and E. D. Lazowska, "Trident: Scientific workflow workbench for oceanography," in *IEEE Congress on Services*, 2008.
- [18] Z. G. Ives, T. J. Green, G. Karvounarakis, N. E. Taylor, V. Tannen, P. P. Talukdar, M. Jacob, and F. Pereira, "The ORCHESTRA collaborative data sharing system," *SIGMOD Record*, no. 3, 2008.
- [19] "Orchestra : Managing the collaborative sharing of evolving data," <http://www.cis.upenn.edu/~zives/orchestra/>.
- [20] M. R. Berthold, N. Cebron, F. Dill, T. R. Gabriel, T. Kötter, T. Meinl, P. Ohl, K. Thiel, and B. Wiswedel, "KNIME - the Konstanz information miner: version 2.0 and beyond," *SIGKDD Explorations*, vol. 11, no. 1, 2009.
- [21] E. Frank, M. A. Hall, G. Holmes, R. Kirkby, and B. Pfahringer, "Weka - a machine learning workbench for data mining," in *The Data Mining and Knowledge Discovery Handbook*, 2005.
- [22] "Rapidminer," <http://rapid-i.com/content/view/181/190/>.
- [23] J. Demsar, B. Zupan, G. Leban, and T. Curk, "Orange: From experimental machine learning to interactive data mining," in *PKDD*, 2004.
- [24] "Orange," <http://www.ailab.si/orange/>.
- [25] J.-D. Fekete and C. Plaisant, "Interactive information visualization of a million items," in *InfoVis*, 2002.
- [26] F. Chevalier, S. Huot, and J.-D. Fekete, "Wikipediaviz: Conveying article quality for casual wikipedia readers," in *PaciViz*, 2010.
- [27] F. Viégas, M. Wattenberg, and K. Dave, "Studying cooperation and conflict between authors with history flow visualizations," in *SIGCHI*, 2004.
- [28] V. Tannen, "Provenance for database transformations," in *EDBT*, 2010, p. 1.
- [29] "The workflow management coalition reference model," <http://www.wfmc.org/reference-model.html>.
- [30] A. Gupta, I. S. Mumick, and V. S. Subrahmanian, "Maintaining views incrementally," in *ACM SIGMOD*, 1993.
- [31] "Wild: Wall-sized interaction with large datasets," <http://insitu.lri.fr/Projects/WILD>.
- [32] "XML Process Definition Language," <http://www.wfmc.org/xpdl.html>.
- [33] O. Alliance, *OSGi Service Platform, Core Specification, Release 4, Version 4.2*, 2009.
- [34] A. Noack, "An energy model for visual graph clustering," in *Proceedings of the 11th International Symposium on Graph Drawing*, ser. LNCS, vol. 2912, 2004.