



- (51) **International Patent Classification:**
G06F 17/30 (2006.01) *H04W 4/00* (2009.01)
- (21) **International Application Number:**
PCT/EP2015/070245
- (22) **International Filing Date:**
4 September 2015 (04.09.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
15177551.7 20 July 2015 (20.07.2015) EP
- (71) **Applicant:** NEC EUROPE LTD. [DE/DE]; Kurfürsten-Anlage 36, 69115 Heidelberg (DE).
- (72) **Inventor:** BAUER, Martin; Seestraße 55/1, 69214 Eppelheim (DE).
- (74) **Agent:** ULLRICH & NAUMANN; Schneidmühlstraße 21, 69115 Heidelberg (DE).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

- (54) **Title:** METHOD AND SYSTEM FOR QUERYING SEMANTIC INFORMATION STORED ACROSS SEVERAL SEMANTICALLY ENHANCED RESOURCES OF A RESOURCE STRUCTURE




Fig. 4

- (57) **Abstract:** A method for querying semantic information stored across several semantically enhanced resources of a resource structure, particularly for a machine-to-machine application, is disclosed, wherein a semantic description is associated with each of the semantically enhanced resources. This method comprises the steps of: accessing a first semantically enhanced resource of the resource structure by a querying entity, wherein a first semantic description is associated with the first semantically enhanced resource, loading the first semantic description into a memory, analyzing the first semantic description, and loading into the memory a second semantic description associated with a second semantically enhanced resource of the resource structure, if a semantic element of the first semantic description contains a resource link, said resource link pointing to the second semantically enhanced resource. Additionally an according system for querying semantic information, a querying entity, a method for generating a resource structure and a database for enabling queries regarding semantic information stored across several semantically enhances resources are disclosed.



METHOD AND SYSTEM FOR QUERYING SEMANTIC INFORMATION STORED ACROSS SEVERAL SEMANTICALLY ENHANCED RESOURCES OF A RESOURCE STRUCTURE

5 The present invention claims priority of European patent application 15 177 551.9. The content of this European patent application is incorporated herein by reference.

10 The present invention relates to a method for querying semantic information stored across several semantically enhanced resources of a resource structure, particularly for a machine-to-machine application, wherein a semantic description is associated with each of the semantically enhanced resources.

15 The present invention further relates to a querying entity for querying semantic information stored across several semantically enhanced resources of a resource structure.

20 The present invention further relates to a method for generating a resources structure.

25 The present invention additionally relates to a database for enabling queries regarding semantic information stored across several semantically enhanced resources of a resource structure, particularly at a machine-to-machine application.

30 The present invention also relates to a non-transitory computer readable medium storing a program causing a computer to execute a process.

Although applicable to any kind of machine-to-machine application, at least one of the embodiments of this invention will be described with regard to oneM2M, particularly to Release 2 of oneM2M which is currently elaborated.

Machine-to-machine applications are gaining more and more importance nowadays leading to standards like ETSI M2M, oneM2M or the like. Conventional

machine-to-machine systems provide well defined machine-to-machine resources to represent things of the real world in electronic form. A resource structure is formed by several to many resources and interconnections between them and describes the structure of a system, e.g. a sensor system or an air conditioning system.

oneM2M is a REST (Representational State Transfer)-based architecture. oneM2M enables applications to create resource structures for exchanging information using a number of pre-defined types, e.g. <AE> for representing oneM2M Application Entities, <Container> for storing content and sub-containers, <contentInstance> for storing content, e.g. measurements. Access rights are handled on a resource basis, i.e. fine-grained policies are possible

In order to provide consistent data translation and data interoperability to heterogeneous machine-to-machine applications, some standards provide the possibility to add semantic information to resources. This document refers to such machine-to-machine resources with associated semantic description as semantically enhanced resources.

At oneM2M, semantic descriptions will be supported in the upcoming Release 2. In the current draft version, support for semantic descriptions is provided for <AE> (Application Entity), <Container> and <contentInstance> resources as well as <node> and <group>. The semantic description is stored in <SemanticDescription> child resources. Besides that, semantically enhanced functionality is provided making use of this semantic information, in particular semantic filtering or other semantic querying capabilities.

One drawback of the structure of oneM2M is that semantic queries, particularly semantic filtering, has to be performed on resource basis, i.e. a semantic query can only cover semantic information that is stored within one resource. As a result, a resource has to contain each piece of semantic information that is needed for a potential query/filter operation, even if the semantic information is not directly linked to the respective resource but to another resource of the resource structure. This will result in duplicated semantic information at different resources and

increases the need of storage space considerably. Additionally, consistency of the duplicated semantic information has to be ensured, i.e. each time a piece of semantic information is changed, each semantically enhanced resource has to be checked, whether an update operation is needed or not. This results in high load to the system during an update process. Correct functioning of the update process is particularly critical in multi-user systems, where the changed piece of information may refer to the access rights of a querying entity. If the update process does not reliably change the data a user may access a piece of semantic information although the access right to this piece of semantic information is already revoked.

As an alternative approach to duplication of semantic information, enhanced query/ filter operations can be provided that allows for queries beyond resource borders. Then the enhanced querying entity/filter can be executed on several resources at the same time. In this case, semantic information can be kept to those resources to which the semantic information belongs. However, such an approach will have a deep impact to oneM2M, as these enhanced queries completely contradict the querying structure of oneM2M.

Although these problems are described in detail with regard to oneM2M, it is also present at other machine-to-machine systems or other resource-based systems that use semantic information.

It is therefore an object of the present invention to provide a method and a system for querying semantic information as well as a querying entity and a database that may be used in connection with this method and/or system, where the semantic information is stored across several semantically enhanced resources of a resource structure, where the required storage space should be kept low and where the impact on a machine-to-machine system should be minimized.

In an embodiment, the invention provides a method for querying information stored across several semantically enhanced resources of a resource structure, particularly for a machine-to-machine application, wherein a semantic description

is associated with each of the semantically enhanced resources, said method comprising the steps of:

accessing a first semantically enhanced resource of the resource structure by a querying entity, wherein a first semantic description is associated with the first semantically enhanced resource,

loading the first semantic description into a memory,

analyzing the first semantic description, and

loading into the memory a second semantic description associated with a second semantically enhanced resource of the resource structure, if a semantic element of the first semantic description contains a resource link, said resource link pointing to the second semantically enhanced resource.

In a further embodiment, the invention provides a querying entity for querying semantic information stored across several semantically enhanced resources of a resource structure, comprising:

a memory for storing semantic descriptions,

a control unit operable to access a first semantically enhanced resource and to load into the memory a first semantic description associated with the first semantically enhanced resource, and

an analyzing unit operable to analyze the first semantic descriptions stored in the memory,

wherein the control unit is additionally operable to access and load into the memory a second semantically enhanced resource, wherein a pointer to said second semantically enhanced resource – resource link – is contained in an element of the first semantic description.

In another embodiment, the invention provides a method for generating a resource structure comprising the steps of:

analyzing a semantic structure representing a behavior of a system,

extracting semantically enhanced resources contained in the semantic structure,

generating, for each semantically enhanced resource, a semantic description as a subset of the semantic structure, wherein each semantic

description contains nodes of the semantic structures associated with the respective semantically enhanced resource,

generating a resource link between a first semantically enhanced resource and a second semantically enhanced resource, if a node of the semantic structure is contained in the first semantically enhanced resource and in the second semantically enhanced resource, said resource link preferably being implemented as a URL (Uniform Resource Locator) or a URI (Uniform Resource Identifier).

In yet another embodiment, the invention provides a database for enabling queries regarding semantic information stored across several semantically enhanced resources of a resource structure, particularly at machine-to-machine application, wherein the database comprises hierarchically structured resources, wherein the resources comprise at least two semantically enhanced resources, wherein a semantic description is associated to each of said semantically enhanced resources, respectively, and wherein at least one of the semantically enhanced resources comprises a resource link, said resource link pointing to another semantically enhanced resource.

In yet another embodiment, the invention provides a system for querying semantic information stored across several semantically enhanced resources of a resource structure, comprising:

storage space storing a database according to one of the embodiments of this application,

a querying entity and

a user device operable to generate a query for the querying entity.

In yet another embodiment, the invention provides a non-transitory computer readable medium that stores a program causing a computer to execute a process comprising:

accessing a first semantically enhanced resource of the resource structure, wherein a first semantic description is associated with the first semantically enhanced resource,

loading the first semantic description,

analyzing the first semantic description, and

loading a second semantic description associated with a second semantically enhanced resource of the resource structure, if a semantic element of the first semantic description contains a resource link, said resource link pointing to the second semantically enhanced resource.

5

According to the invention, it has first been recognized that semantic information may be distributed across several resources. To this end, each semantically enhanced resource is associated to a semantic description which contains one or several pieces of semantic information. The semantic information contained in a semantic description is substantially related to the associated resource. Association between a semantically enhanced resource and a semantic description may be established in different ways. Each kind of association that forms a kind of link between a semantically enhanced resource and a semantic description is suitable. The semantic description may be part of the semantically enhanced resource. It may be contained in a variable or a field of the resource. The semantic description may be incorporated into a sub-container or a child-resource that is linked to a superior container or a parent resource. The link may also be provided by the concept of features, where the semantic description is a feature that is linked to the resource.

20

According to the invention, it has further been realized that it is not necessary to duplicate substantial parts of the semantic information or to rework querying capabilities completely. Instead a resource link is provided at each semantically enhanced resource which refers to semantic information which is not fully defined within its semantic description. A resource link is a kind of pointer that points from one resource description to another resource description. By following a resource link, the semantic information that is needed for a query can be retrieved without providing each needed piece of information within one resource description. By providing a resource link, a query may still be resource based, i.e. the query refers just to one resource. However, by following a resource link, semantic information may be distributed across several semantically enhanced resources. Therefore, not the complete querying concept has to be reworked. It is sufficient that the querying entities may follow a resource link and retrieve the required information from the linked semantic description.

25

30

The resource link may be provided by different elements. It might be a variable of the resource or of the semantic description. It might be a node of a tree that forms the semantic description. With regard to oneM2M, the resource link might be implemented as an oneM2M annotation.

5

According to one embodiment of the invention, a method for querying semantic information stored across several semantically enhanced resources may comprise the step of accessing a first semantically enhanced resource of a resource structure by a querying entity, wherein a first semantic description is associated with the first semantically enhanced description. The first semantically enhanced resource might be each semantically enhanced resource of the resource structure. The adjective "first" is only used to distinguish a semantically enhanced resource that contains a resource link (or the associated semantic description) from a semantically enhanced resource/semantic description to which the resource link points to – second semantically enhanced resource/second semantic description. Even each second semantically enhanced resource might form a first semantically enhanced resource for a next step of the method.

10

15

20

25

30

The querying entity may be implemented by different means. It may be implemented by software and/or hardware. Preferably it is implemented by a program that runs on a processor, like a microprocessor, a DSP (Digital Signal Processor), a suitably configured FPGA (Field Programmable Gate Array), a vCPU (virtual Central Processing Unit) of a virtualization system or other processing units. The processor may be part of a computer, a server, a smartphone, a PDA (Personal Digital Assistant) or the like. The querying entity may be run on the same unit that stores the resource structure and/or the semantic descriptions. The querying entity may also run on a unit that does not store the resource structure and/or the semantic descriptions. In such a case, the query may be sent to the unit that stores the required data via a network, like the internet or a LAN (Local Area Network). The unit that stores the resource structure and/or the semantic descriptions may be implemented by a server or the like. However, the resource structure may also be stored on a distributed system with different nodes that are interconnected by a network, like the internet or a LAN. Each node of such a distributed system may store one or several resources of the

resource structure and/or one or several semantic descriptions. In such a distributed system, the querying entity may access the first semantic description directly (if running at the node which stores the first semantic description) or via the network (if running on a different node than the node which stores the first semantic description). Such a distributed system may be formed by several servers or by a cloud system.

After accessing the first semantically enhanced resource, the first semantic description may be loaded into a memory. The memory may be implemented by various memory technologies. A rewriteable non-transitory memory may be used as well as a volatile memory. The non-transitory memory can be, but not limited to, a hard disk drive, a flash memory, a solid state disk, a USB memory stick or the like. The volatile memory may be a RAM (Random Access Memory). The memory may be part of or may be associated with the querying entity.

After loading the first semantic description, the first semantic description may be analyzed. The step of analyzing may search for a resource link that is contained in the first semantic description, wherein the resource link points to a second semantic description/a second semantically enhanced resource. If a resource link is found, the second semantic description may be loaded into the memory.

In one embodiment, the step of analyzing the first semantic description comprises the step of executing a query stored at the querying entity, wherein the step of loading a second semantic description is preformed preferably only, if the second semantic description is needed for the query. During the step of loading the second semantic description, the step of executing the query may be halted and may be resumed after the step of loading the second semantic description is terminated. As already mentioned above, each second semantic description may also form a first semantic description for another step of the method. According to one embodiment, the step of loading a second semantic description may be repeated each time, a resource link is found. I.e. if the second semantic description contains a resources link, this second semantic description will form a new first semantic description for the next step of loading a second semantic description. The new second semantic description for this next step will be the

semantic description to which the resource link of the new first semantic description points to.

5 In another embodiment, the step of analyzing may comprise searching for each element of the first semantic description containing a resource link, wherein the step of loading a second semantic description is performed for each resource link which is found in the first semantic description, and wherein, after loading each second semantic description, a query stored at the querying entity is executed. In this embodiment, the second semantic description will be loaded, irrespective of
10 the usage of the semantic description by a querying entity. Similar to the previously described embodiment, each second semantic description loaded into the memory may be analyzed regarding the presence of a resource link in one element of the semantic description. This resource link will point to another second semantically enhanced resource. Accessing a second semantic description and
15 loading the second description into the memory may be repeated, until the semantic descriptions of each second semantically enhanced resource is loaded into the memory, thus loading a complete tree with each semantic description associated with the first semantic description. This approach has the advantage that a query can be executed on the semantic information, which is loaded from
20 different semantically enhanced resources into the memory, and that the query does not need separate accesses to the semantically enhanced resources. This approach may reduce network load, particularly if the semantic information is needed for different queries regarding different querying criteria.

25 According to one embodiment, a right of the querying entity for accessing the second semantic description is checked before the second semantic description is loaded. This allows for fine-grained access policies. Each time a querying entity has the right to access the data, the second semantic description may be loaded into the memory. Each time the querying entity has not the right to access the
30 data, the second semantic description will not be loaded. In this case, the query may terminate due to the lack of required data or the query may jump to the next element and protocol the terminated loading request.

Each of the semantic descriptions may comprise one or more triples, preferably RDF (Resource Description Framework) triples. Preferably, classes and properties are defined in OWL (Web Ontology Language). Each of these triples comprise a subject, an object and a predicate linking the subject and the object. When using
5 such triples, the step of analyzing a semantic description may comprise following the links which are provided by the triples. In this case, the step of analyzing is more or less a search in a graph formed by the different triples.

In one embodiment, the querying entity may comprise a memory for storing
10 semantic descriptions, a control unit and an analyzing unit. In most cases the control unit and the analyzing unit will be implemented by a program running on a processor as already mentioned before. The control unit may access a first semantically enhanced resource and load into the memory a first semantic description associated with the first semantically enhanced resource. The
15 analyzing unit may analyze the first semantic description stored in the memory. The control unit may additionally access and load into the memory a second semantically enhanced resource/second semantic description, wherein a pointer to said second semantically enhanced resource – resource link – is contained in an element of the first semantic description.

20 In one embodiment, the querying entity may be a filter, wherein the filter preferably searches for semantically enhanced resources that fulfil a filter criteria stored at the querying entity.

25 In one embodiment, the querying entity is based on SPARQL (SPARQL Protocol And RDF Query Language).

In another embodiment, the memory for storing semantic description is operable
30 as a cache memory that stores each loaded semantic description, preferably for a predefined time. The cache memory may not only have the combined information, but may also store the access rights information for subsequent evaluations. Together this may enable subsequent access to be much faster as the process of combining the graphs does not need to be executed. The cache may be regarded as being valid for a predefined time, e.g. for 1 hour, 10 minutes or 1 minute. If the

predefined time is expired, the cache memory or the expired parts of it may be deleted.

5 In a further development of the cache memory, a cache invalidation entity is provided that checks the validity of the cache content and that may delete parts of the cache or the complete cache. There is a high likelihood that such a cache remains valid for a longer period as the semantic description usually captures (rather static) structural information. Cache invalidation can therefore be done on a time base with a suitable time period. Still there is a chance of an inconsistent state. Additional measures can be added to a system to avoid this. The following list provides three option that may be suitable measures and that may be combined with each other. Although described with regard to oneM2M, the measures may also be applied to other resource-based systems:

15 (a) Introducing an additional query parameter which controls whether a query is allowed to be executed on a cached copy. This enables applications to control the use of that cache. If usage of the cached information is not allowed, the query may be executed on the original semantic descriptions and the cache may be refreshed or invalidated.

20 (b) Introducing an additional query parameter which instructs an oneM2M system to keep the cache of the combined semantic description. Potentially, this parameter can contain a timeout or some other values defining when the cache shall be invalidated. These values may override any system internal cache invalidation parameters.

25 (c) Introduction and use of a subscription to the semanticDescription resources and/or to access control information which fires in case the semanticDescription is changed. This can be internally used for invalidating a cached copy and to reach a consistent state fast.

30 If consistency is absolutely necessary, the oneM2M system can introduce a transaction for updating a master copy of the semantic description and all the cached copies where the semantic description is contained in.

According to one embodiment of a method for generating a resource structure, the method comprises the step of analyzing a semantic structure representing a

behavior of a system. This analyzing step may comprise finding different semantically enhanced resources contained in the semantic structure.

5 In a next step semantically enhanced resources contained in the semantic structure may be extracted from the semantic structure. Further for each semantically enhanced resource, a semantic description may be generated as a subset of the semantic structure, wherein each semantic description contains nodes of the semantic structures associated with the respective semantically enhanced resource. If a node of the semantic structure is contained in two
10 semantically enhanced resources, a resource link between a first semantically enhanced resource and a second semantically enhanced resource may be generated. The resource link may be implemented as a URL (Uniform Resource Locator) or a URI (Uniform Resource Identifier). As such a resource structure which is extracted from a semantic structure will contain semantically enhanced
15 resources only, it might be possible to add resources that do not contain semantic descriptions, in order to describe system behavior more precisely. For adding these "standard" resources and/or for extracting the semantically enhanced resources, additional information about the internal structure of the described system may be provided.

20 There are several ways how to design and further develop the teaching of the present invention in an advantageous way. To this end it is to be referred to the patent claims on the one hand and to the following explanation of preferred embodiments of the invention by way of example, illustrated by the figures on the
25 other hand. In connection with the explanation of the preferred embodiments of the invention by the aid of the figure, generally preferred embodiments and further developments of the teaching will be explained. In the drawing

30 Fig. 1 shows a connection between a relation-semantic graph and an oneM2M resource structure,

Fig. 2 shows the semantic graph of Fig. 1 with regard to a possible filter operation,

Fig. 3 shows an example of a filter operation using a resource link according to the invention,

Fig. 4 an enlarged part of Fig. 3 showing the semantic graphs,

Fig. 5 shows another example of a filter operation using a resource link and

Fig. 6 shows a further example of a filter operation using a resource link.

The upper part of Fig. 1 shows an oneM2M resource structure. Each resource is illustrated by a rectangular box. Each link between two resources is illustrated by a line. The base resource is formed by the resource which is illustrated at the uppermost position. The two or more resources connected to this base resource are child resources of the base resource. Each of these child resources forms a parent resource for at least one child resource that is linked to it. The child resources that contain the semantic descriptions are shown as boxes with a dotted surface and are <semanticDescriptor> resources.

The lower part of Fig. 1 shows a semantic graph. Each node is an instance of an ontology and represents a concept or individual. Each edge represents a relation or a property. The semantic tree forms the logical structure of the system. The nodes, which are stored in one semantically enhanced resource, are contained in a dotted line. Three dotted arrows shows the relation between the areas surrounded by a dotted line and the respective semantically enhanced resource. Nodes with an edge crossing the dotted line between two neighboring areas are contained in the two semantically enhanced resources which are represented by the area surrounded by the dotted line. For these nodes, a resource link is used in order to provide a pointer from one semantically enhanced resource to another semantically enhanced resource.

Fig. 2 shows an example for a filter operation which is dependent on the semantic description contained in two semantically enhanced resources. The dotted box

represents the scope of such a filter, i.e. this is the information required for evaluating it.

Fig. 3 shows such an example. The diagram on the upper left hand side shows the oneM2M resource structure. The two diagrams on the right hand side shows the semantic description of the two semantically enhanced resources. The resource link between the two semantic diagrams is represented by a dotted ellipse. The filter request and the SPARQL representation of this request is shown on the lower left hand side of Fig. 3.

The filter request is: "Find all devices that have a service that has an operation whose output quantifies a temperature aspect". Its SPARQL representation is as follows:

```

SELECT ?device
WHERE { ?device rdf:type base:Device .
        ?device base:hasService ?service .
        ?service base:hasOperation ?operation .
        ?operation base:hasOutput ?output .
        ?output base:quantifies
temp:TemperatureAspect
}

```

When executing the filter request, the filter accesses "Device12" which forms a first semantically enhanced resource. The graph contained in "Device12" is a first semantic description. By following the edge "hasService" to node "Service23" and edge "hasOperation" to node "OperationA", the filter reaches a resource link that is contained in node "OperationA" and that points to a second semantically enhanced resource and a second resource description. At this stage, execution of the filter operation will be halted until the required information are retrieved from the second semantically enhanced resource. By following this resource link, the semantic description which contains a definition of "OperationA" can be accessed. The second semantic description is loaded into the memory. After loading is completed, filter operation will be resumed. By following the edge "hasOutput" to

node "OutputX" and edge "quantifies" to node "temperatureAspect", the filter reached the information that "Device12" fulfills the filter request.

Fig. 4 shows an enlarged part of Fig. 3 showing the semantic graphs.

Fig. 5 shows another example for a filter operation. The basic composition of the figure is similar to Fig. 3. The filter request of this example is: "Find all Rooms for which there is an indoorTemperature property of type Temperature whose unit is given in Celsius and whose resolution is within +/-0.5°C". Its SPARQL representation is:

```
SELECT ?room
WHERE { ?room rdf:type building:Room .
        ?room building:hasIndoorTemperature ?temp .
        ?temp rdf:type qk:Temperature .
        ?temp qk:hasUnit qk:Celsius .
        ?temp qk:hasResolution ?resolution .
        FILTER (?resolution <= 0.5)
}
```

Staring at resource "RoomA", the filter will follow the edge "hasIndoor-Temperature" to node "TemperatureYZ". This node contains a resource link to a second semantic description. By following the resource link and edge "hasResolution" (as specified in the filter request), the filter will reach "0.5". As this resolution is <=0.5, this node fulfills the filter request.

Fig. 6 shows a further example. In this example, the filter request is: "Find all irrigation groups with valves whose manufacturer is Grundfos". The SPARQL representation is:

```
SELECT ?group
WHERE { ?group irr:hasGroupMember ?valve .
        ?valve rdf:type irr:Valve .
        ?valve irr:hasManufacturer Grundfos
}
```

In this case, node “Valve987” contains a resource link to a second semantic description. By following the edges of the first and second semantic description, the filter may evaluate, whether this resource fulfills the filter request.

- 5 The following list summarizes some embodiments of the invention as well as some benefits with reference to oneM2M, a SPARQL filtering engine and semantically enhanced resources implemented by semantic descriptor child resources:
- 10 • The **SPARQL filtering engine** may be modified in the following way:
 - The semantic filter formulated as a SPARQL request is executed on the content of the semantic descriptor resource of the candidate resource
 - If in the course of the execution a class instance with one or more resourceDescriptorLink annotations is encountered, the execution is halted
 - 15 – The content of each of the semantic descriptor resources the semanticDescriptorLink references is added to the content on which the SPARQL request is being executed (*lazy evaluation*, alternative: fetch everything before execution, but may result in fetching unnecessary information)
 - 20 – The execution of the SPARQL request is continued on the enlarged content
 - System with semantically described REST resources, where
 - relevant semantic description in semantic descriptor child resources is distributed across REST resources
 - 25 – Annotation links connect the semantic information in semantic descriptor child resources with other semantic descriptor REST resources containing additional semantic information
 - Semantic operations may follow annotation links during their execution.
 - 30 • Specifically operation can be a filtering operation, where resources are filtered according to semantic description.
 - Specifically operation can be a discovery operation, where resources are discovered based on semantic description.

- During execution of semantic operations individual access rights on semantic descriptor resources may be respected.
- Caching of semantic information as well as access rights can be used to make subsequent executions of operations on the same information more efficient.
- Operations can have an explicit parameter to indicate that information should be cached for future use.
- Operations can have an explicit parameter to indicate that cached information not older than a specified time can be used for evaluation.
- Subscriptions on changes to semantic descriptor resources can be used for cache invalidation.
- Subscriptions on changes to access control rights of semantic descriptor resources can be used for cache invalidation.
- Transactions can be used to update semantic descriptions, updating/invalidating all cached copies consistently.

Many modifications and other embodiments of the invention set forth herein will come to mind to the one skilled in the art to which the invention pertains having the benefit of the teachings presented in the foregoing description and the associated drawings. Therefore, it is to be understood that the invention is not to be limited to the specific embodiments disclosed and that modifications and other embodiments are intended to be included within the scope of the appended claims. Although specific terms are employed herein, they are used in a generic and descriptive sense only and not for purposes of limitation.

Claims

1. Method for querying semantic information stored across several
5 semantically enhanced resources of a resource structure, particularly for a machine-to-machine application, wherein a semantic description is associated with each of the semantically enhanced resources, said method comprising the steps of:
- 10 accessing a first semantically enhanced resource of the resource structure by a querying entity, wherein a first semantic description is associated with the first semantically enhanced resource,
 loading the first semantic description into a memory,
 analyzing the first semantic description, and
 loading into the memory a second semantic description associated with a
15 second semantically enhanced resource of the resource structure, if a semantic element of the first semantic description contains a resource link, said resource link pointing to the second semantically enhanced resource.
2. Method according to claim 1, wherein the step of analyzing comprises the
20 step of executing a query stored at the querying entity, wherein the step of loading a second semantic description is preformed preferably only, if the second semantic description is needed for the query, and wherein during the step of loading the second semantic description the step of executing the query is preferably halted and is preferably resumed after the step loading the second
25 semantic description is terminated.
3. Method according to claim 1, wherein the step of analyzing comprises
30 searching for each element of the first semantic description containing a resource link, wherein the step of loading a second semantic description is performed for each resource link which is found in the first semantic description, and wherein, after loading each second semantic description, a query stored at the querying entity is executed.

4. Method according to claim 3, wherein each second semantic description loaded into the memory is analyzed regarding the presence of a resource link in one element of the semantic description, said resource link pointing to another second semantically enhanced resource, and that the semantic descriptions of each second semantically enhanced resource is loaded into the memory, thus loading a complete tree with each semantic description associated with the first semantic description.

5. Method according to any one of claims 1 to 4, wherein, before loading a second semantic description, a right of the querying entity for accessing the second semantic description is checked.

6. Method according to any one of claims 1 to 5, wherein each of the semantic description comprises one or more triples, preferably RDF – Resource Description Framework – triples, each of said triples comprising a subject, an object and a predicate linking the subject and the object, and wherein the step of analyzing a semantic description comprises following the links which are provided by the triples.

7. Querying entity for querying semantic information stored across several semantically enhanced resources of a resource structure, particularly for performing a method according to any one of claims 1 to 6, comprising:

a memory for storing semantic descriptions,

a control unit operable to access a first semantically enhanced resource and to load into the memory a first semantic description associated with the first semantically enhanced resource, and

an analyzing unit operable to analyze the first semantic descriptions stored in the memory,

wherein the control unit is additionally operable to access and load into the memory a second semantically enhanced resource, wherein a pointer to said second semantically enhanced resource – resource link – is contained in an element of the first semantic description.

8. Querying entity according to claim 7, wherein the querying entity is a filter, wherein the filter preferably searches for semantically enhanced resources that fulfil a filter criteria stored at the querying entity.

5 9. Querying entity according to claim 7 or 8, wherein the querying entity is based on SPARQL (SPARQL Protocol And RDF Query Language).

10 10. Querying entity according to one of claims 7 to 9, wherein the memory is operable as a cache memory that stores each loaded semantic description, preferably for a predefined time.

11. Method for generating a resource structure, particularly a resource structure for the use in a method according to any one of claims 1 to 6, comprising the steps of

15 analyzing a semantic structure representing a behavior of a system, extracting semantically enhanced resources contained in the semantic structure,

20 generating, for each semantically enhanced resource, a semantic description as a subset of the semantic structure, wherein each semantic description contains nodes of the semantic structures associated with the respective semantically enhanced resource,

25 generating a resource link between a first semantically enhanced resource and a second semantically enhanced resource, if a node of the semantic structure is contained in the first semantically enhanced resource and in the second semantically enhanced resource, said resource link preferably being implemented as a URL (Uniform Resource Locator) or a URI (Uniform Resource Identifier).

30 12. Method according to claim 11, wherein the semantic description is generated as an element, preferably an annotation, of the respective semantically enhanced resource, or wherein the semantic description is generated as a property which is linked to the respective semantically enhanced resource.

13. Database for enabling queries regarding semantic information stored across several semantically enhanced resources of a resource structure, particularly at a

machine-to-machine application, preferably being used in connection with a method according to any one of claims 1 to 6, wherein the database comprises hierarchically structured resources, wherein the resources comprise at least two semantically enhanced resources, wherein a semantic description is associated to
5 each of said semantically enhanced resources, respectively, and wherein at least one of the semantically enhanced resources comprises a resource link, said resource link pointing from one semantically enhanced resource to another semantically enhanced resource.

10 14. System for querying semantic information stored across several semantically enhanced resources of a resource structure, comprising:
storage space storing a database according to claim 13,
a querying entity, preferably according to any one of claims 7 to 10, and
a user device operable to generate a query for the querying entity.

15 15. Non-transitory computer readable medium storing a program causing a computer to execute a process, said process comprising:
accessing a first semantically enhanced resource of the resource structure,
wherein a first semantic description is associated with the first semantically
20 enhanced resource,
loading the first semantic description,
analyzing the first semantic description, and
loading a second semantic description associated with a second
semantically enhanced resource of the resource structure, if a semantic element of
25 the first semantic description contains a resource link, said resource link pointing to the second semantically enhanced resource.





Fig. 1



Dotted box represents the scope of a semantic filter, i.e. this is the information required for evaluating it

Fig. 2




Fig. 3






Fig. 4

oneM2M Resource Structure



Semantic Graphs



Filter request: "Find all Rooms for which there is an indoorTemperature property of type Temperature whose unit is given in Celsius and whose resolution is within +/-0.5°C"


SPARQL representation:

```
...
SELECT ?room
WHERE {
  ?room rdf:type building:Room .
  ?room building:hasIndoorTemperature ?temp .
  ?temp rdf:type qk:Temperature .
  ?temp qk:hasUnit qk:Celsius .
  ?temp qk:hasResolution ?resolution .
  FILTER (?resolution <= 0.5) }
}
```

On encountering the TemperatureYZ instance with resourceDescriptorLink referencing the semantic descriptor of the TempContainer resource, its content is added and the SPARQL execution is executed on the extended content

Fig. 5

oneM2M Resource Structure




Filter request: "Find all irrigation groups with valves whose manufacturer is Grundfos"

SPARQL representation:

```
...  
SELECT ?group  
WHERE { ?group irr:hasGroupMember ?valve .  
        ?valve rdf:type irr:Valve .  
        ?valve irr:hasManufacturer Grundfos  
}
```

Semantic Graphs



On encountering the Valve987 instance with resourceDescriptorLink referencing the semantic descriptor of the Device resource, its content is added and the SPARQL execution is executed on the extended content

Fig. 6

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2015/070245

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F17/30 H04W4/00
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F H04W

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	OLAF HARTIG ET AL: "Executing SPARQL Queries over the Web of Linked Data", 25 October 2009 (2009-10-25), THE SEMANTIC WEB - ISWC 2009, SPRINGER BERLIN HEIDELBERG, BERLIN, HEIDELBERG, PAGE(S) 293 - 309, XP019132809, ISBN: 978-3-642-04929-3 abstract page 294, paragraph 2 - paragraph 3 page 295, paragraph 2 - last paragraph page 296, paragraph 1; figure 1 page 304, paragraph 3 - paragraph 4 page 305, paragraph 2 ----- -/-	1-15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

13 April 2016

Date of mailing of the international search report

22/04/2016

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Herri, Edmond

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2015/070245

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	DENNIS PFISTERER ET AL: "SPITFIRE: toward a semantic web of things", IEEE COMMUNICATIONS MAGAZINE, IEEE SERVICE CENTER, PISCATAWAY, US, vol. 49, no. 11, 1 November 2011 (2011-11-01), pages 40-48, XP011364903, ISSN: 0163-6804, DOI: 10.1109/MCOM.2011.6069708 abstract page 40 page 42 - page 45 -----	1-15
A	WO 2014/125119 A1 (NEC EUROPE LTD [DE]) 21 August 2014 (2014-08-21) page 10 - page 12 -----	1-15
A	"Machine to Machine Communications (M2M); Study on Semantic support for M2M Data;TR 101 584 v050 Clean", ETSI DRAFT; TR 101 584 V050 CLEAN, EUROPEAN TELECOMMUNICATIONS STANDARDS INSTITUTE (ETSI), 650, ROUTE DES LUCIOLES ; F-06921 SOPHIA-ANTIPOLIS ; FRANCE, vol. SmartM2M, no. V0.5.0, 7 December 2012 (2012-12-07), pages 1-36, XP014167435, [retrieved on 2012-12-07] page 13 - page 14 page 16 - page 21, paragraph 1 page 9 - page 13 -----	1-15
A	Dominique Guinard ET AL: "From the Internet of Things to the Web of Things: Resource-oriented Architecture and Best Practices" In: "Architecting the Internet of Things", 12 April 2011 (2011-04-12), Springer Berlin Heidelberg, Berlin, Heidelberg, XP055257476, ISBN: 978-3-642-19157-2 pages 97-129, DOI: 10.1007/978-3-642-19157-2_5, abstract page 4 - page 9 page 25 - page 26 ----- -/--	1-15

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2015/070245

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	RAMPARANY FANO ET AL: "Handling smart environment devices, data and services at the semantic level with the FI-WARE core platform", 2014 IEEE INTERNATIONAL CONFERENCE ON BIG DATA (BIG DATA), IEEE, 27 October 2014 (2014-10-27), pages 14-20, XP032717260, DOI: 10.1109/BIGDATA.2014.7004417 [retrieved on 2015-01-07] abstract page 16 -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2015/070245

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2014125119 A1	21-08-2014	US 2015370900 A1 WO 2014125119 A1	24-12-2015 21-08-2014
