

# Evolving, Growing, and Gardening Cyber-Physical Systems

Susan Stepney<sup>1</sup>, Ada Diaconescu<sup>2</sup>, René Doursat<sup>3</sup>,  
Jean-Louis Giavitto<sup>4</sup>, Julian F. Miller<sup>1</sup>, Antoine Spicher<sup>5</sup>

<sup>1</sup>York Cross-disciplinary Centre for Systems Analysis, University of York, UK

<sup>2</sup>CNRS LTCI, Télécom-ParisTech, France

<sup>3</sup>School of Computing, Mathematics and Digital Technology,  
Manchester Metropolitan University, UK

<sup>4</sup>UMR STMS 9912, IRCAM - CNRS, France

<sup>5</sup>LACL, Université Paris-Est Créteil, France

## 1 Introduction

Our artefacts, from small devices to buildings and cities, are, or are becoming, cyber-physical socio-technical systems, with tightly interwoven physical material and computational parts. We build such systems, laboriously placing material components, laboriously programming computational ones, laboriously integrating the parts, laboriously maintaining the resulting structures. In contrast, trees grow, adapting their form and function to the environmental conditions, and trees self-repair, using the same mechanisms as for growth. These properties allow trees to be gardened—planted, fed, pruned, trained—to meet human needs.

Many authors have expressed the desire to replace our current laborious classical engineering with a more agile “self-\*” approach, for example (Abelson, Allen, Coore, Hanson, Homsy, Knight, Nagpal, Rauch, Sussman & Weiss, 2000; Doursat, Sayama & Michel, 2012; Lalanda, McCann & Diaconescu, 2013; Würtz, 2008). Our vision here is of *construction by directed growth*, through gardening macroscopic cyber-physical artefacts formed from a growing, integrated combination of material and virtual subsystems. This is a novel approach to designing, implementing, and maintaining the wealth of cyber-physical artefacts that comprise our “built environment”, including: housing, schools, hospitals, shops, and factories; transport, power, and communication infrastructures. The objective is an unconventional embodied computational process (Stepney, Rasmussen & Amos, 2018) that will produce autonomous, adaptive, robust systems in a controllable and cost-effective manner.

The overall objective of computational growth is to be able to “grow” macroscopic objects that comprise both physical and computational aspects (cyber-physical systems, “smart” materials, etc). The related field of physical self-assembly focuses on the material side: evolving or designing nano- and micro-scale assemblers, or micro-scale artificial cells, that can manipulate and position matter. The computational control of such devices (how and where to position

---

Susan Stepney, Ada Diaconescu, René Doursat, Jean-Louis Giavitto, Julian F. Miller, Antoine Spicher. Evolving, Growing, and Gardening Cyber-Physical Systems. Rachel Armstrong (ed) *Experimental Architecture*, pp.89–101, Routledge, 2019

the relevant material; how and when to divide and replicate) is typically assumed to be a relatively trivial matter of programming. Little consideration is given to scalability, adaptability, correctness, or other computer science issues relating to the macro scale of cyber-physical artefacts. Here we consider the matters in the context of an unconventional computational framework.

The structure of rest of the chapter is as follows. In §2 we introduce our proposed GRO-CYPHY architecture for designing, growing, and gardening complex cyber-physical systems. In §3 we outline how this process might work for a substantial application: a skyscraper. In §4 we conclude with a discussion of how the GRO-CYPHY approach changes our ideas about artefact construction.

## 2 GRO-CYPHY Architecture

### 2.1 Overview


There are several visions of assembling or growing macroscopic artefacts from the bottom up. These include nanoscale robots manipulating material at the molecular level (Drexler, 1986), which are an artificial analogue of termite and other social insect *construction* processes, and synthetic protocells dividing and differentiating (Armstrong, 2014; Solé, Munteanu, Rodriguez-Caso & Macía, 2007), which are an artificial analogue of plant and animal *growth* processes.

Past work on exploiting growth metaphors has tended to concentrate on either purely virtual or purely physical structures. Concentrating on only the virtual loses the advantages of embodied physical reality (Stepney, 2007, 2008), and allows issues of reality constraints, such as the cost of replication and other such operations, to be neglected. Concentrating on only the physical ignores the computational issues of programming the desired growth, and the necessity of integrating physical and virtual subsystems.

Here we take the software perspective, and outline a computational growth process that, given the necessary physical assembly devices (nanobots or protocells), can program these to construct macroscopic cyber-physical structures. Additionally, the resulting growth process technology will be able to grow purely virtual computational structures in virtual environments, such as those running on computer networks and clouds.

In the current absence of such physical assembly devices, initial work is necessarily in simulation. However, the aim is for a process where a given program could either be embodied in a physical constructor that would provide the growth mechanisms, or placed in a virtual “growth engine” that would simulate the growth process. This dual approach also allows close integration and co-growth of purely virtual software components such as software control systems.

The focus of our approach is on the design and growth of a complex artefact with several interrelated subsystems. These subsystems must co-develop to produce an integrated whole. Additionally, we want our cyber-physical artefacts to be responsive and adaptive to their environments. For a growth approach, it makes sense to consider parallels with plants rather than with animals: plants



**Fig. 1.** Overview of the conceptual GRO-CYPHY architecture in application (here, a skyscraper): application-specific subsystem seed specifications are developed (wall, windows, plumbing, etc); the *Seed Factory* develops the relevant seeds (subsystem genomes); *Growth Engines* grow these seeds in the *Computational Garden* where the application specific gardening takes place, delivering the cyber-physical system end product. See following text for a full explanation of the various components and their relationships.

are more “plastic” in their morphological responses. There is evidence that this sort of plasticity can be exploited in artificial developmental systems (Kowaliw & Banzhaf, 2009; Kowaliw, Grogono & Kharma, 2007a; Tufte & Haddow, 2007). This plasticity admits a new possibility: deliberately and selectively applying environmental stimuli to direct growth in a desired direction—or *gardening* the growing artefact. To this end, we have designed the conceptual GRO-CYPHY architecture, which comprises three major components (figure 1):

- a *Seed Factory*, a process for evolving and designing specific computational seeds to meet cyber-physical system requirements
- a *Growth Engine*, providing the computational processes that grow physical seeds in simulation, and grow virtual seeds into software
- a *Computational Garden*, where multiple seeds can be planted and grown in concert, where virtual seeds can be interfaced with embodied growth processes, and where a high-level gardener can shape the whole into complex cyber-physical systems.

The gardening metaphor guiding this approach is an attempt at drastically changing the way we build cyber-physical systems. Traditionally, such systems

are produced by a tightly coupled team of engineers following a life cycle that separates phases of design, production and use. In contrast, plants develop continuously. As a “product”, a plant already has an innate ability to seamlessly adapt to its environmental conditions (within reasonable limits). Additionally, the shape and metabolism of a plant—its “purpose”—can be artificially adjusted and tailored at various stages:

1. before the growth phase, by choosing an adequate genome: this genome is typically programmed or designed by directed evolution (in the Seed Factory) which does not require an understanding of the lower-level processes (implemented in the Growth Engine)
2. at the beginning of the growth phase, by configuring the initial conditions of the seed, including the internal “metabolism” (growth preferences and parameter values), and possibly by introducing a pre-configured collection of initial components (a “seedling”)
3. during the growth phase, by taking into account the various actions of the gardener and the interactions with the environment.

Even if the organism reaches a steady state, the growth phase never stops: there is no distinction between the “building phase” of the artefact, the “functional phase”, or the “repair and maintenance” phase.

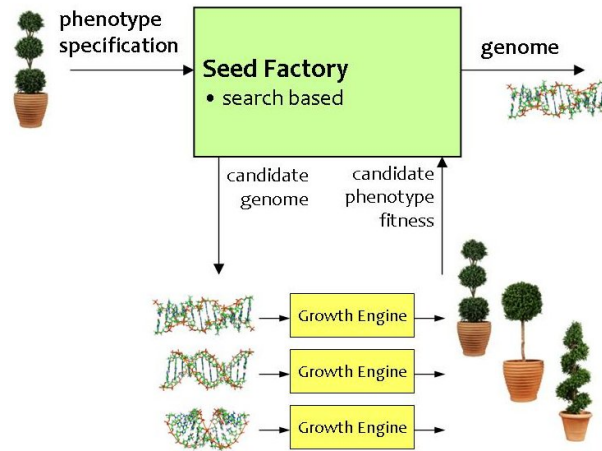
## 2.2 Seed Factory

The Seed Factory develops the seeds (genomes, programs) that are to be implanted either in physical devices (nanobots, protocells) or in virtual Growth Engines, to grow into the desired artefacts.

Seeds are tightly packed data structures that contain the necessary instructions (genome  $G$ , comprising the local rules of growth and self-assembly followed by the components at the microscopic level) and initial conditions (internal state  $S$ ) needed to grow a given target structure (the organism phenotype  $P$ , being the global structure and function of the emergent system at the macroscopic level). By “tightly packed” we mean that the actual phenotypic shape adopted after growth is not explicitly represented in the seed, but will unfold and reveal itself as the seed develops (Stanley & Miikkulainen, 2003).

Several underlying technologies for defining and implementing seeds could be supported simultaneously in a GRO-CYPHY system, and should be chosen as appropriate for growing different kinds of structures. Those for physical technologies such as nanobots or protocells need to be tailored for the physical capabilities of the embodying systems. Example of virtual technologies include

- morphogenetic engineering (ME) (Debbabi, Diaconescu & Lalanda, 2012; Diaconescu, Tomforde & Müller-Schloer, 2018; Doursat, 2008; Doursat, Sayama & Michel, 2012; Sayama, 2009);
- developmental cellular automata (DCA) (Miller, 2004; Miller & Thomson, 2004) ;



**Fig. 2.** Conceptual overview of the search process in the Seed Factory: high-level phenotype (grown) specifications are input; the search process develops the relevant seeds (subsystem genomes); it uses the Growth Engine to grow candidate seeds into phenotypes, which it evaluates against the specification, and feeds the information back into its search process.

- generative systems (GS) (Hornby & Pollack, 2001; Jacob, 1994, 1996; Mock, 1998; Monks, Oh & Dorsey, 2000; Shea & Smith, 2006);
- self-modifying Cartesian Genetic Programming (SMCGP) (Harding, Miller & Banzhaf, 2010; Miller, 2011);
- spatial computing (SC) (DeHon, Giavitto & Gruau, 2007; Maignan & Spicher, 2015; Spicher & Giavitto, 2017; Spicher & Michel, 2005; Spicher, Michel & Giavitto, 2010).

Several seed design methodologies are applicable to various types of computational cyber-physical structures. For example, simple to moderately sophisticated structures could grow out of simple seeds produced by a composition of elementary “rules of thumb” for seed design (Beal, 2011; Werfel, 2010). More elaborate organisms could be difficult to engineer “by hand” in this way, and would require automatic search and optimisation methods to sweep a potentially vast space of possible seeds, in search of good ones able to grow into the desired structures (figure 2).

When “planted” in the computational garden and grown by the relevant Growth Engine, seeds should reliably provide the intended functionality, in a sufficiently flexible and fault-tolerant way such that a gardener can steer their growth towards customised organisms. The resulting organisms should have the features of genuinely adaptable components, which can be used at a higher level in large-scale systems, such as a building’s physical architecture, technical networks, or distributed software.

Designing seeds deals with the mapping from a genome to a phenotype,  $G \rightarrow P$ . Although the growth of a given genotype into a fixed phenotype is already a

challenge in itself, a greater goal is to design genomes that can support *adaptable* growth on multiple different levels, corresponding to the forms of interaction present in the growing models.

One form of adaptability involves modification of the phenotype via environmental influence; for an environment  $E_i$  we have  $(G, E_i) \rightarrow P_i$ . That is, on the ontogenetic timescale, and under a given genotype-phenotype mapping, the development process should be sensitive to, and modifiable by, environmental conditions in the garden (Měch & Prusinkiewicz, 1996; Prusinkiewicz, James & Měch, 1994). External conditions or stimuli encountered by one individual during its growth, whether of a mechanical or signalling type, should be able to influence the outcome. This is the level of  $P$ , or rather the genotype-phenotype mapping, for which natural analogies can be found more readily in the plant kingdom than the animal kingdom. Plants and trees can be pruned, bent, arranged, sculpted, etc., whether intentionally by a human gardener (bonsais, espaliers, topiaries) or spontaneously when faced with adverse or favourable conditions (wind strength, rocky obstacles, soil composition, light intensity). This is precisely what computational gardeners expect to be doing. Therefore, the seeds that they “order” from the Seed Factory’s species engineers should provide appropriately workable material.

A second form of adaptability involves *symbiosis*, or the growing of different phenotypes or individuals together. This includes both individuals from the same “species”, but also individuals from different species as well (ecosystemic interactions). We may denote this  $(G_i + G_j, E) \rightarrow (P_i, P_j)$ .

To qualify as true “components”, grown organisms must also be literally “symbiotic” with each other. Whether they belong to the same breed or species, or to different species, their shape and physiology should offer certain anchor points to be able to attach and function together into a larger architecture. Think of interfaces and method-calling in software objects; prefabricated walls and rooms in trailer homes; clusters and routing in networks. A complex cyber-physical system such as a skyscraper or spaceship (Armstrong, 2017) needs several different species of seeds to grow together: walls, plumbing, power, control systems, etc.

### 2.3 Growth Engine

A Growth Engine provides the computational mechanisms to grow a seed. This might be required to grow in simulation a seed intended for a physical device, or to grow the seed of a virtual component such as a software control system. There are several types of growth engine, corresponding to the several types of seeds, based on different underlying physical technologies such as nanobots or protocells, or on the virtual technologies such as ME, DCA, GS, SMCGP, SC mentioned earlier. For example, a seed defined as a set of rewrite grammar rules would require a growth engine that could implement the grammar rewriting system (among other things, discussed below).

The GRO-CYPHY approach contains a conceptual switch from the notion of “machine” to the notion of “organism”. An organism is an organised ensemble

of components. This ensemble is derived from one generic component, the seed (or fertilised egg in animals), by duplication and differentiation, or alternatively by the self-assembly of components produced elsewhere. Each component has some autonomy, but it is also coupled and interacts with the other components and with the environment. Contrary to the relations between parts of a machine, the coupling between components of an organism is not the result of deliberate design: two components interact because they are neighbours (e.g. in physical space), not because these components have been pre-made to interact. This new type of relationship, natural and spatially explicit, requires the components to be generic: they are not designed for one specific and dedicated purpose but must be able to react to a wide variety of interactions. Furthermore, a complete artefact, classically several “machines”, is now seen as several “organisms”, grown together as a cooperative ecosystem in a garden, from several seeds.

The Growth Engine executes a programmable and reproducible indirect mapping from seed genomes  $G$  to organism phenotypes  $P$ , in the context of an environment  $E$ . Calculating the transformation from a given  $G$  to a resulting  $P$  corresponds to developing an organism. (Solving the inverse problem, of finding an appropriate  $G$  (or family of  $G$ s) given a desired  $P$ , is the challenge of searching and designing useful species, which process occurs in the Seed Factory.) The growing organism interacts with co-growing subsystems and environmental influences, in the Computational Garden.

The generic component that grows, reacts, and interacts has two parts:

1. an internal state  $S$  that develops over time
2. a program (genome)  $G$  that specifies the internal dynamics and the results of the potential interactions with the other components and with the environment.

Program  $G$  is inherited from the (components in the) seed through its successive division steps, or during self-assembly. State  $S$  includes the physical characteristics of the component (size, shape, position, etc.) as well as a set of internal variables corresponding to some gene-regulatory and metabolic activity.

The development of the whole organism is implemented by the Growth Engine. This engine:

- takes care of the interactions between components and with the environment
- updates the state of each component of an organism in accordance with their Genome
- manages the organism as a whole.

Several organisms co-develop in interaction in a garden, achieving an ecosystem of developmental artefacts. Each organism is under the control of its own Growth Engine instance, and these individual Growth Engines interact with each other and with the wider environment.

## 2.4 Computational Garden

The computational garden (Miller, 2018) is where the various seeds are planted and grow together, responding to their environment, into the resultant artefact.

The idea of several “plastic” organisms growing together and adapting to their environment is one path to our gardening metaphor. Another path that also leads us to the garden, springs from a disappointment with progress in software development. Programming still essentially occurs at the individual statement level: there is the necessity of managing an overwhelming amount of low-level detail. Moreover, the resulting software is very *brittle*: the smallest change or error can have a devastating effect. The garden provides us with a higher-level metaphor: high-level guiding of a robust complex growing system, rather than low-level engineering of the precise placement of every cell or particle.

The majority of the computational components in a garden will be growing seeds; the remainder will be supplied by external software such as physics engines (simulating the relevant internal and environmental physical conditions), and (interfaces to) growing embodied physical systems. The gardener needs to manipulate these components. Keeping with the gardening metaphor, we require operations such as:

- Planting seeds (starting the growth of a seed at a particular location)
- Moving growing organisms (transplanting)
- Feeding: applying nutrients, energy, and their virtual analogues, to affect growth rates (enhance growth in a particular direction)
- Pruning growing organisms (stop growth in a particular direction)
- Training along a given direction (change the direction of growth)
- Grafting physical and/or virtual organisms together
- Manipulating environmental variables and other parameter values
- Recording and rerunning gardening “scripts”.

In order to raise the level of the programming task, these operations need to be controlled via some high-level, intuitive visual metaphor, including facilities for:

- Various graphical views of the garden
  - 3D graphical visualisation (for physical systems)
  - Other standard graphical visualisations (e.g. graphs), for virtual systems (and for physical systems if desired)
  - Visual widgets indicating the state or desirability of the computational output of the garden (since computation underlies the Garden, the graphical elements need not be too sophisticated, as it is the computation they perform that is most important)
  - Ability to create customised views, using “drag and drop” from the visual widget library
- User navigation through the garden, to explore, view and manipulate the various components growing there.

It should also be possible to embed user-defined entities into the garden, that is, to place non-“living” but functional objects into the garden. This allows a combination of traditionally engineered and growing parts within the final artefact. For example, in some circumstances one might want to provide a “pre-cast” shell for a building (analogous to a trellis in a real garden), and grow only the utilities and control systems within it.



### 3 Illustrative Application: a skyscraper

There is a sub-discipline of “morphogenetic architecture” (Hensel, Menges & Weinstock, 2004, 2006; Roudavski, 2009), where researchers attempt to “grow” the form of a building. However, these researchers tend to focus on the aesthetic architecture of the physical shell, and tend not to consider the simultaneous growth of the other physical subsystems such as plumbing, or the virtual components such as the control systems.

The generation of useful architectural and related sub-system design is often realized via parametrised geometries; these realisations are sometimes used in real-world construction (Shepherd, 2010; von Buelow, Falk & Turrin, 2010). There have been several successful approaches to the generation of structural designs via developmental techniques, usually cellular growth procedures (Kicinger, Arciszewski & Jong, 2005; Kowaliw, Grogono & Kharma, 2007b; Steiner, Jin & Sendhoff, 2008; Yogev, Shapiro & Antonsson, 2010). Jin & Sendhoff (2009) argue in favour of such life-like representations as a means of incorporating features such as scalability, robustness, evolvability, and self-adaptation. One such example of robustness has been demonstrated in structural design: Kowaliw, Grogono & Kharma (2007a) show the use of a developmental procedure in the creation of truss designs which led to self-adaptation to differing scales and geometric environments, implying a sort of artificial polymorphism.

A skyscraper provides a good exemplar application, being a non-trivial cyber-physical system with multiple physical (structural) and virtual (control) components, existing in a complex environment (including social use, weather, and other external influences). There are several issues to address in such an application. These include:

1. How various subsystems can be grown in concert such that they integrate effectively, and influence each other’s growth, for example: plumbing adapting to the shell; plumbing and wiring adapting so as not to interfere
2. How various subsystems can be grown in the context of a physical environment, allowing some of a system to be built conventionally and the remainder to be grown through it, and an analysis of the tradeoffs
3. How much of the morphological requirements can be placed in the seed, and how much in the subsequent gardening environment and regime (genericity *v* specificity tradeoffs)
4. How to stop the systems growing where necessary (Miller, 2004), potentially based on the mechanisms plants use to stop growing, which process can be triggered either by environmental conditions, such as winter, or by internally regulated developmental programmes (Leyser & Day, 2003):
  - A plant (or parts of it) can go into *dormancy*, for example, over winter, or as a seed, or axillary buds.
  - A plant (or parts of it) can *senesce* and die: leaves in winter, whole plant for annual plants that only last a year. This involves active “suicide” of tissues in a very regulated process, for example, remobilising all the nutrients to the roots before shedding leaves.

- Plants grow mostly from meristems (root and shoot tips). These are usually indeterminate: they can go on for ever maintaining themselves while producing new cells to build the plant, but they can also become *determinate* and use themselves up making other components. For example, the shoot apical meristem becomes a terminal flower, so the shoot stops growing.

Such an exemplar skyscraper application could be used to drive the implementation of a prototype GRO-CYPHY system conforming to the architectural design described above.

## 4 Discussion and conclusions

Currently most of the research effort in growing software has been concentrated either on growing structures or shapes with no obvious computational purpose, or on growing programs from (grammar) rules. GRO-CYPHY is a process to grow both aspects together in an integrated fashion: the shapes will be needed to perform the functions, and programs will be embedded in the structure to make it adaptable and responsive.

Visual programming has been around for a long time. For example, most graphical user interfaces (GUIs) have some development toolkit that allows programmers to create GUI layouts in a visual and intuitive way. Furthermore, complex programs can be specified with the help of UML (Unified Modelling Language) diagrams. However, such visual programming has mainly been used to build software in a static way, following a traditional software development lifecycle: software is rarely manipulated in this way while executing (model-driven software adaptation (Bencomo, Blair & France, 2008) is one exception); growing software in this way is not yet a reality. GRO-CYPHY is a radically new way of conceiving of software: as something that grows, and therefore modifies itself during execution. Such software will naturally be able to withstand external operations on itself during execution, such as pruning and grafting by a “gardener”. The gardener will be a high-level programmer, and ultimately can even be the user directly.

GRO-CYPHY is a vision of “programmed organisms”, software intensive embodied systems or cyber-physical systems that are grown in a garden, where they are autonomous and yet responsive and programmable at a high level. Such programmed organisms help blur the artificial distinction between the abstract software programs and the substrate where the software actuates (execution hardware plus supporting physical structures).

### Acknowledgements

We thank Taras Kowaliw, Ottoline Leyser, Bruce MacLennan, Olivier Michel, Igor Nikolic, Christof Teuscher, Gunnar Tufte, Francisco J. Vico, and Lidia Yamamoto, for many discussions and critiques that helped develop some of the ideas presented here.

## Bibliography

- Abelson, H., Allen, D., Coore, D., Hanson, C., Homsy, G., Knight, Jr., T. F., Nagpal, R., Rauch, E., Sussman, G. J., & Weiss, R. (2000). Amorphous computing. *Communications of the ACM*, 43(5), 74–82.
- Armstrong, R. (2014). Designing with protocells: Applications of a novel technical platform. *Life*, 4(3), 457–490.
- Armstrong, R. (Ed.). (2017). *Star Ark: A Living, Self-Sustaining Spaceship*. Springer.
- Beal, J. (2011). Functional blueprints: an approach to modularity in grown systems. *Swarm Intelligence*, 5, 257–281.
- Bencomo, N., Blair, G., & France, R. (2008). Model-driven software adaptation. In *ECOOP 2007*, volume 4906 of *LNCS* (pp. 132–141). Springer.
- Debbabi, B., Diaconescu, A., & Lalanda, P. (2012). Controlling self-organising software applications with archetypes. In *Sixth IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2012, Lyon, France, September 10-14, 2012*, (pp. 69–78). IEEE Press.
- DeHon, A., Giavitto, J.-L., & Gruau, F. (2007). 06361 executive report – computing media languages for space-oriented computation. In DeHon, A., Giavitto, J.-L., & Gruau, F. (Eds.), *Computing Media and Languages for Space-Oriented Computation*, number 06361 in Dagstuhl Seminar Proceedings.
- Diaconescu, A., Tomforde, S., & Müller-Schloer, C. (2018). Holonic cellular automata: Modelling multi-level self-organisation of structure and behaviour. In *The 2018 Conference on Artificial Life (ALife 2018), Tokyo, Japan*, (pp. 186–193). MIT Press.
- Doursat, R. (2008). Organically grown architectures: Creating decentralized, autonomous systems by embryomorphic engineering. In Würtz, R. P. (Ed.), *Organic Computing*, (pp. 167–200). Springer.
- Doursat, R., Sayama, H., & Michel, O. (Eds.). (2012). *Morphogenetic Engineering: Toward Programmable Complex Systems*. Springer.
- Drexler, K. E. (1986). *Engines of Creation*. Doubleday.
- Harding, S., Miller, J. F., & Banzhaf, W. (2010). Developments in cartesian genetic programming: self-modifying CGP. *Genetic Programming and Evolvable Machines*, 11(3-4), 397–439.
- Hensel, M., Menges, A., & Weinstock, M. (Eds.). (2004). *Emergence: Morphogenetic Design Strategies*. Wiley.
- Hensel, M., Menges, A., & Weinstock, M. (Eds.). (2006). *Techniques and Technologies in Morphogenetic Design*. Wiley.
- Hornby, G. S. & Pollack, J. B. (2001). The advantages of generative grammatical encodings for physical design. In *CEC 2001, vol. 1*, (pp. 600–607). IEEE.
- Jacob, C. (1994). Genetic L-system programming. In *PPSN III*, volume 866 of *LNCS*, (pp. 333–343). Springer.

- Jacob, C. (1996). Evolving evolution programs: Genetic programming and L-systems. In *Proceedings of the First Annual Conference on Genetic Programming*, (pp. 107–115). MIT Press.
- Jin, Y. & Sendhoff, B. (2009). A systems approach to evolutionary multiobjective structural optimization and beyond. *IEEE Computational Intelligence Magazine*, 4, 62–76.
- Kicinger, R., Arciszewski, T., & Jong, K. D. (2005). Evolutionary design of steel structures in tall buildings. *Journal of Computing in Civil Engineering*, 19(3), 223–238.
- Kowaliw, T. & Banzhaf, W. (2009). Augmenting artificial development with local fitness. In *CEC 2009*, (pp. 316–323). IEEE.
- Kowaliw, T., Grogono, P., & Kharm, N. (2007a). Environment as a spatial constraint on the growth of structural form. In *GECCO '07*, (pp. 1037–1044). ACM.
- Kowaliw, T., Grogono, P., & Kharm, N. (2007b). The evolution of structural form through artificial embryogeny. In *IEEE Symposium Series on Computational Intelligence (IEEE-ALIFE '07)*, (pp. 425–432).
- Lalanda, P., McCann, J. A., & Diaconescu, A. (Eds.). (2013). *Autonomic Computing: principles, design and implementation*. Springer.
- Leyser, O. & Day, S. (2003). *Mechanisms in Plant Development*. Blackwell.
- Maignan, L. & Spicher, A. (2015). Global graph transformations. In Plump, D. (Ed.), *GCM 2015, satellite of ICGT 2015*, (pp. 34–49).
- Miller, J. F. (2004). Evolving a self-repairing, self-regulating, French flag organism. In *Proc. Genetic and Evolutionary Computation Conference*, volume 3102 of *LNCS*, (pp. 129–139). Springer.
- Miller, J. F. (Ed.). (2011). *Cartesian Genetic Programming*. Springer.
- Miller, J. F. (2018). The software garden. In R. Walsh & S. Stepney (Eds.), *Narrating Complexity* chapter 15, (pp. 201–212). Springer.
- Miller, J. F. & Thomson, P. (2004). Beyond the complexity ceiling: Evolution, emergence and regeneration. In *Proc. GECCO 2004 Workshop on Regeneration and Learning in Developmental Systems*.
- Mock, K. J. (1998). Wildwood: The evolution of L-system plants for virtual environments. In *Proceedings of the 1998 IEEE World Congress on Computational Intelligence*, (pp. 476–480). IEEE Press.
- Monks, M., Oh, B. M., & Dorsey, J. (2000). Audiooptimization: goal-based acoustic design. *IEEE Computer Graphics and Applications*, 20(3), 76–90.
- Měch, R. & Prusinkiewicz, P. (1996). Visual models of plants interacting with their environment. In *Proc SIGGRAPH '96*, (pp. 397–410). ACM.
- Prusinkiewicz, P., James, M., & Měch, R. (1994). Synthetic Topiary. In *Proc SIGGRAPH '94*, (pp. 351–358). ACM.
- Roudavski, S. (2009). Towards morphogenesis in architecture. *International Journal of Architectural Computing*, 7(3), 345–374.
- Sayama, H. (2009). Swarm chemistry. *Artificial Life*, 15, 105–114.
- Shea, K. & Smith, I. F. C. (2006). Improving full-scale transmission tower design through topology and shape optimization. *Journal of Structural Engineering*, 132(5), 781–790.

- Shepherd, P. (2010). Aviva stadium – the use of parametric modelling in structural design. *The Structural Engineer*, 89(23), 28–34.
- Solé, R. V., Munteanu, A., Rodriguez-Caso, C., & Macía, J. (2007). Synthetic protocell biology: from reproduction to computation. *Philos Trans R Soc Lond B Biol Sci*, 362(1486), 1727–1739.
- Spicher, A. & Giavitto, J.-L. (2017). Interaction-based programming in MGS. In A. Adamatzky (Ed.), *Advances in Unconventional Computing*, volume 1 (pp. 305–342). Springer.
- Spicher, A. & Michel, O. (2005). Using rewriting techniques in the simulation of dynamical systems: Application to the modeling of sperm crawling. In *ICCS'05, part I*, volume 3514 of *LNCS*, (pp. 820–827). Springer.
- Spicher, A., Michel, O., & Giavitto, J.-L. (2010). Declarative mesh subdivision using topological rewriting in MGS. In *ICGT 2010*, volume 6372 of *LNCS*, (pp. 298–313). Springer.
- Stanley, K. O. & Miikkulainen, R. (2003). A taxonomy for artificial embryogeny. *Artificial Life*, 9, 93–130.
- Steiner, T., Jin, Y., & Sendhoff, B. (2008). A cellular model for the evolutionary development of lightweight material with an inner structure. In *GECCO '08*, (pp. 851–858). ACM.
- Stepney, S. (2007). Embodiment. In D. Flower & J. Timmis (Eds.), *In Silico Immunology* chapter 12, (pp. 265–288). Springer.
- Stepney, S. (2008). The neglected pillar of material computation. *Physica D: Nonlinear Phenomena*, 237(9), 1157–1164.
- Stepney, S., Rasmussen, S., & Amos, M. (Eds.). (2018). *Computational Matter*. Springer.
- Tufte, G. & Haddow, P. C. (2007). Extending artificial development: Exploiting environmental information for the achievement of phenotypic plasticity. In *ICES 2007*, volume 4684 of *LNCS*, (pp. 297–308). Springer.
- von Buelow, P., Falk, A., & Turrin, M. (2010). Optimization of structural form using a genetic algorithm to search associative parametric geometry. In *Proceedings of the International Conference on Structures Architecture (ICSA 2010)*.
- Werfel, J. (2010). Biologically realistic primitives for engineered morphogenesis. In *ANTS 2010*, volume 6234 of *LNCS*, (pp. 131–142). Springer.
- Würtz, R. P. (Ed.). (2008). *Organic Computing*. Springer.
- Yogev, O., Shapiro, A. A., & Antonsson, E. K. (2010). Computational evolutionary embryogeny. *IEEE Transactions on Evolutionary Computation*, 14(2), 301–325.