CONSIGLIO NAZIONALE DELLE RICERCHE

RELAZIONE PROGRAMMA SHORT - TERM MOBILITY - Anno 2011

Fruitore:	*Dott.ssa Angela Bonifati, ricercatrice a tempo indeterminato Icar-CNR Rende*
Istituto di afferenza:	*Istituto di Calcolo e Reti ad Alte Prestazioni (ICAR) CNR*
Titolo del programma: Obiettivi (3-4 righe)	*Title: Incremental View Maintenance and View Update Problem for XML Views.* *Goals*: Recently, the W3C has proposed XUpdate, an update extension to the XQuery [XQuery] language, which is gradually being implemented in XML data management platforms. The objective of our research program is double-fold. First, we address the *view incremental maintenance problem*. Our focus is on incrementally modifying an XML view with XML updates, without recomputing the view. Secondly, we address the *view update problem* in the same setting. The latter problem can be considered as the inverse of incremental maintenance, when XML updates are applied to the view and need to be propagated to the document.

Relazione Finale del Programma svolto dalla Dott.ssa Angela Bonifati presso INRIA Saclay dal 1 Maggio 2011 al 22 Maggio 2011:

The project aimed at addressing the problem of XML view maintenance and its opposite direction, namely the view update problem for XML views. This problem is very relevant to Information and Communication Technology, as XML data is becoming pervasive in Web 2.0.

During the program, myself, Dr. Manolescu and our students have been discussing many issues, regarding the development of algorithms and their implementation. More precisely, we have realized the following activities: (i) design of algorithms for incremental view maintenance, in order to propagate the updates from the document to the view; (ii) design of algorithms for view update problem, in order to propagate the updates from the view back to the document; (ii) study of *bidirectional* XML materialized views in order to unify the above algorithms.

The view maintenance problem amounts to evaluating the view *v* over the XML document *d* leading to materialize *v(d)*. An XML update transforms *d* into *d'*, and correspondingly the affected view *v* should be transformed into *v(d')*. One possibility is to evaluate *v* from scratch on the modified document *d'*. Instead, our focus is on incrementally modifying *v* by adding, removing, or modifying data as needed, to transform it into *v(d')*, without re-computing it.

In our work, we addressed the incremental maintenance of XML views in the presence of statement-level XML updates. Our view language corresponds to a core useful conjunctive XQuery subset. This language supports the child and descendant axis, value and branch predicates, and moreover allows returning data from more than one node, unlike the XPath dialects studied in [Wim09,Sawi06,Sawi05]. Our approach is designed to take advantage of advanced artifacts of current XML query processors, such as structural joins and smart identifiers[XuLWB09]. Employing such efficient tools allows our algorithms to outperform node-level update propagation techniques in the frequent case where more than one node is added/removed at the same time. Moreover, our approach integrates smoothly in the process of updating the source document itself, by re-using some partial results of the update process. These features make it a good candidate to be integrated within a persistent XML database.

The view update problem is, given a database instance *t* and a view *v* of *t*, and given an update *u* to be applied on *v*, how to 'correctly' propagate on *t* the changes of *v* made by *u*, where the meaning of 'correctly' has to be precisely defined.
The user only has access to the view, with no knowledge of the source document. This is useful for security reasons in order to limit access to a document or simply for commodity purposes. However, there are problems with having no knowledge on the source document. Undesirable side effects on the hidden part of the document could occur, as for example hidden nodes could be inserted or deleted. Work on view update translation for XML[BonevaEtAl2011] has studied two settings for the problem: with source constraints and without source constraints. Without source constraints all updates are allowed, conversely with source constraints there is a restricted set of authorised source updates.

We were interested in extending our algebraic view maintenance method[BonifatiEtAl2011] to apply to the view update problem. This would currently only be able to support deletions. The current work[BonevaEtAl2011] supports deletions and also rename operations within views. Our current method doesn't support renaming operations so would need to be extended to handle this problem. The reason insertions aren't handled is that should there be an ancestor-descendant relationship within a view then we will be unable to know the details of the missing nodes to be inserted in the source document. For example given a view //a//b, if we have an insertion of a new b node then we need to reflect this in the source document. The b has to be simply added to the view. However, we need to know how to insert the nodes between a and b in the source. Insertions would only be possible if all nodes above the inserted node were related by parent-child relationships.

A comprehensive and more detailed version of the present final report, which provides a precise technical summary of the activities conducted during the short-term mobility program 2011, is in [*BonifatiEtAl2011*, *BGM2011*], which are available on demand.

References

[*BonevaEtAl2011*] Iovka Boneva, Anne-Cecile Caron, Benoit Groz, Yves Roos, Sophie Tison, and Slawek Staworko. View update translation for XML. In ICDT, pages 42-53, 2011.
[*BonifatiEtAl2011*] Angela Bonifati, Martin Goodfellow, Ioana Manolescu, and Domenica Sileo. Algebraic incremental maintenance of XML views. Technical Report (extended version of the EDBT2011 paper, pages 177-188).
[*BGM2011*] Angela Bonifati, Martin Goodfellow, Ioana Manolescu,. The View Update Problem Algebraically. Technical Report, pages 1-14, 2011.
[*Wim09*] H. Björklund, W. Gelade, M. Marquardt, and W. Martens. Incremental XPath evaluation. In ICDT, 2009.
[*Sawi05*] A. Sawires, J. Tatemura, O. Po, D. Agrawal, and K. S. Candan. Incremental maintenance of path-expression views.In SIGMOD, 2005.
[*Sawi06*] A. Sawires, J. Tatemura, O. Po, D. Agrawal, A. E. Abbadi, and K. S. Candan. Maintaining XPath views in loosely coupled systems. In VLDB, 2006.
[*XulWB09*] Liang Xu, Tok Wang Ling, Huayu Wu, and Zhifeng Bao. DDE: from Dewey to a fully dynamic XML labeling scheme. In SIGMOD, 2009.
[*XML*] http://www.w3.org/xml/
[*XQuery*] http://www.w3.org/xmlquery/

Firma del Fruitore
Dott. ssa Angela Bonifati