

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Version accepted for publication in IEEE Systems, Man, and Cybernetics Magazine.

# Self-Organizing Software Models for the Internet of Things

Damian Arellanes

**Abstract**—The Internet of Things (IoT) envisions the integration of physical objects into software systems for automating crucial aspects of our lives, such as healthcare, security, agriculture, and city management. Although the vision is promising, with the rapid advancement of hardware and communication technologies, IoT systems are becoming increasingly dynamic, large, and complex to the extent that manual management becomes infeasible. Thus, it is of paramount importance to provide software engineering foundations for constructing autonomic IoT systems. In this paper, we introduce a novel paradigm referred to as *self-organizing software models* in which IoT software systems are not explicitly programmed, but emerge in a decentralized manner during system operation, with minimal or without human intervention. We particularly present an overview of these models by including their definition, motivation, research challenges, and potential directions.

**Index Terms**—Self-organization, Internet of Things, Cyber-Physical Systems, Self-Composition of Software, Autonomous Systems.

## I. INTRODUCTION

THE Internet of Things (IoT) is a novel paradigm, considered the next Industrial Revolution, that promises the integration of every physical object for the automation of essential concerns of our modern lives such as healthcare, security, agriculture, and city management. Unlike traditional enterprise systems, IoT systems are moving towards environments full of complex interactions, as a consequence of the overwhelming number of objects available worldwide [1]–[3]. Currently, there are over 17 billion connected objects, and it is estimated that this number will increase exponentially in the coming years [4]. Hence, scalability and complexity become a significant challenge for the full realization of IoT.

Most IoT research has extensively focused on hardware and network issues so that early IoT systems operate in closed environments and integrate relatively few static software components. Contrastingly, future software-intensive IoT systems will be deployed in open environments (i.e., software ecosystems [5]) where billions of (off-the-shelf) components will abstract the functionality of an immense number of connected physical objects [4], [6]. Such environments will be highly dynamic and uncertain due to disturbances caused by external perturbations (e.g., change in requirements and increasing workloads) and unforeseeable internal situations (e.g., system failures and sub-optimal behaviors) [4], [7].

Autonomicity represents the most viable solution to manage complex IoT systems that both integrate an ultra-large number of software components and operate in highly dynamic,

uncertain environments. This is because that property enables the adaptation of computational behaviors with minimal or no human intervention. Research on autonomic software has produced significant results, especially in the area of self-adaptive software [7]. However, existing solutions are mainly centralized, and it has been proven over many years that centralized approaches do not scale and are therefore unsuitable to tackle complexity [8].

In this paper, we present *self-organizing software models* which are a new kind of abstractions that allow the construction of autonomous software systems, in which computational behaviors are not predefined but emerge during system operation to dynamically accommodate a given context. Emergence is achieved from the individual interactions of the constituent software components (not hardware devices), without the need of a central authority. The main role of these models is to remove or reduce the role of a programmer in the composition, maintenance and evolution processes of a software system.

## II. WHAT IS A SELF-ORGANIZING SOFTWARE MODEL?

Self-organization is the bottom-up process by which complex behaviors emerge from the decentralized interactions of participant components (e.g., molecules or insects), in order to collectively achieve a global system goal (e.g., foraging). In contrast to top-down processes, self-organization is a well-known technique to deal with uncertainty, scale, dynamism, and complexity [8], [9]. Self-organization is not a new concept. It has been studied in diverse areas from distinct points of view, from biological systems (e.g., flock of birds, school of fish, and ant colonies) to artificial systems (e.g., traffic light ensembles, networking, and swarm robotics). Just recently, self-organization has captured the attention of the software engineering community to study it as an inherent property of software models.

A software model is an abstract system representation that describes software components and their composition [10]. A software component (e.g., a web service or a generic port-based component) is a self-contained unit of composition that provides some computational functionality. In traditional software engineering, composition is the design-time process of combining the functionality of two or more software components, and it is performed manually by system engineers [10], [11]. As manual composition is unsuitable to tackle the imminent challenges that future IoT systems pose (i.e., dynamicity, complexity, and scale), we envision that such an approach will eventually be obsolete. Instead, composition will be a run-time process performed by the model itself (i.e., autonomously), in which complex computational behaviors emerge in the form of

D. Arellanes is with the School of Computing and Communications, Lancaster University, Lancaster LA1 4YW, UK (e-mail: damian.arellanes@lancaster.ac.uk).

complex composite components. These emergent composites can further self-organize to define even more complex composites. Emergence occurs from the decentralized interactions of the available autonomous components (potentially developed independently by different stakeholders), with no or minimal human intervention and according to a set of self-organization rules. As emergent composites cannot be expressed as a simple summation of the composed components, emergent computational behaviors cannot be predicted just by knowing the available components.

Our vision of self-organizing software models is depicted in Figure 1. Although we use a port-based composite for illustrative purposes, other software compositions can emerge as a result of self-organizing component interactions, such as service-oriented workflows [12], [13] or algebraic compositions [14], [15]. In any case, emergent composites lie on top of a three-layer IoT view, physical objects (known as *things*) are situated at the bottom and self-organizing interactions occur in the middle. This conceptual three-layer view shall be referred to as *self-organizing IoT*.

The idea of self-organizing software contrasts with that of self-adaptive systems. This is because the latter often require adaptation managers to control the entire adaptation process outside a software model [7]. We refer to this process as exogenous adaptation. By contrast, in our vision, adaptation is achieved through self-organization, and it therefore occurs in a pure decentralized manner without the need of any central authority, external controllers, or leaders. We refer to this process

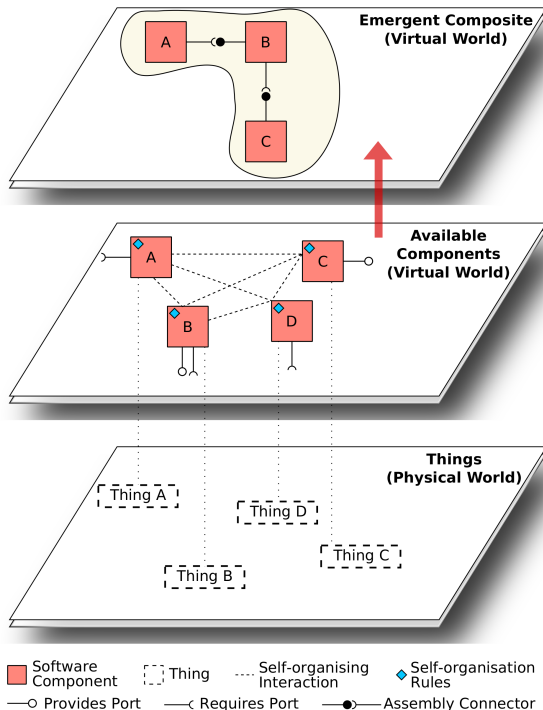


Fig. 1. Self-organizing IoT.

*“A self-organizing software model is a computational abstraction whose software components comply with self-organization rules, so complex composite components are not explicitly programmed but (autonomously) emerge from the decentralized interactions of the available, independent software components. Its overall goal is to accommodate perturbations in a system operating environment.”*

as endogenous adaptation. Unlike self-adaptive software, in self-organizing software models there is no notion of *managed system* or *managing system* since components, which belong to the system itself, collaboratively realize adaptation by the emergence of computational behaviors. Like self-adaptive

software, there are adaptation triggers to initiate the emergence of composite components upon detecting perturbations in the internal or the external system operating environment. Perturbations in the internal environment include system failures and sub-optimal performance. Perturbations in the external environment include change in stakeholder requirements, run-time scaling (i.e., component addition/removal), environmental changes, and increasing workloads [4], [7].

To illustrate how self-organizing software components deal with open environments, let us consider the example depicted in Figure 2. At time  $t_0$ , four components interact to meet the requirements  $R1$ ,  $R2$ , and  $R3$ . In this case, the result of self-organization is a composite assembling  $A$ ,  $B$ , and  $C$ . Suddenly, at time  $t_1$ , the requirement  $R3$  is no longer needed, and component  $E$  becomes available. As requirements changed and component  $E$  offers better performance than  $B$ , the existing components self-organize to connect  $C$  with  $D$  and replace  $B$  with  $E$ . Note that component  $C$  has mutated by adding a new port, in order to fulfill the requirements  $R1$  and  $R2$ . Finally, in the last time window, the requirements  $R1$ ,  $R2$ , and  $R3$  become obsolete since a new specification is defined. So, components self-organize once again to compose an entirely new computational structure. In any time window, composite components emerge without any central controller or leader.

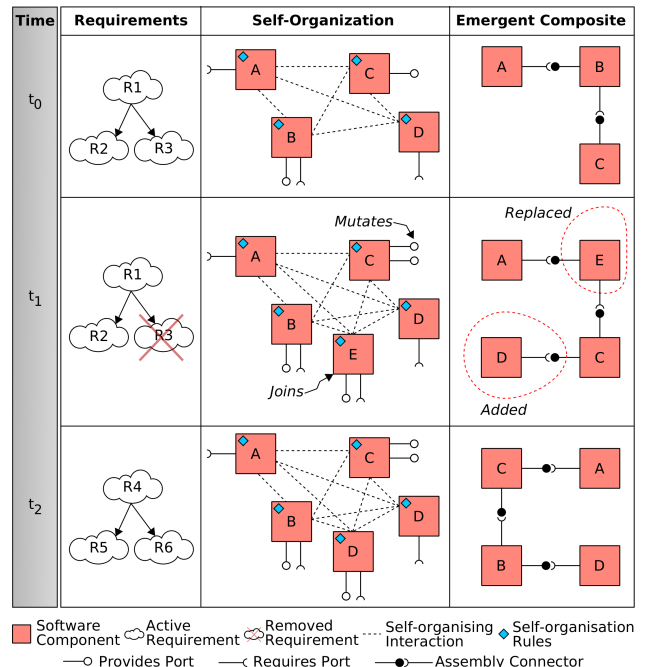


Fig. 2. An example of a self-organizing software model.

### III. WHY SELF-ORGANIZING SOFTWARE MODELS?

With the rapid advancement of hardware and communication technologies, IoT systems are becoming increasingly dynamic, functionally large, and extremely complex to the extent that manual management becomes infeasible. Dynamism, or churn [4], makes *things* (and their software components) to constantly appear and disappear in completely uncertain network environments [13]. This can be a consequence of mobility, failures, or poor network connections. Functional scalability [4] is another problem, which arises from the fact that the functionality of one *thing* can be virtualized by more than one software component, and there are plenty of *things* available. In fact, the more components available, the more possible computational behaviors (potentially leading to a combinatorial explosion problem [14], [15]). Last but not least, complexity refers to the number of interactions between software components [2], [8], [11], and it is closely related to functional scalability since the more software components composed, the more complex a system is. To tackle these imminent challenges, it is therefore of paramount importance to provide software engineering theory that facilitates the construction of autonomic IoT software systems.

Most of the research done in the field of autonomic computing is built upon centralized adaptation managers (e.g., MAPE-K control loop [16]). However, due to the law of requisite variety [17], which refers to an explosion in the number of system states, such top-down autonomic solutions are unsuitable for dealing with realistic, open environments like IoT [18], [19].

Since self-organization is a well-known bottom-up approach that deals with precisely the challenges that IoT faces [9], some preliminary endeavors have been done to apply self-organization principles in the IoT realm. However, most of this work has been done from the perspective of general systems engineering (e.g., [18]) rather than from a software engineering viewpoint. Just a few works have been devised in the context of self-organizing software models [12], [13], [19]–[22].

Because IoT systems are becoming increasingly software-intensive, we need to define self-organization rules in the semantics of software components. Otherwise, self-organization does not occur among software components but among something else (e.g., network nodes [23] or robots [24]). We need to leverage the flexibility and manageability that software offers versus hardware (cf., software-defined networks [25]) in order to adapt IoT systems to different operation contexts. In the end, it is software that controls embedded systems in *things*, and it is because of this that self-organization always originates in the virtual world. Embedded systems are just interfaces to the physical reality.

As they provide the mechanisms to compose/emerge complex computational behaviors on the fly without any central adaptation entity, we see self-organizing software models as an important contribution to the field of autonomic computing.

### IV. RESEARCH CHALLENGES AND POTENTIAL DIRECTIONS

Due to its non-deterministic nature, the main challenge of self-organization is to emerge meaningful functionality from decentralized interactions. As the concept of meaningful varies

from one domain to another, goal models (e.g., RELAX [26] and KAOS [27]) can be used to specify expected system-wide behaviors in the form of *requirements@runtime* (which must be met without any central reasoner). In addition to goal models, run-time verification/testing techniques can be integrated to prove the correctness of

emergent computational behaviors. All of this without sacrificing systems' operation. Apart from dealing with *non-determinism* and *correctness*, other challenges are as follows:

- *Incompleteness*. Due to their decentralized nature, software components cannot always have a complete, consistent view of their operating environment (e.g., knowledge about other available components or knowledge about running composition structures). This is especially true in IoT ecosystems with a large number of IoT software components. So, How to efficiently disseminate knowledge to maintain an accurate representation of the world under highly dynamic environments?
- *Self-explanation*. Transparency of decision-making is a must in approaches lacking human intervention, since users may require an explanation of the emergence of certain computational behavior. So, How to provide self-explanation in self-organizing software models, especially when the participant components are completely autonomous and operate in different administrative domains?
- *Measurement*. To date, there is a huge body of research on self-organizing systems in which plenty of different evaluation metrics have been proposed, e.g., entropy, fragility, and stability [18]. Are these metrics suitable to evaluate computational emergence in self-organizing software models? If not, what are the most suitable metrics for the software engineering domain?
- *Uncertainty quantification*. Since IoT systems operate in environments full of *aleatoric* and *epistemic* uncertainty, it is impossible to predict the space of complex programs that arise from dynamically appearing and disappearing software components. To control the construction of such spaces, uncertainty quantification approaches can be used for defining trustworthy decision-making mechanisms that proactively adjust the way components interact. What are the semantic constructs that allow software components to collaboratively solve uncertainty quantification problems?
- *Dynamic evolution of self-organization rules*. Self-organization rules can be predefined at design-time. However, in certain scenarios, the rules might not be enough to achieve global system goals. So, Is it possible to define decentralized learning techniques (e.g., collaborative/multi-agent reinforcement learning) for dy-

namically evolving self-organization rules without human intervention?

- *Evolution reasoning.* Self-organizing software models are always evolving at run-time, so it becomes necessary to reason about dynamics. Modeling evolution in traditional self-organizing systems can be done using process algebra, temporal logic, or Petri Nets. But are these techniques suitable to reason about self-evolving software? If not, which reasoning techniques could be more appropriate?
- *Software semantics and self-organization rules.* Defining generic self-organization rules for autonomously composing software in different IoT domains is challenging. This raises the question of what are the most suitable self-organization rules for IoT software models? Can we embed those rules in the semantics of a (universal) component model? Can we take inspiration from self-organizing biological systems to define bio-inspired component models?

## V. CONCLUSION

Although the dream of removing the tasks performed by software engineers is far from reality, self-organization offers an encouraging route towards the next generation of software systems in which software is not explicitly programmed, but emerges without any central controller to meet the needs of a given context at run-time. We refer to this class of abstractions as *self-organizing software models*. In this paper, we presented the definition, motivation, challenges and future directions of these abstractions.

To date, self-organizing software models are still in their infancy and their challenges hinder the development of a self-organizing software solution applicable to a wide variety of IoT domains. We envision that this nascent field will be of great relevance in the coming years to deal with the inherent scale, uncertainty, dynamism, and complexity that IoT systems are increasingly posing.

## REFERENCES

- [1] C. Janiesch, A. Koschmider, M. Mecella, B. Weber, A. Burattin, C. Di Ciccio, G. Fortino, A. Gal, U. Kannengiesser, F. Leotta, F. Mannhardt, A. Marrella, J. Mendling, A. Oberweis, M. Reichert, S. Rinderle-Ma, E. Serral, W. Song, J. Su, V. Torres, M. Weidlich, M. Weske, and L. Zhang, "The Internet of Things Meets Business Process Management: A Manifesto," *IEEE Systems, Man, and Cybernetics Magazine*, vol. 6, no. 4, pp. 34–44, 2020.
- [2] D. Arellanes and K.-K. Lau, "Analysis and Classification of Service Interactions for the Scalability of the Internet of Things," in *IEEE ICIOT*, 2018, pp. 80–87.
- [3] G. Fortino, "Agents Meet the IoT: Toward Ecosystems of Networked Smart Objects," *IEEE Systems, Man, and Cybernetics Magazine*, vol. 2, no. 2, pp. 43–47, 2016.
- [4] D. Arellanes and K.-K. Lau, "Evaluating IoT service composition mechanisms for the scalability of IoT systems," *Future Generation Computer Systems*, vol. 108, pp. 827–848, 2020.
- [5] A. Bröring, S. Schmid, C.-K. Schindhelm, A. Khelil, S. Käbisch, D. Kramer, D. Le Phuoc, J. Mitic, D. Anicic, and E. Teniente, "Enabling IoT Ecosystems through Platform Interoperability," *IEEE Software*, vol. 34, no. 1, pp. 54–61, 2017.
- [6] R. Want, B. N. Schilit, and S. Jenson, "Enabling the Internet of Things," *Computer*, vol. 48, no. 1, pp. 28–35, 2015.
- [7] D. Weyns, "Software Engineering of Self-adaptive Systems," in *Handbook of Software Engineering*, S. Cha, R. N. Taylor, and K. Kang, Eds. Cham: Springer, 2019, pp. 399–443.
- [8] F. Heylighen and C. Gershenson, "The Meaning of Self-organization in Computing," *IEEE Intelligent Systems*, pp. 1–6, 2003.
- [9] W. R. Ashby, "Principles of the Self-Organizing System," in *Principles of Self-Organization*, H. Von Foerster and G. W. Zopf, Eds. London, UK: Pergamon Press, 1962, pp. 255–278.
- [10] K.-K. Lau and S. Di Cola, *An Introduction to Component-based Software Development*, 1st ed. Singapore: World Scientific, 2017.
- [11] M. Lehman, "Software's future: managing evolution," *IEEE Software*, vol. 15, no. 1, pp. 40–44, 1998.
- [12] H. Ben Mahfoudh, G. Di Marzo Serugendo, N. Naja, and N. Abdennadher, "Learning-based coordination model for spontaneous self-composition of reliable services in a distributed system," *International Journal on Software Tools for Technology Transfer*, vol. 22, no. 4, pp. 417–436, 2020.
- [13] N. Chen, N. Cardozo, and S. Clarke, "Goal-Driven Service Composition in Mobile and Pervasive Computing," *IEEE Transactions on Services Computing*, vol. 11, no. 1, pp. 49–62, 2018.
- [14] D. Arellanes and K.-K. Lau, "Exogenous Connectors for Hierarchical Service Composition," in *International Conference on Service-Oriented Computing and Applications (SOCA)*. IEEE, 2017, pp. 125–132.
- [15] D. Arellanes and K.-K. Lau, "Workflow Variability for Autonomic IoT Systems," in *International Conference on Autonomic Computing (ICAC)*. IEEE, 2019, pp. 24–30.
- [16] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [17] W. R. Ashby, "Requisite Variety and Its Implications for the Control of Complex Systems," *Cybernetica*, vol. 1, no. 2, pp. 83–99, 1958.
- [18] C. Gershenson, "Guiding the Self-Organization of Cyber-Physical Systems," *Frontiers in Robotics and AI*, vol. 7, pp. 1–41, 2020.
- [19] I. Georgiadis, J. Magee, and J. Kramer, "Self-organising software architectures for distributed systems," in *ACM SIGSOFT Workshop on Self-Healing Systems (WOSS)*. ACM, 2002, pp. 33–38.
- [20] T. Bures, F. Plasil, M. Kit, P. Tuma, and N. Hoch, "Software Abstractions for Component Interaction in the Internet of Things," *Computer*, vol. 49, no. 12, pp. 50–59, 2016.
- [21] N. Cardozo, "Emergent software services," in *International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!)*. ACM, 2016, pp. 15–28.
- [22] J. Dowling and V. Cahill, "Self-managed decentralised systems using K-components and collaborative reinforcement learning," in *ACM SIGSOFT Workshop on Self-Managed Systems (WOSS)*. ACM, 2004, pp. 39–43.
- [23] M. Mamei, A. Roli, and F. Zambonelli, "Emergence and control of macro-spatial structures in perturbed cellular automata, and implications for pervasive computing systems," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 35, no. 3, pp. 337–348, 2005.
- [24] G. Baldassarre, V. Trianni, M. Bonani, F. Mondada, M. Dorigo, and S. Nolfi, "Self-Organized Coordinated Motion in Groups of Physically Connected Robots," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 37, no. 1, pp. 224–239, 2007.
- [25] E. E. Hayek, I. G. Ben Yahia, D. Arellanes, and K.-K. Lau, "Analysis of component-based approaches toward componentized 5G," in *International Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*. IEEE, 2018, pp. 1–5.
- [26] J. Whittle, P. Sawyer, N. Bencomo, B. H. C. Cheng, and J.-M. Bruel, "RELAX: a language to address uncertainty in self-adaptive systems requirement," *Requirements Engineering*, vol. 15, no. 2, pp. 177–196, 2010.
- [27] A. van Lamsweerde, R. Darimont, and E. Letier, "Managing conflicts in goal-driven requirements engineering," *IEEE Transactions on Software Engineering*, vol. 24, no. 11, pp. 908–926, 1998.

**Damian Arellanes** (damian.arellanes@lancaster.ac.uk) is an Assistant Professor in Computer Science at the School of Computing and Communications, Lancaster University, UK. He received his PhD degree in Computer Science from the University of Manchester, UK. His research interests lie at the intersection of models of computation, Cyber-Physical Systems, and autonomic computing, with a particular emphasis on self-organizing software models.