

The Typed Graph Model – a Supermodel for Model Management and Data Integration

Fritz Laux
 Fakultät Informatik
 Reutlingen University
 D-72762 Reutlingen, Germany
 email: fritz.laux@fh-reutlingen.de

Abstract—In recent years, the Graph Model has become increasingly popular, especially in the application domain of social networks. The model has been semantically augmented with properties and labels attached to the graph elements. It is difficult to ensure data quality for the properties and the data structure because the model does not need a schema. In this paper, we propose a schema bound Typed Graph Model with properties and labels. These enhancements improve not only data quality but also the quality of graph analysis. The power of this model is provided by using hyper-nodes and hyper-edges, which allows to present data structures on different abstraction levels. We prove that the model is at least equivalent in expressive power to most popular data models. Therefore, it can be used as a supermodel for model management and data integration. We illustrate by example the superiority of this model over the property graph data model of Hidders and other prevalent data models, namely the relational, object-oriented, XML model, and RDF Schema.

Keywords—typed hyper-graph model; semantic enhancement; data quality.

I. INTRODUCTION

The popularity of the Graph Model (GM) stems primarily from its application to social networks, medicine, scientific literature analysis, drug analysis, power and telephone networks. The flexibility of the GM contributes to its popularity, but its schema-less implementations are prone to data quality problems. This was pointed out in our DBKDA paper [1] in which we introduced the Typed Graph Model (TGM) with schema support. The present work extends our findings about the TGM and provides a proof of the expressive power and demonstrates its superiority over most popular data models.

Commercial graph database products like Neo4J [2], ArangoDB [3], JanusGraph [4], Amazon Neptune [5], and others have been successfully applied to many domains. Advocates of the GM like Robinson et al. of Neo4J recommend in their book [6] to use specification by example, which builds on example objects. But this reaches not far enough as the following example taken from Robinson's book shows. It is depicted in Figure 1 and shows a *User* named Billy with its 5-star *Review* on a *Performance* dated 2012/7/29. From this example we cannot know if Billy is allowed to have multiple reviews (on the same performance). For good data quality, a review should depend on the existence of a user and a performance. But this cannot be derived from one example. This means that we have to deal with class things (like a generic Person) and not only with real objects (like Billy) and specify if a relationship is mandatory or optional.

In order to express structural information, it is necessary to abstract from a particular situation and specify integrity constraints. The use of a schema would help to ensure data

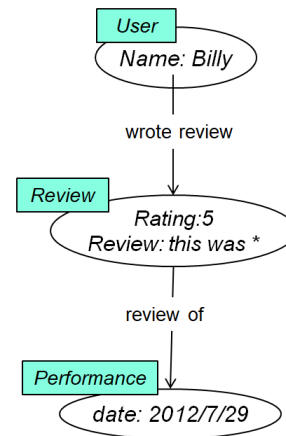


Figure 1. Example graph taken partially from [6], p. 42.

integrity and would clarify the intended situation of the example. Daniel et al. [7] complain that "there are only few solutions that target conceptual modeling for NoSQL databases and even less focusing on graph databases". They also point out the importance of a schema for data consistency and efficient implementation of a graph database and propose a framework, that translates an UML schema definition into a graph representation, and generate database-level queries from business rules and invariants.

Another weakness of the GM is that it has no notation to support different levels of detail and abstraction, which is apparently important for modeling large and complex data structures.

A. Contribution

To overcome these limitations we introduce in this paper a new typed graph model allowing hyper-nodes with complex structured properties (even sub-graphs) and hyper-edges connecting (recursively) one, two or more hyper-nodes. The graph schema provides data types, which allow type checking for instance elements. This ensures a formal data quality. We prove that the model is at least equivalent in expressive power to most popular data models. It can be used on both, the instance and schema level. Its semantic power makes it suitable for a loss-less model management and high quality data integration. Our model has a higher semantic expressiveness and precision than the prevalent data models, namely the relational, object oriented, XML data model, and RDF Schema. This will be demonstrated with typical modeling patterns.

B. Structure of the Paper

With the following overview of Related Work (Section II) the context for our new typed graph model will be settled. Section III introduces and defines formally the Typed Graph Model (TGM) consisting of a typed schema and a hyper-graph instance connected to the schema. We present a compact and easy to read visualization of the model using UML. The definitions are illustrated by some examples and the abstraction capability is demonstrated by the data model of a commercial enterprise.

In Section IV we prove a semantic preserving schema translation for prevalent data models to the TGM. Properties of the TGM are explained and the translation process is illustrated for an Entity Relationship example. In the next Section V our TGM is compared to the Graph Data Model (GDM) of J. Hidders [8]. Then, the semantic expressiveness of the TGM is demonstrated with typical data structures and compared with the prevalent data models, namely the relational, object oriented, and XML data model. The paper ends with a summary of our findings and gives an outlook on ideas for future work.

II. RELATED WORK

Since the beginning of 1980 many papers on the GM have been published. DBLP [9] alone retrieves 920 matches (retrieved July 30, 2021) for the key words "graph data model". If we ignore the papers that present specific applications for the GM incl. XML or Hypertext applications a few dozen of relevant papers remain. In the following, we discuss only papers that present the GM and its extensions (e. g., the Property Graph Model (PGM)) with a formal foundation or papers that use a graph schema:

The notion of PGM was informally introduced by Rodriguez and Neubauer [10]. Spyratos and Sugibuchi [11] use property graphs with hyper-nodes and hyper-edges for their graph data model. The main difference to our approach is that no schema is used and properties have no predefined data type. Another approach with hyper-edges is presented by Bu et al. [12] who treats a label like a node connecting a set of nodes, which he calls hyper-edge. The nodes itself can be of different types. In this case Bu calls the graph a unified hyper-graph. The unified hyper-graph model is then applied to problems of ranking music content and combining it with social media information. Compared to our TGM the unified hyper-graph of Bu is only defined for graph instances. It is not clear if the nodes have any type checking and if the whole graph is ruled by a schema.

Ghrab et al. [13] present GRAB, a schemaless graph database based on the PGM. It supports integrity constraints but cannot ensure data quality because of missing data types for properties and labels. Neo4J [6] has similar foundations and features. It has optional support for integrity constraints and comes with a powerful and easy to use graph query language, called *Cypher*.

All these PGM variants originate as instance graphs and no special attention is given to the graph schema. No attempt is made to specify the different types of edges and the multiplicity of connections (edges) between different node types. Nodes are not typed and labels are not a proper substitute for

data types. Therefore it is important to combine the PGM with a schema.

Amann and Scholl [14] seem to be the first authors who connect a graph schema with its graph database instance. Nodes and edges do not have properties but both must conform to the schema. Their model is used for an algebra (hyperwalk algebra) for traversing the graph.

Marc Gyssens et al. [15] and Jan Hidders [8] use a labeled GM to represent a database schema where each property of an object is modeled as a node in the graph. Labels are used to name node classes and edges. The models become confusing because a node represents either an object, a property or a data type. Still, it is not possible to restrict the cardinality of schema edges (relationships). Hidders' model is explained in more detail and compared to our TGM in Section V.

Similar to Amann and Scholl the paper of Pabón et al. [16] uses a graph schema to query the graph database. They distinguish different node types, which they call "sort". The supported types are: *object class nodes* (complex objects), *composite-value class nodes* (for aggregate values), and *basic-value class nodes* (primitive data types). This model seems to be equivalent to (complex) nodes with properties governed by a schema. A mechanism to abstract and group sub-graphs is missing, but would help to make the model easier to communicate.

Pokorný [17] uses a binary ER-Model as graph conceptual schema. For the graphical rendering he uses a compact entity representation for the nodes with attribute names inside the entity box. This solves the problem using the same node symbol for entities and attributes (properties) as it is the case with Gyssens [15] and Hidders [8] models. The edge cardinality is represented in a form of crow-foot notation.

In order to make the GM usable for real life scenarios with hundreds of schema elements, it is necessary to group or combine graph elements to higher abstracted objects. This would make the model easier to handle.

The need for grouping graph elements is addressed by Junghanns et al. [18]. Their model allows to form logical sub-graphs (graph collections) with heterogeneous nodes and edges. With this it is possible to aggregate sub-graphs, e. g., user communities. The authors use UML-like graphical rendering of nodes to make the model better readable but their model fails to specify the cardinality of schema edges.

A step toward to complex composite nodes as an alternative approach to aggregation presents Levene [19] by allowing the graph vertices to be recursively defined as a finite set of graphs. These hyper-nodes do not form a well-founded set as a node may contain itself, which violates the foundation axiom for the Zermelo-Fraenkel set theory.

A relatively new formal definition including integrity constraints was given by Angles [20]. However, his model does not allow structured objects and grouping or aggregation. In the following section, we simplify his definitions and use it as basis for our TGM.

A. Comparison with Ontology Languages

Ontology languages like the Resource Description Framework Schema (RDFS) [21] and the Web Object Language

(OWL) [22] are designed to specify ontologies and have their strength in allowing reasoning over instances of it. They are often used to semantically describe Linked Open Data (LOD) and the statement triples are usually visualized as graph structures. RDFS and OWL provide a general type system that could be used to form user defined types. This would allow to use it as basis for a graph schema language. But if we look at the W3C OWL 2 Structural Specification [23] it seems difficult to define user specific classes and W3C itself uses UML class diagrams to illustrate OWL structures.

Most approaches that map RDF to property graphs only support instance graphs. This is the case for the papers of Chiba et al., Schätzle et al., and Nguyen et al. The paper of Chiba et al. [24] uses G2GML, a graph-to-graph mapping language where the source graph is a RDF-graph and the destination graph is a Property Graph (PG). With G2GML RDF patterns are specified in SPARQL syntax and the corresponding patterns of the PG are described in openCypher [25]. These patterns are mapped through *node maps* and *edge maps*. The main benefit is that the pattern specifications are domain-specific and declarative. There is no support for schema mappings.

Schätzle et al. [26] define a mapping from RDF to the property graph model of GraphX. Their aim is to provide better analysis performance with S2X, a SPARQL implementation on top of GraphX and Spark, despite the schema-free PGM. Again, no schema mapping is supported.

The paper of Nguyen et al. [27] represents each RDF triple element as a separate node, which justifies the name Labeled Directed Multigraph with Triple Nodes (LDM-3N). While other models represent predicates as labeled arcs, Nguyen et al. map them to nodes. Assertions about RDF statements (triples) are modeled with the singleton property and not by reification. This approach adds an extra computation step and doubles the number of triples, which bloats the graph model. The main application domain seems to be the analysis of RDF triples by mapping it to the LDM-3N graph model allowing the use of graph analysis algorithms. Finally, it does not support schema mapping.

The specification of data structures is not the core intention of RDF. In RDFS for instance it is not possible to define the cardinality of relationships. Likewise, OWL Lite has strong limitations on allowing only 0 or 1 as multiplicity of properties. Simple unique requirements and relations like one-to-one, one-to-many and many-to-one are cumbersome to define even in OWL Full. Complex data structures need a modeling language that allows to define different levels of abstraction, which is not the strength of these ontology languages. Most examples of RDFS or OWL do not care about the multiplicity of relationships (cardinalities may be guessed via property names) and grouping of attributes seem to be on the same level as objects or subjects.

All these arguments and examples make it clear that we need schema support to ensure high data quality when using graph databases. This can be achieved with the Typed Graph Model (TGM), which we develop in the next section.

III. THE TYPED GRAPH MODEL

Our TGM informally constitutes a directed property hypergraph that conforms to a schema. In the following definitions our notation uses small letters for elements (nodes, edges, data types, etc.) and capital letters for sets of elements. Sets of sets are printed as bold capital letters. A typical example would be $n \in N \in \mathbf{N} \subseteq \wp(N)$, where N is any set and $\wp(N)$ is the power-set of N .

A. Graph Schema

Let T denote a set of simple or structured (complex) data types. A data type $t := (l, d) \in T$ has a name l and a definition d . Examples of simple (predefined) types are (int, \mathbb{Z}) , $(char, ASCII)$, $(\%, [0..100])$ etc. It is also possible to define complex data types like an order line (*OrderLine*, $(posNo, partNo, partDescription, quantity)$). The components need to be defined in T as well, e. g., $(posNo, int > 0)$. Recursion is allowed as long as the defined structure has a finite number of components.

Definition 1 (Typed Graph Schema). A typed graph schema is a tuple $TGS = (N_S, E_S, \rho, T, \tau, C)$ where:

- N_S is the set of named (labeled) objects (nodes) n with properties of data type $t := (l, d) \in T$, where l is the label and d the data type definition.
- E_S is the set of named (labeled) edges e with a structured property $p := (l, d) \in T$, where l is the label and d the data type definition.
- ρ is a function that associates each edge e to a pair of object sets (O, A) , i. e., $\rho(e) := (O_e, A_e)$ with $O_e, A_e \in \wp(N_S)$. O_e is called the tail and A_e is called the head of an edge e .
- τ is a function that assigns for each node n of an edge e a pair of positive integers (i_n, k_n) , i. e., $\tau_e(n) := (i_n, k_n)$ with $i_n \in \mathbb{N}_0$ and $k_n \in \mathbb{N}$. The function τ defines the min-max multiplicity of an edge connection. If the min-value i_n is 0 then the connection is optional.
- C is a set of integrity constraints, which the graph database must obey. The constraint language may be freely chosen.

The notation for defining data types T , which are used for node types N_S and edge types E_S , can be freely chosen. The integrity constraints C restrict the model beyond the structural limitations of the multiplicity τ of edge connections. Typical constraints of C are semantic restrictions of the content of an instance graph. This makes the expressiveness of the TGS at least as strong as the models to which it is compared in Section V.

B. Typed Graph Model

Definition 2 (Typed Graph Model). A typed graph Model is a tuple $TGM = (N, E, TGS, \phi)$ where:

- N is the set of named (labeled) nodes n with data types from N_S of schema TGS.
- E is the set of named (labeled) edges e with properties of types from E_S of schema TGS.

- TGS is a typed graph schema as defined in Subsection III-A.
- ϕ is a homomorphism that maps each node n and edge e of TGM to the corresponding type element of TGS, formally:

$$\begin{aligned} \phi : TGM &\rightarrow TGS \\ n &\mapsto \phi(n) := n_S (\in N_S) \\ e &\mapsto \phi(e) := e_S (\in E_S) \end{aligned}$$

The fact that ϕ maps each element (node or edge) to exactly one data type implies that each element of the graph model has a well defined data type. The homomorphism is structure preserving. This means that the cardinality of the edge types are enforced, too. Data type and constraint checking is applied for all nodes and edges before any insert, update, or delete action can be committed. If no single type can be defined, union type or *anyType* (sometimes called *variant*) may be applied. Usually this is an indication for a weak data model and it should be clear that this could affect data quality and processing.

As graphical representation for the TGS we adopt the UML-notation for nodes and include the properties as attributes including their data types. Labels are written in the top compartment of the UML-class. Edges of the TGS are represented by UML associations. For the label and properties of an edge we use the UML-association class, which has the same rendering as an ordinary class but its existence depends on an association (edge), which is indicated by a dotted line from the association class to the edge. This not only allows to label an edge but to define user defined edge types.

The correspondence between the UML notation and the TGS definition is shown in Table I.

TABLE I. TGS correspondence with UML notation

TGS	UML
$n \in N_S$	class
$e \in E_S$	association
$t = (l, d) \in T$	l = name of n resp. e ; d = type of n resp. e
$\rho(e)$	all ends of e
$\tau_e(n)$	(min,max)-cardinality of e at n
C	constraints in $[]$ or $\{ \}$

The use of hyper-nodes $n \in N_S$ and hyper-edges $e \in E_S$ instead of simple nodes resp. edges allow to group nodes and edges to higher abstracted complex model aggregates. This is particularly useful to keep large models clearly represented and manageable. Large graph models may then be grouped into sub-graphs like in Junghanns et al. [18]. Each sub-graph can be rendered as a hyper-node. If the division is disjoint these hyper-nodes are connected via hyper-edges forming a higher abstraction level schema (see Figure 3 (b)).

C. Examples

Lets recall the example graph from Figure 1 and model its corresponding schema. We want to make clear that a user may write as many reviews as he likes, but only one for a particular performance. A rating needs to refer exactly to one performance and one user. This is reflected in Figure 2 by the

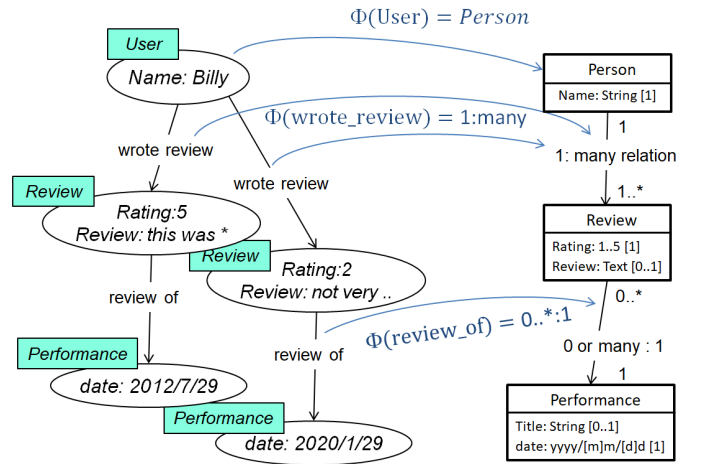


Figure 2. Example graph with schema in UML notation.

”1:many” and ”0 or many:1” relationships. We use the UML-notation for the schema and keep the notation from Figure 1 for the instance graph for clarity.

The homomorphic mapping ϕ guaranties that the instance graph obeys the schema, i. e., type, cardinality, and constraint checking. Now, it is clear from the schema that a user must have at least one review. The review is existence dependent on the user and a performance. The ”wrote review” edge is a 1:many relation and ”review of” is an optional many:1 relation. This has the consequence that a review needs a person and a performance. But, a performance may exist without any review.

In the next example we present a commercial enterprise that sells products and parts to customers. The enterprise assembles products from parts and if the stock level is not sufficient it purchases parts from different suppliers. Figure 3 models this situation using UML rendering. It demonstrates the abstraction power of the TGM showing two schema abstraction levels. The upper part (a) shows the TGM on a detailed level. The properties are suppressed in the diagram for simplicity except for *Customer* and *CustOrder*. The schema is grouped into 3 disjoint sub-graphs depicted with dashed lines.

In the lower part (b) these sub-graphs are shown as hyper-nodes of the graph schema. This allows a simplified and more abstracted view of the model. Also, some aggregate properties (e. g., #orders) are shown to illustrate the modeling capabilities. The coloring of the (hyper-)edges helps the reader to identify which edges have been aggregated. The hyper-edges connecting these abstracted nodes must use the most general multiplicity of the multiple edges it combines. In the example the edge *orders/from* combines two edges, i. e., *orders* with 0..1 - 1 multiplicity and *from* with 0..* - 0..* multiplicity, which leads to the most general multiplicity.

IV. PROPERTIES OF THE TYPED GRAPH MODEL

The TGM has some valuable properties and can be regarded as a **supermodel** because its semantic expressiveness proves to be at least as powerful as prevalent data models. In order to prove this statement we need the following definitions.

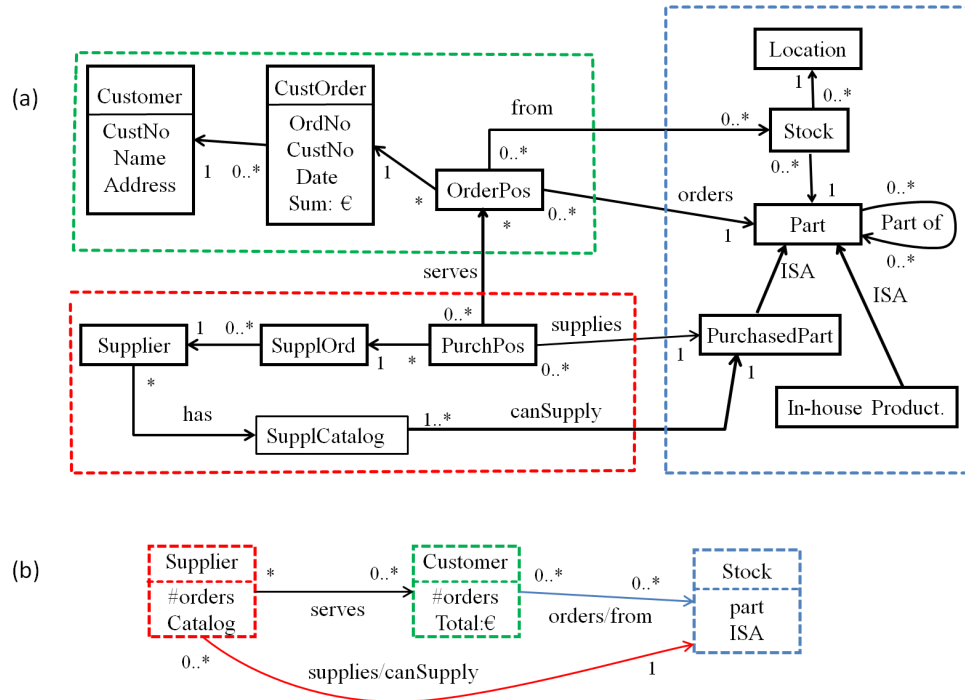


Figure 3. Example TGM of a commercial enterprise showing two levels of detail.

As we are only interested in databases that satisfy the constraints of its schema, we modify the definitions from Angles et al. [28] and restrict it to schema bound databases. More precisely, let M be a database model and S^M be a database schema. A database is then an ordered pair $D^M := (S^M, I^M)$, where I^M is an instance obeying all constraints of S^M .

Definition 3 (Schema Mapping). Let M_1 and M_2 be two data models. A schema mapping from M_1 to M_2 is a total function SM from the set of all database schemas in M_1 , to the set of all database schemas in M_2 .

$$SM : M_1 \rightarrow M_2 \\ s_1 \mapsto s_2$$

where s_1 and s_2 are schemas of M_1 resp. M_2 .

A. Some Property Definitions

Every data model allows to structure the data according to its modeling elements. These conceptual elements determine the representational power of the data model. A model M_2 subsumes the information capacity of M_1 if and only if every schema in M_1 can be translated to a schema in M_2 without loss of information. Two database models can be evaluated in terms of its information capacity considering the following properties:

Definition 4 (Computable Mapping). A schema mapping CM from $s_1 \in M_1$ to $s_2 \in M_2$ is computable if there exists an algorithm that translates schema s_1 into s_2 .

If any schema mapping from M_1 to M_2 is computable then this implies that M_2 subsumes M_1 . A computable mapping may still result in a schema that allows an invalid database instance.

Definition 5 (Semantics Preservation). A computable schema mapping SP is semantics preserving if for every valid database d_1 of schema s_1 , there is a valid database d_2 obeying schema s_2 where s_2 is produced by the mapping SP , i. e., $s_2 = SP(s_1)$.

This property guarantees that the result of the instance mapping will always be a valid database according to s_2 .

Definition 6 (Information Preservation). A semantics preserving schema mapping IP is information preserving if there exists an inverse computable schema mapping IP^{-1} from M_2 to M_1 such that for every $s_1 \in M_1$ it holds $s_1 = IP^{-1}(IP(s_1))$. Such a schema mapping IP is alternatively called "schema translation".

This definition indicates that, for some schema mapping IP , there exists an "inverse" mapping which allows recovering the original schema previously transformed. In general, the inverse mapping IP^{-1} is only a partial function because it is only defined on s_2 and may not be defined on all elements of M_2 , i. e., the image $SM(M_1)$ can be a proper subset of M_2 . Information preservation is an important property because it guarantees that a schema mapping results in a new schema capable to not lose any information. Moreover, it implies that the target database model M_2 subsumes the information capacity of the source database model M_1 . If M_1 subsumes M_2 as well, then both data models are equivalent and IP^{-1} and IP are total functions. McBrien and Poulouvasilis [29][30] describe equivalence-preserving mapping of schema constructs based on a Hypergraph Data Model (HDM). The ideas and problems with schema translation have been reviewed and revised in the light of model management by Bernstein and Melnik [31].

B. Schema Mapping by means of a Meta-model

When considering database mappings two types of mappings can be distinguished: (1) schema independent and (2) schema dependent. We are only interested in schema dependent mappings to always ensure high data fidelity. The above schema mapping of Definition 3 is model independent [32]. A meta-model that is general enough to capture all popular models in the literature would suffice. Hull [33] as well as Atzeni and Torlone [34][35] describe such a framework for heterogeneous data models. It was developed towards a model independent **supermodel** [36] and later implemented as a general tool, called MIDST [37]–[39]. It consists of the following meta-constructs: *lexical*, *abstract*, *aggregation*, *generalization*, and *function*. For example, the Entity-Relationship Model (ERM) involves (i) abstracts (the entities), (ii) aggregations of abstracts (relationships), and (iii) lexicals (attributes) with functions (to entities or relationships). This means that the ERM is a specialization of the supermodel, i. e., a schema in any (sub)model is also a schema in the supermodel, only the names of the model elements differ. Hull and Atzeni give more examples and claim that this supermodel subsumes Relational Model (RM), ERM, XML, Object Oriented Model (OOM), Object Relational (OR), and XSD.

C. Information Preserving Schema Translation to TGM

There are works that show information preserving mappings from ERM to Graph Database Schema [40], RM to RDF and OWL [41], and RM, CSV, XML JSON to RDF using RML, the RDF Mapping Language [42]. Angles et al. [28] have shown that a RDF database can be mapped to a property graph database. This includes information preservation for the schema translation and the instance mapping. Taking all together it seems possible to translate most popular data models to an enhanced graph data model, i. e., a PGM. It is evident that the TGM subsumes the PGM as any model element of PGM can be mapped 1:1 to the corresponding TGM model element. The more general approaches of Arenas et al. [43] and McBrien/Poulovassilis [30] provide criteria for model independent schema mappings that are information preserving. Hull [33] and Atzeni/Torlone [34][35] propose basic meta-constructs for the supermodel that covers all relevant data models. This puts us now in a position to state our main Theorem.

Theorem 1 (Information Preserving Schema Translation to TGM). *Let M be any data model that can be subsumed by the **supermodel** of Hull and let T be the Typed Graph Model. For any schema $s \in M$ there exists an information preserving mapping (translation) M :*

$$M : M \rightarrow T \\ s \mapsto t$$

where t is a TGS and s is a schema from M .

It would be easy to proof the theorem by contradiction. But such a proofs gives no constructive idea how a concrete mapping would look like. Therefore, we proof the theorem by constructing a generic mapping M taking the meta-constructs of the supermodel and assign uniquely model elements from T . These model element pairs allow us to map any schema $s \in M$ to a schema $t \in T$. Then it will be shown that we can

construct an inverse mapping that leads to the original schema s again.

Proof of the Translation Theorem: Following Atzeni and Torlone we define the schema mapping $M(M) = T$ by the following elementary 1:1 transformations for each model element $\sigma \in s$:

- 1) lexical \rightarrow property
- 2) abstract \rightarrow node
- 3) aggregation \rightarrow edge with aggregation type
- 4) generalization \rightarrow edge with generalization type
- 5) function \rightarrow edge with a single target (multiplicity 1)

The above transformation can be chained or composed to form a directed acyclic graph where each transformation assignment represents a node and the directed edges represent the sequence of assignments. Each path from a leaf node ends at the root node representing the mapping M . Given a schema s the mapping M translates the input s into a TGS $t \in T$. All transformation steps are 1:1 such that the resulting function M is a injective mapping that can be reversed. The inverse mapping M^{-1} uses just the opposite transformation assignments listed above. It should be noted that M^{-1} may be a partial function only, i. e., there might exist model elements in T (e. g., a composition edge type with existential dependency of its components) that may have no corresponding model element in the source model M .

Let σ w.l.o.g. be any model element from schema s . M translates σ into an element $\tau \in T$ in the following way:

$$\tau = M(\sigma) = \mathbf{m}_n(\mathbf{m}_{n-1}(\dots(\mathbf{m}_1(\sigma)\dots))$$

where \mathbf{m}_i are the elementary transformations from above. The data type σ_t of an element σ is carried over to the same data type of the translated element τ . Applying the inverse transformations \mathbf{m}_i^{-1} in opposite order results in the identity function I which proofs the information preservation property.

$$\sigma = \mathbf{m}_1^{-1}(\dots(\mathbf{m}_{n-1}^{-1}(\mathbf{m}_n^{-1}(\mathbf{m}_n(\mathbf{m}_{n-1}(\dots(\mathbf{m}_1(\sigma)\dots)) \\ = M^{-1}(M(\sigma)) = I(\sigma)$$

■

D. Example Translation from extended ERM to TGM

To demonstrate the translation algorithm we use an example from Hidders' GDM [8], which will be discussed and compared to the TGM in detail in the next section. At the moment, we concentrate on the translation process taking Hidders' example but using the ERM enhanced with generalization (IS-A) relationship. The original visualization of Hidders is depicted in Figure 5 (a). The extended Entity-Relationship (ER) diagram is shown in Figure 4 where the blue numbers next to the ER-symbols refer to the elementary transformations from Theorem 1.

The complete translation has to execute the elementary transformation steps from left to right (in order of the blue arcs) until all ERM elements have been converted. For instance, before the relationship *Contract* can be mapped to an edge of general aggregate type as indicated by transformation 3), the relationship attributes need to be mapped first. The *salary* is considered as a literal (lexical 1) that transforms directly

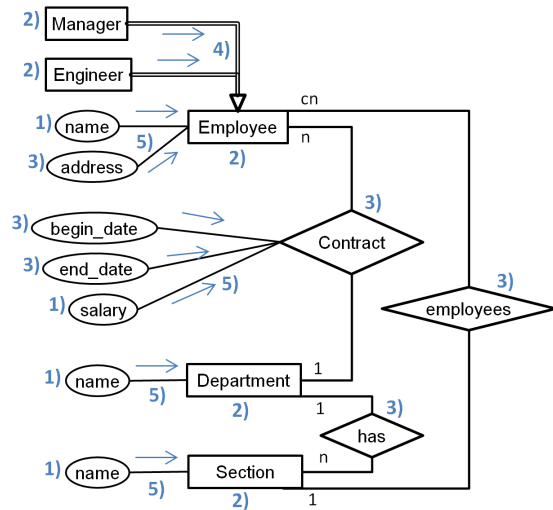


Figure 4. Hidders' example [8] transformed to a TGS.

to a *Contract* property. The *begin_date* and *end_date* are considered as aggregation model elements because of its date structure. All three attributes are single valued (indicated by the function 5) and mapped as properties of *Contract*.

It is also possible to specify more precisely the type of aggregation in the TGM by defining an edge type "Contract" that already includes the properties *begin_date*, *end_date*, and *salary*. This would result into the same semantics but with a user defined edge type. It should be pointed out that this possibility results from TGM's capability to support different abstraction levels.

The result of the translation is shown in the lower part of Figure 5. The visualization of the TGS uses the UML rendering listed in Table I.

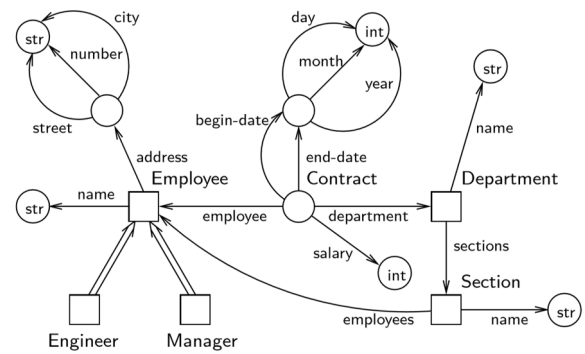
V. COMPARISON WITH OTHER DATA MODELS

In the following, we compare our TGM to other models with respect to structural differences and schema support. We point out modeling restrictions of these models and show how such situations are modeled with TGM. Query and manipulation languages are beyond the scope of this paper.


A. Comparison with GDM of Jan Hidders

Jan Hidders' [8] model added labels and properties together with their data types to nodes and edges (relationships). Property names are modeled as edges in the schema. This allows to model labeled relationships with complex properties. Structured and base data types share the same graphical representation, which makes it difficult to distinguish both. The ISA-relationship is rendered as a double line arrow similar to the extended ERM. Hidders' model does not allow to restrict the cardinality of relationships. This restriction limits its modeling power compared to the TGM, which provides a min-max notation for the cardinality.

The example in Figure 5 is from the publication of Hidders [8]. The schema shows *Employee* and *Department* classes linked by a *Contract*. The relationship *Contract* is existence



(a) Example schema from Hidders' GDM



(b) The same example modeled with TGM

Figure 5. Comparison by example with Hidders' GDM.

dependent on the connected nodes. The properties of *Contract* are salary of type *int*, begin-date and end-date of structure-type *date = (day, month, year)*. In Hidders' model these dates are modeled on the element level using data type *int*. Hidders' schema elements, i. e., nodes (objects), edges (properties) and data types appear on the same visual level, which makes it difficult to read and obscures semantics. The modeling power of complex data types provide a clear advantage for the TGM.

B. Comparison with the Relational Model (RM)

There is a 1:1 correspondence between attributes and properties and any relation can be modeled as a node with properties. The min-max notation for relationship multiplicity can model any link cardinality. The TGM can therefore easily represent tabular structures, foreign key constraints (many-to-one relationships), and join-tables as the building blocks of the RM. Beyond this, the TGM is able to directly model many-to-many relationships of any min-max multiplicity. This makes the TGM strictly stronger than the relational model.

Figure 6 (a) shows a typical FK relationship between Table 1 and Table 2. In part (b) of Figure 6 three tables are linked via a join-table RST with foreign keys fk_1, fk_2, fk_3 representing a many-to-many ternary relationship with attribute col_4 .

Another difference between TGM and RM is that foreign keys (FK) are not necessary because their function is taken over by an edge linking the Tab_1-node (Table 1 without FK) with the referenced node Table 2. This can be seen in Figure 7 (a).

A join-table in the RM is existence dependent on the tables it refers to by FKs. The FKs forming the primary key (PK) of the join table are not necessary in the TGM because of the same reason as mentioned above.

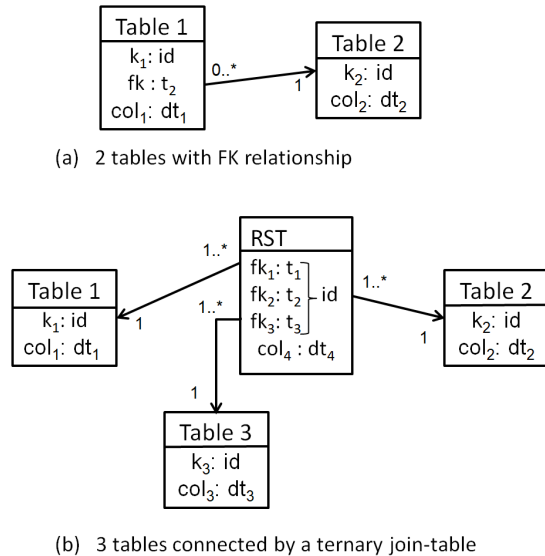


Figure 6. Modeling a many-to-one relationship (FK) and a ternary join-table with the RM.

In Figure 7 (b) the join-table RST maps directly to an hyper-edge labeled RST with property col_4 . The hyper-edge and the join-table in the relational case represent a many-to-many connection linking 3 nodes that correspond to the respective tables in the RM. In the relational model it is not possible to restrict this ternary relationship to $(*,1,1)$. For instance, if many connections of Table 1 should link to exactly *one* connection of Table 2 and *one* connection of Table 3. For the TGM this would be simple; only the cardinality at Table 2 and Table 3 would have to be changed from $1..*$ to 1.


To make the above ternary relationship example less abstract the RST could be an offer of products from Table 1 from one supplier of Table 3 to the client of Table 2. With this in mind it is clear that an offer depends on the product(s), one supplier, and one client.

The TGM can also represent non-normalized tables because the model supports complex structured data types having multivalued or array data. It is only necessary to define the appropriate data types in the set of available data types T .


C. Comparison with XML Schema

XML documents represent hierarchical hypertext documents. The document structure can be defined by an XML schema. The hierarchy of XML-documents is directly supported by the TGM using directed edges. XLink provides references (arcs) between elements of internal or external XML-documents. Extended XLinks can connect to more than one element, but the references are always instance based, i. e., the target elements must be listed by URI. The TGM is more abstract and expressive allowing the definition of non-hierarchical references on the schema level.

As example for the comparison serves a bookstore offering an unlimited number of books. A simple XML-schema for the bookstore is given by w3schools.com. The schema defines books with elements like "title", "author", etc. and its



(a) 2 nodes (tables) with many-to-one link (relationship)




(b) 3 nodes (tables) connected by a ternary hyper-edge (join-table)

Figure 7. Modeling a many-to-one relationship (FK) and a ternary join-table with TGM.

corresponding data types. Some data types are not as precise as they could, e. g., the data type `xs:double` for the price element. We will replace `xs:double` in our TGM by the money-type *euro* to be more precise. Some elements have attributes attached like the language ("lang") of a book title. The attribute `minOccurs="1"` of `xs:sequence` requires the bookstore to have a least one book.

```
<?xml version="1.0" encoding="utf-8"?>
<xs:schema ... >
  <xs:element name="bookstore" >
    <xs:complexType >
      <xs:sequence minOccurs="1"
        maxOccurs="unbounded" >
        <xs:element name="book" >
          <xs:complexType>
            <xs:sequence>
              <xs:element name="title" >
                <xs:complexType>
                  <xs:simpleContent>
                    <xs:extension base="xs:string">
                      <xs:attribute name="lang"
                        type="xs:string" />
                    </xs:extension>
                  </xs:simpleContent>
                </xs:complexType>
              </xs:element>
              <xs:element name="author"
                type="xs:string"/>
              <xs:element name="year"
                type="xs:integer"/>
              <xs:element name="price"
                type="xs:double"/>
            </xs:sequence>
            <xs:attribute name="category"
              type="xs:string"/>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
</xs:schema>
```

If we model the XML-elements as nodes in TGM then



(a) Compact representation (b) Detailed representation of each XML-element

Figure 8. Comparing an XML Schema example with TGM.

XML-attributes and the element values should be represented as properties. The name of an XML-element is mapped to a node label. The order of the XML-elements cannot be represented with this approach and XML-element values can be distinguished from XML-attributes by convention only.


An alternative TGM model represents the complete book structure as one node. In this case the XML-elements and their attributes are modeled as structured properties of the book. The order of the elements and their associated attributes can be preserved. In fact, if XML Schema is used for specifying the data types N_S and E_S (see Subsection III-A) all the flexibility and semantics provided by XML Schema can be represented with the TGS. This argument shows that the TGM is at least as powerful as the XML model.

The example bookstore is depicted in Figure 8 where the left part (a) shows the compact version with the whole book modeled as one node and the right part (b) shows the version where each XML-element is modeled as node. We see from this example another possibility to use sub-graphs for higher abstracted models.


D. Comparison with the Object-Oriented Model

Because we already use the UML for rendering the TGM, it is easy to see that classes correspond one-to-one with typed hyper-nodes. Any methods are simply ignored as we only deal with the network structure of OOM. Any complex internal class structure can be directly modeled by appropriate data types $t \in T$. The type set T is defined beforehand but can contain any user defined structures. In contrast to the OOM the TGM allows different levels of abstraction in the modeling depending whether a structure is modeled by a detailed graph with simple types or a more compact graph using complex data types. This shows the same semantic expressiveness for structures, but a higher flexibility of the TGM. Considering the operations on data the OOM has the advantage to specify the allowed operations by methods.

The UML provides a rich set of association types, which need to be mapped to the label of the edges. Our TGM provides types not only for nodes but also for edges (called associations in UML). With this information it is possible to model different association types like aggregation, generalization, etc. Even user defined associations are possible, e. g., an aggregate could be further qualified as un-detachable or detachable composition



(a) Two models for a bill of material (BOM)



(b) An instance of a product (table) structure conforming to the above schemas

Figure 9. Comparison by example with the OOM.

or a loose containment. The arrow of the edge only indicates the reading direction of the association but does not limit the navigation of the TGM.

It is also possible to model recursive structures as the examples from Figure 9 illustrates. The bill of material (BOM) is an important example for a recursive structure used in production planning and control. It defines recursively a (compound) part with its components until a single part is reached. As example instance a *table* is given in Figure 9 (b) consisting of 4 table legs and a tabletop consisting of a drawer and a mounting.

If the edge of *contains* in Figure 9 (a) is followed against the arrow direction it is possible to find the component where an individual part is built-in (used). A complete *where-used list* for a generic (not an individual) part may be obtained with a small schema modification. Only the from-end of the *contains*-edge needs to change its multiplicity from 0..1 to 0..*. With this small modification all components can be identified where a generic part is used.

E. Comparison with RDF Schema

RDF is a data format for expressing statements about resources with emphasis on Web data. A *resource* can be anything, including data, objects (incl. people), and (abstract) concepts. RDF Schema uses a semantic extension (meta-constructs) of the basic RDF vocabulary to model RDF data. Simple statements about a resource are expressed in the form of *subject-predicate-object* triples. It is clear that the underlying structure of such RDF-triples is graph based where a node represents a *subject* or an *object* and an edge represents the *predicate*. This implies that RDF Schema triples can be expressed as property graph.

Angles et al. [28] describe a schema-dependent computable schema mapping from an RDF Schema to a Property Graph Schema (PGS) that is semantics and information preserving. They prove in their paper that "the PGM subsumes the information capacity of the RDF data model". The mapping is illustrated in their paper with an example RDF Schema depicted in Figure 10. Because of the asymmetrical definition

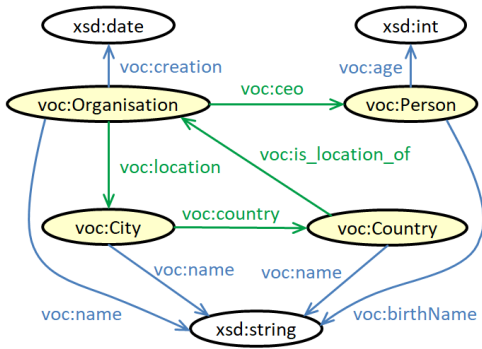


Figure 10. Graphical Illustration of the Example RDFS from Angles [28].

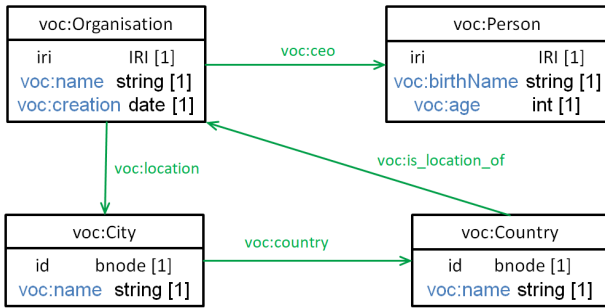


Figure 11. Resulting property graph schema of example Figure 10.

of subject (may not be a literal) and object (can be a literal) the mapping distinguishes objects that are also subjects and objects that are "only objects". The yellow colored subjects are mapped to nodes in the PGS and the "only objects" (the data types xsd:date, xsd:int, and xsd:string in their example) are mapped to PGS properties. The green colored predicates are RDF predicates which are mapped to PGS edges, whereas the blue colored predicates are properties of the subject with a data type specified by an "only object". The resulting property graph schema has 4 nodes and 4 edges and is visualized in Figure 11 as TGS using the same coloring scheme as in the RDF schema.

When comparing the RDF Schema with the TGS we see that both share most characteristics. In particular, nodes and edges have unique labels. These labels serve as global identifiers, called International Resource Identifier (IRI) in the RDF syntax.

There are some limitations of the PGM that the TGM overcomes:

- Properties of RDF graphs support multi-value properties, whereas the PGM usually supports only single value properties [28]. The TGM, however, uses the appropriate data-type to model multi-value properties. This can be array, set, list, or bag data types.
- In RDF edges can have edges, i. e., edges hold meta-data. This is realized with (edge, predicate, object) triples. The simplest situation for this is if a predicate has a label. The TGM handles this and more complex situations by edge properties (see Figure 12 for an example) or hyper-edges,

which is graphically rendered by an UML association class.

- RDF has three types of nodes (IRI, blank node, and literal) which need to be mapped to only one node type in the PGM. Because the TGM supports typed nodes any RDF node-type can be handled directly.
- The RDF model has some special semantics like reification and subclassing (rdfs:subClassOf and rdfs:subPropertyOf) that are not supported in the paper of Angles et al. [28]. The TGM also supports these special structures. This is achieved with the appropriate node or edge types.
- Each node or edge in an RDF graph contains one single IRI, a Literal or nothing (in case of a blank node), whereas each node or edge in a TGM could contain multiple properties, depending on its type. This includes multi-value properties and properties for hyper-edges as explained before.

With the TGM any kind of node type can be supported because of its capability to define user defined types for the nodes and edges. RDF allows to identify an RDF triple (statement) with an IRI. Such an aggregate construct is possible in TGM with a ternary edge type.

For a general schema mapping we have the following rules:

- 1) Abstract elements (subjects) are mapped to TGM nodes.
- 2) Aggregation (predicates) and generalization (rdfs:subClassOf, rdfs:subPropertyOf) elements are mapped to TGM edges with the appropriate type.
- 3) Functions (predicate cardinality is always 1) are mapped to TGM edge ends.
- 4) Lexical elements (literals, IRIs, "only objects") are mapped to TGM properties.

In order to illustrate the mapping to TGM we amend the example from Angles et al. [28] with more properties for the voc:ceo predicate. The predicate now holds the date when the CEO was appointed and the profit dependent bonus payments. For each voc:profit predicate a voc:bonus is associated. The profit is of type xsd:int and the bonus has a user defined percent data type with a range from 0 to 100. RDF auxiliary nodes (bnodes) are needed to add this information pair without any ambiguity. We need a bnode of type rdf:Alt to indicate that the bonuses are alternatives.

Figure 12 shows on the left side the amended RDFS example and the corresponding TGS on the right side. It immediately catches the eye that the TGS is far more readable for humans than the RDFS, which make it better suitable for data modeling. It is cumbersome to add properties to a predicate as it requires blank nodes (bnodes) in order to form legal and unambiguous RDF triples. For the corresponding TGM edge property we simply use an array of size 3 with a structured data type containing profit and the corresponding bonus in percent. The array *bonuses* can have a constraint XOR that clarifies the semantics as an alternative 1 out of 3. The comparison of both schemata clearly shows that even with the coloring the RDFS tends to quickly becomes confusing because properties are not aggregated to objects or predicates. Also data types are expressed on the same level as subjects, which makes the RDF less clear for data modeling purposes.

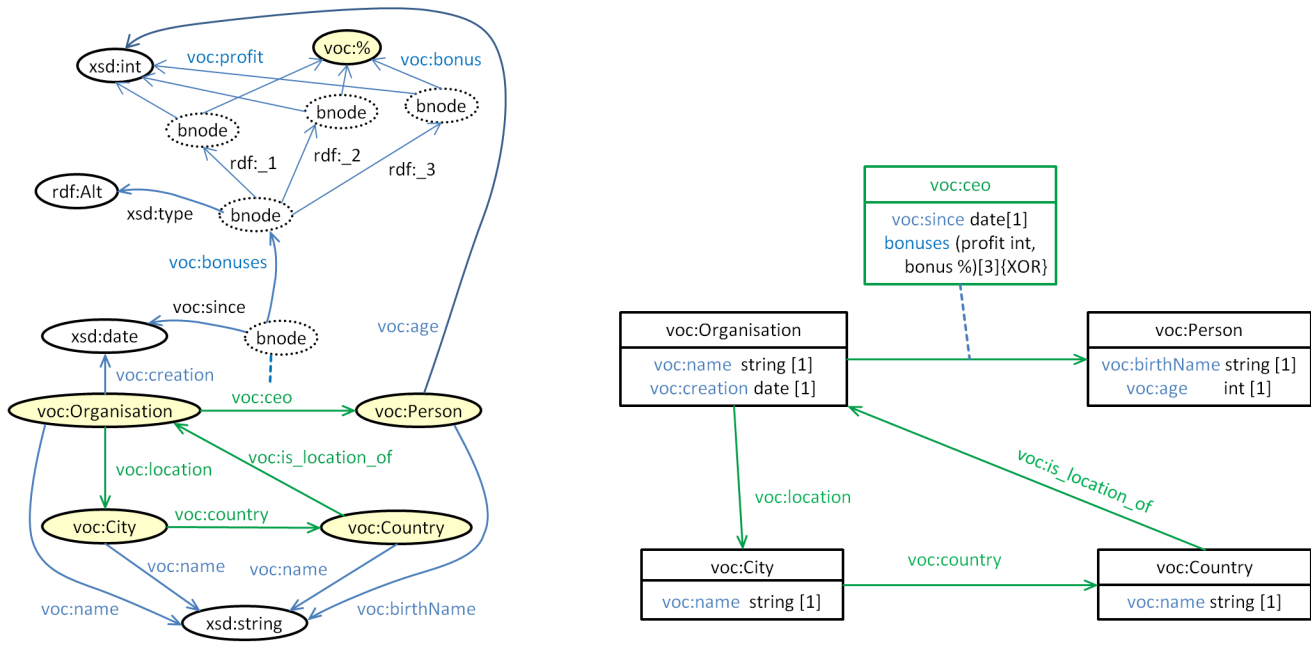


Figure 12. Example RDFS (left part) to TGS (right part) mapping with complex predicates.

VI. CONCLUSION AND FUTURE WORK

In this paper we presented a structure definition of the TGM and an UML-like notation to visualize a graph database and its graph schema. Due to the TGS with predefined and user-defined data types, the TGM improves the formal data quality compared to other graph models. We have demonstrated the superior modeling power in comparison to other graph data models and prevalent data models, namely relational, object oriented, XML model, and RDFS. The model supports built in and user defined complex data types, which allow different abstraction levels. Another possibility for abstraction is to compress a sub-graph into a hyper-node reducing the visible complexity. This capability is especially useful for large and complex data models.

Because of its semantic modeling power the TGM could act as **supermodel** for model-management and serve as a unifying data model that supports data integration from various data sources with different data models. The main challenge for an automated data integration are incompatible data sources where the TGM as **supermodel** could help to solve quality issues and specify information preserving data translations. Details, like a score for the mapping quality and how much of the information was preserved by the mapping still need to be investigated. The development of a manipulation and query language for the TGM is future work. The idea is to combine elements of other graph languages with the dot-notation known from object-oriented languages to support navigation paths.

REFERENCES

- [1] F. Laux, "The Typed Graph Model", DBKDA 2020, pp. 13-19, ISBN: 978-1-61208-790-0, [Online] URL: https://www.thinkmind.org/index.php?view=article&articleid=dbkda_2020_1_30_50016 [retrieved: 2020-11-26].
- [2] Neo4J - Homepage, [Online] URL: <https://neo4j.com> [retrieved: 2020-11-26].
- [3] ArangoDB - Homepage - Graph and Beyond, [Online] URL: <https://www.arangodb.com> [retrieved: 2020-11-26].
- [4] JanusGraph - Homepage, [Online] URL: <https://janusgraph.org> [retrieved: 2020-11-26].
- [5] Amazon Neptune, [Online] URL: <https://aws.amazon.com/de/neptune> [retrieved: 2020-08-01].
- [6] I. Robinson, J. Webber, and E. Eifrem, *Graph Databases*, 2nd ed., O'Reilly Media, 2015.
- [7] G. Daniel, G. Sunyé, and J. Cabot, "UMLtoGraphDB: Mapping Conceptual Schemas to Graph Databases", in Proceedings of Conceptual Modeling - 35th International Conference ER, Gifu, Japan, pp. 430-444, 2016. [Online] URL: <https://hal.archives-ouvertes.fr/hal-01344015/document> [retrieved: 2020-11-26].
- [8] J. Hidders, "Typing Graph-Manipulation Operations", In Proceedings of the 9th International Conference on Database Theory (ICDT), Siena, Italy, pp. 391-406, 2003.
- [9] DBLP computer science bibliography, [Online] URL: <https://dblp.uni-trier.de> [retrieved: 2020-11-26].
- [10] M. A. Rodriguez and P. Neubauer, "Construction from dots and lines", Bulletin of the American Society for Information Science and Technology, Vol. 36(No 6), pp. 35-41, ISSN:1550-8366, 2010. [Online] URL: <https://arxiv.org/pdf/1006.2361.pdf> [retrieved: 2020-11-26].
- [11] N. Spyrtos and T. Sugibuchi, "PROPER - A Graph Data Model Based on Property Graphs", ISIP 10th International Workshop, Communications in Computer and Information Science, vol.622,pp. 23-35, Springer, 2015.
- [12] J. Bu, S. Tan, C. Chen, C. Wang, H. Wu, L. Zhang, and X. He, "Music Recommendation by Unified Hypergraph: Combining Social Media Information and Music Content", In Proceedings of the 18th International Conference on Multimedia (ACM Multimedia 2010), Firenze, Italy, pp. 391-400, 2010.
- [13] A. Ghrab, O. Romero, S. Skhiri, A. Vaisman, and E. Zimányi, "GRAD: On Graph Database Modeling", Cornell University Library, arXiv:1602.00503, 2016. [Online] URL: <https://arxiv.org/ftp/arxiv/papers/1602/1602.00503.pdf> [retrieved: 2020-12-06].
- [14] B. Amann and M. Scholl, "Gram: A Graph Data Model and Query Language", In Proceedings of the ACM Conference on Hypertext (ECHT '92), pp. 201-211, Milan, Italy, 1992.
- [15] M. Gyssens, J. Paredaens, J. Van den Bussche, and D. Van Gucht,

- "A Graph-Oriented Object Database Model", IEEE Transactions on Knowledge and Data Engineering, Vol 6 No 4., pp. 572–586, 1994.
- [16] M. C. Pabón, C. Roncancio, and M. Millán, "Graph Data Transformations and Querying", In Proceedings of the 2014 International C* Conference on Computer Science & Software Engineering (C3S2E '14), Montreal Canada, Article No 20, pp. 1–6, 2014. [Online] URL: <https://doi.org/10.1145/2641483.2641521> [retrieved: 2020-08-01].
- [17] J. Pokorný, "Conceptual and Database Modelling of Graph Databases", In Proceedings of the 20th International Database Engineering & Applications Symposium (IDEAS 2016), Montreal, Canada, pp. 370–377, 2016.
- [18] M. Junghanns, A. Petermann, N. Teichmann, K. Gómez, E. Rahm, "Analyzing extended property graphs with Apache Flink", Proceedings of the 1st ACM SIGMOD Workshop on Network Data Analytics (NDA@SIGMOD 2016), San Francisco, USA, pp. 3:1–3:8 2016.
- [19] M. Levene and A. Poulouvasilis, "The hypernode model and its associated query language", In Proceedings of the 5th Jerusalem Conference on Information Technology, Jerusalem, pp. 520–530, 1990.
- [20] R. Angles, "The Property Graph Database Model", Proceedings of the 12th Alberto Mendelzon International Workshop on Foundations of Data Management, Cali, Colombia, CEUR WS Proc., 2018, [Online] URL: <http://ceur-ws.org/Vol-2100/paper26.pdf> [retrieved: 2020-12-06].
- [21] D. Brickley and R.V. Guha (eds.), RDF Schema 1.1 W3C Recommendation, published 25 February 2014, [Online] URL: <https://www.w3.org/TR/rdf-schema/> [retrieved: 2020-12-06].
- [22] W3C OWL Working Group, OWL 2 Web Ontology Language Document Overview (Second Edition) W3C Recommendation, published 11 December 2012, [Online] URL: <https://www.w3.org/TR/owl2-overview/> [retrieved: 2020-12-06].
- [23] B. Motik, P. F. Patel-Schneider, and B. Parsia (eds.), OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax (Second Edition) W3C Recommendation, published 11 December 2012, [Online] URL: <https://www.w3.org/TR/2012/REC-owl2-syntax-20121211/> [retrieved: 2021-12-06].
- [24] H. Chiba, R. Yamanaka, and S. Matsumoto, "G2GML: Graph to Graph Mapping Language for Bridging RDF and Property Graphs", In proceedings: J. Z. Pan et al. (eds) The Semantic Web 19th International Semantic Web Conference (ISWC 2020), pp. 160–175, Lecture Notes in Computer Science, vol 12507, Springer 2020, DOI: https://doi.org/10.1007/978-3-030-62466-8_11 [retrieved: 2020-12-06].
- [25] N. N., What is openCypher? [Online] URL: <https://opencypher.org/> [retrieved: 2021-12-06].
- [26] A. Schätzle, M. Przyjacieli-Zablocki, T. Berberich, and G. Lausen, "S2X: Graph-parallel querying of RDF with GraphX", in Biomedical Data Management and Graph Online Querying, Big-O(Q) 2015, DMAH 2015, Lecture Notes in Computer Science, vol 9579, pp. 155–168, Springer, Cham. DOI: https://doi.org/10.1007/978-3-319-41576-5_12.
- [27] V. Nguyen, J. Leeka, O. Bodenreider, and A. Sheth, "A formal graph model for RDF and its implementation", 2016, arXiv:1606.00480. [Online]. Available: <http://arxiv.org/abs/1606.00480> [retrieved: 2021-08-01].
- [28] R. Angles, H. Thakkar, and D. Tomaszuk, "Mapping RDF Databases to Property Graph Databases", IEEE Access, vol. 8, pp. 86091–86110, 2020, DOI: 10.1109/ACCESS.2020.2993117 [Online] URL: <https://doi.org/10.1109/ACCESS.2020.2993117> [retrieved: 2021-08-01].
- [29] A. Poulouvasilis and P. McBrien, "A General Formal Framework for Schema Transformation" Data & Knowledge Engineering 28(1), pp. 47–71, 1998, ISSN 0169-023X, DOI: [https://doi.org/10.1016/S0169-023X\(98\)00013-5](https://doi.org/10.1016/S0169-023X(98)00013-5).
- [30] P. McBrien and A. Poulouvasilis, "A Uniform Approach to Inter-model Transformations", CAISE 1999, pp. 333–348, Lecture Notes in Computer Science 1626, Springer 1999, ISBN 3-540-66157-3, DOI: https://doi.org/10.1007/3-540-48738-7_25.
- [31] Ph. A. Bernstein and S. Melnik, "Model management 2.0: manipulating richer mappings", SIGMOD Conference 2007, pp. 1–12, DOI: 10.1145/1247480.1247482, [Online] URL: <https://doi.org/10.1145/1247480.1247482> [retrieved: 2021-11-29].
- [32] M. Lenzerini, "Data Integration: A Theoretical Perspective", ACM PODS 2002, pp. 233–246, [Online] URL: <https://doi.org/10.1145/543613.543644> [retrieved: 2021-11-29].
- [33] R. Hull, "Managing Semantic Heterogeneity in Databases: A Theoretical Perspective", PODS '97: Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, pp. 51–61, ACM Press, 1997, DOI: 10.1145/263661.263668.
- [34] P. Atzeni and R. Torlone, "Schema Translation between Heterogeneous Data Models in a Lattice Framework", DS-6: Proceedings of the Sixth IFIP TC-2 Working Conference on Data Semantics: Database Applications Semantics, pp. 345–364, 1995.
- [35] P. Atzeni and R. Torlone, "Management of Multiple Models in an Extensible Database Design Tool", EDBT 1996, pp. 79–95.
- [36] P. Atzeni, P. Cappellari, and Ph. Bernstein, "Model-Independent Schema and Data Translation", EDBT 2006, pp. 368–385.
- [37] P. Atzeni, P. Cappellari, and G. Gianforme, "MIDST: model independent schema and data translation", SIGMOD Conference 2007, pp. 1134–1136.
- [38] P. Atzeni, L. Bellomarini, F. Bugiotti, and G. Gianforme, "A runtime approach to model-independent schema and data translation", EDBT 2009, pp. 275–286.
- [39] P. Atzeni, P. Cappellari, R. Torlone, Ph. Bernstein, and G. Gianforme, "Model-independent schema translation", VLDB Journal 17(6), pp. 1347–1370, 2008.
- [40] N. Roy-Hubara, L. Rokach, B. Shapira, and P. Shoval, "Modeling Graph Database Schema", IT Professional, Volume: 19, Issue: 6, pp. 34–43, IEEE publ., 2017, DOI: 10.1109/MITP.2017.4241458.
- [41] J. Sequeda, M. Arenas, and D. Miranker, "On directly mapping relational databases to RDF and OWL", WWW '12: Proceedings of the 21st international conference on World Wide Web, pp. 649–658, 2012, DOI: 10.1145/2187836.2187924.
- [42] A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens, and R. Van de Walle, "RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data", LDOW, 23rd International World Wide Web Conference, 2014, [online] URL: http://ceur-ws.org/Vol-1184/ldow2014_paper_01.pdf [retrieved: 2021-12-06].
- [43] M. Arenas, J. Pérez, J. Reutter, and C. Riveros, "Foundations of schema mapping management", ACM PODS 2010, pp. 227–238, DOI: 10.1145/1807085.1807116, [Online] URL: <https://doi.org/10.1145/1807085.1807116> [retrieved 2021-11-29].