

Review



Cite this article: Foley JD, Breiner S, Subrahmanian E, Dusel JM. 2021 Operads for complex system design specification, analysis and synthesis. *Proc. R. Soc. A* **477**: 20210099. <https://doi.org/10.1098/rspa.2021.0099>

Received: 3 February 2021

Accepted: 24 May 2021

Subject Areas:

mathematical modelling, artificial intelligence, category theory

Keywords:

complex systems, system design, automated reasoning, compositionality, applied category theory, operads

Author for correspondence:

John D. Foley

e-mail: foley@metsci.com

Operads for complex system design specification, analysis and synthesis

John D. Foley¹, Spencer Breiner², Eswaran

Subrahmanian^{2,3} and John M. Dusel¹

¹Metron, Inc., 1818 Library St., Reston, VA, USA

²US National Institute of Standards and Technology, Gaithersburg, MD, USA

³Carnegie Mellon University, Pittsburgh, PA, USA

JDF, 0000-0002-0087-5339

As the complexity and heterogeneity of a system grows, the challenge of specifying, documenting and synthesizing correct, machine-readable designs increases dramatically. Separation of the system into manageable parts is needed to support analysis at various levels of granularity so that the system is maintainable and adaptable over its life cycle. In this paper, we argue that operads provide an effective knowledge representation to address these challenges. Formal documentation of a syntactically correct design is built up during design synthesis, guided by semantic reasoning about design effectiveness. Throughout, the ability to decompose the system into parts and reconstitute the whole is maintained. We describe recent progress in effective modelling under this paradigm and directions for future work to systematically address scalability challenges for complex system design.

1. Introduction

We solve complex problems by separating them into manageable parts [1,2]. Human designers do this intuitively, but details can quickly overwhelm intuitive insights. Multiple aspects of a problem may lead to distinct intuitive decompositions and complementary models of a single system—e.g. competing considerations for cyberphysical systems [3,4]—or simulation of behaviour may require many levels of fidelity—e.g. in

© 2021 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.




Figure 1. Separating concerns with operads: (1) Composition separates subsystem designs (boundaries on left); (2) functorial semantics separate abstract systems from computational model instances (single arrows); (3) natural transformations separate and align (double arrow) complementary models (square, diamond).

modelling and simulation [5]—leading to a spectrum of models which are challenging to align. We argue that operads, formal tools developed to compose geometric and algebraic objects, are uniquely suited to separate complex systems into manageable parts and maintain alignment across complementary models.

Operads provide three ways to separate concerns for complex systems: (1) designs for subsystems are separated into composable modules; (2) syntactic designs to compose systems are separated from the semantic data that model them; and (3) separate semantic models can be aligned to evaluate systems in different ways. The three relationships are illustrated in figure 1.

Hierarchical decomposition (figure 1, (1)) is nothing new. Both products and processes are broken down to solve problems from design to daily maintenance. Operads provide a precise language to manage complex modelling details that the intuitive—and highly beneficial—practice of decomposition uncovers, e.g. managing multiple, complementary decompositions and models.

Operads separate the syntax to compose subsystems from the semantic data modelling them (figure 1, (2)). Syntax consists of abstract ‘operations’ to design the parts and architecture of a system. Semantics define how to interpret and evaluate these abstract blueprints. Operad syntax is typically lightweight and declarative. Operations can often be represented both graphically and algebraically (figure 4), formalizing intuitive design diagrams. Operad semantics model specific aspects of a system and can range from fast to computationally expensive.

The most powerful way operads separate is by aligning complementary models while maintaining compatibility with system decompositions (figure 1, (3)). Reconciling complementary models is a persistent and pervasive issue across domains [3,6–12]. Historically, Eilenberg & Mac Lane [13] invented *natural transformations* to align computational models of topological spaces. Operads use natural transformations to align hierarchical decompositions, which are particularly well suited to system design.

This paper articulates a uniform and systematic foundation for system design and analysis. In essence, the syntax of an operad defines *what can be* put together, which is a prerequisite to decide *what should be* put together. Interfaces define which designs are *syntactically* feasible, but key *semantic* information must be expressed to evaluate candidate designs. Formulating system models within operad theory enforces the intellectual hygiene required to make sure that different concerns stay separated while working together to solve complex design problems.

We note five strengths of this foundation that result from the three ways operads separate a complex problem and key sections of the paper that provide illustrations.

Expressive, unifying meta-language. A meta- or multi-modelling [14] language is needed to express and relate multiple representations. The key feature of operad-based meta-modelling is its focus on coherent mappings between models (figure 1, (2), (3)), as opposed to a universal

modelling context, like UML, OWL, etc., which is inevitably under or over expressive for particular use cases. Unification allows complementary models to work in concert, as we see in §5 for function and control. Network operads—originally developed to design systems—were applied to task behaviour. This power to unify and express becomes especially important when reasoning across domains with largely independent modelling histories; compare, e.g. [4] (§2b, d, 3, 4a, 5).

Minimal data needed for specification. Data needed to set up each representation of a problem are minimal in two ways: (i) any framework must provide similar, generative data; and (ii) each level only needs to specify data relevant to that level. Each representation is self-sufficient and can be crafted to efficiently address a limited portion of the full problem. The modeller can pick and choose relevant representations and extend the meta-model as needed (§4, 6b).

Efficient exploration of formally correct designs. An operad precisely defines how to iteratively construct designs or adapt designs by substituting different subsystems. Constructing syntactically invalid designs is not possible, restricting the relevant design space, and correctness is preserved when moving across models. Semantic reasoning guides synthesis, potentially at several levels of detail. This facilitates lazy evaluation: first syntactic correctness is guaranteed, then multitudes of coarse models are explored before committing to later, more expensive evaluations. The basic moves of iteration, substitution and moving across levels constitute a rich framework for exploration. We obtain not only an effective design but also formal documentation of the models that justify this choice (§2b,c, 6, 7e).

Separates representation from exploitation. Operads and algebras provide structure and representation for a problem. Exploitation of this structure and representation is a separate concern. As Herbert Simon noted during his Nobel Prize speech [15]: ‘... decision makers can satisfice either by finding optimum solutions for a simplified world, or by finding satisfactory solutions for a more realistic world’. This is an either-or proposition for a simple representation. By laying the problem across multiple semantic models, useful data structures for each model—e.g. logical, evolutionary or planning frameworks—can be exploited by algorithms that draw on operad-based iteration and substitution (§6, 7e).

Hierarchical analysis and synthesis. Operads naturally capture options for the hierarchical decomposition of a system, either within a semantic model to break up large-scale problems or across models to gradually increase modelling fidelity (§2a, 5, 6c, 7a).

(a) Contribution to design literature

There are well-known examples of the successful exploitation of separation. For instance, electronic design automation has had decades of success leveraging hierarchical separation of systems and components to achieve very large-scale integration of complex electronic systems [16–18]. We do not argue that operads are needed for extremely modular domains. Instead, operads may help broaden the base of domains that benefit from separation and provide a means to integrate and unify treatment across domains. On the other hand, for highly integral domains the ability to separate in practically useful ways may be limited [19,20]. The recent applications we present help illustrate where operads may prove useful in the near and longer term; see §7c for further discussion.

Compared to related literature, this article is application-driven and outward focused. Interest in applying operads and category theory to systems engineering has surged [21–26] as part of a broader wave applying category theory to design databases, software, proteins, etc. [27–33]. While much of loc. cit. matches applications to existing theoretical tools, the present article describes recent *application-driven* advancements and overviews *specific methods* developed to address challenges presented by domain problems. We introduce operads for design to a general scientific audience by explaining what the operads do relative to broadly applied techniques and how specific domain problems are modelled. Research directions are presented with an eye towards opening up interdisciplinary partnerships and continuing application-driven investigations to build on recent insights.

Table 1. The theory of operads draws on many familiar ideas, establishing a dictionary between contexts.

operads	tree	API	equational	systems
types	edges	data types	variables	boundaries
operations	nodes	methods	operators	architectures
composites	trees	scripts	evaluation	nesting
algebras	labels	implementations	values	models

(b) Organization of the paper

The present article captures an intermediate stage of technical maturity: operad-based design has shown its practicality by lowering barriers of entry for applied practitioners and demonstrating applied examples across many domains. However, it has not realized its full potential as an applied meta-language. Much of this recent progress is not focused solely on the analytic power of operads to separate concerns. Significant progress on explicit specification of domain models and techniques to automatically synthesize designs from basic building blocks has been made. Illustrative use cases and successful applications for design specification, analysis and synthesis organize the exposition; see figure 2.

Section 2 introduces operads for design by analogy to other modelling approaches. Our main examples are introduced in §3. Section 4 describes how concrete domains can be specified with minimal combinatorial data, lowering barriers to apply operads. Section 5 concerns analysis of a system with operads. Automated synthesis is discussed in §6. Future research directions are outlined in §7, which includes a list of open problems.

Notations. Throughout, we maintain the following notional conventions for:

- syntax operads (figure 1, left), capitalized calligraphy: \mathcal{O}
- types (figure 1, edges on left), capitalized teletype: X, Y, Z, \dots
- operations (figure 1, nodes on left), uncapitalized teletype: f, g, h, \dots
- semantic contexts (figure 1, right), capitalized bold: **Sem, Set, Rel, ...**
- functors from syntax to semantics (figure 1, single arrows), capitalized sans serif: Model: $\mathcal{O} \rightarrow \mathbf{Sem}$;
- alignment of semantic models via natural transformations (figure 1, double arrow), uncapitalized sans serif: align: $\text{Model}_1 \Rightarrow \text{Model}_2$;

2. Applying operads to design

We introduce operads by an analogy, explaining what an operad is and motivating its usefulness for systems modelling and analysis. The theory [34–36] pulls together many different intuitions. Here, we highlight four analogies or ‘views’ of an operad: hierarchical representations (tree view), strongly typed programming languages (API¹ view), algebraic equations (equational view) and system cartography (systems view). Each view motivates operad concepts; see table 1.

The paradigm of this paper is based on a typed operad, also known as a ‘coloured operad’ [36] or ‘symmetric multicategory’ ([34], 2.2.21). A typed **operad** \mathcal{O} has


- A set T of **types**.
- Sets of **operations** $\mathcal{O}(X_1, \dots, X_n; Y)$ where $X_i, Y \in T$ and we write $f: \langle X_i \rangle \rightarrow Y$ to indicate that $f \in \mathcal{O}(X_1, \dots, X_n; Y)$.
- A specific way to **compose** any operation $f: \langle Y_i \rangle \rightarrow Z$ with $g_i: \langle X_{ij} \rangle \rightarrow Y_i$ whose output types match the inputs of f to obtain a composite $f \circ (g_1, \dots, g_n) = h: \langle X_{ij} \rangle \rightarrow Z$.

¹Application programming interface.


These data are subject to rules ([36], 11.2) governing permutation of arguments and assuring that iterative composition is coherent, analogous to associativity for ordinary categories ([37], I).

(a) The tree view

Hierarchies are everywhere, from scientific and engineered systems to government, business and everyday life; they help to decompose complex problems into more manageable pieces. The fundamental constituent of an operad, called an *operation*, represents a single step in a hierarchical decomposition. We can think of this as a single branching in a labelled tree, e.g.



Formally, this represents an element $op \in \mathcal{O}(\text{Sub}_1, \text{Sub}_2; \text{System})$. More generally, we can form new operations—trees—by *composition*. Given further refinements for the two subsystems Sub_1 and Sub_2 , by op_1 and op_2 , respectively, we have three composites



Together with the original operation, these represent four views of the same system at different levels of granularity; compare, e.g. ([38], fig. 2). This reveals an important point: an operad provides a collection of interrelated models that fit together to represent a complex system.

The relationship between models is constrained by the *principle of compositionality*: the whole is determined by its parts *and* their organization. Here, the whole is the root, the parts are the leaves, and each tree is an organizational structure. Formally, *associativity axioms*, which generalize those of ordinary categories, enforce compositionality. For example, composing the left-hand tree above with op_2 must give the same result as composing the centre tree with op_1 . Both give the tree on the right, since they are built up from the same operations. In day-to-day modelling, these axioms are mostly invisible, ensuring that everything ‘just works’, but the formal definitions ([36], 11.2) provide explicit requirements and desiderata for modelling languages ‘under the hood’.

Operads encourage principled approaches to emergence by emphasizing the organization of a system. Colloquially speaking, an emergent system is ‘more than the sum of its parts’; operations provide a means to describe these nonlinearities. This does not explain emergent phenomena, which requires detailed semantic modelling, but begins to break up the problem with separate (but related) representation of components and their interactions. The interplay between these elements can be complex and unexpected, even when the individual elements are quite simple.² Compositional models may develop and exhibit emergence as interactions between components are organized, in much the same way as the systems they represent.

(b) The API view

For most applications, trees bear labels: fault trees, decision trees, syntax trees, dependency trees and file directories, to name a few. A tree’s labels indicate its semantics either explicitly with numbers and symbols or implicitly through naming and intention.

²For example, diffusion rates (components) and activation/inhibition (interactions) generate zebra’s stripes in Turing’s model of morphogenesis [39].

In an operad, nodes identify operations while edges—called *types*—restrict the space of valid compositions. This is in analogy to type checking in strongly typed programming languages, where we can only compose operations when types match. In the API view, the operations are abstract method declarations

```
def op(x1 : Sub1, x2 : Sub2) : System,
def op1(y1 : Sub11, y2 : Sub12) : Sub1,
def op2(z1 : Sub21, z2 : Sub22, z3 : Sub23) : Sub2.
```

Composites are essentially scripted methods defined in the API. For example,

```
def treeLeft(y1 : Sub11, y2 : Sub12, x2 : Sub2) : System
  = op(op1(y1, y2), x2),
```

is a script for the left-most tree above. However, the compiler will complain with an invalid syntax error for any script where the types do not match, say

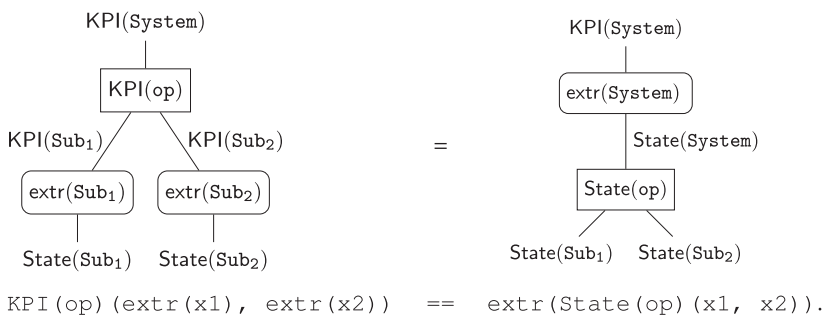
```
def badTree(y1 : Sub11, y2 : Sub12, x2 : Sub2) : System
  = op(x2 ,op1(y1, y2)).
```

If an operad is an API—a collection of abstract types and methods—then an *operad algebra* A is a concrete implementation. An algebra declares: (1) a set of *instances* for each type; and (2) a *function* for each operation, taking instances as arguments and returning a single instance for the composite system. That is, $A: \mathcal{O} \rightarrow \mathbf{Set}$ has

- for each type $X \in T$, a set $A(X)$ of **instances** of type X , and
- for each operation $f: (X_i) \rightarrow Y$, the function $A(f)$ **acts** on input elements $a_i \in A(X_i)$ to obtain a single output element $A(f)(a_1, \dots, a_n) \in A(Y)$.

Required coherence rules ([36], 13.2) are analogous to the definition of a functor into \mathbf{Set} ([37], I.3). For example, we might declare a state space for each subsystem, and a function to calculate the overall system state given subsystem states. Alternatively, we might assign key performance indicators (KPIs) for each level in a system and explicit formulae to aggregate them. The main thing to remember is: just as an abstract method has many implementations, an operad has many algebras. Just like an API, the operad provides a common syntax for a range of specific models, suited for specific purposes.

Unlike a traditional API, an operad provides an explicit framework to express and reason about semantic relationships between *different* implementations. These different implementations are linked by type-indexed mappings between instances called *algebra homomorphisms*. For example, we might like to extract KPIs from the system state. The principle of compositionality places strong conditions on this extraction: the KPIs extracted from the overall system state must agree with the KPIs obtained by aggregating subsystem KPIs. That is, in terms of trees and in pseudocode




For any state instances for Sub_1 and Sub_2 at the base of the tree, the two computations must produce the same KPIs for the overall system at the top of the tree. Here, $\text{KPI}(\text{op})$ and $\text{State}(\text{op})$ implement op in the two algebras, while $\text{extr}(-)$ are *components* of the algebra homomorphism to extract KPIs. Similar to associativity, these compositionality conditions guarantee that extracting KPIs ‘just works’ when decomposing a system hierarchically.

(c) The equational view

We have just seen an equation between trees that represent implementations. Because an operad can be studied without reference to an implementation, we can also define equations between abstract trees. This observation leads to another view of an operad: as a system of equations.

The first thing to note is that equations occur within the sets of operations $\mathcal{O}(X_1, \dots, X_n; Y)$; an equation between two operations only makes sense if the input and output types match. Second, if one side of an equation $f = f'$ occurs as a subtree in a larger operation g , substitution generates a new equation $g = g'$. Two trees are equal if and only if they are connected by a chain of such substitutions (and associativity equations). In general, deciding whether two trees are equal (the word problem) may be intractable. Third, we can often interpret composition of operations as a normal-form computation



We then compare composed operations directly to decide equality. For example, there is an operad whose operations are matrices. Composition computes a normal form for a composite operation by block diagonals and matrix multiplication,

$$\begin{aligned} op &: n \times (m_1 + m_2) \\ op_1 &: m_1 \times (k_{11} + k_{12}) \\ op_2 &: m_2 \times (k_{21} + k_{22} + k_{23}) \end{aligned} \quad \mapsto \quad (op) \cdot \begin{pmatrix} op_1 & 0 \\ 0 & op_2 \end{pmatrix}.$$

Operad axioms constrain composition. For example, the axiom mentioned in §2a corresponds to

$$\begin{aligned} & (op) \cdot \begin{pmatrix} op_1 & 0 \\ 0 & I_{m_2} \end{pmatrix} \cdot \begin{pmatrix} I_{k_{11}} & 0 & 0 \\ 0 & I_{k_{12}} & 0 \\ 0 & 0 & op_2 \end{pmatrix} \\ &= (op) \cdot \begin{pmatrix} I_{m_1} & 0 \\ 0 & op_2 \end{pmatrix} \cdot \begin{pmatrix} op_1 & 0 & 0 & 0 \\ 0 & I_{k_{21}} & 0 & 0 \\ 0 & 0 & I_{k_{22}} & 0 \\ 0 & 0 & 0 & I_{k_{23}} \end{pmatrix}. \end{aligned}$$

The key point is that any algebra that implements the operad must satisfy *all* of the equations that it specifies. Type discipline controls which operations can compose; equations between operations control the resulting composites. Declaring equations between operations provides additional contracts for the API. For instance, any unary operation $f: X \rightarrow X$ (a loop) generates an infinite sequence of composites $\text{id}_X, f, f^2, f^3, \dots$. Sometimes this is a feature of the problem at hand, but in other cases, we can short-circuit the infinite regress with assumptions like idempotence ($f^2 = f$) or cyclicity ($f^n = \text{id}_X$) and ensure that algebras contain no infinite loops.

(d) The systems view

When we apply operads to study systems, we often think of an operation $f: \langle X_i \rangle \rightarrow Y$ as a system architecture. Intuitively, Y is the system and the X_1, \dots, X_n are the components, but this is a bit misleading. It is better to think of types as boundaries or interfaces, rather than systems. Instead, f is the system architecture, with component interfaces X_i and environmental interface Y . Composition formalizes the familiar idea ([38], fig. 2) that one engineer's system is the next's component; it acts by nesting subsystem architectures within higher-level architectures.

Once we establish a system architecture, we would like to use this structure to organize our data and analyses of the system. Moreover, according to the principle of compositionality, we should be able to construct a system-level analysis from an operation by analysing the component-level inputs and synthesizing these descriptions according to the given operations.

The process of extracting computations from operations is called *functorial semantics*, in which a model is represented as a mapping $M: \mathbf{Syntax} \rightarrow \mathbf{Semantics}$. The syntax defines a system-specific architectural design. Semantics are universal and provide a computational context to interpret specific models. Matrices, probabilities, proofs and dynamical equations all have their own rules for composition, corresponding to different semantic operads.

The mapping M encodes, for each operation, the data, assumptions and design artefacts (e.g. geometric models) needed to construct the relevant computational representations for the architecture, its components and the environment. From this, the system model as a whole is determined by composition in the semantic context. The algebras (State, KPI) described in §2b are typical examples, with syntax \mathcal{O} and taking semantic values in sets and functions. The mappings themselves, called *functors*, map types and operations (System, op) to their semantic values, while preserving how composition builds up complex operations.


The functorial perspective allows complementary models—e.g. system state versus KPIs—to be attached to the same design. This includes varying the semantic context as well as modelling details; see §5 for examples of non-deterministic semantics. Though functorial models may be radically different, they describe the *same system*, as reflected by the overlapping syntax.

In many cases, relevant models are *not* independent, like system state versus KPIs. Natural transformations, like the extraction homomorphism in §2b, provide a means to align ostensibly independent representations. Since models are mappings, we often visualize natural transformations as two-dimensional cells

$$\begin{array}{ccc}
 & \text{State} & \\
 & \curvearrowright & \\
 \mathcal{O} & & \text{Set.} \\
 & \curvearrowleft & \\
 & \text{KPI} &
 \end{array} \quad (2.2)$$

Formal conditions guarantee that when moving from syntax to semantics ([36], 13.2) or between representations ([34], 2.3.5), reasoning about how systems decompose hierarchically 'just works.'

Since functors and higher cells assure coherence with hierarchical decomposition, we can use them to build up a desired model in stages, working backwards from simpler models



This is a powerful technique for at least two reasons. First, complexity can be built up in stages by layering on details. Second, complex models built at later stages are partially validated through their coherence with simpler ones. The latter point is the foundation for lazy evaluation: many coarse models can be explored before ever constructing expensive models.

Separating out the different roles within a model encourages efficiency and reuse. An architecture (operation) developed for one analysis can be repurposed with strong coherence

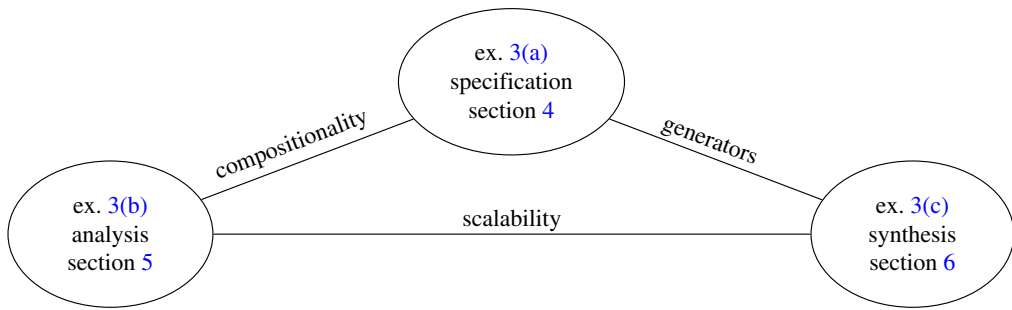


Figure 2. Organization of the paper around applied examples introduced in §3. (Online version in colour.)

between models (algebra instances) indexed by the same conceptual types. The syntax/semantics distinction also helps address some thornier meta-modelling issues. For example, syntactic types can distinguish conceptually distinct entities while still mapping to the same semantic entities. We obtain the flexibility of structural or duck typing in the semantics without sacrificing the type safety provided by the syntax.

3. Main examples

Though operads are general tools [34–36], we focus on two classes of operads applied to system design: wiring diagram operads and network operads. These are complementary. Wiring diagrams provide a top-down view of the system, whereas network operads are bottom-up. This section introduces three examples that help ground the exposition as in figure 2.

(a) Specification

Network operads describe atomic types of systems and ways to link them together with operations. These features enable: (1) specification of atomic building blocks for a domain problem; and (2) bottom-up synthesis of designs from atomic systems and links. A general theory of network operads [40–43] was recently developed under the Defense Advanced Research Projects Agency (DARPA) Complex Adaptive System Composition and Design Environment (CASCADE) program. Minimal data can be used to specify a functor—called a network model ([40], 4.2)—which constructs a network operad ([40], 7.2) customized to a domain problem.

The first example illustrates designs of search and rescue (SAR) architectures. The domain problem was inspired by the 1979 Fastnet Race and the 1998 Sydney to Hobart Yacht Race and we refer to it as the sailboat problem. It illustrates how network operads facilitate the specification of a model with combinatorial data called a network template. For example, figure 3 shows the carrying relationships between different system types to model (e.g. a Boat can carry a UAV (unmanned aerial vehicle) but a Helo cannot). These data specify a network operad $\mathcal{O}_{\text{Sail}}$ whose: (1) objects are lists of atomic system types; (2) operations describe systems carrying other systems; and (3) composition combines carrying instructions. We discuss this example in greater detail in §4.

(b) Analysis

A wiring diagram operad describes the interface each system exposes, making it clear what can be put together [26,44,45]. The designer has to specify precisely how information and physical quantities are shared among components, while respecting their interfaces. The operad facilitates top-down analysis of a design by capturing different ways to decompose a composite system.




Figure 3. Which types are allowed to carry other types—indicated with 1—specify an operad $\mathcal{O}_{\text{Sail}}$; f specifies that a HeLo (red circle) and a QD (blue circle) are carried by a Cut (green circle) and another QD (blue circle) is carried on the HeLo (red circle). (a) Examples of carrying relationships in $\mathcal{O}_{\text{Sail}}$. (b) Operation $f \in \mathcal{O}_{\text{Sail}}$ to specify carrying.

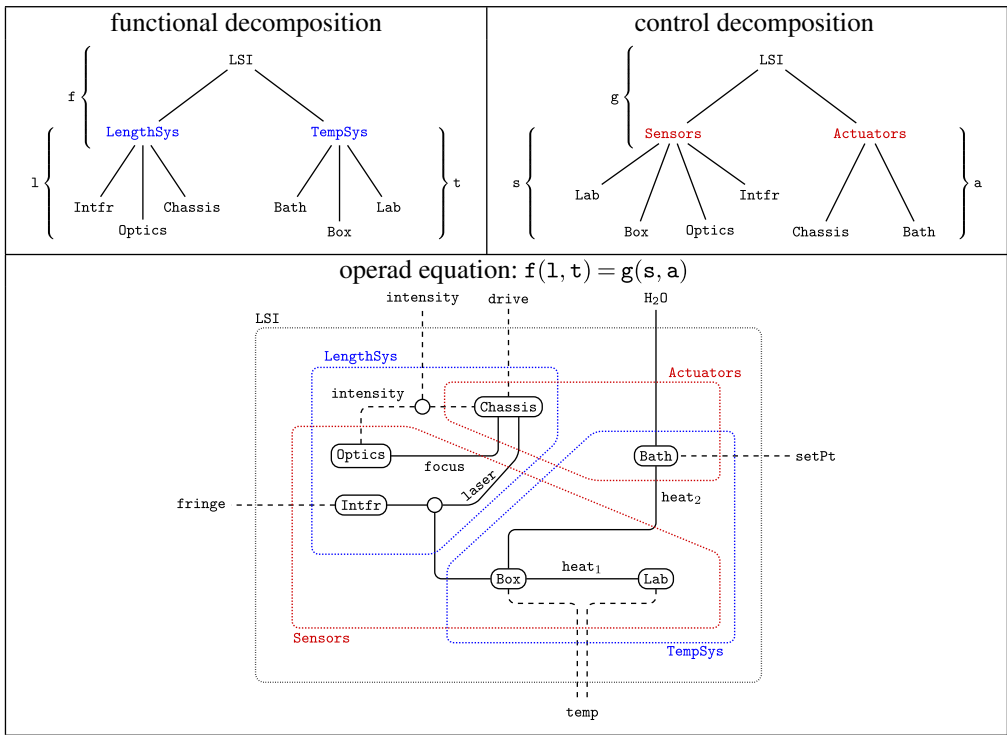


Figure 4. An equation in a wiring diagram operad expresses a common refinement of hierarchies. (Online version in colour.)

The second example analyses a precision-measurement system called the length scale interferometer (LSI) with wiring diagrams. It helps illustrate the qualitative features of operads over and above other modelling approaches and the potential to exploit their analytic power to separate concerns. Figure 4 illustrates joint analysis of the LSI to address different aspects of the design problem: functional roles of subsystems and control of the composite system. This analysis example supports these illustrations in §5.

(c) Synthesis

The third example describes the automated design of mission task plans for SAR using network operads. The SAR-tasking example illustrates the expressive power of applying existing operads and their potential to streamline and automate design synthesis. Figure 5a is analogous to figure 3,

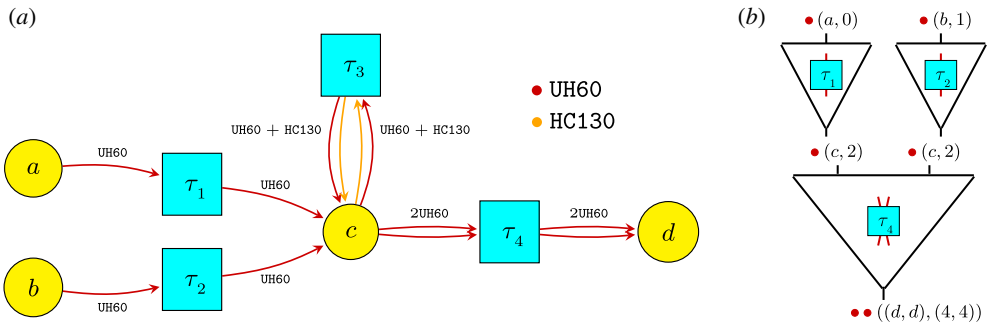


Figure 5. Primitive operations are composed for two UH60s to rendezvous at c and manoeuvre together to d . Each primitive operation is indexed by a transition; types and space–time points must match to compose. (a) Specification of primitive tasks := transitions. (b) Coordinate tasks to compose.

but whereas a sparse matrix specifies an architecture problem, here a Petri net is used to model coordinated groups of agents.

For the SAR-tasking problem, much of the complexity results from agents' need to coordinate in space and time, e.g. when a helicopter is refuelled in the air, as in τ_3 of figure 5a. To facilitate coordination, the types of the network operad are systematically extended via a network model whose target categories add space and time dimensions; compare, e.g. [41]. In this way, task plans are constrained at the level of syntax to enforce these key coordination constraints; e.g. figure 5b where two UH60s at the same space-time point $(c, 2)$ manoeuvre together to d . We describe automated synthesis for this example in §6.

4. Cookbook modelling of domain problems

In this section, we describe some techniques for constructing operads and their algebras, using an example-driven, cookbook-style approach. We emphasize recent developments for network operads and dive deeper into the SAR architecture problem.

(a) Network models

The theory of network models provides a general method to construct an operad \mathcal{O} by mixing combinatorial and compositional structures. Note that this lives one level of abstraction *above* operads; we are interested in *constructing* a language to model systems, e.g. for a specific domain. This provides a powerful alternative to coming up with operads one-by-one. A general construction allows the applied practitioner to cook-up a domain-specific syntax to compose systems by specifying some combinatorial ingredients.

The first step is to specify what the networks to be composed by \mathcal{O} look like. Often this is some sort of graph, but what kind? Are nodes typed (e.g. coloured)? Are edges symmetric or directed? Are loops or parallel edges allowed? What about n -way relationships for $n > 2$ (hyper-edges)? We can mix, match and combine such combinatorial data to define different *network models*, which specify the system types and kinds of relationships between them relevant to some domain problem. The network model describes the operations we need to compose the networks specific to the problem at hand.

Three compositional structures describe the algebra of operations. The *disjoint* or *parallel* structure combines two operations for networks with m and n nodes, respectively, into a single operation for networks with $m + n$ nodes. More restrictively, the *overlay* or *in series* structure superimposes two operations to design networks on n nodes. The former structure combines separate operations to support modular development of designs; the latter supports an incremental design process, either on top of existing designs or from scratch. The last ingredient




Figure 6. Parallel (\sqcup) and in series ($+$) compositional structures define how to combine operations.

permutes nodes in a network, which assures coherence between different ordering of the nodes. This last structure is often straightforward to specify. If it is not, one should consider if symmetry is being respected in a natural way.

We can distill the main idea behind overlay by asking, what happens when we add an edge to a network? It depends on the kind of network being composed by \mathcal{O} :

In a simple graph	but in a labeled graph	and in a multigraph
$x \text{ --- } y$	$x \text{ --- } y$	$x \text{ --- } y$
$+$ $x \text{ --- } y$	$+$ $x \text{ --- } y$	$+$ $x \text{ --- } y$
$x \text{ --- } y,$	$x \xrightarrow{2} y,$	$x \text{ --- } y.$

These differences are controlled by a *monoid*,³ which provides each $+$ shown. Above, the monoids are bitwise OR, addition, and maximum, respectively. As a further example, if edge addition is controlled by $\mathbb{Z}/2\mathbb{Z}$ then $+$ will have a toggling effect.

Consider simple graphs. Given a set of nodes n , write U_n for the set of all undirected pairs $i \neq j$ (a.k.a. simple edges $i-j$), so that $|U_n| = \binom{|n|}{2}$. Then we can represent a simple graph over n as a U_n -indexed vector of bits $\langle b_{i-j} \rangle$ describing which edges to ‘turn on’ for a design. Each bit defines whether or not to add an $i-j$ edge to the network and the overlay compositional structure is given by the monoid $SG(n) := \mathbf{Bit}^{U_n}$, whose $+$ is bitwise OR for the product over simple edges, i.e. adding $i-j$ then adding $i-j$ is the same as adding $i-j$ a single time. The disjoint structure $\sqcup: SG(m) \times SG(n) \rightarrow SG(m+n)$ forms the disjoint sum of the graphs g and h . Finally, permutations act by permuting the nodes of a simple graph. Together, these compositional structures define a network model $SG: \mathcal{S} \rightarrow \mathbf{Mon}$, which determines how operations are composed in the constructed network operad; see figure 6 (or [40], 3.2, 7.2) for complete technical details.

This definition has an analogue for \mathbb{N} -weighted graphs, $lG(n) := (\mathbb{N}, +)^{U_n}$, with overlay given by sum of edge weights and another for multi-graphs, $MG(n) := (\mathbb{N}, \max)^{U_n}$, with overlay equivalent to union of multisets; see ([40], 3.3, 3.4) for details. More generally, we can label edges with the elements of *any* monoid. Many of these examples are strange—binary addition makes edges cancel when they add—but their formal construction is straightforward; see ([40], Thm. 3.1).

Equivalently, we can view the undirected edges in U_n as generators, subject to certain idempotence and commutativity relations: $SG(n) := \langle e \in U_n \mid e \cdot e = e, e \cdot e' = e' \cdot e \rangle$. Here, the idempotence relations come from \mathbf{Bit} while the commutativity relations promote the single copies

³A set with a binary operation, usually written \cdot unless the operation is commutative ($m+n=n+m$). A monoid is always associative, $\ell \cdot (m \cdot n) = (\ell \cdot m) \cdot n$ and has a unit e satisfying $e \cdot m = m = m \cdot e$, e.g. multiplication of $n \times n$ matrices.

of **Bit** for each $i-j$ to a well-defined network model. Similar tricks work for lots of other network templates; we just change the set of generators to allow for new relationships. For example, to allow self-loops, we add loop edge generators $L_n = n + U_n$ to express relationships from a node i to itself. Likewise, network operads for directed graphs can be constructed by using generators $D_n = n \times n$, and one can also introduce higher-arity relationships.

In all cases, the formal definition of a network model assures that all the combinatorial and compositional ingredients work well together; one precise statement of ‘working well together’ is given in ([40], 2.3). Once a *network template*—which expresses minimal data to declare the ingredients for a network model—is codified in a theorem (as in [40], 3.1), it can be reused in a wide variety of domains to set up the specifics of composition.

(b) Cooking with operads

The prototype for network operads is a simple network operad, which models only one kind of thing, such as aircraft. The types of a simple network operad are natural numbers, which serve to indicate how many aircraft are in a design. Operations of the simple network operad are simple graphs on some number of vertices. For example, figure 6 above shows a simple network operad to describe a design for point-to-point communication between aircraft.

Structural network operads extend this prototype in two directions: (1) a greater diversity of things-to-be-modelled is supported by an expanded collection of types; and (2) more sorts of links or relationships between things are expressed via operations. To illustrate the impact of network templates, suppose we are modelling heterogeneous system types with multiple kinds of interactions. For simplicity, we consider simple interactions, which can be undirected or directed.

A *network template* need only declare the *primitive* ways system types can interact to define a network model, e.g. a list of tuples (directed : carrying, HeLo, Cut). These data are minimal in two ways: (1) *any* framework must provide data to specify potentially valid interactions; and (2) this approach allows *only* those interactions that make sense upon looking at the types of the systems involved. Thus, interactions must be syntactically correct when constructing system designs.

Presently, we will consider an example from the DARPA CASCADE program: the sailboat problem introduced in §3a. This SAR application problem was inspired by the 1979 Fastnet Race and the 1998 Sydney to Hobart Yacht Race, in which severe weather conditions resulted in many damaged vessels distributed over a large area. Both events were tragic, with 19 and six deaths, respectively, and remain beyond the scale of current SAR planning. Various larger assets—e.g. ships, aeroplanes, helicopters—could be based at ports and ferry smaller SAR units—e.g. small boats, quadcopters—to the search area.

Specifically, there were eight atomic types to model: $P = \{\text{Port}, \text{Cut}, \text{Boat}, \text{FW}, \text{FSAR}, \text{HeLo}, \text{UAV}, \text{QD}\}$. The primary relationship to specify a structural design is various assets carrying another types, so only one kind of interaction is needed: carrying. This relationship is directed; e.g. a cutter (Cut) can carry a helicopter (HeLo) but not the other way around.

Specifying allowed relationships amounts to specifying pairs of type $(p, p') \in P \times P$ such that type p' can carry type p ; see figure 3 for examples. Figure 3 data are extended to: (1) specify that Port can carry all types other than Port, UAV and QD; and (2) conform to an input file format to declare simple directed or undirected interactions, e.g. the JSON format in figure 7.

If another type of system or kind of interaction is needed, then the file is appropriately extended. For example, we can include buoys by appending Buoy to the array of colours and augmenting the relationships in the carrying node. Or, we can model the undirected (symmetric) relationship of communication by including an entry such as ‘undirected’: {‘communication’: {‘port’: [‘cut’: ...], ...}}. Moreover, modifications to network templates—such as ignoring (undirected : communication) or combining QD and UAV into a single type—naturally induce mappings between the associated operads ([40], 5.8).




Figure 7. After specifying $\mathcal{O}_{\text{Sail}}$, \mathbb{f} places a QD (blue circle) on a Cut (green circle) and another QD (blue circle) on a Helo (red circle). (a) Network template data to specify the operad $\mathcal{O}_{\text{Sail}}$, (b) Example operation $\mathbb{f} \in \mathcal{O}_{\text{Sail}}$.

Table 2. Example properties captured in algebra for sailboat problem including time on station (ToS), speed for search (S) and max speed (R) and sweep widths measuring search efficiency for target types person in water (PIW), crew in raft (CIR) and demasted sailboats (DS) adrift.

type	cost (\$)	ToS (hr)	speed (kn)		sweep width (nmi)		
			S	R	PIW	CIR	DS
Cut	200 M	∞	11	28	0.5	4.7	8.5
Boat	500 K	6	22	35	0.4	4.2	7.5
FW	60 M	9	180	220	0.1	2.2	7.6
FSAR	72 M	10	180	235	0.5	12.1	16.6
Helo	9 M	4	90	180	0.5	1.5	4.8
UAV	250 K	3	30	45	0.5	1.8	4.5
QD	15 K	4	35	52	0.5	1.5	4.8

(c) Cooking with algebras

Because all designs are generated from primitive operations to add edges, it is sufficient to define how primitive operations act in order to define an algebra. For the sailboat problem, semantics are oriented to enable the delivery of a high capacity for search—known in the literature as search effort ([46], 3.1)—in a timely manner. Given key parameters for each asset—e.g. speed, endurance, search efficiency across kinds of target and conditions, parent platform, initial locations—and descriptions of the search environment—e.g. expected search distribution, its approximate evolution over time—the expected number of surviving crew members found by the system can be estimated ([46], ch. 3).

Among these data, the parent platform and initial locations vary within a scenario and the rest describe the semantics of a given scenario. In fact, we assume all platforms must trace their geographical location to one of a small number of base locations, so that the system responds from bases, but is organized to support rapid search. Once bases are selected, the decision problem is a choice of operation: what to bring (type of the composite system) and how to organize it (operation to carry atomic systems). Data for the operational context specify a particular algebra; e.g. table 2. Just as for the operad, these data are lightweight and configurable.

(d) Related cookbook approaches

Though we emphasized network operads, the generators approach is often studied and lends itself to encoding such combinatorially data with a ‘template’, in a cookbook fashion. The

generators approach to ‘wiring’ has been developed into a theory of hypergraph categories [22,47], which induce wiring diagram operads. Explicit presentations for various wiring diagram operads are given in [45]. Augmenting monoidal categories with combinatorially specified data has also been investigated, e.g. in [48].

5. Functorial systems analysis

In this section, we demonstrate the use of functorial semantics in systems analysis. As in §2d, a functor establishes a relationship between a syntactic or combinatorial model of a system (components, architecture) and some computational refinement of that description. This provides a means to consider a given system from different perspectives, and also to relate those viewpoints to one another. To drive the discussion, we will focus on the LSI and its wiring diagram model introduced in §3b.

(a) Wiring diagrams

Operads can be applied to organize both qualitative and quantitative descriptions of hierarchical systems. Because operations can be built up iteratively from simpler ones to specify a complete design, different ways to build up a given design provide distinct avenues for analysis.

Figure 4 shows a wiring diagram representation of a precision measurement instrument called the LSI designed and operated by the US National Institute of Standards and Technology (NIST). Object types are system or component boundaries; figure 4 has: six components, the exterior, and four interior boundaries. Each boundary has an interface specifying its possible interactions, which are implicit in figure 4, but define explicit types in the operad.

An operation in this context represents one step in a hierarchical decomposition, as in §2a. For example, the blue boxes in figure 4 represent a functional decomposition of the LSI into length-measurement and temperature-regulation subsystems: $f: \text{LengthSys}, \text{TempSys} \rightarrow \text{LSI}$. These are coupled via (the index of refraction of) a laser interaction and linked to interactions at the system boundary. The operation f specifies the connections between blue and black boundaries.

Composition in a wiring diagram operad is defined by nesting. For this functional decomposition, two further decompositions l and t describe the components and interactions within LengthSys and TempSys , respectively. The wiring diagram in figure 4 is the composite $f(l, t)$.

This approach cleanly handles multiple decompositions. Here, the red boxes define a second, control-theoretic decomposition $g: \text{Sensors}, \text{Actuators} \rightarrow \text{LSI}$. Unsurprisingly, the system is tightly coupled from this viewpoint, with heat flow to maintain the desired temperature, mechanical action to modify the path of the laser and a feedback loop to maintain the position of the optical focus based on measured intensity. The fact that these two viewpoints specify the *same* system design is expressed by the equation: $f(l, t) = g(s, a)$; see §2c for related discussion.

(b) A probabilistic functor

Wiring diagrams can be applied to document, organize and validate a wide variety of system-specific analytic models. Each model is codified as an algebra, a functor from syntax to semantics (§2d). For the example of this section, all models have the same source (syntax), indicating that we are considering the same system, but the target semantics vary by application. We have already seen some functorial models: the algebras in §4c. These can be interpreted as functors from the carrying operad $\mathcal{O}_{\text{Sail}}$ to the operad of sets and functions **Set**. Though **Set** is the ‘default’ target for operad algebras, there are many alternative semantic contexts tailored to different types of analysis. Here, we target an operad of probabilities **Prob**, providing a simple model of non-deterministic component failure.

Table 3. Failure probabilities form an operad algebra for LSI component failure.

P_f	LengthSys	\mapsto	40%	P_g	Sensors	\mapsto	28%
	TempSys	\mapsto	60%		Actuators	\mapsto	72%
P_l	Interfer	\mapsto	10%	P_s	Lab	\mapsto	21.4%
	Optics	\mapsto	30%		Box	\mapsto	21.4%
	Chassis	\mapsto	60%		Optics	\mapsto	42.9%
					Interfer	\mapsto	14.3%
P_t	Bath	\mapsto	80%	P_a			
	Box	\mapsto	10%		Chassis	\mapsto	33.3%
	Lab	\mapsto	10%		Bath	\mapsto	66.7%

The data for the functor are shown in table 3. Model data are indexed by operations⁴ in the domain, an operad \mathcal{W} extracted from the wiring diagram in figure 4. The functor assigns each operation to a probability distribution that specifies the chance of a failure in each subsystem, assuming some error within the super-system. For example, the length measurement and temperature regulation subsystems are responsible for 40% and 60% of errors in the LSI, respectively. This defines a Bernoulli distribution P_f . Similarly, the decomposition t of the temperature system defines a categorical distribution with three outcomes: Box, Bath and Lab.

Relative probabilities compose by multiplication. This allows us to compute more complex distributions for nested diagrams. For the operation shown in figure 4, this indicates that the bath leads to nearly half of all errors ($60\% \times 80\% = 48\%$) in the system.

Operad equations must be preserved in the semantics. Since $f(1, t) = g(s, a)$, failure probabilities of source components do not depend on whether we think of them in terms of functionality or control. For the bath, this relative failure probability is

$$\underbrace{P_f}_{60\%} \times \underbrace{P_t}_{80\%} = 48\% = \underbrace{P_g}_{72\%} \times \underbrace{P_a}_{66.7\%},$$

and five analogous equations hold for the other source components.

Functorial semantics separates concerns: different operad algebras answer different questions. Here, we considered *if* a component will fail. The LSI example is developed further in ([49], 4) by a second algebra describing *how* a component might fail, with Boolean causal models to propagate failures. The two perspectives are complementary, and loc. cit. explores integrating them with algebra homomorphisms (§2d).

(c) Interacting semantics

Its toy-example simplicity aside, the formulation of a failure model $\mathcal{W} \rightarrow \mathbf{Prob}$, as in table 3 is limited in at least two respects. First, it tells us *which* components fail, but not *how* or *why*. Second, the model is static, but system diagnosis is nearly always a dynamic process. We give a high-level sketch of an extended analysis to illustrate the integration of overlapping functorial models.

The first step is to characterize some additional information about the types in \mathcal{W} (i.e. system boundaries). We start with the dual notions of *requirements* and *failure modes*. For example, in the temperature regulation subsystem of the LSI, we have

$$\begin{aligned} T_{\text{laser}} \leq 20.02^\circ\text{C} &\leftrightarrow T_{\text{laser too high}} \\ 19.98^\circ\text{C} \leq T_{\text{laser}} &\leftrightarrow T_{\text{laser too low}} \\ &\vdots \qquad \qquad \qquad \vdots \end{aligned}$$

⁴Types and operations, more generally, but the types carry no data in this simple example.

Requirements at different levels of decomposition are linked by traceability relations. These subsystem requirements trace up to the measurement uncertainty for the LSI as a whole. Dually, an out-of-band temperature at the subsystem level can be traced back to a bad measurement in the Box enclosure, a short in the Bath heater or fluctuations in the Lab environment.

Traceability is compositional: requirements decompose and failures bubble up. This defines an operad algebra⁵ $\text{Req} : \mathcal{W} \rightarrow \mathbf{Rel}^+$. Functoriality expresses the composition of traceability requirement across levels. See ([49], 5) for discussion of how to link these relations with table 3 data.

For dynamics, we need *state*. We start with a state space for each interaction among components. For example, consider the `laser` interaction coupling `Chassis`, `Intfr` and `Box`. The most relevant features of the laser are its vacuum wavelength λ_0 and the ambient temperature, pressure and humidity (needed to correct for refraction). This corresponds to a four-dimensional state-space (or a subset thereof)

$$\text{State}(\text{laser}) \cong \overbrace{[-273.15, \infty]}^{T_{\text{laser}}} \times \overbrace{[0, \infty]}^{P_{\text{laser}}} \times \overbrace{[0, 1]}^{RH_{\text{laser}}} \times \overbrace{[0, \infty]}^{\lambda_0} \subseteq \mathbb{R}^4.$$

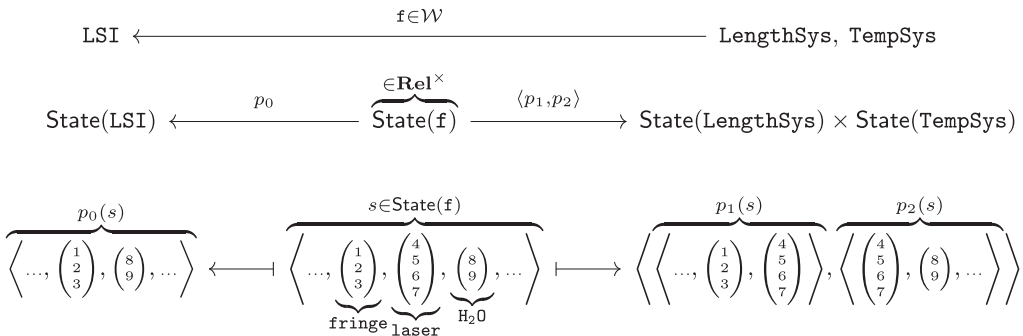
A larger product defines an *external state space* at each system boundary

$$\begin{aligned} \text{State}(\text{TempSys}) &= \text{State}(\text{laser}) \times \text{State}(\text{temp})^2 \times \text{State}(\text{setPt}) \times \text{State}(\text{H}_2\text{O}) \\ \text{State}(\text{Box}) &= \text{State}(\text{laser}) \times \text{State}(\text{temp}) \times \text{State}(\text{heat})^2 \\ &\vdots \end{aligned}$$

Similarly, we can define an *internal state space* for each operation by taking the product over all the interactions that appear in that diagram. We can decompose the internal state space in terms of either the system boundary or the components⁶

$$\begin{aligned} \text{State}(f) &\cong \text{State}(\text{LSI}) \times \overbrace{\text{State}(\text{laser})}^{\text{hidden variable}} \\ &\cong \text{State}(\text{LengthSys}) \times \underbrace{\text{State}(\text{TempSys})}_{\substack{\text{State}(\text{laser}) \\ \text{coupled variable}}} \end{aligned}$$

The projections from these (partial) products form a relation, and these compose to define a functor $\mathcal{W} \rightarrow \mathbf{Rel}^x$:



⁵Many operads are defined from ordinary categories using a symmetric monoidal products ([50], 7). If a category carries more than one product, we use a superscript to indicate which is in use. The disjoint union (+) corresponds to the disjunctive composition 'a failure in one component or another; soon we will use the Cartesian product \times to consider the conjunctive relationship between the state of one component *and* the other'.

⁶Coupled variables are formalized through a partial product called the pullback, a common generalization of the Cartesian product, subset intersection and inverse image constructions.




Figure 8. Requirement specification expressed as lifting problems. (a) Free specification, (b) top-down requirement and (c) bottom-up requirement.

Each requirement $R \in \text{Req}(X)$ defines a subset $|R| \subseteq \text{State}(X)$, and a state is *valid* if it satisfies all the requirements: $\text{Val}(X) = \bigcap_R |R|$. Using pullbacks (inverse image) we can translate validity to internal state spaces in two different ways. External validity (left square) checks that a system satisfies its contracts; joint validity (right square) couples component requirements to define the allowed joint states.

$$\begin{array}{ccccccc}
 \text{Val}(\text{LSI}) & \longleftarrow & \text{XVal}(f) & \longleftarrow \text{---} & \text{JVal}(f) & \longrightarrow & \text{Val}(\text{LengthSys}) \times \text{Val}(\text{TempSys}) \\
 \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 \text{State}(\text{LSI}) & \longleftarrow & \text{State}(f) & \longleftarrow \text{---} & \text{State}(f) & \longrightarrow & \text{State}(\text{LengthSys}) \times \text{State}(\text{TempSys}) \\
 & & p_0 & & \langle p_1, p_2 \rangle & &
 \end{array}$$

A requirement model is *sound* if joint validity entails external validity, corresponding to the dashed arrow above. With some work, one can show that these diagrams form the operations in an operad of entailments **Ent**; see ([21], 6) for a similar construction. The intuition is quite clear

$$\begin{array}{l}
 \text{component reqs.} \Rightarrow \text{subsystem reqs.} \\
 + \text{ subsystem reqs.} \Rightarrow \text{system reqs.} \\
 \hline
 \text{component reqs.} \Rightarrow \text{system reqs.}
 \end{array}$$


There is a functor $\text{Context} : \mathbf{Ent} \rightarrow \mathbf{Rel}^\times$, which extracts the relation across the bottom row of each entailment. Noting that the State relations occur in the validity entailment, we can reformulate requirement specification as a *lifting problem* (figure 8a): given functors State and Context, find a factorization Val making the triangle commute. The second and third diagrams (figure 8b,c) show how to extend the lifting problem with prior knowledge, in this case a top-level requirement and a known (e.g. off the shelf) component capability.

Finally, we are ready to admit dynamics, but it turns out that we have already done most of the work. All that is needed is to modify the spaces attached to our interactions. In particular, we can distinguish between static and dynamic state variables; for the *laser*, T, P and RH are dynamic while λ_0 is static. Now we replace the static values $T, P, RH \in \mathbb{R}$ by functions $T(t), P(t), RH(t) \in \mathbb{R}^\tau$, thought of as *trajectories* through the state space over a timeline $t \in \tau$. For example, we have

$$\text{Traj}(\text{laser}) \subseteq \overbrace{(\mathbb{R}^\tau)^3}^{T, P, RH} \times \overbrace{\mathbb{R}}^{\lambda_0} .$$

From this, we construct $\text{Traj} : \mathcal{W} \rightarrow \mathbf{Rel}^\times$ using exactly the same recipe as above. Trajectories and states are related by a pair of algebra homomorphisms *inst* and *const*. The first picks out a instantaneous state for each point in time, while the second identifies constant functions, which

describe fixed points of the dynamics:



The problem is that the state space explodes; function spaces are very large. Nonetheless, all of the system integration logic is identical, and using the entailment operad **Ent**, we can build in additional restrictions to limit the search space. In particular, we can restrict attention to the subset of functions that satisfies a particular differential equation or state-transition relationship. This drastically limits the set of valid trajectories, though the resulting set may be difficult to characterize and the methods for exploring it will vary by context.

(i) Related analytic applications

Wiring diagrams have an established applied literature for system design problems, e.g. [26,29,31, 44,51,52]. More broadly, the analytic strength of category theory to express compositionality and functorial semantics is explored in numerous recent applied works, e.g. engineering diagrams [29,53–58], Markov processes [59,60], database integration [27,29,31,33,61–63], behavioural logic [29,64–66], natural language processing [67–69], machine learning [70,71], cybersecurity [51,72–74], quantum computation [75–77] and open games [78–80].

6. Automated synthesis with network operads

An operad acting on an algebra provides a starting point to automatically generate and evaluate candidate designs. Formally correct designs (operations in some operad) combine basic systems (elements of some algebra of that operad) into a composite system.

(a) Sailboat example

Consider the sailboat problem introduced in §3a and revisited in §4b,c. Network operads describe assets and ports carrying each other while algebra-based semantics guided the search for effective designs by capturing the impact of available search effort.

To apply this model to automate design synthesis, algorithms explored designs within budget constraints based on costs in table 2. Exploration iteratively composed up to budget constraints and operational limits on carrying.⁷ With these analytic models, greater sophistication was not needed; other combinatorial search algorithms—e.g. simulated annealing—are readily applied to large search spaces. The most effective designs could ferry a large number of low-cost SAR units—e.g. quadcopters (QD)—quickly to the scene, e.g. via helicopters (HelO).

(b) Tasking example

Surprisingly, network operads—originally developed to design systems—can also be applied to ‘task’ them: in other words, declare their behaviour. An elegant example of this approach is given in [41] where ‘catalyst’ agents enable behavioural options for a system.

⁷Though not used for this application, it turns out that degree limits—e.g. how many quadcopters a helicopter can carry—can be encoded directly into operad operations; the relevant mathematics was worked out in [42].

(i) The SAR-tasking problem

The sailboat problem is limited by search: once sailboat crew members are found, their recovery is relatively straightforward. In hostile environments, recovery of isolated personnel (IP) can become very complex. The challenge is balancing the time criticality of recovery with the risk to the rescuers by judiciously orchestrating recovery teams.⁸ Consider the potential challenges of a large-scale earthquake during severe drought conditions that precipitates multiple wildfires over a large area. The 2020 Creek Fire near Fresno, CA required multiple mass rescue operations (MROs) to rescue over 100 people in each case by pulling in National Guard, Navy and Marine assets to serve as search and rescue units (SRUs) [82,83]. Though MRO scenarios are actively considered by US SAR organizations, the additional challenge of concurrent MROs distributed over a large area is not typically studied.

In this SAR-tasking example, multiple, geographically distributed IP groups compete for limited SRUs. The potential of coordinating multiple agent types—e.g. fire-fighting airplanes together with helicopters—to jointly overcome environment risks is considered as well as aerial refuelling options for SRUs to extend their range. Depending on available assets, recovery demands and risks, a mission plan may need to work around some key agent types—e.g. refuelling assets—and maximize the impact of others, e.g. moving protective assets between recovery teams.

Under CASCADE, tasking operations were built up from primitive tasks that coordinate multiple agent types to form a composite task plan. Novel concepts to coordinate teams of SRUs are readily modelled with full representation of the diversity of potential mission plan solutions.

(ii) Network models for tasking

A network model for tasking defines atomic agent types C and possible task plans for each list of agent types. Whereas a network model to design structure $\Gamma: S(C) \rightarrow \mathbf{Mon}$ has values that are possible graphical designs, a network model to task behaviour $\Lambda: S(C) \rightarrow \mathbf{Cat}$ has values that are categories whose morphisms index possible task plans for the assembled types; compare, e.g. ([41], Thm. 9). Each morphism declares a sequence of tasks for each agent, many of which will be coordinated with other agents.

If the system is comprised of only a single UH-60 helicopter, its possible tasks are captured in $\Lambda(\text{UH60})$. In this application, these tasks are paths in a graph describing ‘safe manoeuvres’. For unsafe manoeuvres, UH-60s travel in pairs, or perhaps with escorts such as a HC-130 or CH-47 equipped with a Modular Airborne Fire Fighting System (MAFFS). Anything one UH-60 can do, so can two, but not vice versa. Thus there is a proper inclusion $\Lambda(\text{UH60}) \times \Lambda(\text{UH60}) \subsetneq \Lambda(\text{UH60} \otimes \text{UH60})$. Similarly, $\Lambda(\text{UH60}) \times \Lambda(\text{HC130}) \subsetneq \Lambda(\text{UH60} \otimes \text{HC130})$ since once both a UH-60 and HC-130 are present, a joint behaviour of midair refuelling of the UH-60 by the HC-130 becomes possible. Formally, these inclusions are lax structure maps, e.g. $\Phi_{(\text{UH60}, \text{UH60})}: \Lambda(\text{UH60}) \times \Lambda(\text{UH60}) \rightarrow \Lambda(\text{UH60} \otimes \text{UH60})$, which specifies: given tasks for a single UH-60 (left coordinate) and tasks for another UH-60 (right coordinate), define the corresponding joint tasking of the pair. Here, the joint tasking is: each UH-60 operates independently within the safe graph. On the other hand, tasks in $\Lambda(\text{UH60} \otimes \text{UH60})$ to manoeuvre in unsafe regions cannot be constructed from independent taskings of each UH-60. Such tasks must be set for some pair or other allowed team, e.g. a CH-47 teamed with an UH-60.

(iii) Applying the cookbook: operads

While the above discussion sketches how to specify a network model for tasking, which constructs a network operad [40], these precise details [84] need not concern the applied practitioner.⁹ It

⁸The recovery of downed airman Gene Hambleton, call sign Bat 21 Bravo, during the Vietnam War is a historical example of ill-fated SAR risk management. Hambleton’s eventual recovery cost five additional aircraft being shot down and 11 deaths; for comparison, a total of 71 rescuers and 45 aircraft were lost to save 3883 lives during Vietnam War SAR [81].

⁹That is, a Petri net specifies the network model $\Lambda: S(C) \rightarrow \mathbf{Cat}$ to task behaviour. The construction of Λ [84] is similar to the construction described in ([41], Thm. 9), but adapted to coloured Petri nets whose transitions preserve the number of tokens

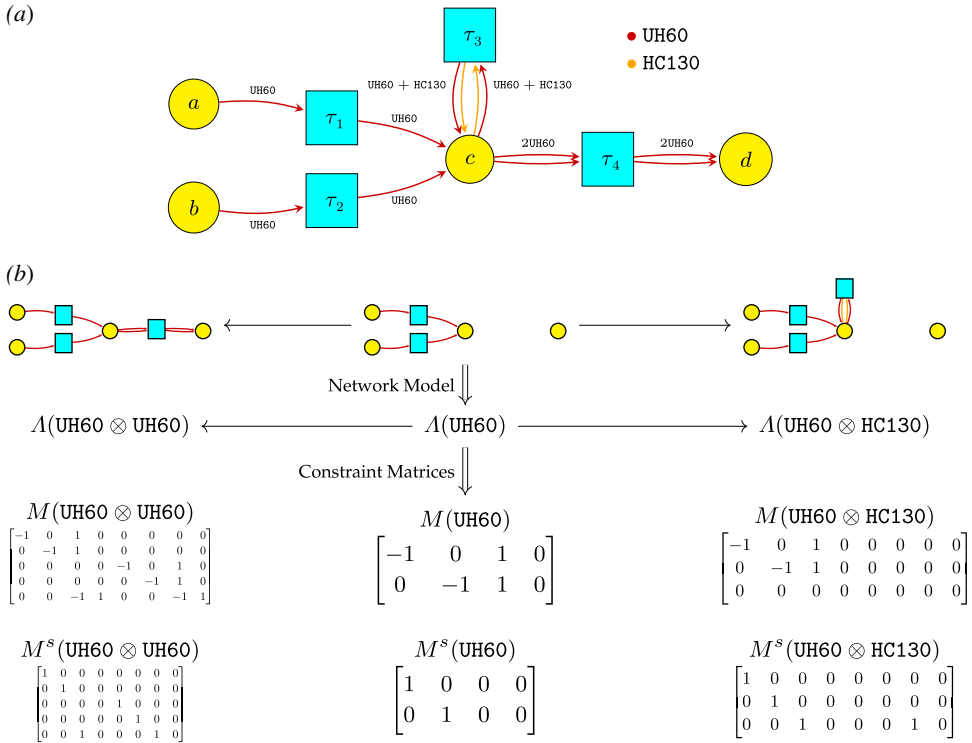


Figure 9. Specified primitive tasks determine an operad \mathcal{O}_{SAR} and a constraint program to explore operations. (a) Four primitive tasks specified in a Petri net; arcs indicate types involved in each task, (b) more primitive tasks become possible as available agent types increase. Type update matrices $M(-)$ and target to source constraint matrices $M^s(-)$ translate type changing and matching, resp.

is sufficient to provide a Petri net as a template, from which a network operad is constructed. Whereas a template to design structures defines the basic ways system types can interact, a template to task behaviour defines the primitive tasks for agent types C , which are token colours in the Petri net.

No specification of ‘staying put’ tasks are needed; these are implicit. All other primitive tasks are (sparsely) declared. For example, each edge of the ‘safe graph’ for a solo UH-60 declares: (1) a single agent of type UH60 participates in this ‘traverse edge’ task; and (2) participation is possible if a UH60 is available at the source of the edge. Likewise, each edge of the ‘unsafe graph’ for pairs of UH-60s should declare similar information for pairs, but what about operations to refuel an UH-60 with a HC-130? It turns out that transitions in a Petri net carry sufficient data [41,84] and have a successful history of specifying generators for a monoidal category [55,85,86]. The Petri net figure 9a shows examples where, for simplicity, tasks to traverse edges are only shown in the left to right direction. This sparse declaration is readily extended, e.g. to add recovery focused CH-47s, which tested their operational limits to rescue as many as 46 people during the 2020 Creek Fire-C and the set of transitions are augmented to encode the new options for primitive tasks.

This specification of syntax is almost sufficient for the SAR-tasking problem and would be for situations where only the sequence of tasks for each agent needs to be planned. When

of each colour; see, e.g. figure 9a. Compared with ([41], Thm. 9), C corresponds to token colours, rather than catalysts ([41], Def. 6), and species index discrete coordination locations. Target categories encode allowed paths for each atomic agent type, e.g. for figure 9a $\Lambda(\text{UH60})$ is (freely) generated by objects $\{a, b, c, d\}$ and morphisms $\tau_1: a \rightarrow c$ and $\tau_2: b \rightarrow c$, whereas $\Lambda(\text{HC130})$ is just by generated $\{a, b, c, d\}$ since no transition involves a single HC130. By describing each target category as an appropriate subcategory of a product of path categories, the symmetric group action is given permuting coordinates, which allows the role of each atomic agent in a task to be specified.

tasking SAR agents, *when* tasks are performed is semantically important because where and how long air-based agents ‘stay put’ impacts success: (1) fuel burned varies dramatically for ground versus air locations; and (2) risk incurred varies dramatically for safe versus unsafe locations. For comparison, in a ground-based domain without environmental costs, these considerations might be approximately invariant relative to the time tasks occur, and therefore, can be omitted from tasking syntax.

Timing information creates little added burden for building a template—transitions declaring primitive tasks need only be given durations derivable from scenario data—and it is technically straightforward to add a time dimension to the network model.

(iv) Constraints from syntax

A direct translation of primitive tasks to decision variables for a constraint program is possible. For syntax, the idea is very simple: enforce type-matching constraints on composing operad morphisms. Here, we will briefly indicate the original mixed integer linear program developed for SAR tasking; later this formulation was reworked to leverage the scheduling toolkit of the CPLEX optimization software package.

To illustrate the concept, let us first consider the constraint program for an operad to plan tasks without time and then add the time dimension.¹⁰ Operad types are translated to Boolean vectors m_j , whose entries capture individual agents at discrete coordination locations. Parallel composition of primitive operations is expressed with Boolean vectors Σ_j indexed over primitive tasks for specific agents. Type vectors m_j indicate the coordination location of each agent with value one; operation vectors Σ_j indicate which tasks are planned in parallel.

Assuming an operation with task vector Σ_j and source vector m_j , the target is $m_{j+1} = m_j + M\Sigma_j$, where M describes the relationship between source and target for primitive tasks. Rows of M correspond to primitive tasks while columns correspond to individual agents. The target to source constraint for a single step of in-series composition is $m_{j+1} \geq M^s \Sigma_{j+1}$ where M^s has rows that give requirements for each primitive task. Here, the LHS describes the target and the RHS describes the source. The inequality appears to allow for implicit identities for agents without tasking, e.g. if Σ_j is a zero vector, then $m_{j+1} = m_j$. This constraint prevents an individual agent from being assigned conflicting tasks or ‘teleporting’ to begin a task.

As seen in figure 9b, additional agents: (1) enable more primitive tasks, indexed by Petri net transitions (top two rows); and (2) expand the type vector/matrix column dimension to account for new agent-location pairs and increase the matrix row dimension to account for new tasks (bottom two rows). For example, the first four rows of $M(\text{UH60} \otimes \text{UH60})$ correspond to the image of $\Lambda(\text{UH60}) \times \Lambda(\text{UH60})$ in $\Lambda(\text{UH60} \otimes \text{UH60})$. The last row corresponds to a new task, τ_4 , for the available pair of UH-60s. During implementation, the constraints can be declared task-by-task/row-by-row to sparsely couple the involved agents. Once a limit on the number of steps of in series composition is set—i.e. a bound for the index j is given—a finite constraint program is determined.


Time is readily modelled discretely with tasks given integer durations. This corresponds to a more detailed network model, Λ_t , whose types include a discrete time index; see figure 5b for example operations. Under these assumptions, one simply replaces the abstract steps of in series composition with a time index and decomposes M and Σ_j by the duration d of primitive tasks

$$m_t + \sum_{d=1}^{d_{\max}} M_d \Sigma_{t-d,d} = m_{t+1}; \quad m_{t+1} \geq \sum_{d=1}^{d_{\max}} M_d^s \Sigma_{t+1,d}$$

so that $\Sigma_{t,d}$ describes tasks beginning at time t ; the inequality allows for ‘waiting’ operations. One can also model tasks more coarsely—with $\Lambda_\bullet: \mathbb{N}(C) \rightarrow \mathbf{Cat}$ —to construct an operad to task counts of agents without individual identity. Then, the type vectors m_j (resp., operation vectors Σ_j)

¹⁰Simply increasing dimensionality is not computationally wise—which was the point of exploring the CPLEX scheduling toolkit to address the time dimension—but this model still serves as a conceptual reference point.

have integer entries to express agent counts (resp., counts of planned tasks) with corresponding reductions in dimensionality. These three levels of network models



naturally induce morphisms of network operads¹¹ ([40], 6.18) and encode mappings of syntactic variables that preserve feasibility. In particular, the top two levels describe a precise mapping from task scheduling (highest) to task planning (middle). The lowest level Λ_\bullet forgets the individual identity of agents, providing a coarser level for planning.

This very simple idea of enforcing type-matching constraints is inherently natural.¹² However, further research is needed to determine if this natural hierarchical structure can be exploited by algorithms—e.g. by branching over pre-images of solutions to coarser levels—perhaps for domains where operational constraints coming from algebras are merely a nuisance, as opposed to being a central challenge for SAR planning. For instance, a precise meta-model for planning and scheduling provides a common jumping off point to apply algorithms from those two disciplines.

(v) Applying the cookbook: algebras

Because the operad template defines generating operations, specifying algebras involves: (1) capturing the salient features of each agent type as its internal state; and (2) specifying how these states update under generating morphisms, including, for operads with time, the implicit ‘waiting’ operations. For the SAR-tasking problem, the salient features are fuel level and cumulative probability of survival throughout the mission. Typical primitive operations will not increase these values; fuel is expended or some risk is incurred. The notable exception is refuelling operations, which return the fuel level of the receiver to maximum. By specifying the non-increasing rate for each agent–location pair, the action of ‘waiting’ operations are specified. In practice, these data are derivable from environmental data for a scenario so that end users can manipulate them indirectly.

(vi) Operational constraints from algebras

Salient features of each agent type are captured as auxiliary variables determined by syntactic decision variables. The values of algebra variables are constrained by update equations, e.g. to update fuel levels for agents with $\max(f_j + F\Sigma_j, f_{\max}) = f_{j+1}$, where f_{\max} specifies max fuel capacities. Having expressed the semantics for generating operations, one can enforce additional operational constraints, e.g. safe fuel levels: $f_{j+1} \geq f_{\min}$.

(vii) Extending the domain of application

As noted above, this sparse declaration of a tasking domain is readily extended, e.g. to add a new atomic type or new ways for agents to coordinate. Syntactically, this amounts to new elements of C or transitions to define primitive tasks. Semantics must capture the impact of primitive operations on state, which can be roughly estimated initially and later refined. This flexibility is especially useful for rapid prototyping of ‘what if’ options for asset types and behaviours, as the wildfire SAR-tasking problem illustrates.

Suppose, for example, that we wanted to model a joint SAR and fire-fighting problem. Both domains are naturally expressed with network operads to task behaviour. Even if the specification formats were independently developed: (1) each format must encode the essential combinatorial

¹¹Strictly speaking, the coarsest (lowest) level is not a network model; its domain is a free commutative monoidal category. Nevertheless, a completely analogous construction produces a typed operad fitting into this diagram.

¹²i.e. operad morphisms push forward feasible assignments variables in the domain to feasible assignments in the co-domain.

data for each domain; and (2) category theory provides a method to integrate domain data: construct a pushout. Analogous to taking the union of two sets along a common intersection, one identifies the part of the problem common to both domains—e.g. MAFFS-equipped HC-130s and their associated tasks appearing in both domains—to construct a cross-domain model

$$\begin{array}{ccc} \text{Spec}_{\cap} & \longrightarrow & \text{Spec}_{\text{SAR}} \\ \downarrow & & \downarrow \\ \text{Spec}_{\text{FF}} & \longrightarrow & \text{Spec}_{\cup} \end{array}$$

The arrows in this diagram account for translating the file format for the overlap into each domain-specific format and choosing a specific output format for cross-domain data.

On the other hand, suppose that the machine readable representation of each domain was tightly coupled to algorithms, e.g. mathematical programming for SAR and a planning framework for fire fighting. There is no artefact suitable for integrating these domains since expression was prematurely optimized. We describe a general workflow to separate specification from representation and exploitable data structures and algorithms in §7e.

(c) Other examples of automated synthesis

Though network templates facilitate exploration from atoms, how to explore valid designs is a largely distinct concern from defining the space of designs, as discussed in §1.

(i) Novel search strategies via substitution

For example, in the DARPA Fundamentals of Design (FUN Design) program, composition of designs employed a genetic algorithm (GA). FUN Design focused on generating novel conceptual designs for mechanical systems, e.g. catapults to launch a projectile. Formulating this problem with network operads followed the cookbook approach: there were atomic types of mechanical components and basic operations to link them.

The operad-based representation provided guarantees of design feasibility and informed how to generalize the GA implementation details. Specifically, composition for atomic algebra elements defined genetic data; crossover produced child data to compose from atoms; and mutation modified parameters of input algebra elements. Crafting a crossover step is typically handled case-by-case while this strategy generalizes to other problems that mix combinatorial and continuously varying data, provided these data are packaged as an operad acting on an algebra. Guarantees of feasibility dramatically reduced the number unfit offspring evaluated by simulation against multiple fitness metrics. Moreover, computational gains from feasibility guarantees increase as the design population becomes more combinatorially complex.

(ii) Integrated structure and behaviour

Large classes of engineering problems compose components to form an ‘optimized’ network, e.g. in chemical process synthesis, supply chains and water purification networks [87–90]. Given a set of inputs, outputs and available operations (process equipment with input and output specification), the goal is to identify the optimal state equipment networks for behavioural flows of materials and energy. A given production target for outputs is evaluated in terms of multiple objectives such as environmental impact and cost. For example, the chemical industry considers the supply chain, production and distribution network problem [90] systematically as three superstructure optimization problems that can be composed to optimize enterprise level, multi-subsystem structures. Each sub-network structure is further optimized for low cost and other metrics including waste, environmental impact and energy costs. The operadic paradigm would provide a lens to generalize and refine existing techniques to jointly explore structure and behaviour.

CASCADE prototyped integrated composition of structure and behaviour for distributed logistics applications. Here, an explicit resupply plan to task agents was desired. Structural composition was needed to account for the resupply capacity for heterogeneous delivery vehicles and the positioning of distributed resupply depots. Probabilistic models estimated steady state resupply capacities of delivery fleet mixes to serve estimates of demand. First, positioning resupply locations applied hill climbing to minimize the expected disruption of delivery routes when returning to and departing from resupply locations. Second, this disruption estimate was used to adjust the resupply capacity estimate of each delivery asset type. Third, promising designs were evaluated using a heuristic task planning algorithm. At each stage, algorithms focused on finding satisficing solutions which allowed broad and rapid explorations of the design and tasking search space.

(iii) Synthesis with applied operads and categories

Research activity to apply operads and monoidal categories to automated design synthesis is increasing. Wiring diagrams have been applied to automate protein design [30,32] and collaborative design ([29], ch. 4) of physical systems employing practical semantic models and algorithms [91–94]. Software tools are increasingly focused on scaling up computation, e.g. [95–98], as opposed to software to augment human calculation, as in [77,99,100], and managing complex domains with commercial-grade tools [27,61–63]. Recent work to optimize quantum circuits [75,76] leverages such developments. The use of wiring diagrams to improve computational efficiency via normal forms is explored in [101].

In the next section, we discuss research directions to develop the meta-modelling potential of applied operads to: (1) decompose a problem within a semantic model to divide and conquer; and (2) move between models to fill in details from coarse descriptions. We also discuss how the flow of representations used for SAR—network template, operad model of composition, exploitation data structures and algorithms—could be systematized into a reusable software framework.

7. Towards practical automated analysis and synthesis

In this section, we describe lessons learned from practical experiences with applying operads to automated synthesis. We frame separation of concerns in the language of operads to describe strategies to work around issues raised by this experience. This gives not only a clean formulation of separation but also a principled means to integrate and exploit concerns.

(a) Lessons from automated synthesis

The direct, network template approach facilitates correct and transparent modelling for complex tasking problems. However, computational tractability is limited to small problems, relative to the demands of applications. More research is needed to develop efficient algorithms that break up the search into manageable parts, leveraging the power of operads to separate concerns.

Under CASCADE, we experimented with the CPLEX scheduling toolkit to informally model across levels of abstraction and exploit domain-specific information. In particular, generating options to plan, but not schedule, key manoeuvres with traditional routing algorithms helped factor the problem effectively. These applied experiments were not systematized into a formal meta-modelling approach, although our prototype results were promising. Specification of these levels—as in §4—and controlling the navigation of levels using domain-specifics would be ideal.

The FUN DESIGN genetic algorithm approach illustrates the potential operads have to: (1) generalize case-by-case methods¹³; (2) separate concerns, in this case by leveraging the operad

¹³In fact, applying genetic algorithms to explore network structures was inspired by the success of NeuroEvolution of Augmenting Topologies (NEAT) [102] to generate novel neural network architectures.

syntax for combinatorial crossover and algebra parameters for continuous mutation; and (3) guarantee correctness as complexity grows. Distributed logistics applications in CASCADE show the flexibility afforded by multiple stage exploration for more efficient search.

(b) Formal separation of concerns

We begin by distinguishing *focus* from *filter*, which are two ways operads separate. Focus selects *what* we look at, while filter captures *how* we look at it. These are questions of syntax and semantics, respectively. To be useful, the *what* of our focus must align with the *how* of the filter.

Separation of focus occurs within the syntax operad of system maps. In §2a, four trees correspond to different views on the same system. We can zoom into one part of the system while leaving other portions black-boxed at a high level. Varying the target type of an operation changes the scope for system composition, such as restricting attention to a subsystem.

Filtering, on the other hand, is semantic; we choose which salient features to model and which to suppress, controlled by the semantic context used to ‘implement’ the operations. As described in §5b, the default semantic context is **Set** where: (1) each type in the operad is mapped to a set of possible instances for that type; and (2) each operation is mapped to a function to compose instances. Instances or algebra elements for the sailboat problem (§4) describe the key features of structural system designs. For SAR tasking (§6), mission plan instances track the key internal states of agents—notably fuel and risk—throughout its execution. Section 5 illustrates alternative semantic contexts as such probability **Prob** or relations between sets **Rel**.

Focus and filter come together to solve particular problems. The analysis of the LSI system in §5 tightly focuses the syntax operad \mathcal{W} to include only the types and operations from figure 4. Formally, this is accomplished by considering the image of the generating types and operations in the operad of port-graphs ([49], 3). This tight focus means semantics need only be defined for LSI components. In each SAR-tasking problem of §6, an initial, source configuration of agent types is given, narrowing the focus of each problem. The SAR focus is much broader because an operation to define the mission plan must be constructed. Even so, semantics filter down to just the key features of the problem and how to update them when generating operations act.

Functorial semantics, as realized by an operad algebra $A: \mathcal{O} \rightarrow \mathbf{Sem}$, helps factor the overall problem model to facilitate its construction and exploitation. For example, we can construct the probabilistic failure model in table 3 by normalizing historical failures. First, we limit focus from all port-graphs \mathcal{P} to \mathcal{W} then semantics for counts in \mathbb{N}^+ , an operad of counts and sums, are normalized to obtain probabilities in **Prob**:

$$\begin{array}{ccc} \mathcal{W} & \longrightarrow & \mathbb{N}^+ \\ \downarrow & \searrow A & \downarrow \\ \mathcal{P} & & \mathbf{Prob}. \end{array}$$

The power to focus and filter is amplified because we are not limited by a single choice of how to filter. In addition to *limiting* focus with the source of an operad algebra, we can *simplify* filters. Such natural transformations between functors are ‘filters of filters’ that align different compositional models precisely, e.g. requirements over state (§5c) or timed scheduling over two levels of planning (§6b). In this first case, the syntax operad \mathcal{W} stays the same and semantics are linked by an algebra homomorphism (§2d). In the second case, both the operad and algebra must change to determine simpler semantics, e.g. to neglect the impact of waiting operations, which bound performance. Such precision supports automation to explore design space across semantic models and aligns the ability to focus within each model. By working backward relative to the construction process, we can lift partial solutions to gradually increase model fidelity, e.g. exploring schedules over effective plans. This gives a foundation for lazy evaluation during deep exploration of design space, which we revisit in §7e.

For a simple but rich example of these concepts working together, consider the functional decomposition $f(1, \tau)$ in figure 4. We could model the length system 1 using rigid-body

dynamics, the temperature system \mathfrak{t} as a lumped-element model and super-system \mathfrak{f} as a computation (Edlén equation) that corrects the observed fringe count based on the measured temperature:

$$\begin{array}{ccccc}
 \{1\} & \longrightarrow & \mathcal{N}_{\text{mech}} & \xrightarrow{\text{Rigid}} & \text{Dyn} \\
 \downarrow & & \Downarrow & & \downarrow \text{Impl} \\
 \mathcal{W} & \longrightarrow & \mathcal{N}_{\text{comp}} & \xrightarrow{\text{Edlén}} & \text{Type} \\
 \uparrow & & \Uparrow & & \uparrow \text{Impl} \\
 \{\mathfrak{t}\} & \longrightarrow & \mathcal{N}_{\text{therm}} & \xrightarrow{\text{Lump}} & \text{Dyn}
 \end{array} \tag{7.1}$$

The upper and lower paths construct implementations of dynamical models based on the aforementioned formalisms. The centre path implements a correction on the data stream coming from the interferometer, based on a stream of temperature data. The two natural transformations indicate extraction of one representation, a stream of state values, from the implementation of the dynamical models. Composition in \mathcal{W} then constructs the two data streams and applies the correction.

A key strength of the operadic paradigm is its genericity: the same principles of model construction, integration and exploitation developed for measurement and SAR apply to all kinds of systems. In principle, we could use the same tools and methodology to assemble transistors into circuits, unit processes into chemical factories and genes into genomes. The syntax and semantics change with the application, but the functorial perspective remains the same. For the rest of this section, we describe research directions to realize such a general purpose vision to decompose a complex design problem into subproblems and support rapid, broad exploration of design space.

(c) Recent advancements, future prospects and limits

Progress driven by applications. Section 4 describes how cookbook-style approaches enable practitioners to put operads to work. Generative data define a domain and compositionality combines it into operads and algebras to separate concerns. Network operads [40–42] were developed *in response to the demands of applications* to construct operads from generative data. Section 5 describes rich design analysis by leveraging multiple decompositions of complex systems and working across levels of abstraction. Focusing on a specific applied problem—the LSI at NIST—provided further opportunities for analysis since *model semantics need only be defined for the problem at hand*; see also equation (7.1). Progress in streamlining automated synthesis from building blocks is recounted in §6 where the domain drives coordination requirements to task behaviour.

Prospects. If interactions between systems are well-understood (specification) and can be usefully modelled by compositional semantics (analysis), then automated design synthesis leveraging separation for scalability becomes possible. For instance, most references from the end of §5 correspond to domains that are studied with diagrams that indicate interactions and have associated compositional models. This allows intricate interactions to be modelled—compare, e.g. classical [18] versus quantum [75,76,103] computing—while unlocking separation of concerns. Cookbook and focused approaches guide practitioners to seek out the minimal data needed for a domain problem—as in the examples presented—but operads for design *require* compositional models.

Limitations. We note three issues limiting when operads apply: (1) key interactions among systems and components are *inputs*; (2) not all design problems become tractable via decomposition and hierarchy; and (3) there is no guarantee of compositional semantics to exploit. For instance, though the interactions for the n -body problem are understood (1), this does not lend itself to decomposition (2) or exploitable compositional semantics (3).

Whitney [20] notes that integral mechanical system design must address safety issues at high power levels due to challenging, long-range interactions. Some aspects of mechanical system design may yield to operad analysis—e.g. bond graphs [57] or other sufficiently ‘diagrammatic’ models—but others may not.

Both examples illustrate how overnumerous or long-range interaction can lead to (2). Operads can work at the system rather than component level if system properties can be extracted into compositional models. However, operads do not provide a means to extract such properties or understand problems that are truly inseparable theoretically or practically.

(d) Research directions for applied operads

We now briefly overview research directions toward automated analysis and synthesis.

Operad-based decomposition and adaptation. Decomposition, ways a complex operation can be broken down into simpler operations, is a dual concept to the composition of operations. Any subsystem designed by a simpler operation can be adapted: precisely which operations can be substituted is known, providing a general perspective to craft algorithms. To be practical, the analytic questions of *how* to decompose and *when* to adapt subsystems must be answered.

One research direction applies the lens of operad composition to abstract and generalize existing algorithms that exploit decomposition, e.g. to: (1) generalize superstructure optimization techniques discussed in §6c; (2) extend the crossover and mutation steps for the FUN DESIGN work §7a, which are global in the sense that they manipulate full designs, to local steps which adapt parts of a design, perhaps driven by analysis to re-work specific subsystems; and (3) explore routing as a proxy for tasking planning, analysing foundational algorithms like Ford & Fulkerson [104] and decomposition techniques such as contraction hierarchies [105]. An intriguing, but speculative, avenue is to attempt to learn how to decompose a system or select subsystems to adapt in a data-driven way, so that the operad syntax constrains otherwise lightly supervised learning. A theoretical direction is to seriously consider the dual role of decomposition, analogous to Hopf and Frobenius algebra [106], and attempt to gain deeper understanding of the interplay of composition and decomposition, eventually distilling any results into algorithms.¹⁴

Multiple levels of modelling. The LSI example shows how a system model can be analysed to address different considerations. This sets the stage to adapt a design—e.g. bolster functional risk points and improve control in back and forth fashion—until both considerations are acceptable. Applied demonstrations for SAR tasking suggest a multi-level framework: (1) encoding operational concepts; (2) planning options for key manoeuvres; and (3) multi-stage planning and scheduling to support these manoeuvres.

Unifying top-down and bottom-up points of view. We have laid out the analytic—exemplified by wiring diagrams—and synthetic—exemplified by network operads—points of view for complex systems. Even if the goal is practical automated synthesis, scalability issues promote analytic decomposition and abstraction to efficiently reason towards satisficing solutions. Two approaches to unification include: (1) create a combined syntax for analysis and synthesis, a ‘super operad’ combining both features; and (2) act by an analytic operad on the synthetic syntax, extending composition of operations. While the former approach is arguably more unified, the latter more clearly separates analysis and synthesis and may provide a constructive approach to the former.

(e) Functorial programming with operads

At this point, experience implementing operads for design suggests a software framework. While conceptually simple, this sketch helps clarify the practical role of a precise meta-model.

Rather than working directly with operads to form a core meta-modelling language, cf. [14], a workflow akin to popular frameworks for JavaScript development would put developers in the driver’s seat: adopters focus on controlling the flow of data and contribute to an ecosystem

¹⁴For example, Bellman’s principle of optimality is *decompositional*, i.e. parts of an optimal solution are optimal.

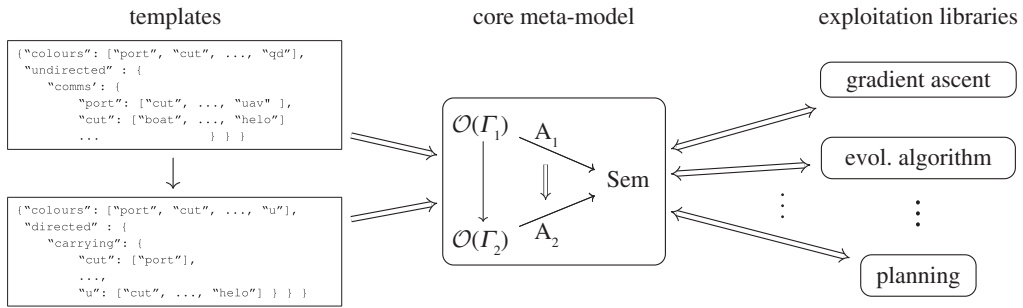


Figure 10. A software framework to leverage a meta-model: templates define each level and how to move between, libraries exploit each level, and core meta-model facilitates control across levels.

of libraries for lower-level data processing. Achieving this requires work before and after the meta-model. First, transferable methods get an applied problem into operads (figure 10, left). As in §4, these data construct operads and algebras to form the core meta-model. Core data feed explicitly exploitable data structures and algorithms to analyse (§5) and automatically construct (§6) complex systems (figure 10, right). On the far left, end user tools convert intent to domain inputs. Rightmost, libraries access exploitation data structures and algorithms, including those exploiting the syntax and semantics separation or substitution and adaptation. At the centre, the core meta-model guarantees that the scruffier ends of the framework exposed to end users and developers are correctly aligned and coherently navigated.

This framework provides significant opportunities to separate concerns compared with other approaches. Foremost, the core model separates syntax from semantics. As noted in §1, applied methods tend to conflate syntax and semantics. For instance, aggregate programming [107] provides: (1) semantics for networked components with spatial and temporal extent; and (2) interactions are proximity-based. The former feature is powerful but limiting: by choosing a single kind of semantics, modelling is wedded to the scales it abstracts well. The individual component scale is not modelled, even syntactically, which would complicate any attempt to align with other models. The latter precludes syntactic declaration of interactions—e.g. to construct architectures not purely based on proximity—and the absolute clarity about what can be put together provided by the operad syntax. Relative to computational efforts to apply operads or monoidal categories, e.g. [95,96,98], this sketch places greater emphasis on specification and exploitation: specification of a domain is possible without exposing the meta-model, algorithms searching within each model are treated as black boxes that produce valid designs. Separate specification greatly facilitates set up by experts in the domain, but not the meta-model. Separate exploitation encourages importing existing data structures and algorithms to exploit each model.

(f) Open problems

The software framework just sketched separates out the issues of practical specification, meta-modelling and fast data structures and algorithms. We organize our discussion of open problems around concrete steps to advance these issues. In our problem statements, ‘multiple’ means at least three to assure demonstration of the genericity of the operadic paradigm.

Practical specification. The overarching question is whether the minimal combinatorial data, which can specify operads, their algebras and algebra homomorphisms in theory, can be practically implemented in software. We propose the following problems to advance the state-of-the-art for network template specification of operads described in §4:

- (i) Demonstrate a specification software package for operad algebras for multiple domains.
- (ii) Develop specification software for algebra homomorphisms to demonstrate correctly aligned navigation between multiple models for a single domain.

- (iii) Develop and implement composition of specifications to combine multiple parts of a domain problem or integrate multiple domains.

This last point is in line with the discussion of extending a domain in §6b and motivates a need to reconcile independently developed specification formats.

- (iv) Demonstrate automatic translation across specification formats.

Core meta-model. As a practical matter, state-of-the-art examples exercise general principles of the paradigm but do not leverage general purpose software to encode the meta-model.

- (v) Develop and demonstrate reusable middleware to explicitly encode multiple semantic models and maps between them which (a) takes inputs from specification packages; and (b) serves as a platform to navigate models.

We have seen rich examples of focused analysis with wiring diagrams in §5 and automated composition from building blocks in §6. Theoretically, there is the question of integrating the top-down and bottom-up perspectives:

- (vi) Develop unified foundations to integrate: (a) analytic and synthetic styles of operads; and (b) composition with decomposition.

Potential starting points for these theoretical advancements are described in §7d. Developing understanding of limitations overviewed in §7c requires engagement with a range of applications:

- (vii) Investigate limits of operads for design to (a) identify domains or specific aspects of domains lacking minimal data; (b) demonstrate the failure of compositionality for potentially useful semantics; and (c) characterize complexity barriers due to integrality.

Navigation of effective data structures and algorithms. Lastly, there is the question of whether coherent navigation of models can be made practical. This requires explicit control of data across models and fast data structures and algorithms within specific models. The general-purpose evolutionary algorithms discussed in §7c motivate:

- (viii) Develop reusable libraries that exploit (a) substitution of operations and instances to adapt designs and (b) separation of semantics from syntax.

SAR-tasking experience and prototype explorations for distributed logistics illustrate the need to exploit moving *across* models:

- (ix) Develop and demonstrate general purpose strategies to exploit separation across models via hierarchical representation of model fidelity, e.g.: (a) structure over behaviour; and (b) planning over scheduling.
- (x) Quantify the impact of separation of concerns on: (a) computational complexity; and (b) practical computation time.

For this last point, *isolating* the impact of each way to separate concerns is of particular interest to lay groundwork to systematically analyse complex domain problems. Finally, there is the question of demonstrating an end-to-end system to exploit the operadic, meta-modelling paradigm.

- (xi) Demonstrate systematic, high-level control of iteration, substitution and moving across multiple models to solve a complex domain problem.

- (xii) Develop high-level control framework—similar to JavaScript frameworks for UI—or programming language—similar to probabilistic programming—to systematically control iteration, substitution and movement across multiple models.

8. Conclusion

Operads provide a powerful meta-language to unite complementary system models within a single framework. They express multiple options for decomposition and hierarchy for complex designs, both within and across models. Diverse concerns needed to solve the full design problem are coherently separated by functorial semantics, maintaining compositionality of subsystems. Each semantic model can trade-off precision and accuracy to achieve an elegant abstraction, while algorithms exploit the specifics of each model to analyse and synthesize designs.

The basic moves of iteration, substitution and moving across multiple models form a rich framework to explore design space. The trade-off is that the technical infrastructure needed to fully exploit this paradigm is daunting. Recent progress has lowered barriers to specify domain models and streamline tool chains to automatically synthesize designs from basic building blocks. Key parts of relevant theory and its implementation in software have been prototyped for example applications. Further research is needed to integrate advancements in automatic specification and synthesis with the analytic power of operads to separate concerns. To help focus efforts, we described research directions and proposed some concrete open problems.

Ethics. This article does not present research with ethical considerations.

Data accessibility. This article has no additional data.

Authors' contributions. All authors contributed to the development of the review, its text and provided feedback and detailed edits on the work as a whole. J.F. coordinated the effort and led the writing of §6. S.B. led §5 and co-led §§2, 3 and 4 with J.F. E.S. focused on connections to applications while J.D. focused on assuring accessible mathematics discussion.

Competing interests. We declare that we have no competing interests.

Funding. J.F. and J.D. were supported by the DARPA Complex Adaptive System Composition and Design Environment (CASCADE) project under contract no. N66001-16-C-4048.

Acknowledgements. The authors thank John Baez, Tony Falcone, Ben Long, Tom Mifflin, John Paschkewitz, Ram Sriram and Blake Pollard for helpful discussions and two anonymous referees for comments that significantly improved the presentation.

Disclaimer. Official contribution of the National Institute of Standards and Technology; not subject to copyright in the USA. Any commercial equipment, instruments or materials identified in this paper are used to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

References

1. Alexander C. 1964 *Notes on the synthesis of form*. Cambridge, MA: Harvard University Press.
2. Simon HA. 1991 The architecture of complexity. *Facets of systems science*. pp. 457–476. New York, NY: Springer.
3. Lee EA. 2015 The past, present and future of cyber-physical systems: a focus on models. *Sensors* **15**, 4837–4869. (doi:10.3390/s150304837)
4. Sirjani M, Lee EA, Khamespanah E. 2020 Verification of cyberphysical systems. *Mathematics* **8**, 1068. (doi:10.3390/math8071068)
5. Sokolowski JA, Banks CM. 2010 *Modeling and simulation fundamentals: theoretical underpinnings and practical domains*. Hoboken, NJ: John Wiley & Sons.
6. Chechik M, Nejati S, Sabetzadeh M. 2012 A relationship-based approach to model integration. *Innov. Syst. Softw. Eng.* **8**, 3–18. (doi:10.1007/s11334-011-0155-2)
7. Chungoora N, Young RI. 2011 Semantic reconciliation across design and manufacturing knowledge models: a logic-based approach. *Appl. Ontol.* **6**, 295–295. (doi:10.3233/AO-2011-0092)

8. Choi H *et al.* 2021 On the use of simulation in robotics: opportunities, challenges, and suggestions for moving forward. *Proc. Natl Acad. Sci. USA* **118**, e1907856118.
9. Lisciandra C, Korbmacher J. 2021 Multiple models, one explanation. *J. Econ. Methodol.* **28**, 1–21.
10. Marder E, Taylor AL. 2011 Multiple models to capture the variability in biological neurons and networks. *Nat. Neurosci.* **14**, 133–138. (doi:10.1038/nn.2735)
11. Quarteroni A. 2009 Mathematical models in science and engineering. *Not. Am. Math. Soc.* **56**, 10–19.
12. Simon G, Levendovszky T, Neema S, Jackson E, Bapty T, Porter J, Sztipanovits J. 2012 Foundation for model integration: Semantic backplane. *Int. Design Engineering Technical Conf. and Computers and Information in Engineering Conf.*, vol. 45011, pp. 1077–1086. American Society of Mechanical Engineers.
13. Eilenberg S, Mac Lane S. 1945 General theory of natural equivalences. *Trans. AMS* **58**, 231–294. (doi:10.1090/S0002-9947-1945-0013131-6)
14. Boronat A, Knapps A, Meseguer J, Wirsing M. 2008 What is a multi-modeling language? *International Workshop on Algebraic Development Techniques*, pp. 71–87. Berlin, Heidelberg: Springer.
15. Simon HA. 1979 Rational decision making in business organizations. *Am. Econ. Rev.* **69**, 493–513.
16. Brayton RK, Hachtel GD, McMullen C, Sangiovanni-Vincentelli A. 1984 *Logic minimization algorithms for VLSI synthesis*, vol. 2. Springer Science & Business Media.
17. Ferrari A, Sangiovanni-Vincentelli A. 1999 System design: traditional concepts and new paradigms. *Proceedings 1999 IEEE Int. Conf. on Computer Design: VLSI in Computers and Processors*, pp. 2–12.
18. Sangiovanni-Vincentelli A. 2003 The tides of EDA. *IEEE Design & Test of Comput.* **20**, 59–75. (doi:10.1109/MDT.2003.1246165)
19. Sosa ME, Eppinger SD, Rowles CM. 2003 Identifying modular and integrative systems and their impact on design team interactions. *J. Mech. Des.* **125**, 240–252. (doi:10.1115/1.1564074)
20. Whitney DE. 2002 Physical limits to modularity. Working paper, ESD-WP-2003-01.03-ESD. Cambridge, MA: MIT. (<https://dspace.mit.edu/handle/1721.1/102731>)
21. Breiner S, Sriram RD, Subrahmanian E. 2019 Compositional models for complex systems. IN *Artificial Intelligence for the Internet of Everything* eds. pp. 241–270. Cambridge, MA: Academic Press.
22. Fong B. 2016 The algebra of open and interconnected systems. DPhil thesis. Oxford, UK: University of Oxford.
23. Gürlebeck K, Hofmann D, Legatiuk D. 2017 Categorical approach to modelling and to coupling of models. *Math. Methods Appl. Sci.* **40**, 523–534. (doi:10.1002/mma.3978)
24. Mabrok MA, Ryan MJ. 2017 Category theory as a formal mathematical foundation for model-based systems engineering. *Appl. Math. Inf. Sci.* **11**, 43–51. (doi:10.18576/amis/110106)
25. Master J, Patterson E, Canedo A. 2020 String diagrams for assembly planning. *DIAGRAMS 2020 11th Int. Conf. on the Theory and Application of Diagrams*.
26. Spivak DI, Tan J. 2017 Nesting of dynamical systems and mode dependent networks. *J. Complex Netw.* **5**, 389–408.
27. Brown KS, Spivak DI, Wisnesky R. 2019 Categorical data integration for computational science. *Comput. Mater. Sci.* **164**, 127–132. (doi:10.1016/j.commatsci.2019.04.002)
28. Diskin Z, Maibaum T. 2014 Category theory and model-driven engineering: from formal semantics to design patterns and beyond. In *Model-Driven Engineering of Information Systems: Principles, Techniques, and Practice*, p. 173. Boca Raton, FL: CRC Press.
29. Fong B, Spivak DI. 2019 *An invitation to applied category theory: seven sketches in compositionality*. Cambridge, UK: Cambridge University Press.
30. Giesa T, Jagadeesan R, Spivak DI, Buehler MJ. 2015 Matriarch: a python library for materials architecture. *ACS Biomater. Sci. Eng.* **1**, 1009–1015. (doi:10.1021/acsbmaterials.5b00251)
31. Spivak DI. 2014 *Category theory for the sciences*. Cambridge, MA: MIT Press.
32. Spivak DI, Giesa T, Wood E, Buehler MJ. 2011 Category theoretic analysis of hierarchical protein materials and social networks. *PLoS ONE* **6**, e23911. (doi:10.1371/journal.pone.0023911)
33. Spivak DI, Kent RE. 2012 Ologs: a categorical framework for knowledge representation. *PLoS ONE* **7**, e24274. (doi:10.1371/journal.pone.0024274)

34. Leinster T. 2003 *Higher operads, higher categories*, London Math. Soc. Lec. Note Series, 298.
35. Markl M, Shnider S, Stasheff JD. 2002 *Operads in algebra, topology and physics*. Providence, RI: AMS.
36. Yau D. 2016 *Colored operads*. Providence, RI: AMS.
37. Mac Lane S. 1998 *Categories for the working mathematician*, 2nd edn, vol. 5. New York, NY: Springer.
38. Leveson N. 2013 The Drawbacks in using the term 'System of Systems'. *Biomed. Instrum. Technol.* **47**, 115–118. (doi:10.2345/0899-8205-47.2.115)
39. Turing AM. 1952 The chemical basis of Morphogenesis. *Phil. Trans. R. Soc. Lond. B* **237**, 37–72. (doi:10.1098/rstb.1952.0012)
40. Baez JC, Foley J, Moeller J, Pollard B. 2020 Network models. *Theor. Appl. Categ.* **35**, 700–744.
41. Baez JC, Foley J, Moeller J. 2019 Network models from Petri nets with catalysts. *Compositionality* **1**. (doi:10.32408/compositionality-1-4)
42. Moeller J. 2020 Noncommutative network models. *Math. Struct. Comput. Sci.* **30**, 14–32. (doi:10.1017/S0960129519000161)
43. Moeller J, Vasilakopoulou C. 2020 Monoidal Grothendieck construction. *Theor. Appl. Categ.* **35**, 1159–1207.
44. Spivak DI. 2013 The operad of wiring diagrams: formalizing a graphical language for databases, recursion, and plug-and-play circuits. Preprint. (<http://arxiv.org/abs/1305.0297>)
45. Yau D. 2018 *Operads of wiring diagrams*, Vol. 2192. New York, NY: Springer.
46. Stone LD, Royset JO, Wasburn AL. 2016 *Optimal search for moving targets*, vol. 237. New York, NY: Springer.
47. Fong B, Spivak DI. 2019 Hypergraph categories. *J. Pure Appl. Algebra* **223**, 4746–4777. (doi:10.1016/j.jpaa.2019.02.014)
48. Fong B, Spivak DI. 2020 Supplying bells and whistles in symmetric monoidal categories. Preprint. (<http://arxiv.org/abs/1908.02633>)
49. Breiner S, Pollard B, Subrahmanian E, Marie-Rose O. 2020 Modeling Hierarchical System with Operads. *Proc. of the 2019 Applied Category Theory Conf.*, pp. 72–83.
50. Hermida C. 2000 Representable multicategories. *Adv. Math.* **151**, 164–225. (doi:10.1006/aima.1999.1877)
51. Bakirtzis G, Vasilakopoulou C, Fleming CH. 2020 Compositional cyber-physical systems modeling. *Proc. 3rd Annual Int. Applied Category Theory Conf. 2020 (ACT 2020)*.
52. Vagner D, Spivak DI, Lerman E. 2015 Algebras of open dynamical systems on the operad of wiring diagrams. *Theor. Appl. Categ.* **30**, 1793–1822.
53. Baez JC, Coya B, Rebro F. 2018 Props in network theory. *Theor. Appl. Categ.* **33**, 727–783.
54. Baez JC, Fong B. 2018 A compositional framework for passive linear networks. *Theor. Appl. Categ.* **33**, 1158–1222.
55. Baez JC, Master J. 2020 Open Petri nets. *Math. Struct. Comp. Sci.* **30**, 314–341. (doi:10.1017/S0960129520000043)
56. Baez JC, Weisbart D, Yassine A. 2021 Open systems in classical mechanics. *J. Math. Physics.* **62**, 042902. (doi:10.1063/5.0029885)
57. Coya B. 2018 *Circuits, Bond Graphs, and Signal-Flow Diagrams: A Categorical Perspective*. PhD thesis. University of California–Riverside.
58. Ghica DR, Jung A, Lopez A. 2017 Diagrammatic semantics for digital circuits. *26th EACSL Annual Conf. on Computer Science Logic (CSL 2017)*, Leibniz Int. Proc. in Informatics, vol. 82, pp. 24:1–24:16.
59. Baez JC, Fong B, Pollard BS. 2016 A compositional framework for Markov processes. *J. Math. Phys.* **57**, 033301. (doi:10.1063/1.4941578)
60. Pollard BS. 2016 Open Markov processes: a compositional perspective on non-equilibrium steady states in biologys. *Entropy* **18**, 140. (doi:10.3390/e18040140)
61. Nolan JS, Pollard BS, Breiner S, Anand D, Subrahmanian E. 2020 Compositional models for power systems. *Proc. of the 2019 Applied Category Theory Conf.*, pp. 72–83.
62. Spivak DI, Wisnesky R. 2020 Fast Left-Kan extensions using the chase. Preprint. Available at www.categoricaldata.net.
63. Wisnesky R, Breiner S, Jones A, Spivak DI, Subrahmanian E. 2017 Using category theory to facilitate multiple manufacturing service database integration. *J. Comput. Inf. Sci. Eng.* **17**, 021011. (doi:10.1115/1.4034268)
64. Johnson-Freyd P, Aytac J, Hulett G. 2020 Topos semantics for a higher-order temporal logic of actions. *Proc. of the 2019 Applied Category Theory Conf.*, pp. 161–171.

65. Schultz P, Spivak DI. 2019 *Temporal type theory*. Cham, Switzerland: Springer International Publishing.
66. Zardini G, Spivak DI, Censi A, Frazzoli E. 2020 A compositional sheaf-theoretic framework for event-based systems (extended version). Preprint. (<http://arxiv.org/abs/2005.04715>)
67. Coecke B, Sadrzadeh M, Clark S. 2010 Mathematical foundations for a compositional distributional model of meaning. *Linguist. Anal.* **36**, 345–384.
68. Grefenstette E, Sadrzadeh M, Clark S, Coecke B, Pulman S. 2018 Concrete sentence spaces for compositional distributional models of meaning. *Computing meaning*, pp. 71–86. Dordrecht, The Netherlands: Springer.
69. Hedges J, Sadrzadeh M. 2019 A generalised quantifier theory of natural language in categorical compositional distributional semantics with bialgebras. *Math. Struct. Comput. Sci.* **29**, 783–809. (doi:10.1017/S0960129518000324)
70. Fong B, Johnson M. 2019 Lenses and learners. *Proc. of the Eighth Int. Workshop on Bidirectional Transformations (Bx 2019)*.
71. Fong B, Spivak DI, Tuyéras R. 2019 Backprop as functor: a compositional perspective on supervised learning. *2019 34th Annual ACM/IEEE Symp. on Logic in Computer Science (LICS)*, pp. 1–13.
72. Bakirtzis G, Genovese F, Fleming CH. 2021 *Yoneda hacking: the algebra of attacker actions*. Preprint. (<http://arxiv.org/abs/2103.00044>)
73. Bakirtzis G, Fleming CH, Vasilakopoulou C. 2021 Categorical semantics of cyber-physical systems theory. Preprint. (<http://arxiv.org/abs/2010.08003>)
74. Bakirtzis G. 2021 *Compositional cyber-physical systems theory*. PhD thesis. Charlottesville, VA: University of Virginia. (<https://doi.org/10.18130/xn8v-5d89>)
75. Fagan A, Duncan R. 2019 Optimising clifford circuits with quantomatic. *Proc. 15th Int. Conf. on Quantum Physics and Logic (QPL 2018)*, pp. 85–105.
76. Kissinger A, van de Wetering J. 2020 Reducing the number of non-Clifford gates in quantum circuits. *Phys. Rev. A* **102**, 022406. (doi:10.1103/PhysRevA.102.022406)
77. Kissinger A, Zamdzhiev V. 2015 Reducing the number of non-Clifford gates in quantum circuits. *Int. Conf. on Automated Deduction*, pp. 326–336. Cham, Switzerland: Springer.
78. Di Lavore E, Hedges J, Sobociński P. 2021 Compositional modelling of network games. *Computer Science Logic 2021, Leibniz Int. Proc. in Informatics 183*.
79. Ghani N, Kupke C, Lambert A, Forsberg FN. 2020 Compositional game theory with mixed strategies: probabilistic open games using a distributive law. *Proc. of the 2019 Applied Category Theory Conf.*, pp. 95–105.
80. Ghani N, Winschel J, Zahn P. 2018 Compositional game theory. *Proc. of the 33rd Annual ACM/IEEE Symp. on Logic in Computer Science*, pp. 472–481.
81. Busboom S. 1990 *Bat 21: a case study*. Carlisle Barracks, PA: US Army War College.
82. Guy J. 2020 142 Evacuated From Creek Fire by Military Helicopters; Body of Deceased Man Flown to Fresno. *Fresno Bee*, Sept. 8.
83. Kuwada R, Guy J, Cooper D. 2020 Creek Fire roars toward mountain resort towns, after airlift rescues hundreds trapped by flames. *Fresno Bee*, Sept. 6.
84. Foley JD. 2018 An example of exploring coordinated SoS behavior with an operad and algebra integrated with a constraint program, CASCADE tech report.
85. Baez JC, Genovese F, Master J, Shulman M. 2021 Categories of nets. Preprint. (<http://arxiv.org/abs/2101.04238>)
86. Meseguer J, Montanari U. 1990 Petri nets are monoids. *Inf. Comput.* **88**, 105–155. (doi:10.1016/0890-5401(90)90013-8)
87. Khor CS, Chachuat B, Shah N. 2012 A superstructure optimization approach for water network synthesis with membrane separation-based regenerators. *Comput. Chem. Eng.* **42**, 48–63. (doi:10.1016/j.compchemeng.2012.02.020)
88. Neuro SM, Pinto JM. 2003 Supply chain optimization of petroleum refinery complexes. *Proc. of the 4th Int. Conf. on Foundations of Computer-Aided Process Operations*, pp. 59–72.
89. Raman R, Grossmann IE. 1992 Integration of logic and heuristic knowledge in MINLP optimization for process synthesis. *Comput. Chem. Eng.* **16**, 155–171. (doi:10.1016/0098-1354(92)85003-Q)
90. Yeomans H, Grossmann IE. 1999 A systematic modeling framework of superstructure optimization in process synthesis. *Comput. Chem. Eng.* **23**, 709–731. (doi:10.1016/S0098-1354(99)00003-4)

91. Censi A. 2015 *A mathematical theory of co-design*. Preprint. (<http://arxiv.org/abs/1512.08055>).
92. Censi A. 2017 Uncertainty in monotone codesign problems. *IEEE Rob. Autom. Lett.* **2**, 1556–1563. (doi:10.1109/LRA.2017.2674970)
93. 2020 Zardini, N. Lanzetti, M. Salazar, A. Censi, E. Frazzoli, and M. Pavone. Towards a co-design framework for future mobility systems. *Annual Meeting of the Transportation Research Board*.
94. Zardini G, Milojevic D, Censi A, Frazzoli E. 2020 Co-design of embodied intelligence: a structured approach. Preprint. (<http://arxiv.org/abs/2011.10756>)
95. de Felice G, Toumi A, Coecke B. 2020 DisCoPy: Monoidal Categories in Python. *Proc. 3rd Annual Int. Applied Category Theory Conf. 2020 (ACT 2020)*.
96. Halter M, Patterson E, Baas A, Fairbanks J. 2020 Compositional scientific computing with Catlab and SemanticModels. Preprint. (<http://arxiv.org/abs/2005.04831>)
97. A. Kissinger and J. van de Wetering: PyZX: Large Scale Automated Diagrammatic Reasoning. *Proc. 16th Int. Conf. on Quantum Physics and Logic (QPL 2019)*, (2020) 229–241.
98. Breiner S, Pollard B, Subrahmanian E. 2020 Workshop on applied category theory: bridging theory and practice. *Special Publication (NIST SP)* 1249.
99. Bar K, Kissinger A, Vicary J. 2018 Globular: an online proof assistant for higher-dimensional rewriting. *Log. Methods Comput. Sci.* **14**, 1–16.
100. Reutter D, Vicary J. 2019 High-level methods for homotopy construction in associative n-categories. *2019 34th Annual ACM/IEEE Symp. on Logic in Computer Science (LICS)*, pp. 1–13. IEEE.
101. Patterson SM, Spivak DI, Vagner D. 2020 Wiring diagrams as normal forms for computing in symmetric monoidal categories. *Proc. of the 2020 Applied Category Theory Conf. (ACT 2020)*.
102. Stanley KO, Miiikkulainen R. 2002 Evolving neural networks through augmenting topologies. *Evol. Comput.* **10**, 99–127. (doi:10.1162/106365602320169811)
103. Abramsky S, Coecke B. 2009 Categorical quantum mechanics. *Handb. Quantum Logic Quantum Struct.* **2**, 261–325.
104. Ford LR, Fulkerson DR. 1956 Maximal flow through a network. *Can. J. Math.* **8**, 399–404. (doi:10.4153/CJM-1956-045-5)
105. Geisberger R, Sanders P, Schultes D, Delling D. 2008 Contraction hierarchies: faster and simpler hierarchical routing in road networks. *Int. Workshop on Experimental and Efficient Algorithms*, pp. 319–333. Berlin, Heidelberg: Springer.
106. Dusko P. 2013 Monoidal computer I: basic computability by string diagrams. *Info. Comput.* **226**, 94–116. (doi:10.1016/j.ic.2013.03.007)
107. Beal J, Pianini D, Viroli M. 2015 Aggregate programming for the internet of things. *Computer* **48**, 22–30. (doi:10.1109/MC.2015.261)