WoT Graph as Multiscale Digital-Twin for Cyber-Physical Systems-of-Systems

Gilles Privat
Orange Labs Services,
Grenoble, France
gilles.privat@orange.com

Thierry Coupaye
Orange Labs Research
Grenoble, France
thierry.coupaye@orange.com

Sebastien Bolle
Orange Labs Services
Grenoble, France
sebastien.bolle@orange.com

Philippe Raipin-Parvedy
Orange Labs Services
Rennes, France
philippe.raipin@orange.com

ABSTRACT

The Web of Things need not be restrained to a mere software-engineering construct or interoperability enabler. Viewed as a graph, the WoT may become the scaffolding of a comprehensive model of physical environments, capturing relevant aspects of their intertwined structural, spatial and behavioral dependencies As such, it can support context-rich mediation of data for network-enabled monitoring and control of these environments, augmenting them into multiscale and multilevel cyber-physical systems (CPS) and, more generally, systems of systems (CPSS).

Keywords

Internet of Things, Web of Things, Cyber-Physical System, Property Graph, System of Systems, Digital Twin

1. The Web of Things as a Graph

The Web of Things has been originally promoted [1] as addressing IoT interoperability issues by shifting the "narrow waist" upwards in the protocol stack. Significant as this protocol-centric view may be, it falls short of achieving the full potential of the WoT. The Web itself has evolved from solving a basic interoperability problems for exchanging documents, to being viewed, in the vision of web science [2][3][4] and network science[5], as a full-strength graph, scaffolding an information universe that stands on its own, liberated from its original hyperlink semantics.

Viewed as a graph, the original Web and its extension as the Linked Data cloud provide structure and meaning to the mass of inchoate information that get interlinked through it.

Viewed as a graph, the Web of Things provides a similar informational scaffold to physical environments that get interlinked and interfaced through the original "Internet of connected Devices" and its extensions to non-connected Things [6]. These environments make up Cyber-Physical Systems and, more generally, systems of systems [7]. Information models and interfaces for these systems have, so far, been highly heterogeneous, piecemeal and siloed. Extending from present-day models and interfaces for individual connected devices, the Web of Things bears the promise of unifying Cyber-Physical Systems of Systems under the umbrella of a comprehensive graph representation, as we will explain in the following.

2. Graphs as comprehensive information models for IoT/CPS environments

IoT/WoT environments (e.g. buildings, industrial compounds cities) bring together multiple data sources (e.g. sensors) and data users (IoT/CPS applications), for which this environment is both a physical anchor and a common denominator. A comprehensive model of this environment is indispensable to support the horizontal sharing, consolidation and enrichment of this data and to avoid that each of these applications remain custom-designed and vertically integrated, within its own silo, using its own exclusive data sources.

2.1 Starting point 1: "pure" distributed WoT

In its mainstream view, the Web of Things revolves, much like the original web, around an implicit and 100% distributed graph of hyperlinks, not between web pages or text anchors, but between resources exposed by REST interfaces to various functions of individual connected devices.

In a "pure" ROA interface, (corresponding to the third level, Hypermedia As The Engine of Application State, "HATEOAS", of the Richardson Maturity Model), no global functional description is provided, all resources are self-descriptive and provide their own URI that can be interpreted by applications without requiring any "out-of-band" information. In this approach, a graph representation of the overall system comprising these devices does exist, but only implicitly through these hyperlinks. All the information that may be required from applications is in principle available by traversing the graph made up of these hyperlinks, providing the equivalent of interface introspection, discovery and dynamic service composition from more traditional approaches. Just as the original public Web itself, the "pure" distributed Web of Things requires external or third party tools to provide the equivalent of the functionalities that are natively provided by more traditional cloud-based IoT platforms, such as database-like queries.

2.2 Starting point 2 : IoT-enabled digital twins of physical systems

Digital twins started from a very different premise, in a very different community (industrial manufacturing): they aim at providing a comprehensive and accurate informational/numerical replica of some advanced contraption (e.g. a satellite), to optimize its operation and maintenance. From its origin in the aerospace industry, the digital twin concept has evolved and widened, to integrate remote sensing and actuation capabilities through the IoT [9][10]. In this view, the twin could be seen as a mere informational proxy, maintaining key information about the physical system it stands for and serving as a monitoring/actuation interface to this system. We allow ourselves to retain the phrase even if the corresponding model is strictly minimal and stripped of all the more advanced trappings of a "genuine" digital twin, such as multiphysics simulation

In this downgraded acceptation, the WoT interface to a connected device could be seen as the bare-bones digital twin of this device.

2.3 Starting point 3: Cyber-Physical Systems and Systems of Systems

Cyber-Physical Systems could be seen as the evolution of traditional industrial control in the era of open networks. They started from classical industrial systems that were not only closed, but also designed in a 100% top-down fashion, with every component and subsystem exactly fitted as a part of the overall tightly-coupled system they made up. The integration of CPS with open networks and platforms inherited form the IoT is the stage we are in now, which will also widen the narrow range of present-day IoT as it started from consumer devices and mostly upwards (sensor-originating) data collection, towards more integration with actuators and real-time control.

Yet the next stage of CPS evolution is just as important: moving from traditional top-down-designed systems to bottom-up-assembled, open and loosely coupled "systems of systems". In a proper sense, systems of systems [8] such as cities or large buildings are the engineering counterpart of "complex systems" that have emerged since the 1980s as a far-reaching transdisciplinary concept from physical, biological and social sciences. And this is where the graph idea is creeping in again! Graphs have been, since the consolidation of the "Network Science"[5] domain in the 1990s, a foundational model for the analysis of complex systems. They should also become a pivotal model for the analysis and operation of systems of systems. These models are not, by any means, intended to replicate the whole system in a digital twin fashion: they only capture relevant interdependencies between subsystems to support the analysis of potentially unforeseen emergent phenomena that may result from the interworking of systems that had not been designed to operate jointly.

2.4 Integrating device/thing twins and system twins under a single graph

The coming-together of the previous three viewpoints opens up a fresh and promising understanding of expanded WoT graphs as representations of CPSS environments.

Networked IoT devices, as exposed by the current Web of Things are but the basis subset of primary nodes to be included in this graph. The graph grows orders-of-magnitude richer by linking, through these primary nodes, all kinds of non-directly-connected "Things"[11], be they pieces of furniture in a building, trees in a park or bicycles in a city. The only requirement is that these secondary things are individually monitored by networked sensors, or possibly controlled by networked actuators (both of which are these primary nodes of the graph) through what we proposed to call a "phenotropic" or "stigmergic"[12] link in this extended WoT. Both these non-connected things and devices may be seen as atomic or black-box systems, represented by one and only one node of the graph and described only though their interfaces, their properties, and possibly their state, as attached to this node. They might be decomposed further down with component subsystems that may become nodes of the graph in turn, but they remain, normally, top-down self-contained systems. We propose to call these nodes of the graph "Thing Twins" (TT for short) to capture their nature

Representing only such atomic "things" as nodes of what would be a "flat" graph is insufficient to capture the complete structure that corresponds to all nested and intertwined subsystems of the overall referent system, such as e.g. security or HVAC systems within a building, water distribution or traffic management systems within a city. We need to match these systems to a different species of nodes, within an overlay "graph of graphs", as will be explained next. These nodes will be defined as "System Twins" (ST for short).

Just as for device twins and thing twins, these "graphs as twins" are not fully-fledged digital twins (in the original sense of the phrase) of the systems they stand for. They capture only a few key structural features of these systems as the inter-relationships and properties of their components, together with, crucially, their own inter-relationships with peer systems at the upper level of nesting.

A graph comprising nested "System Twins" and underlying "Thing Twins" could be seen, at its own level and scale, as an overarching "System of System" twin. Unlike the digital twin of a top-down-manufactured system, it cannot start its life as one instance of a common blueprint, because there is, by definition, no such blueprint for a system of systems. The system of systems twin graph has to evolve incrementally and organically, from the bottom-up, as the aggregation and interconnection of partial information from its "Thing Twin" and "System Twin" nodes.

3. Property Graphs as a Meta-model for persistent structural representation of Systems of Systems

An adequate graph model for the representation of Cyber-Physical Systems and IoT/WoT environments at large should have a sufficient level of expressivity to match the structure of these environments. These CPS graphs have in this regard a kind of similarity-based semantics when they are meant to mirror the structure of a physical network such as e.g. a transportation network. These semantics apply to the graphs as a whole and are not reducible to the kind of "per-resource" semantics, which RDF is meant to describe. As previously expressed, our position [13] , endorsed by the ETSI CIM group for the NGSI-LD Context information management data model [17], is that property graphs are the best existing meta-model for capturing these CPS graphs. Property Graphs [14][15][16] (PG for short in the following) are a class of directed, labeled & attributed multigraphs, informally defined as the common denominator model of graph databases. They have so far lacked a strong theoretical grounding as they have emerged from the use of database practitioners as a compromise to retain familiar key-value or object primitives within a graph. Property Graphs make it possible, crucially, to single out as relationships (and thus first class citizens of the meta-model) those arcs that represent actual physical linkages between physical entities, themselves represented as nodes. Mere properties (corresponding to OWL datatype properties) are directly attached to both entities and relationships as attributes would be in an object-based model. Paradoxical as it may seem, in a property graph, properties are usually not represented as arcs of the graph proper! This keeps the graph uncluttered and "clean" to feature saliently what matters the most: relationships as representations of physical connections between entities that make up the structural scaffolding of a system. This makes it possible to run graph-theoretical algorithms that rely on this structure: these can be extremely classical algorithms like maximal flow computation or complexity analysis tools such as evaluations of average path-lengths or degree distributions, or more sophisticated tools from spectral graph theory.

By contrast, the RDF metamodel is meant to capture "pure" knowledge as predicates and is not adapted to properly capture the structure of physical systems. : if an arc/edge of these graphs stands for a physical connection between two nodes, such as e.g. a power line in an electrical grid, or a pipe in a water distribution system, it should be entitled to have properties of its own to fully describe the underlying piece of physical plant, which an RDF graph cannot natively support without either a cumbersome reification process.

Besides this, RDF graphs are unsuited to an analysis by classical graph-theoretical tools because they obfuscate, flatten and dissolve the physically-matched graph structure in the mixing of structural arcs with mere property values and typing arcs that have no structural relevance whatsoever. This obfuscation is compounded by the need for reification or conversion of arcs into vertices as mentioned before, hiding the actual connectivity structure underneath an additional layer of transformation.

4. Conjoining "Thing Twins" (TT) and "System Twins" (ST) in a multiscale Property Graph

The example from Figure 1 shows how "TT" nodes that stand for individual physical entities, i.e. things and devices (black rectangles) are captured in a property graph with their relationships (black diamonds) and their properties (black ovals for the predicate/key of the property, hexagon for the target value).

Subgraphs of the overall TT (black) graph corresponding to relevant sub-systems are matched to "ST" nodes of the overlay red graph through the "*isNodeOfGraph*" special relationship (red diamond). This relationship applies to all nodes in the corresponding subgraph, yet is shown only once to avoid cluttering the diagram. ST nodes are themselves caught in graph relationships, such as the relationship expressing that the city traffic management system may impose constraints on the traffic-light management system, or the inclusion of these (distributed) systems into the overarching graph capturing the Smart City as a system of system through the "isSubGraphOf" relationship.

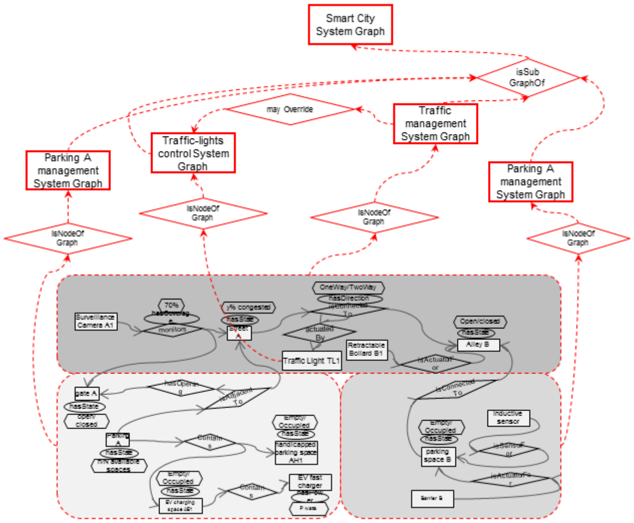

Figure 1 : Smart City Graph example

5. System clustering as PG overlay on CPS graph

Even though System Twins and Thing Twins are proposed to be included in a joint Property Graph, they are of a different nature when interpreted in view of their matching to the actual physical "stuff" (the *plant,* in control theory parlance). Figure 2 gives a partial, tentatively perspective view of our previous smart city example. Thing Twins stand for physical entities as their direct informational representation, on an information-object-instance per physical-entity-instance basis (one-to-one mapping). By contrast, System Twins do not directly stand for Physical Entities, but for technical groupings of the informational representatives of such entities (their Thing Twins) that provide a joint functionality as a system. The mapping of ST nodes to physical entities is many-to-one and one-to-many. An example is the city traffic management system illustrated above: this is a distributed cyber-physical system which is "more cyber than physical", and does not "include" the streets as components, but may still let traffic flow into them, by means of relevant actuators (e.g. retractable bollards or barriers) that are attached to both the street entity (Alley B in the example above) and the traffic management system that controls them. The same street entities will also be associated with other technical systems (e.g. lighting management) with which they will also share their state as an entity.

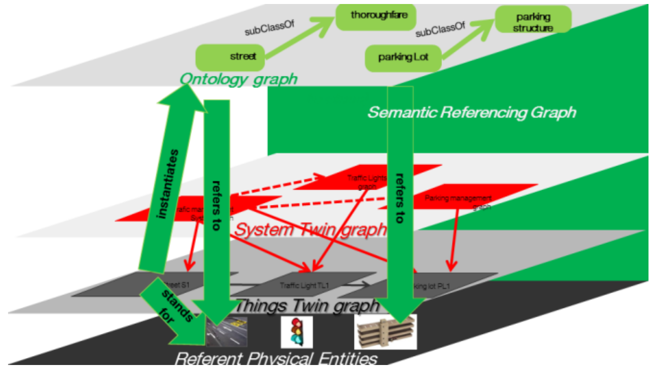

Figure 2 : A tale of three graphs : System Twin graph, Semantic referencing & Ontology graphs as complementary overlays on Things Twin graph

This view of (mostly distributed) cyber-physical systems as informational abstractions overlaid above the physical plant might seem to contradict a more classical view of systems as self-contained physical enclosures. These two views are actually complementary and may coexist as different descriptions of the same system in the proposed paradigm, or as complementary subsystems of a larger overall system of systems. There is *not*, in the environments we address with the Web of Things and Cyber-Physical Systems, a sharp binary distinction between :

1. self-contained, top-down, tightly-coupled, "mostly physical" CPS
2. distributed, bottom-up, loosely-coupled "mostly informational" CPS

The reality of current Cyber-Physical Systems and Systems of Systems is that of a mixture, at different scales and levels, of systems that are somewhere in a spectrum between these two extremes! Obviously, the larger we go in scale, the more we tend towards the second alternative, but there do still exist large-scale systems that remain squarely within the confines of the top-down-designed genre, (such as e.g., airplanes, or nuclear power plants…) , and they are the primary target of digital twin modeling, in its classical acceptation.

If it has long since been accepted that most physical systems evolve towards becoming cyber-physical at different scales and levels. Our view is that most cyber-physical systems are also evolving towards becoming systems of systems, at different scales and levels. Our proposed model is intended to account for this momentous evolution.

6. Semantic referencing as RDF-based overlay on PG-based CPS graphs

We stated previously [13] that RDF graphs should be used complementarily to CPS graphs for which we propose to use property graphs: semantics à la RDF may still apply to individual resources (nodes, relationships and properties) of the CPS graph, even if this graph has a distinct global semantics of its own, as previously stated. Figure 2 shows how this pans out, with ontology classes (pictured as solid green, rounded rectangles) playing a similar role to system groupings (red ST nodes), within a graph overlay of their own. The similarities between the two:

- both graphs are overlays upon the "TT" graph
- ontology classes have a many-to-one and one-to-many mapping to TT instances

are just as interesting as are the differences :

- the ontology graph and the semantic referencing links are defined are RDF graphs typically using *rdf:type* and *rdfs:subClassOf* properties, while the System Twin graph we propose uses the PG model
- ontology classes that are *instantiated* by TT graph nodes *refer to* the actual physical thing that the TT graph node *stands for,* NOT to the TT node itself, The TT node is but an informational representation (a *signifier*, in semiotics parlance) whereas the actual *referent* is the piece of physical stuff it represents: this is the crucial idea of the semiotic triangle, picture as solid green arrows in figure 2.
- ST nodes do, by contrast, stand on their own as informational objects : if they instantiate classes (not shown in figure 2), these classes refer to the system as an informational abstraction. This is also the case for properties and relationships of the ST graph.

7. Perspective

The TT & ST graph as put forward here do not include persons, their roles or legal entities. This would obviously be useful, if not indispensable, for technical reasons (associating administrative roles to various stakeholders or contractors operating the corresponding subsystems). A coupling to regular social graphs (with e.g. links to passers-by in the streets) would raise daunting privacy issues and is, if only for this, hard to envision, even if the hybrid triplicate graph resulting from the combination of cyber-physical graphs, knowledge graphs and social graphs could be an even more far-reaching incarnation of the "Giant Global Graph" contemplated by Tim Berners-Lee…

Acknowledgments

We acknowledge the invaluable contributions of all our colleagues from the Orange Labs "Thing'in" (http://www.thinginthefuture.com) integrative platform project who have implemented the graph concept presented here and its accompanying enablers. We acknowledge as well the contributions of the partners of the ETSI CIM (Context Information Management) Industry Specification Group with whom we have evolved the NGSI-LD property graph information model used for this project, especially our colleague Dr. A. Abbas who was a key author of the formal RDF specification and conversion of this model. The FIWARE project and open source community have also been a breeding ground and early promotion vehicle for some of the ideas presented here

REFERENCES

[1] Guinard, D., & Trifa, V. (2009, April). Towards the web of things: Web mashups for embedded devices. In *Workshop on Mashups, Enterprise Mashups and Lightweight Composition on the Web (MEM 2009), in proceedings of WWW (International World Wide Web Conferences), Madrid, Spain* (Vol. 15).

[2] Hendler, J., Shadbolt, N., Hall, W., Berners-Lee, T., & Weitzner, D. (2008). Web science: an interdisciplinary approach to understanding the web. *Communications of the ACM*, 51(7), 60-69.

[3] Kleinberg, J. M., Kumar, R., Raghavan, P., Rajagopalan, S., & Tomkins, A. S. (1999, July). The web as a graph: measurements, models, and methods. In *International Computing and Combinatorics Conference* (pp. 1-17). Springer, Berlin, Heidelberg.

[4] Albert, R., Jeong, H., & Barabási, A. L. (1999). Internet: Diameter of the world-wide web. *Nature*, *401*(6749), 130.

[5] Brandes, U., Robins, G., McCranie, A., & Wasserman, S. (2013). What is network science?. *Network science*, *1*(1), 1-15.

[6] G. Privat, "Extending the Internet of Things," *Communications & Strategies, Digiworld Economic Journal,* pp. 101-119, 2012.

[7] G. Privat "The "Systems of Systems" viewpoint in telecommunications, https://hellofuture.orange.com/en/the-systems-of-systems-viewpoint-in-telecommunications/

[8] Maier, M. W. (1998). Architecting principles for systems-of-systems. *Systems Engineering: The Journal of the International Council on Systems Engineering*, *1*(4), 267-284.

[9] Boschert, S., Rosen, R., & Heinrich, C. (2018). Next generation digital Twin. In *Proceedings of the 12th International Symposium on Tools and Methods of Competitive Engineering TMCE* (pp. 209-218).

[10] Siemens Mindsphere Whitepaper

[11] G. Privat, S. Bolle, T. Coupaye, P. Raipin : "Thing'in, the Things Graph Platform", *Orange Research Blog,* November 2018

[12] G. Privat, (2012). Phenotropic and stigmergic webs: the new reach of networks. *Universal Access in the Information Society*, 11(3), 323-335.

[13] G. Privat & A. Abbas, "Cyber-Physical Graphs" vs. RDF graphs" Position paper for W3C Workshop on Web standardization for Graph Data, Berlin, March 2019

[14] M. A. Rodriguez and P. Neubauer, "Constructions from dots and lines," *Bulletin of the Association for Information Science and Technology*, vol. 36, no. 6, pp. 35–41, 2010.

[15] Margitus, M., Tauer, G., & Sudit, M. (2015, July). RDF versus attributed graphs: The war for the best graph representation. In *Information Fusion (Fusion), 2015 18th International Conference on* (pp. 200-206). IEEE.

[16] Das, S., Srinivasan, J., Perry, M., Chong, E. I., & Banerjee, J. (2014, March). A Tale of Two Graphs: Property Graphs as RDF in Oracle. In *EDBT* (pp. 762-773).

[17] CIM/NGSI-LD Data model Specification Preliminary Draft https://docbox.etsi.org/ISG/CIM/70-Draft/006-MOD0/CIM-006-MOD0v002.doc