

Phenotropic & stigmergic webs: the new reach of networks

GILLES PRIVAT

*Orange Labs R&D, Technologies
28 Chemin du vieux Chêne, F-38240 Meylan, France*

gilles.privat@orange-ftgroup.com

Keywords: Internet of Things, stigmergy, phenotropics, web of services

Abstract

This article proposes a conceptual extension of sensor-actuator networks, taking in all "things" that can be sensed by sensors, or acted upon by actuators, in various physical modalities. These things become nodes of a web, graph or virtual network overlaid on the existing sensor-actuator networks that make up the "Internet of Things". We explain how the broader concepts of phenotropics and stigmergy may account for the special kind of connections that these networks entail. Phenotropics refers to a model of communication between these nodes by way of pattern recognition. Stigmergy refers to a model of self-organization that uses communication between entities by modifications of a shared physical environment. Phenotropic-stimergic webs "loop back" sensor-actuator networks by way of the physical world. Graph-based complexity models provide a means of analysing the hybrid systems made up by these networks and the additional nodes attached to them in this way.

We explain how an evolution towards such paradigms in the realm of network-to-environment interfaces draws upon a similar, long-standing evolution in the realm of human to information interfaces. We explore the consequences of these new networking paradigms on the architecture, management and organization of networks. We show how these ideas can expand and enrich present-day applications of pervasive networking, by taking full advantage of the physical nature of the new end-points of digital networks, and how they bear upon human interfaces to networked services, possibly opening up new territories for universal access.

1. Introduction

1.1. Beyond the "internet of things"

Most mainstream visions of the "Internet of Things"[1] come down to an extension of the range of devices that may become attached to networks, usually by means of radio-based technologies such as RFID or Zigbee¹.

The underlying rationale is straightforward: there are trillions of "things" waiting to get connected, when billions of humans already are. If some new-found variant of Metcalfe's law would apply, the promise of these "things to things" connections would appear boundless.

Under such earlier catchphrases as "smart devices", "communicating/cooperating objects", "pervasive networking" or M2M², it is no surprise that the telecom sector had been embracing this

¹ Zigbee is a short-range, low-bit-rate and low-power radio protocol used for connection of sensors or other low-end devices.

² M2M (Machine to Machine) is, in the telecom industry, the favored designation for the new domain of services where mobile terminals are used in conjunction with sensors for remote monitoring or remote control. Viewed originally as a mere extension of the subscription base for cellular services (using embedded SIM cards), M2M services are now understood as potentially using all kinds of special-purpose wireline or wireless access networks, extending their range to low-end devices for which direct connection to regular cellular networks would not, technically or economically, make sense.

evolution as a legitimate extension of its territory, well before the "Internet of Things" (IoT) gained currency as the new buzzword of choice. When incorporated into the lingo of a perennially parochial telecom industry, these early attempts at redefining communication beyond person-to-person have created some confusion, as the distinction between the different categories of new "things" objects or devices that became attached to networks was not always clearly understood, especially when mobile phones and their avatars were added to the mix. Regular IT interface devices and telecom terminals, for which sensors and actuators are used exclusively to support classical human interfaces should normally be excluded from the internet of things proper. IoT devices should be characterized in that *they are endowed with function-specific capabilities to interact with the physical environment through embedded sensors and/or actuators*. Understood in this way, IoT devices are whatever appliances, machinery, fittings, apparatuses, contrivances which have their own physical function in the physical environment, in whatever form factor or outer appearance they may come. They are, in a proper sense, "embedded" in this environment, with information processing and transmission capabilities added on top of it.³

Another defining thread of the Internet of Things originated from RFID and other tagging technologies as used in such canonical applications as supply chain management. Originally distinct from the telecom & M2M vision, this view centred more on the low-end of the spectrum of connected things, towards entirely passive items such as supermarket goods or pieces of hardware.

Compelling and extensive as it may seem, this "things-to-things" vision misses the crux of the broader IoT evolution, which, whatever its name, is not a purely quantitative enlargement. Encompassing the whole spectrum of connected things, from passive items to sensors, actuators and other embedded appliances, it is a quantum leap, portending the liberation of networks from their informational confines. By connecting "things" that are deeply embedded in the physical environment, ICT systems become strongly coupled with all kinds of physical systems, opening up entire new domains that had remained outside the purview of ICT, or for which information systems were entirely disconnected from the corresponding physical plant/system/process, requiring manual data entry to relate the two.

In this view, the outermost border of the digital network is still the sensor or actuator itself, beyond which is the hazy analog world. The revolution of pervasive networking that lead to the multiplication of these connected sensors and actuators [9,11] afforded an order-of-magnitude enlargement in the interaction bandwidth between the analog environment and the digital word. Yet, for all their transformative role, this current generation of sensors and actuators do not correspond to the ultimate possible displacement of the network border. How this border may shift further is precisely the next stage of the evolution that we intend to describe.

Drawing an analogy from the domain of human interfaces, we first provide a view on the possible extensions of networks to things that can be sensed by sensors, and actuated by actuators, and explain how these extensions can be described in a graph-based formalism.

Taking a different viewpoint, we then relate these two ideas, with all due reservations, to the original concepts of "phenotropics" and stigmergy⁴, as introduced respectively by Jaron Lanier and Pierre-Paul Grassé.

³ The distinction is not entirely clear-cut because many IT devices now include new physical interfaces (such as location sensors or NFC tags/readers), that, even if they are not used directly for human interaction, provide contextual information that may be used directly or indirectly for human interaction.

⁴ The latter word has already been widely adopted in the scientific literature on collective intelligence, even though it first appeared in a french-language journal on social insects, whereas the Lanier phenotropics article has had very little following so far.

We then provide concrete examples to show how these ideas can be applied, and conclude by drawing a link to robotics.

1.2. Drawing upon human interfaces

Even the most radical advocates of universal RFID or IPv6 would shy away from enrolling human beings into their systems with a lifelong ID or IP address. Humans do belong, just as physical things, in the analog world outside of digital networks. Yet they are, for good reason, treated in a radically different way. Contrasting the evolution of the border between digital networks and "things" on the one hand, between digital networks and humans on the other hand, provides interesting insights, to be applied from the latter to the former, and not the other way round!

Most grand schemes devised for the Internet of Things (such as the EPCglobal Network⁵, the uIDCenter⁶, or, in a very different vein, the "Internet 0"[3]) boil down to attaching a universally unique, network-ready digital identity to analog things, be it their General ID, ucode or IP address. This amounts to digitizing these "analog things", or, more broadly to making the physical world more digital by letting the digital world encroach upon the world of analog things.

In the realm of human interfaces, *exactly the opposite trend* has been at work, which could be summarized as "making the digital world appear more like the analog/human/physical world". All varieties of human interfaces have been moving in the very same direction: they try not to force human users to meet the digital environment on its own digital terms, or, equivalently, to make the interface for human users look less like the interface between programs or networked entities. The entire agenda of so-called "perceptual interfaces" [10] bears witness to this. The difference between digital data input through a keyboard or command entry through a menu selection and through a voice recognition software should make this clear. A less obvious and more interesting example is the replacement of clicking on a menu item by the grasping of a tangible interface that represents (physically impersonates) the same digital entity. Obviously, the interface is moving much further into the analog world in the latter case, and it requires more sophisticated sensing and perception capabilities on the part of the system.

These concurrent evolutions have each been advocated for valid and widely accepted reasons in their own right.

As for human interfaces, convenience and ease of use are not the only reasons for the un-digitization trend: robustness, graceful degradation and reliability are complementary and equally valid reasons. This is not obvious when, e.g., keyboard input is replaced by imperfect speech recognition. It would become clearer if we could replace a password input by a 100% foolproof and transparent biometric identification software.

In the slow-moving world of software architecture, it certainly is provocative, even preposterous, to propose, as a few visionary authors have done [1][2] [4], that communication between different modules of a large software system should try to move towards this analog interface model. The rationale for this "de-protocolization" of network & software interfaces is not that they have to be accessible to a human user (though this could be also an argument in this case) but that they should overcome the brittleness inherent in syntax-bound protocols and programmatic interfaces, in order to become, if possible, more robust, gracefully degradable and scalable.

⁵ <http://www.epcglobalinc.org>

⁶ <http://www.uidcenter.org>

The main thesis of this paper is that similar arguments of scalability, expandability and robustness can apply for networks of things and, more generally, for networks that are closely coupled to things through sensors and actuators. Communication between these things need not be more digital, it may retain the specific properties of the physical world in which these things belong. Beyond this, the physical world provides both the inspiration and the model for these new paradigms that may percolate back in the digital world.

2.The new web of sense-able/actionable things

Starting from a definition of an internet of embedded devices (i.e. sensor-actuator-equipped devices) as outlined before, we may try to examine how we can actually try to do for this internet of things what has been done for human interfaces. This amounts to try to make the outer interfaces of this network more analog and "thing-friendly", instead of enforcing digitization on these things.

2.1. Integrating the network borderland of "sense-able" things

In this perspective, the range of things that may become *indirectly* part of networks can actually be extended much further than sensor devices themselves, as pictured in figure 1. If we use a sensor (e.g. a camera) and a "thing" recognition software analysing the data acquired by this camera, we can consider that every single "thing", every passive item within the field of view of this camera that can be "recognized" by this software becomes ipso facto a "networked thing", without requiring an RFID tag or even an optical code (such as a 1D or 2D barcode) for this. This is represented by a new kind of network link on figure 1, directed from the passive item in question to the sensor. Much rests on the sense given to "thing recognition" here, and we will elaborate on this in the phenotropics section of this article. For the time being, let's say that this idea goes much beyond the sensing of individual items by individual sensors. If we have a federation of distributed networked sensors available (such as represented in figure 1, as a miniature of a "smart environment"), networked "things" will comprise all "stuff" that can be sensed by data fusion & pattern recognition software operating on top of these federated sensors *working together*, potentially overcoming their individual limitations as single-modality devices.

This is not an evolutionary, incremental and quantitative extension of networks, such as can be obtained by integrating a new radio-based protocol. It is really a qualitative leap in that it makes it possible to integrate all analog "stuff" as it is, discrete or bulk amorphous analog things without any digital identity or without any network interface whatsoever, and without adhering to any kind of standard, at any level, for this network connection.

Not only is there no prior barrier to the integration of new things, this integration is also 100% universal as it requires no prior standardization of any kind of code or interface.

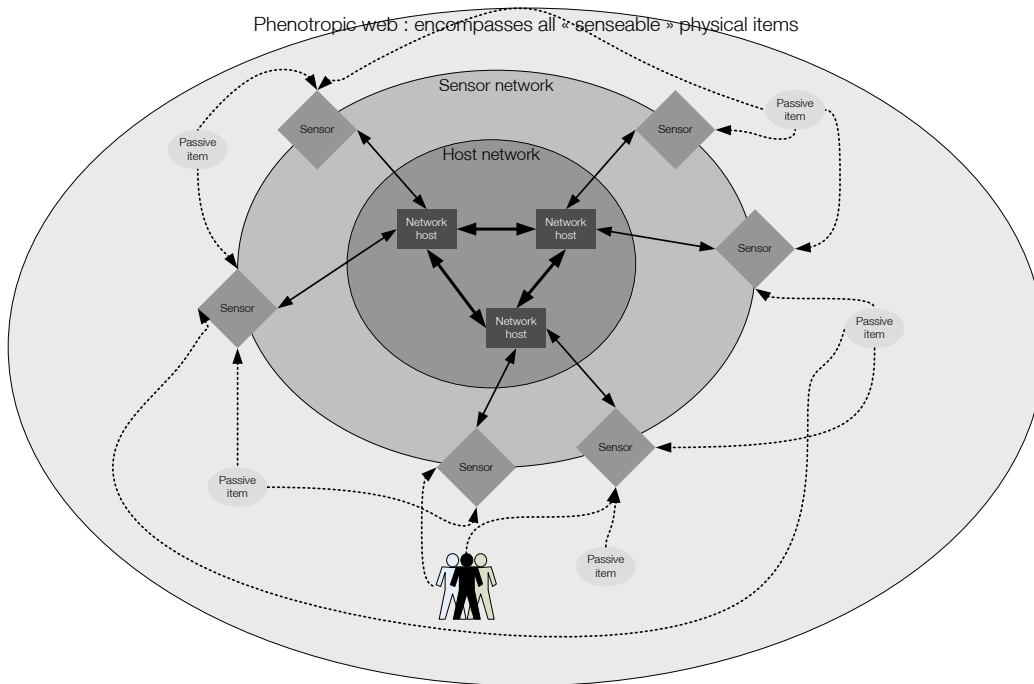


Figure 1 : Enlarged perimeter of networks encompassing all sense-able things

2.2. Integrating the network borderland of "actionable things"

Actuators enact physical modifications of the physical environment and these modifications are sensed by sensors, either directly indirectly, through passive "things" that are modified by the actuators. These new physical links (actuator → environment → sensor) or (actuator → thing(s) → sensor) make up a graph, or virtual network that we may call a stigmergic network, overlaid upon the wireline/wireless data network to which these sensors and actuators are attached to receive or transmit their respective numeric data (figure 2).

In this view, sense-able things that belong to the outer borderland of the digital network described before have a complementary way of becoming integrated into the digital networks, as potentially being acted upon by actuators. The condition for these actionable things to become integrated in the network proper is that effects of their actuation can in turn be picked up by sensors.

What we integrate in the network here is not new nodes, but new links that close loops of sensor-actuator networks in a way that does not use the modalities of classical networks and complements them.

We may distinguish between three cases of "stigmergic links":

Generic actuators acting on generic passive things that are independent from the actuators

These actionable things may be purely passive, and in this case, they may be actuated by external actuators that are independent from the thing itself and may act upon a variety of external things, or loci (e.g. surfaces or volumes) in their environment. Such would be the case for e.g. a robotic arm fetching or moving something, or a projector projecting an image on something, or a foot leaving a footprint on the ground, a pen leaving a mark on a sheet of paper, etc. In this case we consider that the stigmergic network establishes a link between the actuator and the thing or locus being acted upon, which may be considered as a separate node of the stigmergic network.

A very important case corresponds to what has been called [12] "sematectonic" stigmergy, or collaborative physical "work", where such a thing or locus can be acted upon jointly by several

actuators, with physical constraints mediating these coupled effects, corresponding to a node with several incoming links in the stigmergic graph, where the only means of coordinating the collaboration is the thing being acted upon itself. The prototypical example is a track being blazed on soft ground, or a hole being dug, where implicit spontaneous coordination of this kind will lead several agents to plow the same track or dig the same hole, mutually benefiting each other and reinforcing their own work. This is also the case for the kind of coordinated work among social insects (such as termites) that was the focus of Grassé's original definition of stigmergy [4].

Actuators tightly coupled with thing they act upon

It may not be relevant to separate an actuator from the physical thing or locus it acts upon when this relationship is fixed. In this case we may consider that the state of the thing being actuated is the output state of the actuator. Such is the case for e.g. a motor used to open or close a valve, a window blind or a door. The state of this actuated device may in turn influence other things in the environment, such as the state of the room being changed by having doors open, and in this case the link between the two falls into the previous category.

Actuators that are part of a system composed of several actuators and sensors

They may be complex appliances or devices that integrate both sensors and actuators, ranging in complexity from simple control systems or domestic appliance to mechatronic systems such as a robots. In this case, tightly coupled actuator-sensor loops are already part of the design of the device, yet what is interesting is that other sensors available in the environment may come into play to augment the sensing capabilities of the device itself, complementing the internal actuator-sensor loops of the device with external, more loosely coupled ones. For example, a robot is normally able to sense the position of its own joints by using its internal sensors, but other sensors in the environment will also pick up the position of the robot.

It could be considered that humans, be they active users of the target system or passive passers-by, belong to this latter category...As passive entities they are acted upon and sensed, yet they also comprise actuators that can modify the environment, the effect of which is in turn sensed by sensors.

In fact, the distinction between these three cases is only relevant to the degree to which we analyse the system, i.e. whether we dissect subsystems into their constituent parts as actuators and things/system parts being acted upon, or consider them as black boxes. If we could drive down the modeling to the finest granularity, we could model all these cases similarly, with a stigmergic link between actuators and "things" or "system parts". The only difference would be that in some cases these links would represent tight & fixed coupling and in some others a transient one.

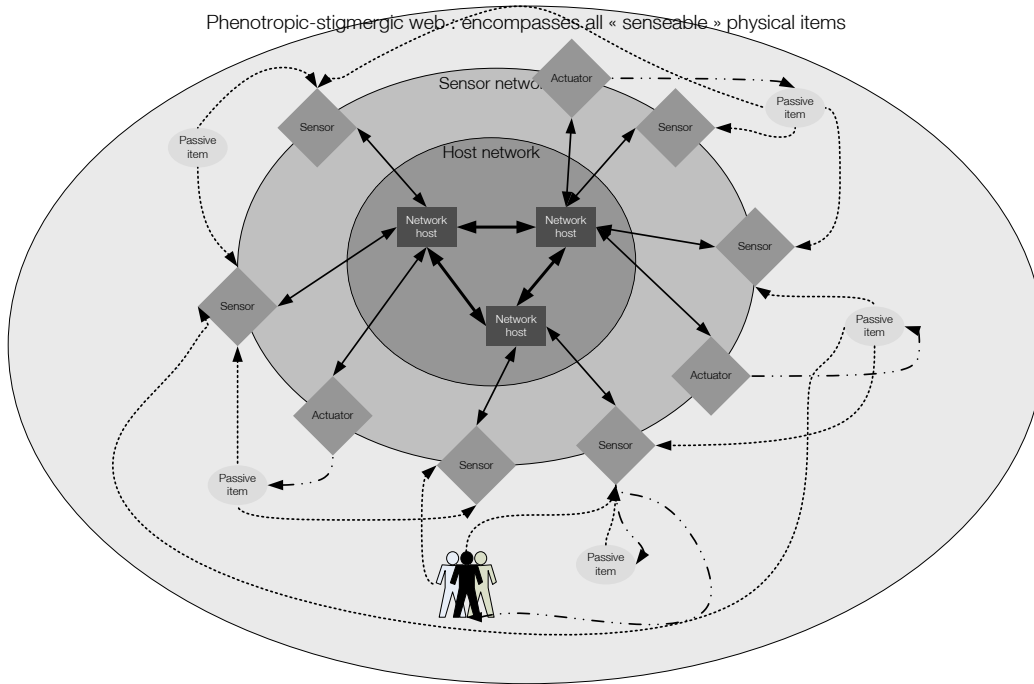


Figure 2 : "Stigmergic" network links between actionable things and the corresponding actuators and sensors

2.3. Internet of Things or Web of Things?

If we called "internet of devices" the network of sensors and actuators (devices that are attached to networks in a traditional sense), could we call the extended phenotropic network as construed above the "internet of things" proper? It should be clear that "internet" is a double misnomer to convey such a vision.

Much as the early World-Wide-Web that we know of was a virtual network of hyperlinked static HTML documents overlaid on top of the internet, the network of sensed things is itself a virtual network overlaid upon an internet of devices and sensors, which is itself an order of magnitude extension of the early internet. As a virtual network (a graph in mathematical terms), the web of things does actually comprise a far larger number of nodes than an IP network ever will and does also correspond to a different topology, as links between things correspond to their mutual "senseability" or "actionability". Another key difference is that, whereas the graph representing an IP network is non-directed, the graphs representing either the classical document-centric web, or the web of things, are *directed*.

Trying to leapfrog the marketing hype that has surrounded the so-called web 2.0, no less an authority than Vinton Cerf has proposed that the "web 3.0" should correspond to the "internet of things". "Web of things" is an alternative name that could fit the bill, hadn't this phrase already been put forward [5] to describe a different idea⁷.

2.4. Alternative Graph models

As we have said earlier, these new kinds of network links are not limited to bilateral connections between sensors and things: potentially several sensors will jointly contribute their data to identify each of these items. A weighted graph could make it possible to represent the "contribution" of a

⁷ Namely the application of lightweight RESTful protocols based on the original web for the internet of things, in lieu of the more cumbersome web services (WS-*) suite

sensor to the identification/attachment of a particular device when these devices are really on the same footing. The estimation of these weights would in practice be difficult and their evaluation arbitrary, and this may not be the best way to make their role clear. Another way to address this is to draw a distinction between

- ♦ a primary sensor device, i.e. one to which an item is virtually attached, much like an RFID tag to is temporarily attached to a reader antenna, because it provides the main identification data for this item (typically a feature that is used to characterize this item)
- ♦ ancillary sensors that do merely provide context that makes it possible to disambiguate the item characterized, but not uniquely identified by this primary feature.

3. The phenotropic web

3.1. Phenotropics : beyond protocols and sequential communication

There is a fundamental difference between classical protocol-based network links and the sensor-things links that make up the "web of sense-able things" as proposed before: these links do not rely on the previous standardization of sequential, syntactically-defined protocols, where the sending and receiving ends of a network link have to fit together like a key in a keyhole.

Jaron Lanier [6][7] has coined a new word : "phenotropics⁸", to conceptualize the special nature of these links, where the matching between both ends of the link is analog, based on global parallel pattern recognition rather than discrete and serial pattern matching as used in classical protocol-based interfaces. Pattern recognition is supposed to play the role of "interface glue" for a phenotropic networks, and to replace discrete syntactic pattern matching used in regular protocols. Phenotropic interfaces are, in Lanier's view, to be preferred because they are adaptable, extensible and bendable, whereas argument-based or protocol-based interfaces are inflexible and brittle, leading to all the problems encountered when scaling up software systems.

Applying this concept to a broader view of networks, new frontiers open up, that question our deep-seated implicit= assumptions about what a network is. This new perspective does also, potentially, give us a lead to an alternative route of evolution for the future of networks and complex software systems.

3.1.1. Phenotropics à la Jaron Lanier : surfaces vs. wires, patterns vs. linear syntaxes :

Jaron Lanier had strongly emphasized, in his original contributions on phenotropics [6] [7], what he believed was a fundamental difference with an implicit and ingrained sequential wire-based model that had molded all traditional communication models: the serial transmission of information along a wire that was originally a physical constraint had, beyond the actual existence of these very wires, gained the status of a universal metaphor that permeated all models of interfaces within a system. He contrasted this with his proposed phenotropic model where communication would be based on the parallel interaction of surfaces rather than the serial attachment of wires. We explain in the following how this parallel vs. serial distinction does actually miss the point.

A further and more significant difference is in the models and representations on which the data being transferred through these interfaces are based, and the way in which correspondence is

⁸ The dual classical greek stem of this work means literally : "appearance", (like in "pheno-type") and "turn" or "direction" like in "iso-tropic"

established between these representations on both sides of the interface. Traditional interfaces usually rely on several layers of symbolic representations and languages defined through a formal syntax, and these languages are matched through discrete and exact symbol-by-symbol pattern matching. Phenotropic interfaces do, by contrast, rely on analog iconic representations (aggrandizingly called "post-symbolic" communication by Lanier) and approximate global matching of patterns.

3.1.2. Phenotropics and the 2D/1D distinction

The surface vs. wire (parallel vs. serial) distinction highlighted by Lanier is not really the most relevant to differentiate phenotropic and classical interfaces. Though any serial interface cannot be parallelized, any parallel interface can obviously be serialized : in this sense serial interfaces are more general than parallel ones, which negates Lanier's assumption that there would be something special about surface-based parallel interfaces.

Besides, there are examples of interfaces like matrix codes (a.k.a. 2D "barcodes"⁹) that have a 2D "syntax"¹⁰ and are thus intrinsically "surface-based", yet do rely on a predefined symbology, are recognized by discrete pattern matching, not global pattern recognition, and are thus a special case of discrete syntactic interfaces, *not* phenotropic interfaces. Conversely, intrinsically serial interfaces can rely on analogue pattern matching and should be considered to be proper phenotropic interfaces. This would be the case for an audio interface that would work by recognizing a sound pattern from a temporal (i.e. 1D) sound waveform.

3.1.3. Beyond APIs and declarative interfaces

It had been advocated that traditional protocols should come to be replaced by APIs and programmatic interfaces [14]¹¹. This technological prophecy has not been entirely vindicated by the evolution of interfaces within distributed software systems : the current trend is more towards declarative interfaces (à la REST or web services), which hide the mechanics of a programmatic interface beneath a more general purpose language (such as WSDL) that is itself based on XML.

From the point of view presented here, this evolution remains within the dominant paradigm of interfaces based on a formal language, whereas phenotropic interfaces forgo this model altogether.

3.2. Extending and clarifying the pattern recognition model of Lanier's phenotropics

The difference between phenotropics and classical interfaces draws upon the theory of semiotics, which does itself subsume the classical theory of formal languages.

The use of pattern recognition does not in itself characterize phenotropic interfaces. An OCR-based interface where the text transcript of a declarative, programmatic or protocol-based interface would be recognized is not a phenotropic interface, because the pattern recognition works at the level of glyphs and the upper levels (the lexical and grammatical levels) are still handled in a classical way. As mentioned before, a 2D graphical code that uses a different kind of formal language with a 2D syntax is not phenotropic either.

⁹ Barcode is a misnomer for these codes as, unlike their 1D counterparts, they are not limited to using bars as their basic symbols. Examples are Aztec codes, cybercodes, „data matrix, QRcodes, shotcodes, semacodes, etc...

¹⁰ Syntax should be taken here to mean the arrangement of individual signs in a language with at least two levels of articulation (where a meaningful sign is made up of a combination of elemental signs). This syntax defines a global sign with semantic mapping (e.g. a morpheme) as a 2D geometric assemblage between lower-level individual signs (e.g. graphemes), rather than merely prescribing a generation/recognition mechanism mapped to a sequential arrangement of alphabet signs, as Chomskyan syntaxes do.

¹¹ The Jini infrastructure[14] was such an attempt at hiding protocols under programmatic interfaces, where the mutual adaptation between both parties was made possible by code mobility

What characterizes a phenotropic interface is that the *upper level of representation of the interface* should be analog and recognized as a global pattern. It should not need to be parsed through a syntactic analyzer! This does not preclude the existence of a lower level syntax-based representation, provided it is not used as such in the recognition process. This would be the case for e.g. recognizing statistical patterns in a text, where the lower level analysis of the text itself would not be relevant for recognition, though it could be used as an input to the recognition process.

As such this recognition should

- lend itself naturally to approximation, which confers non-brittleness and robustness properties
- bridge the semantic gap inherent in symbolic representations, without requiring the previous definition of a language at various levels (an alphabet of symbols and a syntax for the arrangement of these symbols)

3.2.1. Sensing as contextual identification

Things that can be sensed are *implicitly* identified, which means that they are identified on a non-absolute basis, only relatively, *in a given context*. It is the context itself that makes this matching amount to an identification that may become more or less explicit according to the amount and relevance of the context that can be brought to bear.

The key difference between traditional identification and the kind of sensing that we advocate here is that this recognition does not require prior knowledge of a set of codes or protocols used for communication between the sensor and the object, nor does it require a priori registration of the object in a database or its one-to-one mapping with some universally unique identifier.

An example is given in the table below through the differences between four possible means to identify an appliance : with an RFID tag, a matrix code, by recognizing its shape and texture, or by recognizing its sound patterns.

	<i>RFID</i>	<i>Matrix code</i>	<i>Shape + texture</i>	<i>Sound pattern</i>
<i>Modality</i>	Radio	optical	Optical	Audio
<i>Representation</i>	Symbolic	Symbolic	Analog/ Iconic	Analog
<i>Layers of code</i>	3 (object ID, binary sequence, modulation)	>2, depending on particular code (object ID or reference of object ID, 2D code)	0 (1 if digitized)	0 (1 if digitized)

Table 1 Differentiating non-phenotropic (columns 2-3) and phenotropic (columns 4-5) interfaces for identification of an appliance

Only the rightmost two columns of this table correspond to phenotropic interfaces proper. We will elaborate on this example in the last section of this article.

4. The stigmergic web

4.1. Stigmergic networks as distributed physical memories

4.1.1 Actuating as leaving signs to be picked up by sensors

This meaning of stigmergy extends the original concept proposed by P.P. Grassé [4] as it had already been adopted in such fields as swarm robotics and collective intelligence in general since then.

Modifications of the physical environment enacted by actuators can be of any kind, either transient (such as the emission of sound waves, electromagnetic waves, etc.) or remanent (such as leaving a mark on a surface, moving an object, etc.) and the corresponding environmental variations get sensed in turn. Remanent effects are in principle the only ones that correspond to stigmergic communication in the strict original sense, whereas signals transmitted by sound or electromagnetic waves require implicit synchronous coupling between the actuator and the sensors that detect the change. Stigmergic communication is thus, if taken in a strict sense, asynchronous, relying on a change of state of the environment and making possible a temporal decoupling of sensors and actuators that share this environment, much as different processes pass messages or concurrently access a shared memory in traditional computing paradigms.

Common usage has already tended to extend the original sense of stigmergy, especially towards communication through virtual environments or even shared wiki-like web sites. These extensions are not consistent with the original ideas inherent in the concept of stigmergy: we think it is important to restrict the word to its proper meaning of communication that is intrinsically physical, non-symbolic and non-protocol-based, by contrast to data communication (or language-based communication for human agents). Even if it does not rely on protocols or articulated language, stigmergic communication involves different kinds of implicit representations for the information shared through this "channel", which may be classified as either "sematectonic" or sign-based, qualitative or quantitative. They share the property that they are learnt through the operation of the system, relying on pattern recognition rather than on network protocols and predefined data formats.

4.1.2. Mediating access conflicts by physical laws

Coupled phenotropic-stigmergic extensions of networks view the physical world as a scratchpad read-write memory, where the read mechanism is phenotropic/associative and the write mechanism is stigmergic. This raises the question, common in concurrency and database theory, of managing write-write conflicts. As in the case of phenotropic read access, no specific protocol is required for this: physics itself provides a convenient mechanism for avoiding conflicts or arbitrating them if they appear. For example if 2 actuators try to move one and the same thing concurrently (at the same time) by applying a force to it, the resulting motion may be predicted by the law of dynamics as resulting from the vector sum of these two forces.

4.2. Stigmergic network as learning & self-organizing network

4.2.1. Stigmergic links as closing the control loop

Stigmergic links play a key role when the network is viewed from a control theory viewpoint: they are the links that transform the system represented by the sensor-actuator network from a disconnected set of *feed-forward* sensor systems and *open-loop* controllers to an overall *closed-loop*

(feedback) control system. This means that, thanks to their network connection, sensors and actuators that may have been deployed for different applications and were not meant to operate together, will "make up a system" because the sensors pick up the effect of changes to the environment effected by actuators. Of course some sensors and actuators were already configured by construction to operate in closed loop as a tightly coupled system (e.g. a robot). Stigmergic links model either these existing intended links and new, unintended relationships that are both relevant when analysing the behavior of the overall sensor-actuator network as a whole.

4.2.2. Actuating as probing the environment to integrate learning in the overall network

Modelling the overall sensor-actuator network as a closed-loop system brings the possibility of applying general machine learning theory to the network viewed as an overarching system.

Going beyond the most classical machine learning models that could apply, it is interesting to view the system in the light of such transdisciplinary concepts as enaction, embodied cognition and developmental learning, whereby the learning process of either a human (from the natural sciences viewpoint) or a robot (from an engineering viewpoint) is either analysed or engineered as directly based on physical interaction with the environment.

The idea that an infant learns a great deal about its environment by prodding this environment and absorbing the responses to the stimuli it applies, has been applied with astonishing success to robotics, opening up the whole field of "developmental learning" [8] .

Viewing an indoor environment equipped with a sensor-actuator network , i.e. a smart space, as an outside-in robot, the same idea can apply to the learning process where the smart space would probe its own environment, i.e. its inside to learn from its reactions and discover its own sensorimotor "affordances". This idea can be used for dynamic configuration of the smart space, or more mundanely, for mutual calibration of the sensors, by broadening the configuration space under which they operate.

5. Architecture and complexity issues with phenotropic-stigmergic graphs

5.1. The 3 levels of stigmergic-phenotropic webs, sensor-actuator networks and web of things/services

The IoT is potentially too far-reaching and too heterogeneous to be subsumed by a single unified networking protocol, model or architecture, at any level.

Two widespread misconceptions are the following that the IoT will be an all-IPv6 transparent and homogeneous network, and that all "things" in its reach will ultimately be identified through RFID.

The architecture that can be envisioned for the IoT as envisioned here (if we keep that name in spite of its limitations) is not directly similar to classical layered network models. It could be represented by an hourglass, where upper layers mediate as proxy nodes for those of the lower hierarchical level, and lower levels have a wider reach than the upper ones, so the mapping from upper to lower levels is one-to-many.

The uppermost level is a virtual overlay network whose nodes are software entities that can be integrated in a high-level service architecture, making up a new web of virtual entities going beyond the original document-centric web as well as the web of services that grew out of it. Some of these entities will be the digital representatives of "things" in whatever infrastructure is appropriate for this.

The second lower level corresponds more or less to the classical notion of network proper comprising all machines, devices and physical things that are integrated in a classical network, making up the internet of devices (networks hosts in a classical sense, plus embedded machine nodes and smart communicating devices comprising sensors & actuators)

The lowermost level is a virtual network of physical things that extends its reach towards all things that can be identified and sensed through sensors, making up from beneath a capillary overlay on the network of devices, with finer granularity and wider reach, as explained before.

Things are thus "hyperlinked" physically through sensors, forming a graph of sensing links , but these things also have "digital shadows", their representations that are conveyed and made accessible through the network, whatever they are (symbolic IDs such as EPC numbers, services attached to these things or devices, direct or abstracted iconic representations...).

These representations may be linked in a way that is closer to the web of services which itself grew out of the original web of documents, matching outgrowths of the physical web of things in an extended digital web. These digital shadows may correspond at the very least, for passive items, to an entry in a database, or, in the case of more active devices like sensors or actuators, to a service registered in an UDDI (Universal Description, Discovery and Integration)-compatible registry or similar networked service directory.

As the IoT extends its reach beyond the present-day internet of devices and becomes integrated with the internet of services, new applications will emerge beyond existing bread-and-butter M2M.

Present-day M2M applications are mostly one-to-one and ad hoc. Emerging applications rely on the federated use of coupled sensors and actuators in any given environment, and thus make use of federated capabilities of all sensors and actuators available in this environment rather than simple individual sensors and actuators

In this broader vision of the IoT, things or persons can get attached to networks in a dynamic and temporary fashion, relying on context to disambiguate their identity with regard to the network at large, whenever a permanent universal identity of these things or persons is either not available, not needed, or withdrawn for privacy reasons.

This opens up new applications that bring together those that had been previously addressed from the ambient intelligence viewpoint (enriched and contextual user interfaces) and those from M2M (networked sensors &actuators)

5.2. Generalizing network/services directory

A classical network or service directory integrates network nodes or services that have a permanent identity attached to them, and this identity is used as a key to register them in the corresponding registry. In the sense articulated above, a phenotropic network may integrate "things" that are only labelled in a temporary and contextual way, not identified in an absolute way. These may correspond to untagged items in such applications as inventory management, or to physical placeholders used as tangible interfaces. These things need not be matched to an absolute identifier à la EPCglobal, provided they are recognized unambiguously *in a given context of use*.

Such "phenotropic directories" could natively be queried in an associative way instead of being queried exclusively by an exactly matching, ID, key, or digital attribute. This could correspond to querying by location, by shape, or more generally by all kinds of analog pattern-like attributes.

A more far-out application of these ideas in the domain of networking could be towards the possibility to bridge two networks with incompatible protocols with some kind of "phenotropic bridge". Again, this does not mean reverting to analog networking: provided some low-level lowest common denominator digital standard could be shared, only the upper levels need be matched by pattern matching. This is actually closer to what Lanier had in mind when he originally put forward phenotropic interfaces as an alternative glue software systems.

5.3. The complex systems view

Phenotropic-stigmergic networks represent not only an order-of-magnitude quantitative leap in the number of nodes connected to the networks, but a qualitative leap in their complexity, especially due to the intricacy of the feedback loops they entail between sensors and actuators, both through the physical environment and through the network. The tools of spectral graph theory and graph-based complexity theory provide a new basis to study several complexity aspects related to these new kinds of networks, using such tools as clustering coefficients, average path lengths and degree exponents [15]. This complexity analysis is essential to uncover the potential undesirable phenomena that might emerge in these networks, and should also allow us to get a handle on their cognitive properties.

Beyond the first stage of analysing these networks, complexity models open up the possibility to optimize them according to various criteria: robustness, safety, security, adaptability, evolvability

6. Application examples

6.1. Multisensor-based registration and monitoring/control of legacy devices in an energy management system

Spontaneous ("zero-conf") integration of new devices in networks is much more than a convenience short-cut: it is often a prerequisite for systems that have to be deployed at large (e.g. in the homes of non-technical end users) without requiring the costly intervention of a skilled technician. Generalizing PC-centric "plug and play", many distributed discovery protocols have been proposed: whenever they address levels of interoperation above the basic network protocols, as is the case in service-oriented architectures, these solutions usually rest on the fact that the devices to be integrated are "known" in advance by the system in order to be recognized. Be it of a programmatic or declarative nature, the corresponding interface of the device is "recognized" by pattern matching and has to fit the interface of the reciprocating party in an exact fashion. This does not make it possible to interface with legacy devices, or even with devices whose interfaces conform to a high-level standard that is different from the one used by the host system or network. Semantic-level interoperability solutions have been proposed to circumvent this requirement for exact syntax-level matching, but they do mostly push the matching problem upwards by requiring the alignment of (possibly implicit) ontologies under which different syntaxes corresponding to the same semantics can be matched.

Phenotropic integration of non-network-enabled devices in a network or distributed system amounts to this: instead of being identified by syntactic pattern-matching of some service interface through a network, these devices are "recognized" on the basis of patterns of physical features that would be sensed by the system, using available sensors in different modalities. The system matches these observed patterns with its own stored patterns, supposed to be sufficiently generic for this. This does more or less play the role of the semantic matching that has been attempted as a

replacement for the purely syntactic matching, but is much more robust as it does not rely on a predefined standard.

We are currently implementing this analog matching as a replacement zero-conf mechanism for legacy non-network-enabled home appliances that have to be indirectly integrated in a home Local Area Network for the purposes of energy management. Multiple combined sensor modalities are used for this, and pattern recognition is performed on these joint modalities after they have undergone a binding process . We use a basic kit of sensors mounted on a multi-sensor radio mote (figure 3) comprising:

- microphone (detects noise patterns)
- vibration sensor
- temperature sensor
- light sensor
- magnetometer



Figure 3 Multisensor radio module for device monitoring

together with an electrical current sensor working via a plug inserted on the appliances mains connection. Patterns observed jointly through these sensors makes it possible to "recognize these appliances through their characteristic features (like e.g. for electric power consumption an oven showing a fairly steady plateau pattern whereas a washing machine has characteristic peaks and troughs). These appliances are then matched to a category in the system's own ontology. Contrary to a direct protocol-based semantic matching, this process has a property of graceful degradation : if the data provided by sensors is incomplete or ambiguous to match the appliance to a specific category, it is matched to a more generic one, and the system can still make do with this matching.

This multisensory pattern recognition mechanism is also used for network-enabled devices that provide only interfaces in low-level protocols, to provide a replacement for non-existent semantic matching. These legacy devices may have to fall back afterwards on a mode of operation relying on a least common denominator protocol (whatever it is) that can be shared with the rest of the system. However, the integration of these devices in the system will still be enhanced by the fact that they are recognized at a level higher than that afforded by this least common denominator protocol.

After appliances are "recognized" by the system in this way, the operation of our energy management system (requiring the monitoring of the actual state of the appliances in real-time) can still rely on sensor-based interfaces as a replacement for direct network connexion, if needed. Multisensor pattern recognition is used to recognize the instantaneous states of devices and this state can be taken into account by the system.

A similar mechanism is used to monitor rooms (or subsets of the overall target space) as entities that can be integrated in the system as distinct entities and become "peers" in the extended home LAN that also integrates the legacy devices monitored as described above. The multi-sensor module does in this case integrate (figure 4):

- ultrasonic sensors
- passive infrared sensors
- ambient microphones (coupled with voice activity detection system)
- light sensor




Figure 4 Multi-sensor radio mote for room monitoring

6.2. Coupling of an informational system and a physical system

The same idea can be generalized whenever an information system has to automatically identify a model of a physical system composed of distinct entities that may individually be represented by the system as single database entries, software components/agents or full-fledged services. These entities can be contextually recognized by the system through some pattern detected by sensors rather than through an identification system. Though this won't replace RFID or equivalent, it may be sufficient for the purpose at hand.

This may be applied in the following examples :

- ♦ monitoring/managing a fleet of vehicles (supposing they are not identical, or, if they are identical, that their individual identity does not matter to the system)
- ♦ managing an inventory of items that need not be identified individually à la RFID, only by category. If these items need to be identified individually, this could be done by differentiating them by optical codes such as 2D codes without necessarily applying the corresponding standard and without having these items registered in a database through these codes used as an absolute ID
- ♦ monitoring people in a e.g. a public place, a shopping mall, a square, a neighborhood in a way that respects their rights to privacy, i.e. by not matching them with an absolute identity that can be cross-referenced. People are only "labeled" by the system through a temporary contextual ID that makes it possible to differentiate them from other persons in the same environment, not to identify them in an absolute way.

6.3. Tangible and gesture-based interfaces

We have said that the evolution of the internet of things as proposed here is inspired from the evolution of human interfaces. Tangible user interfaces (TUI) bring the two together by recruiting everyday "things" as physical proxies for virtual entities. Tangible interfaces were initially proposed [13] as input interface alternatives to the classical "controls" associated with the WIMP-GUI interface model. They were meant to overcome a cognitive gap between device and function, as tangible controls are supposed to be dedicated (non-multiplexed) and may be directly representational, in an iconic and concrete rather than symbolic or abstract way, of the particular control functionality they support.

Most tangible interfaces proposals (usually at the concept or demo stage) rely on devices that are identified by the system in a very classical way, through either RFID tags or 2D optical codes. The ideas proposed here could very much be applied to these particular things : TUI objects need not at all be matched by the system to a unique ID, they need only be identified in their context of use, and most of them will in fact be used only temporarily.

If we wish to associate a physical item, whatever it is, with a particular functionality, we need only perform a "phenotropic association" with the TUI system. This amounts to have the item registered as a pattern through some combination of sensors (the most obvious for this is a camera, but another "weaker" modality can also be used if sufficient in context, like the weight of the item).

The association need not be limited to the static recognition of the item itself, it can be extended to actions performed with this item that can be matched with particular functionalities. This resonates with the general idea of gesture-based interfaces, which are phenotropic interfaces when the recognized gestures are not limited to a predefined symbolic repertoire. These gestures are then recognized in an analogue way, with a graceful degradation and approximation mechanism inherent in this pattern recognition.

7. Conclusions and perspective

The general ideas that are the subject of this paper have been linked to very concrete engineering examples in the applications we just described, but they may still appear to be, in their previous exposition as a general conceptual framework, very far outside the mainstream of computer science.

Yet, beyond our examples, actual digital systems with a coupled phenotropic-stigmergic interface to their environment do already exist on a very broad basis : robots work in exactly this way!

Seen from a very narrow perspective, robotics could seem to lag behind computer science. Mainstream robotics is still dominated by an archetype of self-contained, stand-alone contraptions, whose connectivity is still at the stage of the pre-internet PC industry, exploiting only marginally, at best, the potential for network-based operation. Robotics software has yet to move beyond closed platforms and to adopt generic high-level models such as service-oriented architectures that would make it possible for them to interoperate in networked environments. Distributed robotics has yet to become something else than an oxymoron, moving beyond basic remote control, as used mostly for industrial robotic equipments, a stage of evolution corresponding to early networked computing of the client-server kind.

Yet for what concerns the discovery of their own environment, advanced robotic systems are way ahead of anything that has been attempted by mainstream computer engineering. They can be set to run in such an environment and “phenotropically” discover it without any previous manual configuration, possibly by “stimergically” prodding it extensively. They do not need to identify obstacles with RFID tags to avoid bumping into them, or to run an association protocol with objects before grasping them.

Computer science has a few things to learn from robotics in these regards, and what we propose is a way to carry this beyond traditional robotics... Robots are not only good at recognizing their environments, we could even suppose that, when they become networked for good, they will be better at recognizing each other than regular parts of a distributed system currently are, and precisely because this recognition process will be phenotropic. Maybe not quite in the sense that Jaron Lanier originally envisioned, where serial 1D wire-based communication would be *replaced* by some parallel 2D surface-based form of communication. Phenotropic and stigmergic communication between two robots is already built-in as communication mediated by the physical environment, supported by the cross-coupling of their respective sensors and actuators. So maybe, contrary to what Lanier thought, it is not that much of a problem if they use standard wireless networks and their cumbersome protocols as a regular means of communication, because the analog pattern-recognition-based robustness and graceful degradation properties will be there nonetheless, supported by the parallel mode of physical communication that is inherent in the nature of robots.

Generalizing these ideas from networked robots to distributed embedded systems, all of which are also composed, deep down, of doubly coupled sensors and actuators, leads us beyond the simple graph-based models from which we started. The conclusive idea we can draw from this is that phenotropic-stigmergic communication, mediated by the physical environment through sensors and actuators, may become just as important as digital communication to allow the self-configuration, analyse the complexity and ensure the robustness of distributed embedded systems. This new research agenda deserves to be addressed with competences from graph-based complexity theory, cognitive sciences, control theory, as well as system architecture.

References

- [1] Floerkemeier, C.; Langheinrich, M.; Fleisch, E.; Mattern, F.; Sarma, S.E. (Eds.), *The Internet of Things*, Proceedings First International Conference, IOT 2008, Zurich, Switzerland, Lecture Notes in Computer Science vol.4952, Springer Verlag, 2008.
- [2] Gabriel, R. P. & Goldman Ron "Conscientious software", *Proceedings OOPSLA'06*, Portland, Oregon, October 2006.
- [3] Gershenfeld, N., Krikorian, R. Cohen, "The internet of Things", *Scientific American*, October 2004
- [4] Grassé, P-P. : "La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs", *Insectes Sociaux*, Volume 6, n°1, mars 1959
- [5] Guinard D. and Trifa V.. "Towards the web of things: Web mashups for embedded devices", MEM 2009 in *Proceedings of WWW 2009*. AM, 2009.
- [6] Lanier, J. : "Why Gordian Software Has Convinced Me To Believe In The Reality Of Cats And Apples" *edge.org*, November 1, 2003
- [7] Lanier, J. : "The Complexity Ceiling," in *The Next Fifty Years : Science In The First Half Of The Twenty-First Century*, John Brockman (Editor,) Vintage, May 2002
- [8] Oudeyer Pierre-Yves, & Kaplan Frédéric and Verena V. Hafner and Andrew Whyte : "The Playground Experiment:Task-Independent Development of a Curious Robot", Lecture Notes in Computers Science, Springer, 2004
- [9] Privat, G. " From Smart Devices to Ambient Communication", Invited conference, *From RFID to the internet of Things*, EU Workshop Brussels, March 3 2006., ftp://ftp.cordis.europa.eu/pub/ist/docs/ka4/au_conf670306_privat_en.pdf
- [10] Reeves, B. and Nass, C. 2000. Perceptual user interfaces: perceptual bandwidth. *Commun. ACM* 43, 3 (Mar. 2000), 65-70.
- [11] Streitz, N., Privat G. "Ambient Intelligence", Chapter 60 of *The Universal Access Handbook*, CRC Press, 2009
- [12] Theraulaz, G., Bonabeau, E., A Brief History of Stigmergy, *Artificial Life*, Spring 1999, Vol. 5, No. 2, Pages 97-116
- [13] Ullmer, B. and Ishii, H.. Emerging Frameworks for Tangible user interfaces. *IBM Systems Journal* 39(3&4) 2000: 915-931.
- [14] Waldo, J. "The End of Protocols or Middleware for the New Millennium", Keynote to Middleware 2000
- [15] Watts, D., *Small Worlds, The dynamics of networks between order and randomness*, Princeton University Press, 1999