Advanced Techniques and Tools for Secure Collaborative Modeling

Csaba Debreceni
Budapest University of Technology and Economics
MTA-BME Lendület Research Group on Cyber-Physical Systems
Hungary
debreceni@mit.bme.hu

Abstract

Model-based systems engineering of critical cyber-physical systems necessitates effective collaboration between different collaborators, teams, stakeholders. Engineering artifacts stored in model repositories are concurrently developed in either offline (checkout-modify-commit) or online (GoogleDoc-style) scenario where the confidentiality and integrity of design artifacts need to be protected by access control policies. Unfortunately, traditional approaches for managing concurrent code development do not naturally extend to collaborative modeling which implies novel challenges.

My research focuses on developing (i) a general secure collaboration scheme that guarantees that high-level access control policies are respected during collaboration and it can be integrated into existing version control systems (e.g. SVN) to support offline scenario; (ii) automated merging and fine-grained locking to enhance the efficiency of conflict resolution and prevention upon concurrent modification of the models; (iii) derivation and incremental maintenance of view models to provide specific focus of the designers by abstracting from unnecessary details of the underlying system model.

Keywords collaborative modeling, fine-grained locking, automated model merge, view model

1 Problem and Motivation

The adoption of model driven engineering (MDE) by system integrators (like airframers or car manufacturers) has been steadily increasing in the recent years [40], since it enables to detect design flaws early and generate various artifacts (source code, documentation, configuration tables, etc.) automatically from high-quality system models.

The use of models also strengthens collaboration between distributed teams of different stakeholders (system integrators, software engineers of component providers/suppliers, hardware engineers, certification authorities, etc.) via model repositories, which significantly enhances productivity and reduces time to market. An emerging industrial practice of system integrators is to outsource the development of various design artifacts to subcontractors in an architecture-driven supply chain.

Collaboration scenarios include traditional *offline collaborations* with asynchronous long transactions (i.e. to check out an artifact from a version control system and commit local changes afterwards) as well as *online collaborations* with short and synchronous transactions (e.g. when a group of collaborators simultaneously edit a model). Even though, various collaborative modeling frameworks (like [18, 21], etc.) exist to support such scenarios, additional challenges arise that cannot be naturally extended from traditional code-based approaches due to the graph-like nature of the artifacts.

1.1 Secure Collaborative Modeling

An increased level of collaboration in a model-driven development process introduces additional confidentiality challenges to sufficiently protect the intellectual property of the collaborating parties, which are either overlooked or significantly underestimated by existing initiatives. Even within a single company, there are often teams with differentiated responsibilities, areas of competence and clearances described by high-level access control policies. Such processes likewise demand confidentiality and integrity of certain modeling artifacts. My first research question is constructed as follows:

RQ-1 How to capture and enforce high-level access control policies during collaborative modeling?

1.2 Conflict Prevention and Resolution

Enabling a high degree of concurrent edits for collaborators is required to make the traditionally rigid development processes more agile. The increasing number of collaborators concurrently developing artifacts increases the probability of introducing conflicts. *Conflict avoidance* techniques such as locks try to prevent conflicts by letting the users request that certain engineering artifacts should be made unmodified by all other participants for a duration of time. But it usually leads to unnecessary preventions (locks) which significantly limits the degree of concurrent development and does not scale with the increasing number of collaborating teams. *Model merging* aims to resolve the conflicts, but, it can be complex tasks as the interdependence within a model makes conflicts easy to introduce and hard to resolve. My second research question is the following.

RQ-2 How to provide fine-grained prevention and automatized resolution strategies of conflicts?

1.3 Bidirectional Synchronization of View Models

Views are key concepts of domain-specific modeling in order to provide specific focus of the system to the engineers with various knowledge and expertise by abstracting the unnecessary details of the underlying model. Usually, these views are represented as models themselves (view models), computed from the source model. On one hand, the efficient forward propagation of changes from the source model to the views is challenging, as recalculating the view from scratch has to be avoided to achieve scalability. On the other hand, the efficient backward propagation of complex changes from one or more abstract view models to the underlying source model is also a challenging task which requires to limit the propagation to a well-defined part of the source model. My third research question is as follows.

RQ-3 How to derive and incrementally maintain view models and trace back complex changes to the underlying source models?

2 Preliminaries

2.1 Related Work

2.1.1 Secure Collaborative Modeling

Traditional version control systems (like [3]) adopt file-level access policies, which are clearly insufficient for fine-grained access control specifications. [18] allows for role-based access control with type-specific (class, package and resource-level) permissions, but disallows instance level access control policy specifications. Access control is not considered in recent collaborative modeling environments like [4, 10, 21, 29, 30, 39], or the tools developed according to [22]. [38] provides fine-grained role-based access control for online collaboration but no offline scenario is supported, though. Both online collaboration and role-based access control with type-specific (class, package and resource-level) permissions

is provided in [18], but no facility for instance level access control policy specifications. However, there is a pluggable access control mechanism that can specify access on the object level.

2.1.2 Locking Support

The state-of-the-art locking techniques are the fragment-based and object-based locks. *Fragment-based locking* requires that models are partitioned into storage fragments, e.g. files or projects and entire fragments can be locked at once. *Object-based locking* locks individual model objects (including their attributes and connections) which requires to inspect the structure of the model.

Existing collaborative modeling tools either lack locking support or implement rigid strategies such as fragment-based locking, or locking subtrees or elements of a specific type, which hinder effective collaboration. Most of *offline collaborative modeling tools* [2, 21, 27], rely on traditional version control systems using file-based (same as fragment-based) locking with contributors committing large deltas of work. *Model repositories* [18, 39], support both implicit and explicit locking of subtrees and sets of elements. These locks can prevent others from modifying elements to avoid conflicts. *Online collaborative modelling frameworks* [4, 10, 29, 38], rely on a short transaction model: a single, shared instance of the model is concurrently edited by multiple users, with all changes propagated to all participants instantaneously. These approaches use timestamped operations to resolve conflicts or provide only lightweight lock mechanisms, e.g., explicit locks to certain elements.

2.1.3 Conflict Resolution in Model Artifacts

Model comparison refers to identifying the differences between models. Based on its result, *model merge* synthesizes a combined model which reconciles the identified differences. My research focuses on *three-way merge*, which uses the common ancestor O of local copy L and remote copy R to derive the merged model M. To determine the changes executed on O, a comparison is conducted between $O \leftrightarrow L$ and $O \leftrightarrow R$. The solution of merge M is obtained by applying a combination of changes performed either on L or R to the original model O.

Most approaches [6, 19, 20, 32, 34] are semi-automated as they use a two-phase process: (i) first, they apply the non-conflicting operations and then (ii) let the user prioritize and select the operation to apply in case of two conflicting changes. This always results in a single solution due to the manual resolution by the user. In comparison, [11, 28] resolve the conflicts automatically in different ways and offer several solutions.

2.1.4 Incremental Maintenance of View Models.

View maintenance by incremental and live QVT transformations is used in [37] to define views from runtime models.

The proposed algorithm operates in two phase, starting in check-only mode before an enforcement run, but its scalability is demonstrated only on models up to 1000 elements. [10] allows the composition of multiple EMF models into a virtual model based on a composition metamodel, and provides both a model virtualization API and a linking API to manage these models. The approach is also able to add virtual links based on composition rules. In [41], an ATL-based method is presented for automatically synchronizing source and target models. In [26], correspondences between models are handled by matching rules defined in the Epsilon Comparison Language, but incremental derivation is not discussed.

2.1.5 Backward Propagation

For the backward propagation of changes, the use of traceability links is a well-accepted approach to define which part of the source model has to be updated upon a change on the target model. In [33], these links are stored as a *correspondence model* where their maintenance is derived from the TGG rules. [31] also specifies *trace* classes to facilitate and maintain traceability links. [23] stores traceability links in Alloy[25] as a bijective mapping. [7] uses a *weaving* model that stores the traces of references between different models in the view, however all objects in the view model act as proxies to an object in the source model.

2.2 Foundational Techniques

2.2.1 Graph Patterns

A *graph pattern* represents structural constraints prescribing the interconnection between nodes and edges of given types extended with algebraic expressions to define attribute constraints. *Pattern parameters* are a subset of nodes and attributes representing the model elements interesting from the perspective of the pattern user. A *match* of a pattern is a tuple of pattern parameters that has the same structure as the pattern and satisfies all structural and attribute constraints.

2.2.2 Design Space Exploration

Design space exploration (DSE) aims to find optimal design candidates of a domain with respect to different objectives where design candidates are constrained by complex structural and numerical restrictions (e.g. described by *graph pattern*) and are reachable from an initial model by applying a sequence of exploration rules.

3 Overview of the Approaches

3.1 General Secure Collaboration Scheme

Approach. In [5], we proposed a query-based approach for modeling fine-grained access control policies, and we defined bidirectional model transformations to (i) derive filtered views (*front models*) for each collaborator from the original model (*gold model*) containing all the information and to (ii) propagate changes introduced into these views


Figure 1. MONDO Offline Collaboration - Architecture

back to a server in both *online* and *offline* scenarios. Access control policies consist of rules that allow, obfuscate or deny read and/or write permissions of model parts identified by graph patterns detailed in [13].

In [14], a collaboration scheme between the clients of multiple collaborators and exactly one server is described to support fine-grained access control in offline scenario. The server stores the gold models and the clients can download their specific front models. Modifications, executed by a clients, are submitted to the server and they are accepted if write permissions are successfully checked. Right after the submission, the changes are propagated to the other front model while read permissions are enforced. Finally, clients can downloaded their updated front models.

The scheme is realized by extending *SVN*[3] using its hooks. The server and clients are realized as a *gold repository* and multiple *front repositories*, respecively. The *gold repository* contains *gold models*, but it is not accessible to collaborators. Each collaborator is assigned to a specific *front repository* containing a full version history of the front models. Change propagations are maintained between the repositories.

My contributions related to the fulfillment of *RQ-1*:

Contribution 1 I proposed a generic modeling language to capture fine-grained access control policies integrated into a provenly secure collaborative architecture.

C1.1 *Access Control Language.* I proposed a rule-based access control language to describe high-level and fine-grained policies in both online and offline scenarios. Rules may allow, obfuscate or deny read and/or write permissions of model parts identified by graph patterns[13, 14].

C1.2 *Read and Write Dependencies.* I analyzed read and write dependencies implied by high-level access control policies as read and write permissions of a model part may depend on other model parts implied by internal consistency rules [13].

C1.3 *Formalization of Transformation Rules.* I formalized transformation rules to derive secure front models with respect to the read and write permissions [14].

C1.4 *Secure Collaboration Scheme.* I formalized a collaboration scheme as *communicating sequential processes* (CSP) to enforce high-level access control policies. I specified correctness criteria and proved the correctness of the scheme [14].

C1.5 *Realization of Secure Collaboration.* I realized the collaboration scheme in case of offline scenarios by extending an existing version control system to enforce fine-grained access control while collaborators can use off-the-shelf tools [12, 14].

C1.6 *Evaluation.* I evaluated the scalability of the collaboration architecture on a case study of offshore wind turbine controllers [5, 12, 14, 24].

The bidirectional transformation and the algorithm to derive effective permission based on the proposed language is the contribution of Gábor Bergmann whereas the concept of the common architecture to support both online and offline scenarios is the contribution of István Ráth.

Uniqueness. Our provenly correct collaboration scheme is able to enforce fine-grained access control policies of modeling artifacts over existing version control system in case of offline scenarios.

3.2 Conflict Reduction and Handling

Approach. In our preliminary work [9], we introduced the concept of property-based locking where collaborators request locks specified as a property of the model which need to be maintained as long as the lock is active. Hence, other collaborators are permitted to carry out any modifications that do not violate the defined property of the lock. In [15], the realization of property-based locking strategy is proposed as a common generalization of existing fragment-based and object-based locking approaches. Complex properties are described as graph patterns to express structural (and attribute) constraints for a model where the result set, i.e. the matches of graph pattern, can be calculated by pattern matchers or query engines. Only those modifications are allowed that do not change the result set of a list of queries as depicted in Fig. 2.

In [17], we proposed DSE-Merge that exploits guided rule-based *design space exploration (DSE)* [1] to automate the three-way model merge with an architecture depicted in Fig. 3. Three-way model merge is applied to DSE problem where the *initial model* consists of the original model O and two difference models (ΔL and ΔR); the *goal* is that there are no executable changes left in ΔL and ΔR; *operations* are defined by change driven transformation rules to process generic composite (domain-specific) operators; and *constraints* may identify inconsistencies and conflicts
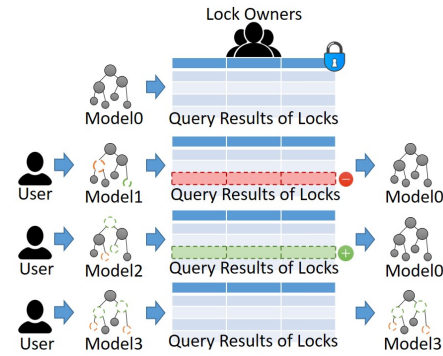

Figure 2. Behavior of Property-Based Locks

Figure 3. Architecture of DSE Merge

to eliminate certain trajectories. The output is a *set of solutions* consisting of (i) the merged model M; (ii) the set of non-executed changes $\Delta L',\Delta R'$; and (iii) the collection of the deleted objects stored in *Cemetery*.

My contributions related to the fulfillment of RQ-2 :

Contribution 2 I proposed a fine-grained property-based locking technique to avoid conflicts and an automated three-way model merge technique to resolve conflicts.

C2.1 *Fine-grained Property-based Locking.* I proposed a property-based locking technique as generalization of traditional fragment-based and object-based locking techniques which captures fine-grained locks as graph patterns and exploits incremental query engines to maintain and evaluate locks [15].

C2.2 *Automated Model Merge using DSE.* I proposed an automated three way model merge technique by adapting rule-based design space exploration to derive consistent and semantically correct merged models [1, 17].

C2.3 *Realization of DSE-merge.* I realized an infrastructure of automated model merge over EMF integrated into the Eclipse IDE [12, 17].

C2.4 *Evaluation.* I evaluated the scalability of the automated model merge and I compared the effectiveness of fine-grained property-based locking and traditional locking strategies for conflict prevention on a case study of offshore wind turbine controllers [15, 17, 24].

Figure 4. Overview of integration architecture

Figure 5. Overview of backward change propagation

The novel concept of property-based locking has been carried out in a collaborative work [9] where my contribution is the first adaption in a practical setting.

Uniqueness. Our property-based approach is general and can be used for both implicit locking of subtrees and set of elements or explicit locking of a certain element and its incoming and outgoing references. In addition it extends these lock types with the definition of properties to provide less restrictive locking for the collaborators.

The closest to our merge approach are [11] and [28], but we rely on state-based comparison, apply a guided local-search strategy (vs. [28]), detect conflicts at runtime and allow complex generic merge operations (vs. [11]). Internally, we uniquely use incremental and change-driven transformations to derive the merged models. Finally, we reported scalability of merge process for models which are at least one order of magnitude larger compared to [11] and [28].

3.3 Synchronization of View Models

Approach. In [16], we introduced an approach where view models are conceptually equivalent to regular models and they are defined using a fully declarative, rule based formalism. *Preconditions* of rules are defined by graph patterns, which identify parts of interest in the source model. *Derivation rules* then use the match set of a graph pattern to define elements of the view model. Informally, when a new match of a query appears then the corresponding derivation rule is fired to create elements of the view model. When an existing match of a query disappears, the inverse of the derivation rule is fired to delete the corresponding view model elements.

View models derived by a unidirectional transformation are read-only representations, and they cannot be changed directly. To tackle this problem, we proposed an approach in [35] to automatically calculate possible source model candidates for a set of changes in different view models as depicted on Fig. 5. First, the possibly impacted partition of the source model is need to be identified by observing traceability links to restrict the impact of a view modification. Then the modified view models and the query-based view specification are transformed into logic formulae. Finally, multiple valid resolutions of the source model are generated using logic solvers corresponding to the changes of view models and the constraints of the source model.

My contributions related to the fulfillment of *RQ-3* :

Contribution 3 I proposed a novel technique of bidirectional synchronization of view models where the forward incremental synchronization is achieved by unidirectional derivation rules while the backward propagation of changes is generated using logic solvers.

C3.1 *Incremental Forward Synchronization.* I formalized a fully forward incremental, unidirectional synchronization technique of view models allowing chaining of views where the object of view model depend on the match set of the precondition of derivation rules [16, 36].

C3.2 *Change Impact Analysis.* I analyzed the impact of changes in underlying source models in case of backward propagation. The impacted part is added to the logic solver as additional constraints to calculate minimally modified source model candidates [35].

C3.3 *Realization of Forward Synchronization.* I realized the incremental and forward view synchronization technique where elementary derivation rules are captured by graph patterns and the reactive synchronization process uses the Viatra Event-driven Virtual Machine (EVM) [16].

C3.4 *Evaluation.* I evaluated the scalability of the proposed approaches on case studies from the avionics and the health-care domain [16, 35, 36].

The transformation of the preconditions described by graph patterns and the impacted parts to first order logic is the contribution of Oszkár Semeráth whereas my contributions are the impact analysis and the concept of using logic solver for backward propagation extended with impacted parts as additional constraints.

Uniqueness. Definition of a view model is *unidirectional*, while the forward propagation of the *operation-based* changes are *live*, *incremental* and executed *automatically* that also maintains *explicit traces*. At backward propagation, using partitioning as an additional input of the logic solver improves scalability issues and limits the impact of changes to a well-defined part of the source model.

Acknowledgement

References

[1] Hani Abdeen, Dániel Varró, Houari A. Sahraoui, András Szabolcs Nagy, Csaba Debreceni, Ábel Hegedüs, and Ákos Horváth. 2014. Multi-objective optimization in rule-based design space exploration. In *ASE '14*. 289–300.

[2] Kerstin Altmanninger, Gerti Kappel, Angelika Kusel, Werner Retschitzegger, Martina Seidl, Wieland Schwinger, and Manuel Wimmer. 2008. AMOR–towards adaptable model versioning. In *MCCM@MoDELS'08*, Vol. 8. 4–50.

[3] Apache. Subversion. https://subversion.apache.org/.

[4] Axellience. GenMyModel. http://genmymodel.com.

[5] Gábor Bergmann, Csaba Debreceni, István Ráth, and Dániel Varró. 2016. Query-based access control for secure collaborative modeling using bidirectional transformations. In *MoDELS'16*. 351–361.

[6] Petra Brosch, Martina Seidl, Konrad Wieland, and Manuel Wimmer. 2009. We can work it out: Collaborative Conflict Resolution in Model Versioning. In *ECSCW'09*. 207–214.

[7] Hugo Brunelière et al. 2015. EMF Views: A View Mechanism for Integrating Heterogeneous Models. In *Conceptual Modeling - ER'15*. 317–325.

[8] Jian Chang, Peter Gebhard, Andreas Haeberlen, Zachary G. Ives, Insup Lee, Oleg Sokolsky, and Krishna K. Venkatasubramanian. 2013. TrustForge: Flexible access control for collaborative crowd-sourced environment. In *PST'13*. 291–300.

[9] Marsha Chechik, Fabiano Dalpiaz, Csaba Debreceni, Jennifer Horkoff, István Ráth, Rick Salay, and Dániel Varró. 2015. Property-Based Methods for Collaborative Model Development. In *GEMOC+MPM@MoDELS'15*. 1–7.

[10] Cauê Clasen, Frédéric Jouault, and Jordi Cabot. 2011. VirtualEMF: A Model Virtualization Tool. In *Advances in Conceptual Modeling. Recent Developments and New Directions*. 332–335.

[11] Hoa Khanh Dam, Alexander Reder, and Alexander Egyed. 2014. Inconsistency Resolution in Merging Versions of Architectural Models. In *WICSA'14*. 153–162.

[12] Csaba Debreceni, Gábor Bergmann, Márton Búr, István Ráth, and Dániel Varró. 2017. The MONDO Collaboration Framework: Secure Collaborative Modeling over existing Version Control Systems. *ESEC/FSE'17* (2017). https://goo.gl/uTsQeg In Press.

[13] Csaba Debreceni, Gábor Bergmann, István Ráth, and Dániel Varró. 2016. Deriving Effective Permissions for Modeling Artifacts from Fine-grained Access Control Rules. In *COMMitMDE@MoDELS'16*. 17–26.

[14] Csaba Debreceni, Gábor Bergmann, István Ráth, and Dániel Varró. 2017. Enforcing Fine-grained Access Control for Secure Collaborative Modeling using Bidirectional Transformations. *Software and System Modeling, MODELS 2016 Special Section* (2017). https://goo.gl/ZAegbo Submitted.

[15] Csaba Debreceni, Gábor Bergmann, István Ráth, and Dániel Varró. 2017. Property-based Locking in Collaborative Modeling. In *MoDELS'17*. In Press.

[16] Csaba Debreceni, Ákos Horváth, Ábel Hegedüs, Zoltán Ujhelyi, István Ráth, and Dániel Varró. 2014. Query-driven incremental synchronization of view models. In *VAO@STAF'14*. 31–38.

[17] Csaba Debreceni, István Ráth, Dániel Varró, Xabier De Carlos, Xabier Mendialdua, and Salvador Trujillo. 2016. Automated Model Merge by Design Space Exploration. In *FASE'16*. 104–121.

[18] Eclipse Foundation. CDO. http://eclipse.org/cdo.

[19] Eclipse Foundation. EMFCompare. http://eclipse.org/emf/compare/.

[20] Eclipse Foundation. EMFDiff/Merge. http://eclipse.org/diffmerge/.

[21] Eclipse Foundation. EMFStore. http://eclipse.org/emfstore.

[22] Jesús Gallardo, Crescencio Bravo, and Miguel A. Redondo. 2012. A model-driven development method for collaborative modeling tools. *J. Network and Computer Applications* 35, 3 (2012), 1086–1105.

[23] Hamid Gholizadeh et al. 2015. Analysis of Source-to-Target Model Transformations in QueST. In *Proceedings of the 4th Workshop on the Analysis of Model Transformations co-located with (MODELS 2015, Ottawa, Canada*. 46–55. http://ceur-ws.org/Vol-1500/paper6.pdf

[24] Abel Gómez, Xabier Mendialdua, Gábor Bergmann, Jordi Cabot, Csaba Debreceni, Antonio Garmendia, Dimitrios S. Kolovos, Juan de Lara, and Salvador Trujillo. 2017. On the Opportunities of Scalable Modeling Technologies: An Experience Report on Wind Turbines Control Applications Development. *ECMFA'17* (2017). https://goo.gl/surozr In Press.

[25] Daniel Jackson. Alloy Analyzer. http://alloy.mit.edu/.

[26] Dimitrios S. Kolovos. 2009. Establishing Correspondences between Models with the Epsilon Comparison Language. In *ECMDA-FA'09*. 146–157.

[27] Gerhard Kramler, Gerti Kappel, Thomas Reiter, Elisabeth Kapsammer, Werner Retschitzegger, and Wieland Schwinger. 2006. Towards a semantic infrastructure supporting model-based tool integration. In *GaMMa@ICSE'06*. ACM, 43–46.

[28] Usman Mansoor, Marouane Kessentini, Philip Langer, Manuel Wimmer, Slim Bechikh, and Kalyanmoy Deb. 2015. MOMM: Multi-objective model merging. *Journal of Systems and Software* 103 (2015), 423–439.

[29] Miklós Maróti, Tamás Kecskés, Róbert Kereskényi, Brian Broll, Péter Völgyesi, László Jurácz, Tihamer Levendovszky, and Ákos Lédeczi. 2014. Next Generation (Meta)Modeling: Web- and Cloud-based Collaborative Tool Infrastructure. In *MPM@MODELS'14*. 41–60.

[30] Obeo. Obeo Designer. https://obeodesigner.com/en/collaborative-features.

[31] OMG. MOF 2.0 QVT. http://www.omg.org/spec/QVT/1.2/.

[32] Julia Rubin and Marsha Chechik. 2013. N-way model merging. In *ACM SIGSOFT Symp@ESEC/FSE'13*. 301–311.

[33] Andy Schürr. 1994. Specification of Graph Translators with Triple Graph Grammars. In *Graph-Theoretic Concepts in Computer Science, WG'94*. 151–163.

[34] Felix Schwägerl, Sabrina Uhrig, and Bernhard Westfechtel. 2013. Model-based tool support for consistent three-way merging of EMF models. In *ACME@ECOOP'13*. 2:1–2:10.

[35] Oszkár Semeráth, Csaba Debreceni, Ákos Horváth, and Dániel Varró. 2016. Change Propagation of View Models by Logic Synthesis using SAT solvers. In *BX@ETAPS'16*. 40–44.

[36] Oszkár Semeráth, Csaba Debreceni, Ákos Horváth, and Dániel Varró. 2016. Incremental backward change propagation of view models by logic solvers. In *MoDELS'16*. 306–316.

[37] Hui Song, Gang Huang, Franck Chauvel, Wei Zhang, Yanchun Sun, Weizhong Shao, and Hong Mei. 2011. Instant and Incremental QVT Transformation for Runtime Models. In *MoDELS'11*. 273–288.

[38] Eugene Syriani, Hans Vangheluwe, Raphael Mannadiar, Conner Hansen, Simon Van Mierlo, and Hüseyin Ergin. 2013. AToMPM: A Web-based Modeling Environment. In *Invited Talks, Demonstration Session, Poster Session, and ACM Student Research Competition@MoDELS'13*. 21–25.

[39] Juha-Pekka Tolvanen. 2016. MetaEdit+ for collaborative language engineering and language use (tool demo). In *Tool Demo@SLE'16*. 41–45.

[40] Jon Whittle, John Hutchinson, and Mark Rouncefield. 2014. The State of Practice in Model-Driven Engineering. *IEEE Software* 31, 3 (2014), 79–85.

[41] Yingfei Xiong, Dongxi Liu, Zhenjiang Hu, Haiyan Zhao, Masato Takeichi, and Hong Mei. 2007. Towards automatic model synchronization from model transformations. In *ASE'07*. 164–173.