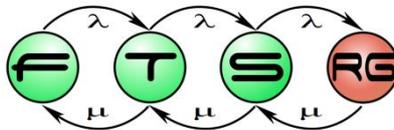


Bi- and Multidirectional Synchronization: Challenges and Approaches

a State of the Art Review by **Gábor Bergmann**
FTSRG Seminar, 2019-02-19



**Hungarian Academy of
Sciences**

Supported by: the MTA-BME Lendület Cyber-Physical Systems Research Group, the János Bolyai Research Scholarship of the Hungarian Academy of Sciences and the ÚNKP-18-4 New National Excellence Program of the Ministry of Human Capacities.

Intro

- Bidirectional / multidirectional synchronization
 - Cloud drive on mobile ↔ computer
 - Data binding, database editing based on view
 - Dentist appointments vs assistant
 - Healthcare, HIPAA
 - Complex MBE workflows
- Is all this really easy?
- How do I come into the picture?
 - MONDO, Dagstuhl 2018

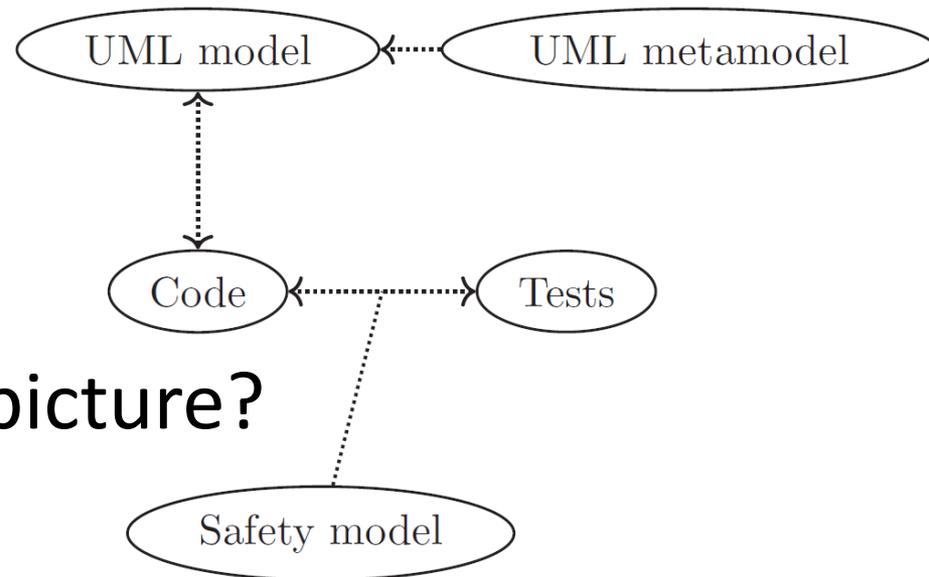


Image source: P. Stevens, "Bidirectional Transformations in the Large," *2017 ACM/IEEE 20th International Conference on Model Driven Engineering Languages and Systems (MODELS)*, Austin, TX, 2017, pp. 1-11. doi: 10.1109/MODELS.2017.8

Topics to cover

Basics

Properties

Constructions

Applications &
Generalizations

Basics

Properties

Constructions

Applications &
Generalizations

Basic Formal Notions

Symmetric / asymmetric BX
Semantic variants

(Symmetric) BX

- Consistency between two models

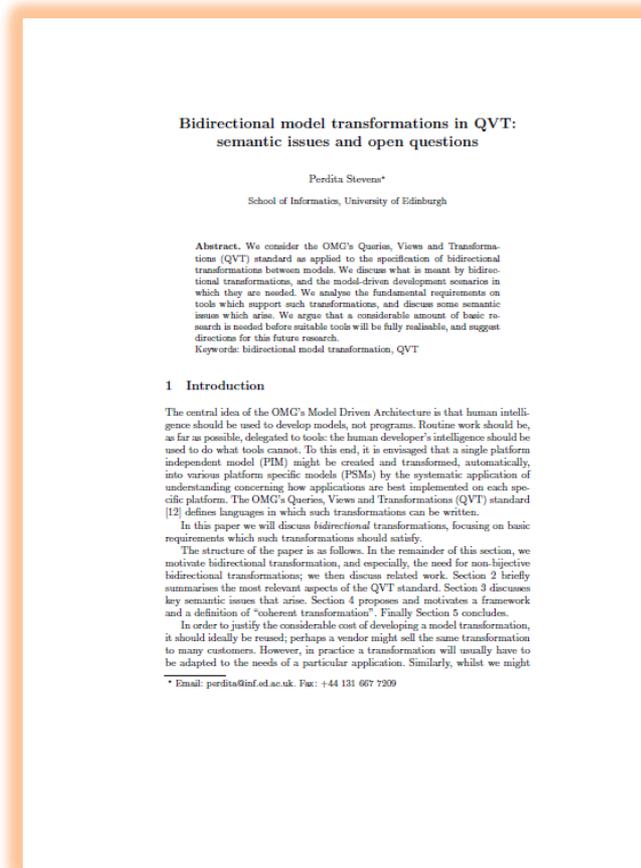
$$R \subseteq M \times N \quad R(m, n)$$

- + a pair of consistency restorers

$$\begin{aligned} \vec{R} &: M \times N \longrightarrow N \\ \overleftarrow{R} &: M \times N \longrightarrow M \end{aligned}$$

- Bidirectional \neq bijective

„bivariate transformations are the exception rather than the rule: the fact that one model contains information not represented in the other is part of the reason for having separate models”



P. Stevens. Bidirectional Model Transformations in QVT: Semantic Issues and Open Questions. Journal of Software and Systems Modeling (SoSyM) 9(1):7–20, 2010.

(Asymmetric) Lenses

■ View-Update Problem @ „Harmony Group”

- GET: view as function of source/master
- PUT(BACK): update master from view

$$B = \text{get}(A) \text{ and } A' = \text{put}(B', A)$$

(asymmetric) lens, **a-lens**

■ Trivial case: $A = \langle B, C \rangle$

- GET/PUT ignores complement (C)
- Again, this is not always applicable



Foster, J.N., Greenwald, M.B., Moore, J.T., Pierce, B.C., Schmitt, A.: Combinators for bi-directional tree transformations: A linguistic approach to the view update problem. ACM Trans. Program. Lang. Syst. 29(3), 17 (2007)

Orig. SIGPLAN paper: 2005 (coinage of „lens”)

Semantic Variants

■ State-based vs. delta-based

- Incrementality
- Preservation of intent

Diskin, Z, Xiong, Y, Czarnecki, K. (2011). From State- to Delta-Based Bidirectional Model Transformations: the Asymmetric Case.. Journal of Object Technology. 10. 6: 1-25.10.5381/jot.2011.10.1.a6.

■ Reflective updates

The best choice may include self-propagation from the updated view to itself

Diskin Z., König H., Lawford M. (2018) Multiple Model Synchronization with Multiary Delta Lenses. In: Russo A., Schürr A. (eds) Fundamental Approaches to Software Engineering. FASE 2018. Lecture Notes in Computer Science, vol 10802. Springer, Cham

■ Partial consistency

- Partially ordered degree of consistency
- E.g. hypocratic = monotonic

Stevens P. (2014) Bidirectionally Tolerating Inconsistency: Partial Transformations. In: Gnesi S., Rensink A. (eds) Fundamental Approaches to Software Engineering. FASE 2014. Lecture Notes in Computer Science, vol 8411. Springer, Berlin, Heidelberg

- High-level consistency applicable only after low-level

Basics

Properties

Constructions

Applications &
Generalizations

Desirable Properties

(Very) well-behaved BX
Least **surprise**

Well-behaved

- Well-behaved BX: correct + hyppocratic

Correct

$$T(m, \overrightarrow{T}(m, n))$$

$$T(\overleftarrow{T}(m, n), n)$$

Hypocratic

$$T(m, n) \implies \overrightarrow{T}(m, n) = n$$

$$T(m, n) \implies \overleftarrow{T}(m, n) = m$$

„First, do no harm.” Hippocrates, 450-355BC

- Special case for a-lens: PUTGET + GETPUT

$$l.get (l.put v s) = v \quad (\text{PUTGET})$$

$$l.put (l.get s) s = s \quad (\text{GETPUT})$$

Very well-behaved

- Very well-behaved (overwriteable):
 - Well-behaved + **history ignorant**

$$\overrightarrow{R}(m, \overrightarrow{R}(m', n)) = \overrightarrow{R}(m, n)$$

$$l.put\ v' (l.put\ v\ s) = l.put\ v' s \quad (\text{PUTPUT})$$

$$\overleftarrow{R}(\overleftarrow{R}(m, n'), n) = \overleftarrow{R}(m, n)$$

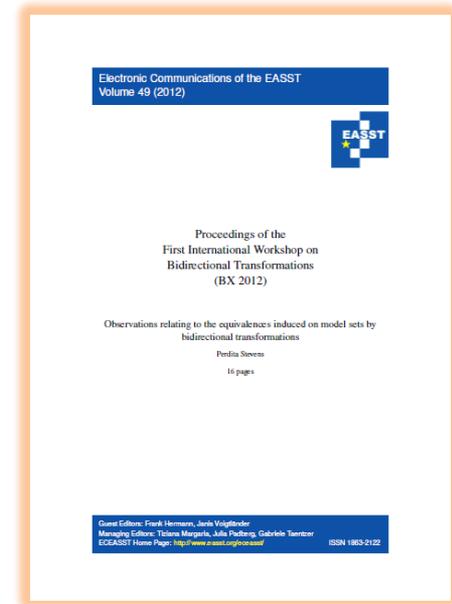
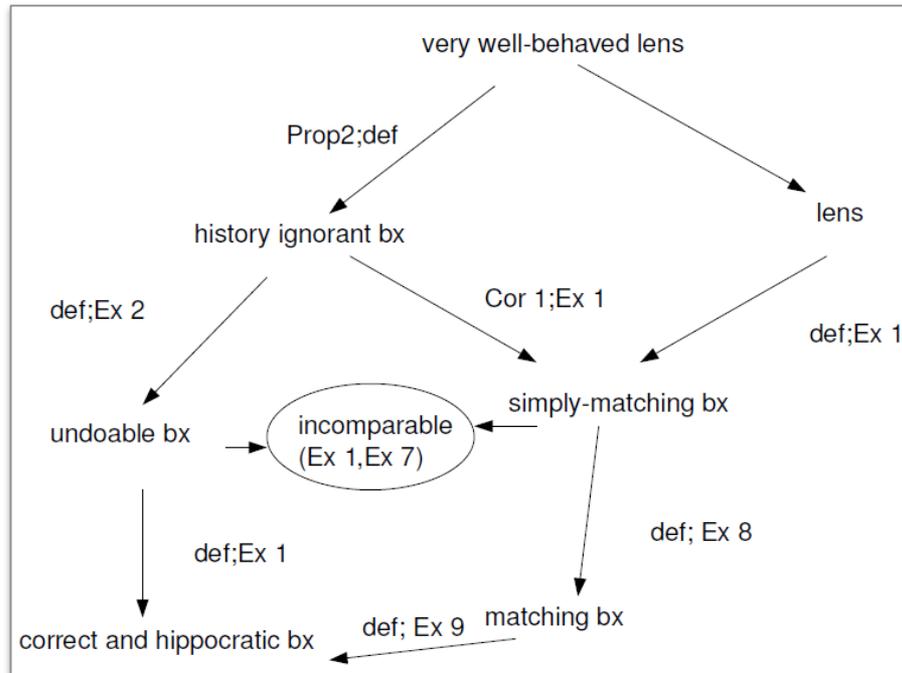
- This is an *optional* requirement

„fail to satisfy it for reasons that seem pragmatically unavoidable”

„a draconian restriction”

Very well-behaved cont'd.

- Very well-behaved:
 - Well-behaved + **history ignorant**
- Refined taxonomy



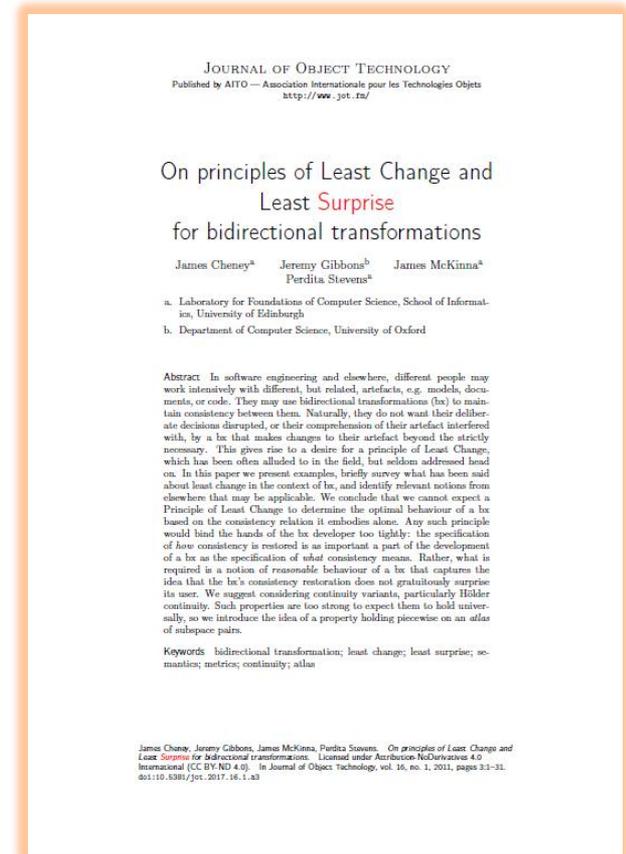
Stevens, Perdita. (2012). Observations relating to the equivalences induced on model sets by bidirectional transformations. EC-EASST 049. 10.14279/tuj.eceasst.49.714.

Least Change / Least Surprise

- Many correct choices for restorer / PUT (if non-bijective)
 - Which is the „best choice“?
 - c.f. „constant complement“ in database view updates

The action taken by the maintainer of a constraint after a violation should change no more than is needed to restore the constraint.

It turns out, however, that there are devils in the details.



James Cheney, Jeremy Gibbons, James McKinna, Perdita Stevens, “On principles of Least Change and Least Surprise for bidirectional transformations”, Journal of Object Technology, Volume 16, no. 1 (February 2017), pp. 3:1-31, doi:10.5381/jot.2017.16.1.a3.

Basics

Properties

Constructions

Applications &
Generalizations

Constructions

Composition and decomposition

Bidirectional Programming

Grammar-based and federation-based synchronization

Composition of Lenses

- How to build large synchronization systems?
- Composition of lenses
 - Sequential composition
 - Span (shared source/master)
 - Co-span (shared target/view)
- Composition laws, e.g.

Proposition 10. The co-targetial composition ($\mu \bowtie \nu$) of lenses is a (very) well-behaved/undoable diagonal system as soon as both lenses are such. In addition, the system is Hippocratic.

Algebra of bidirectional model synchronization

Technical Report CSRG-573
Department of Computer Science,
University of Toronto, 2008

Zinovy Diskin
zdiskin@cs.toronto.edu

Abstract. The paper presents several algebraic models for semantics of bidirectional model synchronization and transformation. Different types of model synchronization are analyzed (view updates, selective and incremental synchronization), and this analysis motivates the formal definitions. Particularly, a new formal model of updates is proposed. Relationships between the formal models are precisely specified and discussed.

1 Introduction

By the very nature of modeling, a snapshot of an MDD-project appears as a diverse collection of interrelated models. If one of these models is modified, other related models must be also modified to maintain the relationships. In other words, updates to a model need to be propagated to other related models to keep the entire collection consistent. We will call this activity *model synchronization*.

Two main classes of synchronization scenarios can be distinguished. One is synchronization of overlapping models. For example, consider a behavioral model m (e.g., a UML sequence diagram) and a structural model n (a class diagram or a database schema) of the same domain. These models overlap in the sense that the structural part of m is normally a part of n , or a view to it. We may say that these two models have a common view/abstraction a and informally write something like $m \succ a \prec n$. If one of the models m, n is updated, its counterpart must be correspondingly updated too. Note also that in both cases the consistency relation between models is not a (single-valued) function: given m , there may be

* Supported by Bell Canada through the Bell University Labs, NSERC, and the Ontario Centres of Excellence.

Diskin Z. (2008) Algebraic Models for Bidirectional Model Synchronization. In: Czarnecki K., Ober I., Bruel JM., Uhl A., Völter M. (eds) Model Driven Engineering Languages and Systems. MODELS 2008. Lecture Notes in Computer Science, vol 5301. Springer, Berlin, Heidelberg

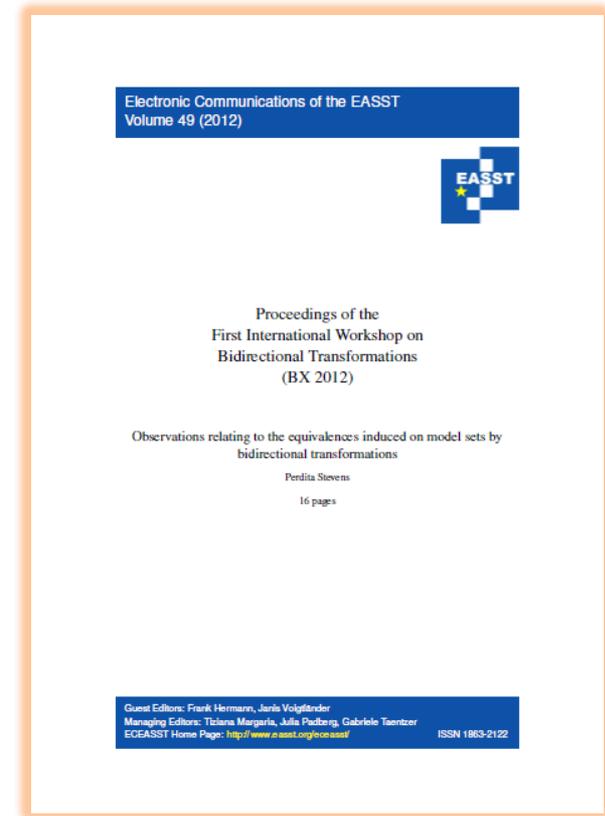
Decomposition of Lenses

- Decomposition: $BX \rightarrow$ a-lenses
 - Programming lenses may be easier!
 - Span of a-lenses: trivial
 - Any WB BX is \sim a span of lenses from a „union“ master model
 - Co-span of a-lenses:

Definition 8 Let $R : M \leftrightarrow N$ be a bidirectional transformation inducing equivalences and a coordinate system as usual. We say that R is *matching* if there is a bijection $f : M_F \rightarrow N_B$ such that $R(m_F, f(m_F))$ for all $m_F \in M_F$. We say that R is *simply matching* if, in addition, $R(m_F, n_B)$ holds *only* when $n_B = f(m_F)$.

Corollary 1 *If R is history ignorant, then R is simply matching.*

Theorem 1 *A bidirectional transformation $R : M \leftrightarrow N$ is simply matching if and only if it can be decomposed into a pair of lenses working “tail to tail”. Any such decomposition gives rise to a choice of matching transversal for R , and vice versa.*



Stevens, Perdita. (2012). Observations relating to the equivalences induced on model sets by bidirectional transformations. EC-EASST 049. 10.14279/tuj.eceasst.49.714.

Bidirectional Programming @ Harmony Group

- „Write one program, execute in 2 directions”
 - based on lens compositions (sequential, conditional...)
- **Harmony: updateable XML views (fork, hoist, etc.)**

Foster, J.N., Greenwald, M.B., Moore, J.T., Pierce, B.C., Schmitt, A.: Combinators for bi-directional tree transformations: A linguistic approach to the view update problem. ACM Trans. Program. Lang. Syst. 29(3), 17 (2007)

- **Relational Lens: updateable RelDB query (with FD)**

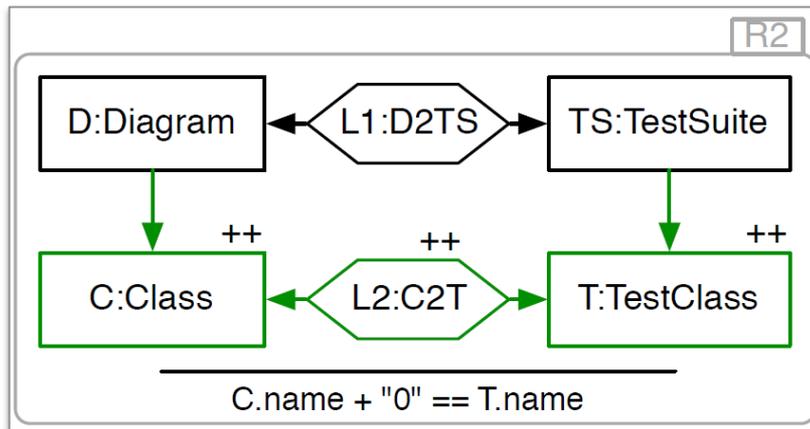
Bohannon, A., Vaughan, J.A., Pierce, B.C.: Relational lenses: a language for updateable views. In: Principles of Database Systems (PODS), Extended Version Available as University of Pennsylvania technical report MS-CIS-05-27 (2006)

- **Boomerang: bidirectional typed ~regex (+dict)**

Bohannon, A., Foster, J.N., Pierce, B.C., Pilkiewicz, A. and Schmitt, A.: Boomerang: Resourceful Lenses for String Data. In ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), San Francisco, California, January 2008.

Triple Graph Grammars

- Rule-based specification for consistency
 - Automatically derived BX restorers
 - Correctness✓, other properties?
- ~ translation in string grammars



Aachener Informatik-Berichte AIB 94-12

Specification of Graph Translators with Triple Graph Grammars

Andy Schürr
Lehrstuhl für Informatik III, RWTH Aachen,
Ahornstr. 55, D-52074 Aachen
e-mail: andy@is.informatik.rwth-aachen.de

Abstract. Data integration is a key issue for any integrated set of software tools where each tool has its own data structures (at least on the conceptual level), but where we have many interdependencies between these private data structures. A typical CASE environment, for instance, offers tools for the manipulation of requirements and software design documents and provides more or less sophisticated assistance for keeping these documents in a consistent state. Up to now almost all of these data consistency observing or preserving integration tools are hand-crafted due to the lack of generic implementation frameworks and the absence of adequate specification formalisms. Triple graph grammars, a proper superset of pair grammars, are intended to fill this gap and to support the specification of interdependencies between graph-like data structures on a very high level. Furthermore, they form a solid fundament of a new machinery for the production of batch-oriented as well as incrementally working data integration tools.

1. Introduction

Graphs play an important role within many application areas of computer science, as e.g. in the form of data flow or control flow graphs in compiler construction, structured analysis and entity relationship diagrams in software engineering, or hypertexts in office automation. Furthermore, *rule-based systems* have proven to be well-suited for the description of complex transformation or inference processes on complex data structures.

Although graphs and rule-based systems are quite popular, their combination in the form of *graph rewriting systems* or *graph grammars* was more or less unknown for a very long time. Nowadays the situation is gradually improving with the appearance of graph rewriting system implementations like PÄGG [6], GraphED [8], AGG [10], and PROGRES [16]. Especially the latter one has quite successfully been used within the project IPSEN for the development of an Integrated Project Support Environment [4, 15, 22], and within the project SUKITS for the development of an (a posteriori) integrated CIM environment [5].

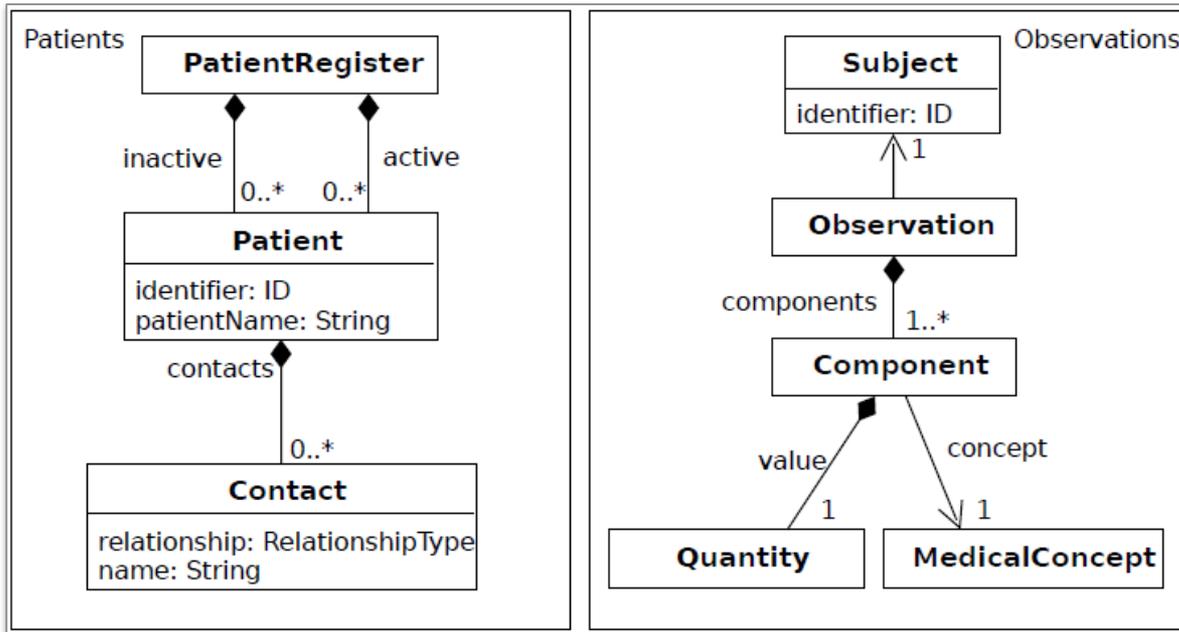
Nevertheless, graph rewriting systems are usually restricted to the specification of processes which perform in-place modifications and transform one instance of a class of graphs into another instance of the same class. Therefore, they are not well-suited for the specification of compilers and integration or traceability tools which

- either take a complex data structure (source graph) as input and translate it into a new, separate data structure (target graph),
- or check consistency between different data structures,
- or propagate small changes of one data structure as incremental updates into another related data structure.

Schürr A. (1995) Specification of graph translators with triple graph grammars. In: Mayr E.W., Schmidt G., Tinhofer G. (eds) Graph-Theoretic Concepts in Computer Science. WG 1994. Lecture Notes in Computer Science, vol 903. Springer, Berlin, Heidelberg

Metamodel composition for federation

- If models are split/merged based on metamodel...
 - Composition as „colimit“



Multimodel Correspondence through Inter-model Constraints

Patrick Stünkel
Western Norway University of Applied Sciences
Bergen, Norway
pst@hvl.no

Yngve Lamo
Western Norway University of Applied Sciences
Bergen, Norway
yla@hvl.no

Harald König
FH DW Hannover
Hannover, Germany
Harald.Koenig@fh-w.de

Adrian Rutle
Western Norway University of Applied Sciences
Bergen, Norway
aru@hvl.no

ABSTRACT
The synchronisation of a 2 heterogeneous typed models requires a thorough understanding of global consistency rules. After having related these models by determining identical entities in them, we express the global rules in terms of diagrammatic predicates imposed on a comprehensive metamodel, which integrates the structural properties of all involved model spaces. If the global rules are violated, a possible subsequent consistency restoration can make use of formal descriptions for the verification of these rules.

The comprehensive metamodel is constructed in the category of directed graphs. If there is an arbitrary number of related models, comprehensive artifacts can formally be represented by colimits, i.e. by universal categorical constructions. The goal of the paper is to establish a practical algorithm for this construction. The main example, a web service integration scenario from the health care domain, also shows that relating graph morphisms may be non-injective, which is incorporated into the algorithm.

CCS CONCEPTS
• Software and its engineering → System modeling languages; Integration frameworks;

KEYWORDS
Bidirectional Transformations, Modeling Languages, Inter-Model Constraints

ACM Reference Format:
Patrick Stünkel, Harald König, Yngve Lamo, and Adrian Rutle. 2018. Multimodel Correspondence through Inter-model Constraints. In *Proceedings of 2nd International Conference on Art, Science, and Engineering of Programming (Programming'18 Companion)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3191697.3191715>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made for distribution for profit or commercial advantage and that copies bear the notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with modest fee is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5174-4, 2018.
<https://doi.org/10.1145/3191697.3191715>

1 INTRODUCTION
Bidirectional Transformations (BX) [5, 14] marks an important field of cross-disciplinary research within computer science, comprising Programming Languages [26], Databases [5], and Graph Transformation [10, 26]. The general content underlying BX are model transformations with respect to a consistency relation between a related models.

Traditionally, research on BX focuses on the binary case $n = 2$. The synchronisation of updatable data sources, however, is not limited to model pairs. The challenge of the general n -ary case ($n \geq 2$) is to treat more complex model inter-relations, which are beyond simple model-to-model transformations. Recently this case has gained more attention [27]. In this spirit we will consequently assume the existence of an arbitrary number of model spaces. Moreover, we assume that all models of these spaces can be formalized as directed multi-graphs. This is no serious limitation, as directed graphs are widely accepted as an appropriate formalism in the field of model-driven engineering (MDE) [10, 18, 25, 26]. In the examples of this paper each model space is represented by a metamodel and elements of such a space are models typed over the metamodel.

Model tuples are connected via a correspondence relation. In this paper correspondences arise from commonalities. If a person is simultaneously contained in a family register A_1 and in a person register A_2 (cf. corresponding metamodels in Fig. 3), then this fact is declaratively specified. The correspondence is refined as a special model A_3 together with mappings from A_3 to $A_{1,2}$, resp. If $n = 2$, this is a binary relation as in the theory of Triple Graph Grammars (TGG) [26]. The general case ($n \geq 2$), however, requires partial commonality mappings as in [16].

Once these correspondences are settled, joint consistency of models may still be violated in the presence of global constraints, e.g. if we require the number of persons in the family register to be the same as the number of persons in the person register. Hence, it is important to formally determine, when corresponding models are jointly consistent. We will show how the Diagrammatic Predicate Framework (DPF) [24, 25] can be used for this. Constraints can be imposed on a comprehensive metamodel, which represents the union (the merge) of all participating metamodels. To achieve effective results, it has turned out, that the computation of merged models and metamodels has to be carried out with the help of a categorical construction, the so-called “colimit” [12].

*We explain this construction in detail in Sec. 4

Stünkel, P., König, H., Lamo, Y., Rutle, A.. 2018. Multimodel correspondence through inter-model constraints. In *Conference Companion of the 2nd International Conference on Art, Science, and Engineering of Programming (Programming'18 Companion)*. ACM, New York, NY, USA, 9-17. DOI: <https://doi.org/10.1145/3191697.3191715>

Basics

Properties

Constructions

**Applications &
Generalizations**

Applications and Generalizations

Functional Programming
Access Control
MX

Monadic BX in Functional Programming

- Provide result under some monad, e.g.
 - Future<T>
 - User choice, e.g. Input<T>
 - State (transformation remembers history)
- Does it compose / commute with other monads?
- Many approaches

Abou-Saleh F., Cheney J., Gibbons J., McKinna J., Stevens P. (2015) Notions of Bidirectional Computation and Entangled State Monads. In: Hinze R., Voigtländer J. (eds) Mathematics of Program Construction. MPC 2015. Lecture Notes in Computer Science, vol 9129. Springer, Cham

Abou-Saleh F., Cheney J., Gibbons J., McKinna J., Stevens P. (2016) Reflections on Monadic Lenses. In: Lindley S., McBride C., Trinder P., Sannella D. (eds) A List of Successes That Can Change the World. Lecture Notes in Computer Science, vol 9600. Springer, Cham

Updateable Security Views

- RW access control as lens
 - Read access control = GET
 - Write access control = PUT...
 - ...that can reject
- Contributions
 - Basic formalism
 - Desirable properties
 - Composition
 - Boomerang-like applications

Updateable Security Views

J. Nathan Foster Benjamin C. Pierce Steve Zdancewic
University of Pennsylvania

Abstract

Security views are a flexible and effective mechanism for controlling access to confidential information. Rather than allowing untrusted users to access source data directly, they are instead provided with a restricted view, from which all confidential information has been removed. The program that generates the view effectively embodies a confidentiality policy for the underlying source data. However, this approach has a significant drawback: it prevents users from updating the data in the view.

To address the "view update problem" in general, a number of bidirectional languages have been proposed. Programs in these languages—often called lenses—can be run in two directions: read from left to right, they map sources to views; from right to left, they map updated views back to updated sources. However, existing bidirectional languages do not deal adequately with security. In particular, they do not provide a way to ensure the integrity of source data as it is manipulated by untrusted users of the view.

We propose a novel framework of secure lenses that addresses these shortcomings. We enrich the types of basic lenses with equivalence relations capturing notions of confidentiality and integrity, and formulate the essential security conditions as non-interference properties. We then instantiate this framework in the domain of string transformations, developing syntax for bidirectional string combinators with security-annotated regular expressions as their types.

1. Introduction

Security views are a widely used mechanism for controlling access to confidential information in databases and other systems that manage structured information. By forcing users to access data via views that only expose public information, data administrators ensure that secrets will not be leaked, even if the users mishandle the data or are malicious. Security views are robust, making it impossible for users to leak the source data hidden by the view, and they are flexible: since they are implemented as arbitrary programs, they can be used to enforce extremely fine-grained access control policies. However, they are not usually updateable—and for good reason! Propagating updates to views made by untrusted users can, in general, alter the source data, including the parts that are hidden by the view.

Still, there are many applications in which having a mechanism for reliably updating security views would be extremely useful. For example, consider IntelliPedia, a collaborative data sharing system based on Wikipedia that is used by members of the intelligence community. The data stored in IntelliPedia is classified at the granularity of whole documents, but many documents actually contain a mixture of highly classified and less-classified data. In order to give users with low clearances access to the portions of documents they have sufficient clearance to see, documents often have to be regraded: i.e., the highly classified parts need to be erased or redacted, leaving behind a residual document—a security view—that can be reclassified at a lower level of clearance. Of course (since we are talking about a wiki), we would like the users of these views to be able to make updates—e.g., to correct errors or add new information—and have their changes be propagated back to the original document.

In general, for a view to be updateable, the program that generates it needs to be bidirectional. That is, it must not only be able to transform sources to views but also to map updated views back to updated sources. In previous work, we and many others have proposed a family of languages for describing bidirectional transformations, often called lenses [19], [8], [7], [21], [37], [44], [26], [9], [24], [35], [17], [23], [30], [28]. Formally, a lens l mapping between a set S of "source" structures and a set V of "views" comprises three functions

$$\begin{aligned} l.get &\in S \rightarrow V \\ l.put &\in V \rightarrow S \rightarrow S \\ l.create &\in V \rightarrow S \end{aligned}$$

that obey "round-tripping" laws for every $s \in S$ and $v \in V$.

$$\begin{aligned} l.get(l.put\ v\ s) &= v && \text{(PUTGET)} \\ l.get(l.create\ v) &= v && \text{(CREATEGET)} \\ l.put(l.get\ s)\ s &= s && \text{(GETPUT)} \end{aligned}$$

The get function defines the view and is a total function from S to V . There are two functions that handle updates: the put function takes an updated V and the original S and weaves them together to yield a correspondingly modified S , while the $create$ function handles the special case where we need to compute an S from a V but have no S to use as the original (it fills in any source data that is not reflected in the view with default values).

J. N. Foster, B. C. Pierce and S. Zdancewic, "Updateable Security Views," 2009 22nd IEEE Computer Security Foundations Symposium, Port Jefferson, NY, 2009, pp. 60-74. doi: 10.1109/CSF.2009.25

Unique MX Challenges

■ Two seminal papers

P. Stevens, "Bidirectional Transformations in the Large," *2017 ACM/IEEE 20th International Conference on Model Driven Engineering Languages and Systems (MODELS)*, Austin, TX, 2017, pp. 1-11. doi: 10.1109/MODELS.2017.8

Diskin Z., König H., Lawford M. (2018) Multiple Model Synchronization with Multiary Delta Lenses. In: Russo A., Schürr A. (eds) *Fundamental Approaches to Software Engineering. FASE 2018. Lecture Notes in Computer Science*, vol 10802. Springer, Cham

■ Dagstuhl Seminar 2018 Dec: many open questions

- Non-local consistency?
- Concurrent updates?
- Whack-a-mole property?
- Composition of BX into MX?
 - Cyclic networks?
- MX type theory etc.

