

Topic Maps Reference Model, 13250-5

Patrick Durusau
patrick@durusau.net

Steve Newcomb
srn@coolheads.com

Robert Barta
rho@devc.at

May 16, 2008

Contents		Page
1	Scope	1
2	Normative References	1
3	Subjects	1
4	Subject Proxies and Maps	1
5	Ontological Commitments	2
6	Navigation	3
7	Merging	4
8	Constraints	5
9	Map Legends	6
10	Conformance	6
	Annex A (normative) Path Language	7
	Annex B (informative) Notation	11

Figures

Figure 1 — Proxy Structure	2
--------------------------------------	---

Tables

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of technical committees is to prepare International Standards. Draft International Standards adopted by the technical committees are circulated to the member bodies for voting. Publication as an International Standard requires approval by at least 75% of the member bodies casting a vote.

Introduction

The Topic Maps family of standards is designed to facilitate the gathering of all the information about a subject at a single location. The information about a subject includes its relationships to other subjects; such relationships may also be treated as subjects (*subject-centric*).

ISO/IEC 13250-2 (TMDM, Topic Maps Data Model) provides a foundation for syntaxes and notations, such as those defined in *ISO/IEC 13250-3* Topic Maps XML Syntax and *ISO/IEC 13250-4* Topic Maps Canonicalization. Of necessity, the TMDM makes ontological commitments in terms of how particular subjects are identified (topics, associations, occurrences), what properties are required, the tests to be used to determine whether two or more proxies represent the same subject, and other matters.

This Standard defines TMRM (Topic Maps Reference Model), which is more abstract and has fewer ontological commitments. Its purpose is to serve as a minimal, conceptual foundation for subject-centric data models such as the TMDM, and to supply ontologically neutral terminology for disclosing these. It defines what is required to enable the mapping of different subject-centric data models together to meet the overall goal of the Topic Maps standards, that each subject has a single location for all the information about it.

TMRM also provides a formal foundation for related Topic Maps standards such as *ISO/IEC 18048* Topic Maps Query Language (TMQL) and *ISO/IEC 19756* Topic Maps Constraint Language (TMCL).

Topic Maps — Reference Model — ISO 13250-5

1 Scope

The following are within the scope of this part of ISO 13250:

- a formal model for subject maps;
- minimal access functionality and information retrieval from subject maps;
- a constraint framework governing the interpretation of subject maps.

The following are outside the scope of this part of ISO 13250:

- a particular formalism to constrain subject maps.

2 Normative References

NOTE Each of the following documents has a unique identifier that is used to cite the document in the text. The unique identifier consists of the part of the reference up to the first comma.

TMDM, ISO/IEC IS 13250-2:2006, *Information Technology – Document Description and Processing Languages – Topic Maps–Data Model*, ISO, 2006.

3 Subjects

A subject is defined in the Topic Maps family of standards as something which '[...] can be anything whatsoever, regardless of whether it exists or has any other specific characteristics, about which anything whatsoever may be asserted by any means whatsoever' (*ISO/IEC 13250-2:2006* 5.3.1). According to the TMRM, there is only one representative for subjects: subject proxies (*proxies*).

4 Subject Proxies and Maps

Proxies consist of properties. These are key/value pairs which—in turn—may contain references to other proxies. This recursive relationship is defined via two postulated sets. One is the finite set of proxy *labels*, \mathcal{L} . The second set is \mathcal{V} , a finite set of values. It contains values (such as numbers, strings, etc.), and all the labels in \mathcal{L} .

A *property* is the pair $\langle k, v \rangle \in \mathcal{L} \times \mathcal{V}$. The first component of this pair is called the *key*, the other the *value* of the property. The set of all such properties is denoted as \mathcal{P} . Keys in properties are always labels, the values in properties may be labels or any other value from the value set \mathcal{V} .

EXAMPLE 1 Given the label `shoesize` and the integer 43, then $\langle \text{shoesize}, 43 \rangle$ is a property.

A *proxy* is a finite set of properties, $\{p_1, \dots, p_n\}$, with $p_i \in \mathcal{P}$ (see Fig. 1). The multiset of all keys of a proxy x is retrievable via the function $keys(x)$, i.e. keys can occur more than once in a proxy with different values. The multiset of all values is $values(x)$. Particular values may appear more than once in one proxy.




Figure 1 – Proxy Structure

EXAMPLE 2 A particular person may be represented by the following proxy: $\{\langle \text{shoesize}, 43 \rangle, \langle \text{beardcolor}, \text{white} \rangle, \langle \text{beardlength}, \text{verylong} \rangle\}$

The set of all proxies \mathcal{X} is the set of all subsets of \mathcal{P} , $\mathcal{X} = 2^{\mathcal{P}}$.

The connection between proxies and their labels is modeled with a partial, injective function $\tilde{\cdot}: \mathcal{X} \mapsto \mathcal{V}$. It returns the label for a given proxy x whereby two different proxies never share the same label. The function is extended to values in that $\tilde{v} = v$.

NOTE 1 A proxy is defined by the totality of its properties. Individual properties can provide a basis for mapping multiple proxies of the same subject to each other.

NOTE 2 One proxy may contain several properties which all share the same key but have different values; or share the same value, but have differing keys.

A *subject map* (*map*) is a finite set of *proxies*. The set of all such maps is denoted as \mathcal{M} . As maps are simply sets of proxies, *generic merging of maps* is achieved via set union, $m \cup m'$.

NOTE 3 The model of subject maps described herein assumes no particular implementation technology or strategy.

5 Ontological Commitments

This Standard deliberately leaves undefined the methods whereby subject proxies are derived or created. No specific mechanism of subject identification is inherent in or mandated by this Standard, nor does it predefine any subject proxies.

NOTE 1 Any subject proxy design choices would be specific to a particular application domain and would exclude equally valid alternatives that might be appropriate or necessary in the contexts of various requirements.

Two types of relationships, *ako* (subclass of) and *isa* (instance of), are defined. These predicates are always interpreted *relative* to a given map m :

- a) Two proxies c, c' can be in a *subclass-superclass* relationship, $\text{ako}_m \subseteq m \times m$. In such a case, the same relationship can be stated either *c is a subclass of c'* or *c' is a superclass of c*.

ako_m is reflexive and transitive. *Reflexive* means that any proxy is a subclass of itself, regardless whether the proxy is used as a class in the map or not: $x \text{ako}_m x$ for all $x \in m$.

Transitive means that if a proxy c is a subclass of another, c' , and that subclasses c'' , then c is also a subclass of c'' , i.e. if $c \text{ako}_m c'$ and $c' \text{ako}_m c''$ then also $c \text{ako}_m c''$ must be true.

NOTE 2 Circular subclass relationships may exist in a map.

- b) Two proxies a, c can be in an *isa* relationship, $\text{isa}_m \subseteq m \times m$. In such a case, the same relationship can be stated either *a is an instance of c* or *c is the type of a*.

The *isa* relationship is non-reflexive, i.e. $x \text{isa}_m x$ for no $x \in m$, so that no proxy can be an instance of itself. Additionally, whenever a proxy a is an instance of another c , then a is an instance of any superclass of c : if $x \text{isa}_m c$ and $c \text{ako}_m c'$, then $x \text{isa}_m c'$ is true.

NOTE 3 Other definitions of the above relationships are possible with different properties than defined by the TMRM. Such definitions would appear in legends and would be distinguished from those defined in the TMRM.

NOTE 4 This Standard does not mandate any particular way of representing such relationships inside a map. One option is to model such a relationship simply with a property using a certain key (say **type**). An alternative way is to provide a proxy for each such relationship. Such relationship proxies could, for example, have properties whose keys are **instance** and **class**, or respectively **subclass** and **superclass**.

6 Navigation

Given a map m and particular proxies $x, y \in m$ in it, the following *primitive navigation operators* are defined:

- a) A postfix operator \downarrow to return the multiset of *all local keys* of a given proxy:

$$x \downarrow = \text{keys}(x) \quad (1)$$

- b) A postfix operator \uparrow_m to retrieve the multiset of *remote keys* of a proxy inside m . These are those where the given proxy (more precisely its label) *is* the value in another proxy:

$$x \uparrow_m = [k \mid \exists y \in m : \langle k, \tilde{x} \rangle \in y] \quad (2)$$

This is easily generalized to an operator for all values:

$$v \uparrow_m = [k \mid \exists y \in m : \langle k, \tilde{v} \rangle \in y] \quad (3)$$

- c) A postfix operator $\rightarrow k$ to retrieve the multiset of *local values for a particular key* $k \in \mathcal{L}$:

$$x \rightarrow k = [v \mid \exists \langle k, v \rangle \in x] \quad (4)$$

Using the predicate ako_m the operator can be generalized to honor subclasses of the key k :

$$x \rightarrow_m k^* = [v \mid \exists \langle k', v \rangle \in x : k' \text{ ako}_m k] \quad (5)$$

- d) A postfix operator $\leftarrow_m k$ which navigates to all proxies in the given map which use a given value v together with a certain key $k \in \mathcal{L}$:

$$v \leftarrow_m k = [x \in m \mid \exists \langle k, \tilde{v} \rangle \in x] \quad (6)$$

Using the predicate ako_m the operator can be generalized to honor subclasses of the key k :

$$v \leftarrow_m k^* = [x \in m \mid \exists \langle k', \tilde{v} \rangle \in x : k' \text{ ako}_m k] \quad (7)$$

It is straightforward to generalize all these navigation operators on individual proxies and values to multisets of them. As a consequence the result of one postfix can be used as startpoint for another postfix, enabling the building of *postfix chains*. This primitive path language is denoted as \mathcal{P}_M .

NOTE 1 \mathcal{P}_M only serves as a minimal baseline for functionality to be provided by conforming implementations. It can be also used as the basis for a formal semantics for higher-level query and constraint languages. Annex A describes one.

7 Merging

Generic merging of maps only combines two (or more) proxy sets. Application-specific merging includes a second aspect. A mechanism has to be found to state whether—in a given map—two proxies are *regarded to be about the same subject*. Then all such equivalent proxies have to be actually merged.

NOTE 1 How *subject sameness* is determined and how the actual *proxy merging* is effectively done is not constrained by this Standard. Such a process may be defined as having inputs that consist only of the proxies to be merged. Alternatively, the inputs may also include other information that may appear either inside the map or elsewhere in the merging process's environment.

NOTE 2 Given the appropriate expressivity of the used constraint language, such equivalence and the consequent merging process can be described with a constraint.

Merging is modeled with a partial function $\bowtie: \mathcal{X} \times \mathcal{X} \times \mathcal{E} \mapsto \mathcal{X}$. It takes two proxies and an—otherwise unconstrained—environment \mathcal{E} as parameters and produces a new proxy. In the special case that the environment has no influence in this process, \bowtie is an infix operator $\mathcal{X} \times \mathcal{X} \mapsto \mathcal{X}$ between proxies.

NOTE 3 The reason for including \mathcal{E} as a term in the definition of merging is to account for the fact that merging criteria may be defined as being dependent on conditions external to the maps. When environmental differences affect the results of merging, a single interchangeable subject map may be realized as different subject maps in different environments.

The fact that \bowtie is *partial* means that it may be applicable only to some pairs of proxies but not to others. Those where the result is defined are supposed to be merged.

The operator \bowtie must be commutative and associative:

$$x \bowtie x' = x' \bowtie x \quad (8)$$

$$(x \bowtie x') \bowtie x'' = x \bowtie (x' \bowtie x'') \quad (9)$$

Additionally, \bowtie must be idempotent, as proxies merged with themselves should not result in a different one:

$$x \bowtie x = x \quad (10)$$

The operator \bowtie factors a given proxy set into equivalence classes: two proxies x, y from a given map $m \in \mathcal{M}$ are then in the same class if $x \bowtie y$ is defined. The set of equivalence classes is written m / \bowtie . Every such class can be merged into a single proxy by combining all its members by applying \bowtie . Given a set of proxies $c = \{x_1, \dots, x_n\}$, the merging of members of an equivalence class is defined as:

$$\bowtie c = x_1 \bowtie x_2 \bowtie \dots \bowtie x_n \quad (11)$$

Implied with this is a *change set*, i.e. a prescription to replace proxies x_1, \dots, x_n in m with $\bowtie c$. We denote the change set of c as Δ . When it is applied to a map m , then every occurrence of a label of x_i in some property is replaced with the label of $\bowtie c$.

Given a map m and a particular \bowtie , the *merged view of the map* $m|_{\bowtie}$ is defined as:

$$m|_{\bowtie} = \left(\bigcup_{i=1}^n \bowtie c_i \right) \Delta_{c_1} \dots \Delta_{c_n} \quad (12)$$

with all $c_i \in m / \bowtie$ and the Δ_{c_i} the corresponding change sets. Obviously the individual change sets have to be applied to all newly merged proxies.

One such merging step may result in proxies being created which again may be mergeable. The process can be repeated until a *fully merged map* is computed, so that $m|_{\bowtie} = m$. Fully merged maps are symbolized as $m|_{\bowtie}^*$. Since the sizes of newly merged maps are weakly monotonously decreasing, the process to produce a fully merged map takes always a finite number of steps.

8 Constraints

Subject maps are structures which are used to encode assertional knowledge. To interpret a map, be it for modelling, retrieval or modification, some background information about the map may be necessary. That information is provided in form of *constraints*, so that a given map m either satisfies a constraint, or not. A constraint language is a formalism which allows the expression of such constraints.

NOTE 1 Constraints may conditionally or unconditionally require the existence of certain proxies in maps, the existence of properties in proxies, and/or values in properties. Constraints may also prohibit the existence of any of the foregoing.

EXAMPLE 1 A constraint language may allow the expression of constraints such as *all instances of the concept person must have at least one shoesize property or any shoesize property must have an integer value between 10 and 50*.

NOTE 2 The precise ways in which constraints may be expressed are not constrained by this Standard. Different constraint languages will differ in expressivity and, consequently, in computational complexity.

This Standard imposes two requirements on any constraint language \mathcal{C} :

- a) \mathcal{C} must define the *application of a constraint to a map* in the form of a binary operator $\otimes : \mathcal{M} \times \mathcal{C} \mapsto \mathcal{M}$. The operator \otimes is used to define the *satisfaction relation* $\models_{\subseteq} \mathcal{M} \times \mathcal{C}$ between a map and a constraint.

13250-5 FCD

A particular map m is said to *satisfy a constraint*, $m \models c$, if the application of the constraint results in a non-empty map:

$$m \models c \iff m \otimes c \neq \emptyset \quad (13)$$

- b) \mathcal{C} must define a *merging operator* $\oplus : \mathcal{M} \times \mathcal{M} \mapsto \mathcal{M}$ as binary operator between two maps. \oplus must be commutative, associative and idempotent.

NOTE 3 The provision of \oplus and \otimes may be done in any manner that is sufficiently expressive. \oplus can but does not need to be defined in terms of \bowtie . Annex A demonstrates one way of defining \otimes .

9 Map Legends

A *map legend (legend)* $G = \{c_1, \dots, c_n\}$ is a finite set of constraints, all from a given constraint language \mathcal{C} . A legend G is said to *govern* a map m if $m \models c_i$ for all $c_i \in G$.

NOTE 1 The legend of a subject map plays a role similar to the legend on more familiar city or road maps. The legend associated to a subject map is one key to its interpretation. Just as the legends of geographic maps describe and define the symbols that appear in them, their scaling rules, etc., subject map legends explain the symbols that appear in them (such as property keys) and other interpretation rules.

NOTE 2 There is no explicit connection between maps and legends, other than whether the map is governed by a legend or not. Any subject map can be simultaneously viewed with multiple legends written for different purposes or users. Such views may be quite different from each other. Maps themselves exist without any legend describing the rules how they may be used.

10 Conformance

An implementation conforms to this Standard if:

- a) Its information structures are homomorphous to the subject map structure in Clause 4;
- b) It implements merging by providing an operator \oplus according to Clause 8;
- c) It implements access methods equivalent to those in Clause 6;
- d) Those methods must honor the predicates from Clause 5;
- e) It implements a constraint language by providing an operator \otimes according to Clause 8;
- f) It implements legends according to Clause 9.

Annex A (normative) Path Language

The primitive path language $\mathcal{P}_{\mathcal{M}}$ (clause 6) can be used to define a more expressive language $\mathcal{T}_{\mathcal{M}}$ which provides navigation, filtering, sorting and general function application.

In the following, first tuples and then tuple sequences are defined. Sequences are the results of applying a $\mathcal{T}_{\mathcal{M}}$ path expression to a map (or any multiset of proxies and values). With the help of that, the path expression language $\mathcal{T}_{\mathcal{M}}$ is characterized.

A.1 Tuples

Path expressions describe patterns to be identified in a map. In order to provide a model for both queries and constraints, the results of path expressions are modeled as tuples of values and organized into tuple sequences.

A single tuple with values from a value set \mathcal{V} is denoted as $\langle v_1, v_2, \dots, v_n \rangle$. Tuples can be concatenated simply by collating their values: $\langle u_1, \dots, u_m \rangle \cdot \langle v_1, \dots, v_n \rangle = \langle u_1, \dots, u_m, v_1, \dots, v_n \rangle$. This enables the representation of tuples as the products of singleton tuples:

$$t = \prod_{i=1}^n \langle v_i \rangle = \langle v_1 \rangle \cdot \langle v_2 \rangle \cdot \dots \cdot \langle v_n \rangle \quad (14)$$

The concatenation symbol \cdot will be omitted from now on. For a tuple slices we use $t(j..k) = \prod_{i=j}^{k-1} \langle v_i \rangle$ and $t(i) = t(i..i+1)$.

Tuples and proxies are closely related. All the values can be taken out of a proxy and arranged into a value tuple. If order is important, an order can be postulated on the keys and the values are sorted according to it. Conversely, a value tuple can be converted into a proxy, provided that the tuple of keys is available (*zipping*):

$$\langle v_1, \dots, v_n \rangle \text{z} \langle k_1, \dots, k_n \rangle = \{ \langle k_1, v_1 \rangle, \dots, \langle k_n, v_n \rangle \} \quad (15)$$

Tuples are identical if all their values in corresponding positions are identical. Tuples can also be compared so that $t \leq t'$ is derived from some ordering over the tuple values $v \leq v'$. To allow to control ascending and descending tuple ordering an *order control tuple* $o = \prod_1^k \langle d_i \rangle$ with the direction $d \in \{\uparrow, \downarrow, -\}$ is introduced. Given that and tuples t and t' , the ordering on tuples \leq_o is defined via $t \leq_o t' \Leftrightarrow t(1..m) \leq_o t'(1..m)$ with m the length of the shorter tuple. For tuple slices we define \leq_o to be:

$$t(i..j) \leq_o t'(i..j) \Leftrightarrow \begin{cases} t(i) \leq t'(i) & \text{if } d_i = \uparrow \\ t'(i) \leq t(i) & \text{if } d_i = \downarrow \\ t(i+1..j) \leq_o t'(i+1..j) & \text{if } d_i = - \text{ or } t(i) = t'(i) \end{cases}$$

A.2 Tuple Sequences

When tuples are organized into sequences, a single sequence is written:

$$s = \sum_{i=1}^m t_i = [t_1, \dots, t_m] \quad (16)$$

13250-5 FCD

For a given set of values $\{v_1, \dots, v_n\}$ the tuple sequence $\sum_{i=1}^n \langle v_i \rangle$ can be built. This conversion is denoted as $[\{v_1, \dots, v_n\}]$. Under this interpretation a map $m = \{x_1, \dots, x_n\}$ can be represented as the tuple sequence $[m]$. Conversely, a tuple sequence can be interpreted as a map when the tuples it contains are singletons with proxies.

Given an order tuple o , a sequence s can be ordered such that $t_i \leq t_j$ iff $i \leq j$. This is written as \vec{s}^o .

Sequences behave like multisets, i.e. operations such as union \cup , intersection \cap and subtraction \setminus are those of multisets. Any ordering of tuples within the operand sequences is lost in the result. Ordering in any of the operand sequences is maintained when tuple sequences are concatenated:

$$\sum_{i=1}^m s_i + \sum_{j=1}^n t_j = [s_1, \dots, s_m, t_1, \dots, t_n] \quad (17)$$

This is achieved by interleaving the tuples of the second operand with those of the first. The indices may be omitted if the range is obvious.

Tuple sequences can also be combined by *multiplying* (joining) them. The product of two tuple sequences is defined recursively:

$$(s) \left(\sum_{j=1}^m \langle v_{1j}, v_{2j}, \dots, v_{lj} \rangle \right) = \left(s \sum_{j=1}^m \langle v_{1j} \rangle \right) \sum_{j=1}^m \langle v_{2j}, \dots, v_{lj} \rangle \quad (18)$$

$$\sum_{i=1}^n t_i \sum_{j=1}^m \langle v_j \rangle = \sum_{i,j=1}^{nm} (t_i \langle v_j \rangle) \quad (19)$$

NOTE 1 Every tuple of the left hand operand sequence is concatenated with every other one of the right-hand operand. The first value of each tuple of the second operand is removed and combined with every tuple of the first operand. This is repeated until the second operand does not have tuples with any values.

The zipping function (section A.1) can be generalized to tuple sequences by repeating the process for every tuple of the sequence. Consequently, every tuple sequence can be interpreted as a sequence of proxies once a set of keys has been chosen. This sequence of proxies can then be interpreted as a subject map.

A.3 Path Expressions

A particular path expression can be interpreted as an *expression of interest*, i.e. as a pattern to be identified in a map. Given a tuple sequence s , a path expression p can be applied to it, with the expectation of another tuple sequence in return. This application can be symbolized as $s \otimes_m p$. This operation is to be understood in the context of a map m . If that is implicit, the index can be dropped.

The set of path expressions $\mathcal{T}_{\mathcal{M}}$ is characterized as the smallest set satisfying the following conditions:

- a) Every value from \mathcal{V} (and consequently every proxy label) is in $\mathcal{T}_{\mathcal{M}}$. If such a value is applied to a sequence, then the sequence itself is discarded. Instead a new sequence with a singleton tuple is created in which the value is used.

- b) The *projection* postfix π_i is in $\mathcal{T}_{\mathcal{M}}$ for any positive integer i . It can be used to extract a certain column from a given tuple sequence:

$$\sum_{i=1}^n \langle v_{1i}, v_{2i}, \dots, v_{li} \rangle \otimes \pi_j = \sum_{i=1}^n \langle v_{ji} \rangle \quad (20)$$

Projection here plays a similar role like in query languages like SQL, except that an index is used for selection instead of names.

To organize values freely into a new tuple sequence, the *tuple projection* $\langle p_1, \dots, p_n \rangle$ is used with p_i being path expressions. For a single tuple it evaluates all the path expressions and builds the product of all partial result sequences:

$$t \otimes \langle p_1, \dots, p_n \rangle = \prod_{i=1}^n t \otimes p_i \quad (21)$$

When applied to a tuple sequence, all applications to its tuples are concatenated. As a special case, the *empty projection* $\mathbf{0} = \langle \rangle$ always returns an empty tuple sequence and the *identity projection* $\mathbf{1} = \prod \langle \pi_i \rangle$ always returns the incoming tuple sequence.

- c) The navigation operators \uparrow , \downarrow , $\leftarrow k$ and $\rightarrow k$ are in $\mathcal{T}_{\mathcal{M}}$. When applied to a tuple sequence these operators are applied to every tuple:

$$\left(\sum_{i=1}^n \langle v_{1i}, v_{2i}, \dots, v_{li} \rangle \right) \otimes \leftarrow k = \sum_{i=1}^n \prod_{j=1}^l \langle v_{ji \leftarrow m k^*} \rangle \quad (22)$$

$$\left(\sum_{i=1}^n \langle v_{1i}, v_{2i}, \dots, v_{li} \rangle \right) \otimes \rightarrow k = \sum_{i=1}^n \prod_{j=1}^l \langle v_{ji \rightarrow m k^*} \rangle \quad (23)$$

These operators simply iterate over each tuple and compute an intermediate result for each tuple. This intermediate result is achieved by applying the navigation to each value in the current tuple. As one such application results in a multiset of values, that is converted into a singleton tuple sequence. All these singleton tuple sequences are multiplied and all these intermediate results are then concatenated into the overall result.

The operators can be naturally extended to sets of keys:

$$s \otimes \leftarrow \{k_1, k_2, \dots, k_n\} = \sum_{i=1}^n (s \otimes \leftarrow k_i) \quad (24)$$

$$s \otimes \rightarrow \{k_1, k_2, \dots, k_n\} = \sum_{i=1}^n (s \otimes \rightarrow k_i) \quad (25)$$

An analogous approach is used for finding keys:

$$\left(\sum_{i=1}^n \langle v_{1i}, v_{2i}, \dots, v_{li} \rangle \right) \otimes \downarrow = \sum_{i=1}^n \prod_{j=1}^l \langle v_{ji \downarrow} \rangle \quad (26)$$

$$\left(\sum_{i=1}^n \langle v_{1i}, v_{2i}, \dots, v_{li} \rangle \right) \otimes \uparrow = \sum_{i=1}^n \prod_{j=1}^l \langle v_{ji \uparrow m} \rangle \quad (27)$$

- d) Given path expressions p_1, \dots, p_n and a function $f : \mathcal{V}^n \mapsto \mathcal{V}$ then also $f(p_1, \dots, p_n)$ is in \mathcal{T}_M . When an n -ary function $f : \mathcal{V}^n \mapsto \mathcal{V}$ is applied to a tuple sequence, it is interpreted as one which takes a value tuple of length n and renders one value from \mathcal{V} . To apply it to a tuple sequence, it is applied to every individual tuple and the results are organized back into a sequence:

$$f\left(\sum t_i\right) = \sum \langle f(t_i) \rangle \quad (28)$$

- e) The *conditional* $p?q:r$ is in \mathcal{T}_M for path expressions p, q and r . When it is applied to a tuple sequence, every tuple is tested whether it produces a result when p is applied to it. If that is the case, the *then* branch is used, i.e. the tuple will be subjected to q and these results will be added to the overall result. Otherwise, the tuple will get r applied:

$$\left(\sum_{i=1}^n t_i\right) \otimes (p?q:r) = \sum_{i=1}^n ([t_i \otimes q \mid t_i \otimes p \neq \emptyset] \cup [t_i \otimes r \mid t_i \otimes p = \emptyset]) \quad (29)$$

- f) For two path expressions p and q the *alternation* $p + q$, the *reduction* $p - q$, and the *comparison* $p = q$ are in \mathcal{T}_M :

$$s \otimes (p + q) = (s \otimes p) \cup (s \otimes q) \quad (30)$$

$$s \otimes (p - q) = (s \otimes p) \setminus (s \otimes q) \quad (31)$$

$$s \otimes (p = q) = (s \otimes p) \cap (s \otimes q) \quad (32)$$

- g) The *slicer* $[i..j]$, the *sorter* sort_o with an order control tuple o and the *duplicate remover* uniq are in \mathcal{T}_M :

$$\sum_{k=1}^n t_k \otimes [i..j] = \sum_{k=i}^{j-1} t_k \quad (33)$$

$$s \otimes \text{sort}_o = \vec{s}^{\vec{o}} \quad (34)$$

$$s \otimes \text{uniq} = \sum_k t_k \quad \text{with } t_k \in s \text{ and } t_i \neq t_j \text{ for } i \neq j \quad (35)$$

- h) For two path expressions p and q the *concatenation* $p \circ q$ is in \mathcal{T}_M :

$$s \otimes (p \circ q) = (s \otimes p) \otimes q \quad (36)$$

If—from the context—it is clear that two path expressions are to be concatenated, the infix is omitted.

Annex B (informative) Notation

Symbol	Used For	ASCII Equivalent
k^*	All subclasses (key or value)	k^*
\bowtie	Merging function	$ >< $
$m _{\bowtie}$	Set of equivalence classes	$ >< $
$m _{\bowtie}^*$	Fully merged subject map	$m >< ^*$
\downarrow	All keys in a proxy	\backslash
\uparrow	Proxy as a key	$/$
\rightarrow	All values for a key	\rightarrow
\leftarrow	Proxies with given value for a key	\leftarrow
\mapsto	Results in	$ \rightarrow$
\models	Satisfaction relationship	$ =$
\otimes	Constraint operator	(x)
\oplus	Merging operator	$(+)$