

The Domain Relationship Diagram: a Mapping Framework for Digital Twin Design in Complex Cyber-Physical Ecosystems

Matt Tipuric^{1*}, Daniel Mitchell^{2†}, Blair Archibald^{2†},
David Flynn^{3†}, Paul Harvey^{2†}, Michele Sevensgani^{2†},
David Wagg^{1†}

^{1*}School of Mechanical, Aerospace and Civil Engineering, The University of Sheffield, Mappin Street, Sheffield, S1 4DT, United Kingdom.

²School of Computing Science, University of Glasgow, Glasgow, G12 8RZ, United Kingdom.

³James Watt School of Engineering, University of Glasgow, Glasgow, G12 8RZ, United Kingdom.

*Corresponding author(s). E-mail(s): m.tipuric@sheffield.ac.uk;

Contributing authors: daniel.mitchell@glasgow.ac.uk;

blair.archibald@glasgow.ac.uk; david.flynn@glasgow.ac.uk;

paul.harvey@glasgow.ac.uk; michele.sevensgani@glasgow.ac.uk;

david.wagg@sheffield.ac.uk;

†These authors contributed equally to this work.

Abstract

The ability to reduce risk and uncertainty in the operation and performance of complex adaptive systems is highly desirable across a range of applications and sectors, especially when the correct operation of those systems is critical. Recently, digital twins have delivered this capability for cyber-physical systems by combining modelling with real/relevant-time information flows. Increasingly, however, critical national infrastructure which were previously isolated networks (such as transport and energy networks) are accelerating their levels of integration, forming cyber-physical *ecosystems* (CPES) and creating new design challenges for digital twins. One fundamental challenge lies in synthesizing information from multiple sources in order to understand the ecosystem as a whole. In this paper, we propose a novel mapping framework to meet this need: the ‘domain

relationship diagram’, which balances flexibility and interpretability with sufficient rigour to support model-based design. After establishing the method, we illustrate its usage through two use cases, both motivated by priority research in critical CPES: (i) flooding resilience in northern England and (ii) energy infrastructure at the Port of Dover. We also tie this approach into other crucial topics in digital twins for CPES, including ownership, ontologies and communication, and trust and resilience.

Keywords: Digital Twins, Cyber-physical ecosystems, Complex adaptive Systems, Model-based design, Concept mapping

1 Introduction

A digital twin, as defined by the UK Government’s National Digital Twin Programme, is a “. . . virtual model of an object, a system, or a process. It is connected to its real-world counterpart (the physical twin) by a two-way flow of right-time data, meaning it mimics it in all aspects. . .” (Department for Business and Trade 2024). This definition contrasts with that of a digital shadow, where data flows only from physical to digital, and a digital model, which lacks live data exchange (Wagg et al. 2020). Digital Twins have found use cases in multiple sectors, including engineering and manufacture (Tao et al. 2019), transport (Department for Transport 2024), communication networks (Almasan et al. 2022), policy (Papyshev and Yarime 2021) and energy (Mitchell et al. 2025a) and they can be used for a wide variety of purposes, such as condition monitoring, visualisation, control, (optimal) planning, prediction/-forecasting, and decision-making support through what-if scenarios (Wagg et al. 2020).

Typically, digital twins and the corresponding physical asset/system (the ‘physical twin’) are owned by a single entity or organisation. As an example, the planned Port of Dover (PoD) digital twin is being developed by the same company that owns the physical port itself (Port of Dover 2024). There are, however, many situations where digital twins are desired for systems and processes which are the responsibility of multiple bodies, entailing a complex cyber-physical ecosystem (CPES). The UK’s Department for Transport (DfT) has expressed an interest in developing a digital twin(s) of the UK transport network, with one use case being freight management at ports (Department for Transport 2024). Such a digital twin would need to include Dover Harbour, along with the other UK ports, road and rail links, and logistics hubs. Other examples exist within such varied fields as healthcare, policy, and crisis response, among others

Model-based development is a powerful tool for designing cyber-physical systems (Rajhans et al. 2014; Tellioglu 2009), with one key benefit being a reduction in costs through identifying errors within early-stage models, rather than as part of the final deployment. Expanding the scope from cyber-physical systems to cyber-physical *ecosystems* increases the expense of deployment and so, correspondingly, the value gained from using model-based design. This argument is supported by recent work

by Tadeja et al. (2026), who conducted a qualitative study on design approaches for digital twins, which they contextualise as part of a seventeen-step process concluding with ‘a system of systems’.

For model-based design to be effective, we require our model of the CPES to accurately represent the physical ecosystems. This requires us to engage with stakeholders (including system users, domain experts, and project funders), whether throughout, as part of collaborative or co-operative design, or as part of the discovery phase of more linear design methodologies, such as waterfall. Equally, Tadeja et al. (2026) highlight that the analysis of existing tools and knowledge is “frequently omitted in the first instance”. These issues are only compounded by the social complexity involved in CPES, where knowledge, tools and responsibility can be split across multiple organisations.

Flexible, unstructured methods – such as brain dumps or open discussions – are well suited to synthesizing knowledge of this kind, without inflicting pre-existing biases or assumptions. However, such methods integrate poorly with the rest of model-based design, which inherently require structure Eng et al. (2017). For the stakeholder engagement to form a coherent part of the design process, the outputs need to be consistent with the models used elsewhere in the process. In other terms, we require a framework to enable concept mapping (Rosas 2012) while remaining interoperable with the entity-relationship models used elsewhere in the design process.

We propose the **Domain Relationship Diagram (DRD)** as a new framework to fulfil this need. The DRD can be used to capture interdependencies between both virtual models and their physical counterparts, and between the twins of different owners. A core feature of this format is the selection of a single perspective (the *view-point owner*) to account for a suitable level of abstraction. Additionally, the division of the ecosystem into physical and virtual ‘domains’, without judgement over whether they constitute part of a ‘true’ digital twin allows for the formulation and discussion of architectures not captured by the existing terminology.

We continue this paper in Section 2 by presenting some background information on digital twins, cyber-physical systems, and model-based design as it relates to both topics. After demonstrating the strong theoretical underpinnings of the DRD, we then present the mapping method in detail in Section 3, including a description of six archetypical relationships, guidance on how to create a DRD, and detail of its four underlying principles. To demonstrate the practical relevance of the framework, as well as its applicability to a range of CPES, we then present two separate case studies, grounded in our previous works. In Section 4.1 we apply the DRD to a CPES at the intersection of transport resilience and crisis response, based on the DTBOC project (Tipuric et al. 2025), while in Section 4.2 we draw upon the work of the TransiT research hub into digital twins for decarbonisation (TransiT 2023) and apply the DRD to the decarbonisation of the Port of Dover. Finally, in Section 5 we discuss some additional aspects of digital twins for CPES, including issues of fragmented ownership; communication and data flow; trust, security and resilience; and some topics for future work, before concluding the paper in Section 6.

2 Background

2.1 Digital twins in complex cyber-physical ecosystems

The initial concept of the digital twin is often ascribed to the National Aeronautics and Space Administration’s Apollo programme in the 1960s (Rosen et al. 2015). While digital twins are prominent in manufacturing (Atalay et al. 2022), they have also found use in transportation (Department for Transport 2024), communication networks (Almasan et al. 2022), policy (Papsyshev and Yarime 2021), construction (Ira 2023) and energy (Podvalny and Vasiljev 2021; Mitchell et al. 2025a). Specific uses include: condition monitoring, visualisation, control (Mitchell et al. 2021), (optimal) planning, prediction/forecasting (Blanche et al. 2020; Tang et al. 2021), and decision-making support through what-if-scenarios (Wagg et al. 2020).

A common representation of digital twins is as a pseudo-block diagram, like those found in the control engineering literature. Diagrams such as in Figure 1a make explicit the common distinction between digital models, shadows, and twins (Thelen et al. 2022; Grieves and Vickers 2017). Modifications to this representation are used to describe other aspects of a digital twin, for example describing the representation of disconnected/semi-connected/fully connected digital twins as in Figure 1b, or capturing the time-evolving aspect of a digital thread as in Figure 1c.

These types of representations reinforce a limiting aspect of our definition of a digital twin (introduced in Section 1) which is the implicit one-to-one correspondence between a single physical twin and its digital twin: “a virtual model of *an* object, system, or process...”. This is compounded by the requirement that the later mimics the former “in all aspects”. Practically speaking, in most cases where this distinction matters it is possible to abstract away through careful definition of the physical twin (thus defining the boundary of what it is the digital twin must represent). A system, after all, can consist of multiple objects, or even sub-systems and so there is no contradiction in referring to the virtual representation as a digital twin, even if it represents multiple physical twins. Equally, pragmatism allows us to be selective with the second constraint: creating a ‘true’ one-to-one representation of anything is philosophically challenging. Instead, this requirement gets recast as ‘all *useful and practical* aspects’. These abstractions serve well in many cases. However, they become less useful and possibly obstructive when digital twins are considered in the context of cyber-physical ecosystems (CPES).

According to Giese et al. (2012), cyber-physical *systems* are “engineered systems of synergistically interacting physical and computational components...[where] functionality and salient system properties are emerging from an intensive interaction of physical and computational components”. By extension, Bujorianu (2023) describes CPES as “ecosystems of networked [Cyber-physical systems]... provided with an interaction activity between them and with their environment.” Mori and Chen (2014) provide a more detailed analysis of *software* ecosystems, which they characterise as representing “interconnected entities that depend upon each other as well as interact with others outside their domain for survivability, growth, and sustainability”. They also discuss why these ecosystems tend to use the concept of ‘autonomous decentralized

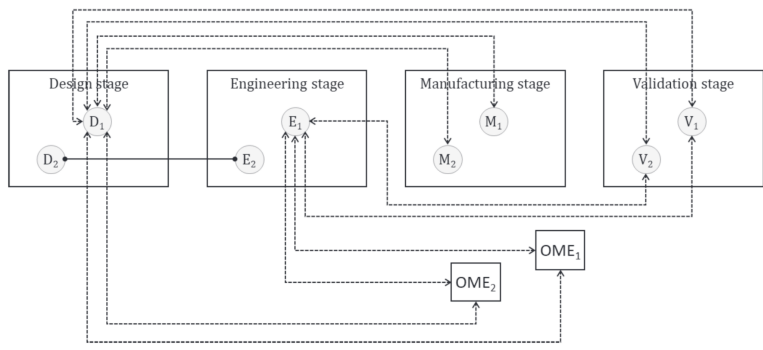
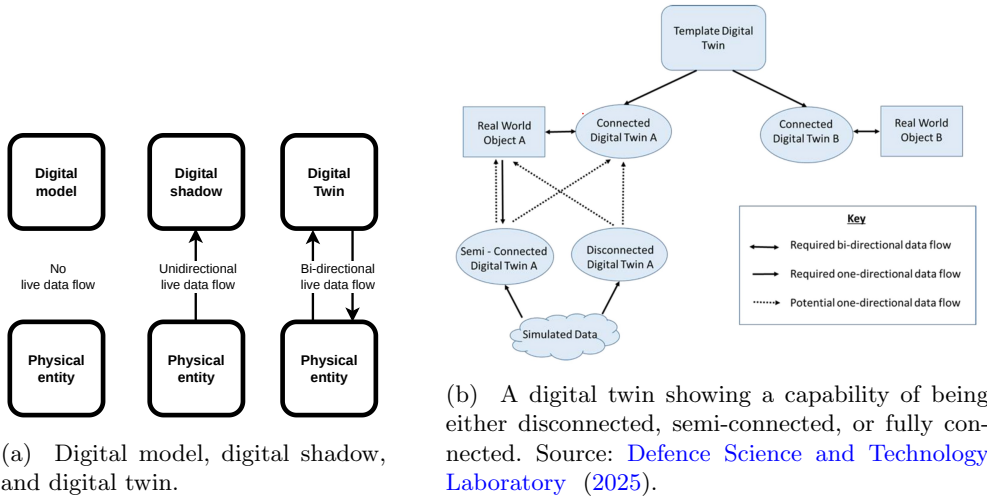


Fig. 1: Three different pseudo-block diagram representations relating to digital twins.

systems’, larger components of the ecosystem are decomposed into smaller subcomponents and managed by largely autonomous organisations. We can easily expand this to CPES by allowing these autonomous organisations to control physical subcomponents in addition to the virtual ones.

Returning to digital twins, it is clear that each instance, by definition, forms a cyber-physical system with its physical twin. Thus, any analysis of digital twins of or in an ecosystem, is inherently an analysis of a cyber-physical ecosystem, even if the only cybernetic component is the digital twin itself. In a world where government and industry often seek to create digital twins of such ecosystems, this distinction is important, as it allows us to more clearly consider what the definition of the physical twin is (or, correspondingly, what the boundaries are on what the digital twin must represent). Crucially, we need a way to identify the different aspects and facets of a

cyber-physical ecosystem: which entities have access to data, which have responsibility over the cybernetic and physical components, and which can affect the outcomes of the digital twin onto the physical.

2.2 Design for cyber-physical systems

In their review paper, [Khaitan and McCalley \(2015\)](#) describe how the continuous interaction of “complex and heterogeneous systems...[necessitate] careful codesign of the overall architecture of CPSs”, thus requiring domain specific methods. They highlight the role played by meta modelling, formal semantic models, and various computational models in laying a theoretical foundation. They also describe multiple specific methodologies employed, including discrete event calculus, ‘cyberizing the physical’/‘physicalizing the cyber’, and the model-based design methodology employed by [Jensen et al. \(2011\)](#), detailed later in this section. Specifically focusing on digital twins, [Tadeja et al. \(2026\)](#) mention approaches based on model based system engineering ([Wang et al. 2023](#)), V-design ([Wagg et al. 2020](#)), and agile ([Aheleroff et al. 2021](#)).

[Rajhans et al. \(2014\)](#) describe model-based design¹ as “...the use of computational and formal models in the system design process”. In the context of CPS, including digital twins they present an architecture-centric approach, with the stated aim of striking a balance between specificity and generality. This results in a consistent style, although the resulting diagrams are still aimed at systems design engineers.

[Jensen et al. \(2011\)](#) take a more general approach, presenting a 10 step design process for CPS. This methodology is sufficiently generalised to be readily expanded to CPES, although step 4 (“derive a control algorithm”) needs to be reframed more broadly. The design process is presented in a linear fashion for clarity, although the authors note in several points that the steps should be iterative in nature, a clarification which we echo. On the occasions where we describe actions as being taken ‘earlier’ or ‘later’ in the design process, it is in the context of these ten steps, with no assumption that they necessarily follow one another linearly. In the context of this design methodology, the remit of the DRD lies in the first step: “state the problem”, which entails describing the problem to be solved in “simple language... without the use of language or technical terminology.”

Through mapping the digital twin landscape – using qualitative approaches to interrogate a pool of experts – [Tadeja et al. \(2026\)](#) identified 17 ‘components’ of the design process. These overlap with the ten steps of [Jensen et al. \(2011\)](#), with much of the discrepancy resulting from choices made by the authors in how to consolidate different sub-steps. For example, the single step ‘10: Verify, and Validate, and Test’ becomes the two steps ‘(10) Verify and Validate Model(s)’ and ‘(11) Test Model(s)’ in the more recent work. Instead of ‘derive a control algorithm’, ([Tadeja et al. 2026](#)) use the terminology of ‘(7) Initial Digital Twin Design’ which, due to the additional sub-categorisation of digital twins² may be more flexibly applied to different CPEs. [Tadeja et al. \(2026\)](#) incorporated the 17 components into a six stage design methodology for digital twins, based on the waterfall methodology. The role of the DRD lies in the first

¹The authors use the term ‘model-based *development*’ but no distinction is made with design.

²simulation, operational, semi-automated, fully automated, and mature

Table 1: Correspondence of the six design steps and 17 components proposed by Tadeja et al. (2026) and the ten steps proposed by Jensen et al. (2011). Items marked with † are within the primary scope of the DRD (see Section 3.2). Components marked with * do not align with any specific step.

Step (Tadeja)	Component (Tadeja)	Step (Jensen)
Plan†	Define Scope† Reuse Knowledge†	1. State the Problem†
Analyse†	Analyze Tools†	2. Model Physical Processes† 3. Characterize the Problem†
Develop	Acquire Data Asset Modeling Develop Model(s)	4. Derive a Control Algorithm 5. Select Models of Computation 6. Specify Hardware
Design	Initial DT Design	7. Simulate 8. Construct 9. Synthesize Software
	Test-bed Asset* Simulation Data*	
Test	V & V Model(s) Test Model(s)	10. Verify, and Validate, and Test
	Human-in-the-Loop*	
Operate	Digital Twin Operate/Maintain Physical Twin Collect Data	
	System of Systems*	

two steps – ‘plan’ and ‘analyse’. For comparison, Table 1 details these steps and the corresponding components, aligning them with those proposed by Jensen et al. (2011).

2.3 Communication, boundary objects, and concept mapping

Regardless of the choice of methodology, there will be certain knowledge that is prerequisite to the actual development and design of the CPS. Generically, this falls into two categories: 1. *what already exists* and 2. *what is required*. Tadeja et al. (2026) captures this knowledge in the ‘plan’ and ‘analyse’ stages (‘define scope’, ‘reuse knowledge’, ‘analyze tools’), Jensen et al. (2011) in the steps 1-3 (‘state the problem’, ‘model physical processes’, ‘characterize the problem’). For model-based design, this knowledge is required in order to create said models, while in other methodologies it underpins the creation of a design specification. Accuracy is key in these early stages to avoid compounding issues later in the process. However, in practice they can receive less focus – Tadeja et al. (2026) notes that the ‘analyze tools’ step in particular is “frequently omitted”.

The difficulties involved in communication during the design process are well-documented and are not unique to the design of cyber-physical systems. In a review of the causes of failure of space systems, Newman (2001) identified the vast majority as involving some form of misunderstanding. Eckert et al. (2005) highlights necessity of active negotiation between participants in the design process and identifies four causes of communication breakdown: ‘Not understanding the big picture’, ‘Missing information provision’, ‘Information distortion’, and ‘Interpretation of representation’. The issue of interface management – the need to negotiate interactions between components designed by teams from different disciplines – is particularly relevant for both CPS and CPES, which inherently involve a great number of such interfaces.

Subrahmanian et al. (2003) describe how ‘boundary objects’³ “can inhabit the space of negotiation at the interfaces and serve to support cooperation” without requiring participants to have a pre-existing shared understanding. Boujut and Blanco (2003) give this concept further weight, documenting the use of simplified representations such as sketches and CAD diagrams for negotiation within the design process. In this context, they highlight the need for boundary objects to lack ambiguity and to remain malleable to avoid creating cognitive traps.

For established areas of design, there is often a clear choice for these boundary objects. A lay person can be expected to interpret a floor plan well enough to communicate with an architect, electrical and control engineers can negotiate over block and circuit diagrams, while the use of Universal Modelling Language (UML) diagrams for requirements capture in software engineering is well established. In the case of CPES, however, we can expect to encounter interfaces between design participants, from entirely different domains, who do not normally interact. There is no clear choice for a general framework to use to create our boundary objects.

One option for those following a model-based methodology is that described by Tellioglu (2009), who focuses on integrating *collaborative and model-based design*. They categorize such models into four groups, based on function. Of interest here is the first type of model (“models to visualize several issues”), for which they describe using an ‘organisational design game’ to capture a company’s employees’ understanding of the state of that company. This is very similar to the kinds of understanding we aim to capture with the DRD.

Looking more broadly, an alternative method is *concept mapping*. This is a tool which has been successfully employed for participatory and collaborative research (Rosas 2012). Davies (2011) contrasts concept maps with mind maps, characterising the former as “more structured and less pictorial in nature”. Concept maps “outline relationships between ideas”. Eng et al. (2017) investigated mapping tools more generally for engineering design, especially for complex systems. They categorise mapping methods as a graphical model sitting part-way between semi-structured (which integrate poorly with later design work), and structured information (which enable detailed design but require long learning curves). These tools are best suited to early-stage design tasks, including aggregating and synthesizing information, and interdisciplinary communication. A key benefit of mapping is “their support for

³also called ‘intermediary objects’ (Boujut and Blanco 2003)

exploring emergent structuring information. This is particularly important when approaching chaotic situations where imposing structure is key”.

Borrowing the terminology used by Eng et al. (2017), Fig. 2 compares the level of structure supported by – and required for – visual methods across some disciplines which relate to design for cyber-physical (eco-)systems. The least structured methods, such as brain dumps and mind maps are the most broadly applicable. In contrast, the most structured methods sit firmly within a discipline. It can be seen that some methods, such as block diagrams (Bolton 2002) and Universal Modelling Language (UML) (Phillips 2018) are flexible enough to support either semi-structured or structured information, through a loosening or tightening of requirements as standard. These methods are still tied to the assumptions of their disciplines and so while they can be used to great effect as boundary objects *within* their disciplines, they serve less well for the boundaries *between* disciplines. At the more structured extreme, some disciplines rely on formal graph-based models with mathematically defined semantics, particularly within software engineering (Ehrig et al. 2006; Taentzer 2000; Archibald et al. 2025). These support rigorous reasoning about structure, interaction, and reconfiguration, but typically impose higher cognitive and tooling overheads.

Setting the DRD against this background serves to clarify our aims: to provide a visual, semi-structured representation for CPEs, with the flexibility to serve as boundary object between the disciplines involved in the design of the components therein.

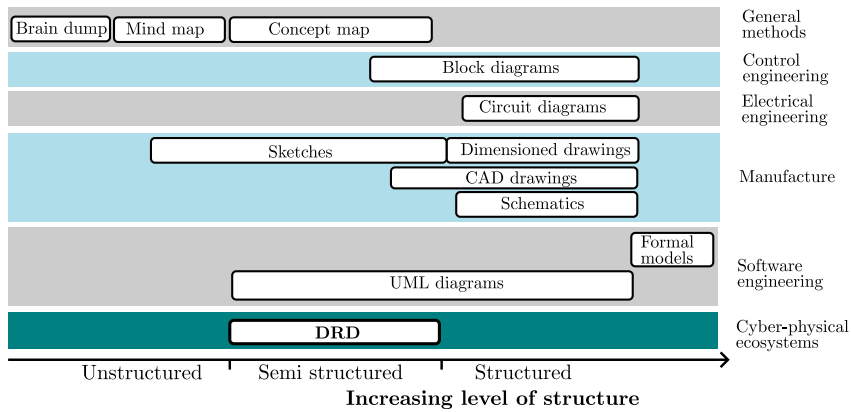


Fig. 2: A selection of visual methods which can be used to create boundary objects with varying levels of structure.

2.4 Summary and problem statement

Digital twins are conventionally used as CPS but there is a growing desire to apply them to CPES. By analogy with software ecosystems (Bujorianu 2023), we can achieve this by capturing the interactions between the autonomous centralised systems which make up the ecosystem within our representation of the ecosystem.

If we seek to expand on the existing model-based design frameworks from CPS, then “given the multidisciplinary nature of cyber-physical systems, [stating the problem] is necessary to effectively communicate design requirements” (Jensen et al. 2011). This requires communication and negotiation with stakeholders in diverse disciplines, leading to interface issues and risk of miscommunication.

Such conversations can be aided through use of diagrams as boundary objects, reducing ambiguity and highlighting potential knowledge gaps. As we seek to represent “loosely structured information in ways that balance usable formalisms and cognitive costs” (Eng et al. 2017), a graphical mapping framework is a sensible choice. This framework should:

- Present an unambiguous representation of the cyber-physical ecosystem
- Provide sufficient structure to be consistent with more formalised models used later in the design process
- Be sufficiently flexible to be applicable to a variety of heterogeneous contexts
- Be sufficiently accessible to enable use by non-expert users with a minimal level of instruction.

It is in this context which we define the Domain Relationship Diagram.

3 The Domain Relationship Diagram

3.1 Definition

A domain relationship diagram is a graphical representation of a cyber-physical ecosystem which represents the ecosystem as a collection of connected subsystems, called ‘domains’, where each domain is a) identified as either physical or virtual in nature (but never both); b) connected to at least one other domain by at least one of i. a flow of data, ii. a flow of information, or iii. a transfer of some physical property or object; and c) attached to at least one ‘owner’, representing the autonomous agent with agency over that domain.

3.2 Aims and scope

A complete, general design methodology for CPES is beyond the scope of this paper. Equally, the DRD is not and does not aim to be either a formal model or ontology of cyber-physical ecosystems. Instead, it is a mapping framework to aid knowledge capture and requirements creation required in order for the development of those models and ontologies specific to a given CPES.

In order to achieve this, the DRD should act as a boundary object to facilitate design conversations both between design teams in different disciplines, designers and system users, and designers and domain experts. Following the terminology established by Eckert et al. (2005), the framework needs to cover the following interface scenarios:

1. Request for information
2. Negotiation of clarity and negotiation of constraints
3. Handover
4. Joint designing

This specific selection of scenarios are those well-suited to the compromise between flexibility and rigour offered by concept mapping, as well as requiring either a holistic view of the whole ecosystem, or an interface between two disciplines which do not commonly cooperate. In contrast, idea generation is better suited to less constrained methods (such as mind mapping), while conflict resolution, decision making, and justification can be expected to require more formal design documents/models.

Similarly, the purview of the DRD is restricted only to those use cases where the fact that the design entity is part of cyber-physical ecosystem is relevant, and to the interface between different disciplines. For other cases, the familiarity of designers with pre-existing discipline specific tools will invariably make these the better choice. For example, two software engineers working within the same domain of the CPES might be expected to use UML. However, if they required a conversation with a geographer in a different part of the ecosystem, they could be expected to fall back to the DRD as a shared representation.

The DRD is primarily intended for design within cyber-physical ecosystems, and especially of digital twins (or systems/ecosystems of digital twins). This is a relatively loose classification – whether something is a system, a system-of-systems, or an ecosystem is largely a choice of representation. The presence of the following factors can be used as a heuristic indicator that the DRD is an appropriate choice:

1. The presence of multiple interested autonomous parties from different disciplines
2. A requirement for data flow between heterogenous systems and components
3. Pre-existing cyber-physical infrastructure with no single owner
4. A requirement for data flow between parties without pre-existing established channels

3.3 Elements of the DRD

In constructing the DRD, we make use of a well established methodology for representing generalised systems, building on the practices of systems engineering. [Willem \(2007\)](#) describes an object-oriented strategy, which makes use of three techniques to convey underlying structure: ‘tearing’, which represents the system as an interconnection of subsystems; ‘zooming’, or modelling of the subsystems; and ‘linking’ or modelling the interconnections of subsystems. The novelty of the DRD comes through two modifications made to this framework to align it to our goal of mapping cyber-physical ecosystems:

1. *Declaration of owners.* The core difference between representing an ecosystem vs a system of systems lies in the autonomy of the agents within the former. As such, construction of the DRD requires each ‘domain’ element (systems and subsystems) to be attributed to an ‘owner’, who acts with autonomy over that domain.
2. *Opinionated restriction on element classes.* We require the elements of the DRD to be represented as one of six classes of element (two classes of domain, three classes of link, and owners). This restriction balance the need to enforce a level of consistency and support interface negotiation, while avoiding the complexity design and interpretation complexity required by a more granular classification.

The DRD maps the cyber-physical ecosystem using three classes of element: *domains* (represented with coloured boxes), *owners* (black surrounding boxes), and *flows* (lines with arrows to show directionality), as shown in Fig. 3. The DRD builds on the established block diagram language, maintaining compatibility with this representation for conventional digital twins.

In terms of representation, a preference is given for use of colour and line type to represent different components. The colours chosen here are representative but where alternative colour palettes are chosen, care should be taken that the diagram meets accessibility requirements. Where possible, position of elements is used to reinforce the representation in line with the standard convention (virtual above physical, data flowing clockwise) but given the complexity of cyber-physical ecosystems, this requirement can be loosened when required.

Domains may be either physical or virtual. *Physical domains* (shown in brown \square) include infrastructure (such as roads or buildings), people (such as local residents or transport passengers), and phenomena such as the weather. Some physical domains are subdomains of other domains (e.g. roads are a subdomain of the transport network).

Virtual domains (green \square) are the ‘cyber’ part of cyber-physical infrastructure. For the most part, these represent one or more of the physical domains. However, some cases of stand-alone virtual domains exist, such as data lakes.

Domains are connected to one another by *flows* denoted with lines-and-arrows. *In-flows* (turquoise $->$) describe the transfer of data (raw signals or knowledge) into a virtual domain; *Out-flows* (purple $->$) describe the transfer of information (processed data) out of a virtual domain or actions taken on a physical domain; and *physical links* (blue \rightarrow) connect physical domains with one another directly. Note that physical links can be either uni-directional or multidirectional.

Owners are shown using solid black lines (\square). They are the entities with *responsibility* for one or more domain of the digital twin architecture. Here, owners are *risk* owners, rather than conveying any implication of legal ownership. For example, The Met Office do not *own* the weather, but they are responsible for accurate forecasting and so are the owners of the physical domain for the purposes of the DRD. Responsibility might be de-facto or de-jure, and in many cases is already defined by a statutory obligation or legal agreement. For example, DfT’s responsibility for the transport network during a crisis is defined by the ‘Lead Government Department’ framework (Cabinet Office 2023). Owners are a subcategory of the stakeholders: rail users, local residents and insurance companies are examples of stakeholders with an interest in resilient transport, but without ownership over any part of the digital twin architecture. Owners may be in hierarchical relationships, for example DfT has responsibility over National Rail (NR) who themselves have responsibility over the rail network.

To allow for simplification of the DRD, one owner is chosen to be the *viewpoint owner*, denoted with a bolded outline. For example, in Fig. 5, DfT is selected. This does not imply a greater or lesser importance overall, but simply that this particular DRD is drawn from their perspective. The areas relating to the viewpoint owner’s interests will be more detailed, while those relating to other owners can be shown with a greater level of abstraction. This is discussed further in Section 3.6.

3.4 Archetypical DRD Relationships

The DRD offers a flexible framework with which we can represent the overall nature of a cyber-physical ecosystem. The combination of the domains being represented, the types of link between them, and the ways in which the responsibilities of different owners overlap (or don't overlap) all affect this representation. Thus the specific DRD is unique to both the ecosystem it represents and the way in which it is represented.

As an informal taxonomy, we propose six archetypical classes of relationship that can form the basis of DRD:

1. **Centralised** (Fig. 3a): A single owner is responsible for both the physical and virtual domains and they are not connected to other domains. Effectively the digital twin is disconnected from the ecosystem as a whole.
2. **Shared Physical Domain** (Fig. 3b): Multiple owners share responsibility for the same physical domain. If the responsibilities of the owners are distinct, clarity can be improved by showing sub-domains within the shared physical domain. For example, the *flooded area* in Fig. 5 is a shared physical domain, with no clear division of responsibility between the two owners.
3. **Interacting Physical Domains** (Fig. 3c): Multiple distinct physical domains which are connected by at least one physical link. For example, in Fig. 5, *weather* and *rivers* are interacting physical-domains, linked by rainfall.
4. **Horizontal Federation** (Fig. 3d): Two virtual domains with at least one **direct** flow between them. For example, the *transport model* owned by DfT in Fig. 5 is horizontally federated with the *area model* owned by the LRF. These represent virtual information sharing between the owners.
5. **Hierarchical Federation** (Fig. 3e): Two virtual domains with at least one **direct** flow between them, with both virtual domains corresponding to physical domains *and* one of those physical domains is contained within the other. For example, in Fig. 5, rail is a subdomain of the strategic transport network and so we show Network Rail's model as being hierarchically federated with DfT's Transport model ⁴.
6. **Brokered Federation** (Fig. 3f): A virtual domain which is only connected to other virtual domains, facilitating connections between multiple virtual domains with different owners. This represents central data exchange infrastructure of various kinds, including data warehouses, data lakes and marketplaces. The Transit DT in Fig. 7 operates in this capacity.

We do not claim this to be an exhaustive list. It is likely that other relationship types exist in context beyond the scope of this paper, or will be created in the future as the requirements of digital twins develop.

3.5 Creating a DRD

The DRD is a flexible tool and can be constructed in a variety of ways depending on the context. A key goal of the DRD is not just to have a final, diagrammatic, overview

⁴A hierarchical relationship does not imply anything about the legal or functional relationship between the owners themselves, just the domains for which they are responsible.

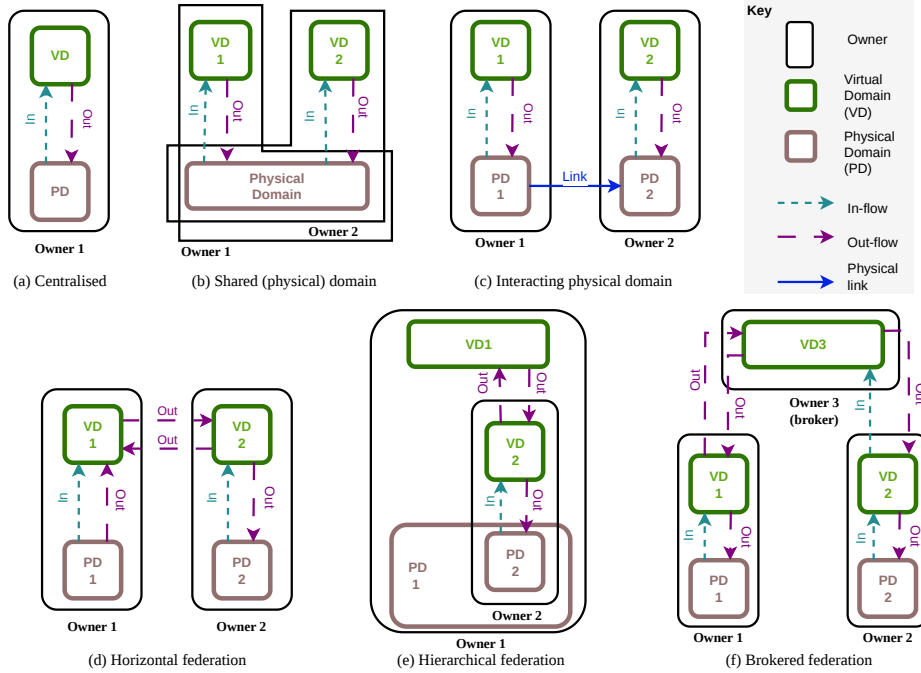


Fig. 3: The six archetypal classes of relationship observed within the DRD.

of the system, but the act of creating the DRD should highlight gaps in knowledge or assumptions. As such, it is best employed within an iterative design process with constant opportunities for review and stakeholder consultation.

In practice we recommend starting by selecting the *viewpoint owner*. This guides the main purpose of the digital twin system and sets the level of abstraction, for example DfT are unlikely to care about the specifics of the flood model implementation, and whether this was outsourced to a hierarchical domain within DEFRA, only the forecast information. Generally the viewpoint owner is made obvious by the context: the transport resilience scenario and the fact it is led by DfT makes them the obvious choice in our scenario above. If there is more than one option for a viewpoint owner, then this indicates that multiple DRDs are required.

Once the viewpoint owner is determined, we recommend building a DRD describing the current state of affairs. It is often easiest to first describe the physical domains and relations between them, as these tend to be more intuitive and existing agreements or regulations often already describe the corresponding owners. This then informs the virtual domains and the in-flows to them. The involvement of domain specialists is invaluable in achieving accuracy at this stage.

The DRD describing the current situation can then be adapted to describe the desired end goal. This can be achieved by adding flows, both to describe the desired actions on the physical domains and the information flow between them as well

as the information exchange between virtual domains. New virtual domains might also be identified in order to enable this. Considering the archetypical relationships (Section 3.4) might aid in shaping the desired architecture – for example is it desirable to have a hierarchical federation or a brokered one.

As with many diagrammatic formats, while the overall structure can be much clearer than textual formats, space for specifics is often limited. We suggest keeping a textual log alongside the DRD to define each element as they are jointly understood, as well as to document questions or assumptions. This allows the text on the DRD itself to be kept to a minimum, increasing clarity, and serving as the basis for other design documentation, such as user stories and hardware specifications.

The figures in this paper were made using `draw.io` (now renamed to `diagrams.net`). The library used is made freely available (Tipuric 2025).

3.6 Principles of the DRD

There is no canonical form of the DRD – the framework is designed to be flexible in order to capture a range of scenarios. However, we propose the following principles in order to maximise the utility of the diagram.

1. The DRD should be easily interpretable by non-technical domain experts.

The primary purpose of the DRD is to support the design and development of digital twins, in part by capturing the insights of non-technical domain experts. As such, it needs to be clear and to the point, with an appropriate level of detail and abstraction for its intended audience.

2. The DRD is iterative.

One of the fundamental goals of the DRD is to expose assumptions or overlooked details. As such issues come to light, the DRD will evolve, until it is in a form that all stakeholders agree with, at which point it can form the basis of more technical specifications.

3. The DRD takes a single point of view.

While the systems represented by the DRD are complex and multifaceted, each stakeholder only typically cares about their own area of responsibility. Each DRD should therefore take the viewpoint of a single owner. Thus the complexity of the diagram is reduced, making it easier to interpret. Where multiple viewpoints are of interest, multiple DRDs should be created – comparison of the similarities and differences between them is then a powerful tool in the design process.

4. The DRD is primarily concerned with the relationships between domains.

The ways in which different owners relate to one another have an impact on how best to design and manage protocols between them. However, the DRD does not seek to represent these inter-stakeholder relationships directly. Notably, just because one domain is in a hierarchical relationship with one another, that does not imply that the owner of one of these domains is subordinate to the other.⁵

⁵In practice, there will often be either a de facto or contractual hierarchical relationship between the owners *within the scope of the specific system being described*. However, in a different context, the same stakeholders might have a different relationship.

4 Case studies

In this section we demonstrate the usage of the DRD in the context of two real-world digital twin proposals. These projects entail designing digital twins within two (different but overlapping) complex cyber-physical ecosystems, with the owners of the domains being autonomous state and private actors. In both cases, we first present the context of each project and then walk through the process of constructing the DRD. We follow this by presenting some feedback on the DRD given by stakeholders familiar with each use case.

4.1 Case Study 1: Transport Resilience

Fig. 4 shows the path of a hypothetical storm hitting the North-East coast of the UK after two weeks of heavy rainfall: leading to significant river flooding across multiple counties. Impacts include landslides, road flooding, flooding of rail and metro stations, wind damage to communication and power infrastructure, sewage overflows, and overflow from reservoirs. We want to determine, and mitigate if possible, the effects on the (critical) transport network including across multiple modalities: specifically rail and road.

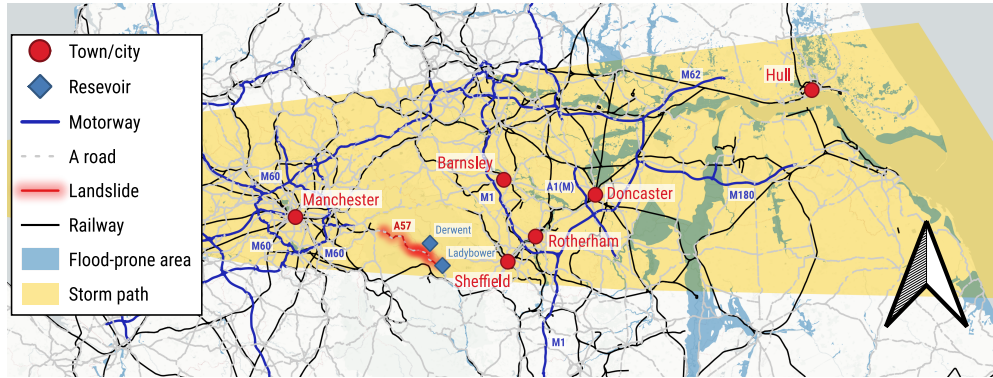


Fig. 4: Map of northern England showing the path and aftermath of a hypothetical storm. Used to support discussion of digital twins for flooding resilience related to transport (Tipuric et al. 2025).

Both the risk to human life and the severe economic impacts make preparing for, responding to, and recovering from disasters such as this a major priority for national governments. As these impacts affect multiple areas, mitigation activities often require collaboration between multiple stakeholders. In the UK, responsibility at the governmental level for both man made and natural disasters is defined by the lead government department model (Cabinet Office 2023). For a storm and fluvial flood event as described here, localised to England alone and impacting the transport network and power grid, the document assigns responsibility for response and recovery to four separate government departments. On a practical level, these agencies will also

need to work with many additional stakeholders, including arms-length bodies, local governments, and the emergency services. In totality, this creates a complex ecosystem of autonomous and semi-autonomous actors.

This risk-scenario was presented at a sandpit event, as part of the Digital Twins Beyond Observed Capabilities (DTBOC) research project (Tipuric et al. 2025) alongside six other risk-scenarios with impacts on the UK transport network. The scenarios covered a wide spectrum of the risks described in the UK National Risk Register (Cabinet Office 2025) and aimed to describe “realistic worst case scenarios”. Groups, consisting of researchers and experts from relevant UK agencies collaboratively identified relevant stakeholders, the needs of those stakeholders, and ways in which digital twins might aid them. The aim of this research was to identify the most promising use cases for crisis response digital twins for transport.

Figure 5 shows a DRD which describes one of the digital twin proposals developed by the sandpit participants (DTBOC 2025) to ‘coordinate emergency service response by utilising meteorological data, river levels, road closures, flood levels, population information, transport flows and logistics information to assign resources and evacuate residents’. The participants judged the potential benefits of this use case to be significant – both in terms of safety and economic impact.

The project highlighted many critical issues in designing digital twins for complex cyber-physical ecosystems, which reinforce our claims in Section 2. The viewpoint from which any given stakeholder approaches a problem affects their initial assumptions and biases in ways that are hard to communicate, or to even identify, until significant work has already been invested in the project. These assumptions and biases often relate to such concepts like the ownership of part of the digital twin ecosystem, the relationships between twins, and the availability of data/models. Equally, definitional issues were common – participants often struggled to agree on the scope and ownership of a given digital twin solution. The DTBOC report (Tipuric et al. 2025) identifies 49 categories of challenge presented by complex cyber-physical ecosystems, such as the intersection of flooding and transport under discussion here. Some indicative examples include:

- Organisational accountability: according to the National Risk Register (Cabinet Office 2025), DfT are responsible for the impacts of flooding on the transport network, while DEFRA is responsible for other impacts.
- Data timeliness: in this example the priorities of DfT exist on a longer timescale than those of the emergency services, for example – the former is more strategic in nature, the latter more tactical.
- Data access: no single stakeholder collects all of the data which would be required, while in many cases the data that might be desired is not collected at all.

Effective design of digital twins that operate in these complex cyber-physical ecosystems requires the ability to easily discuss issues such as these with stakeholders. The lack of an accessible shared ‘language’ for the participants to use was readily apparent. In order to enable digital twin design in these kinds of complex cyber-physical ecosystems, we need a framework which can be used to:

- Map stakeholder responsibilities for both physical and virtual infrastructure;
- Describe existing desired linkages between digital models, shadows and twins;

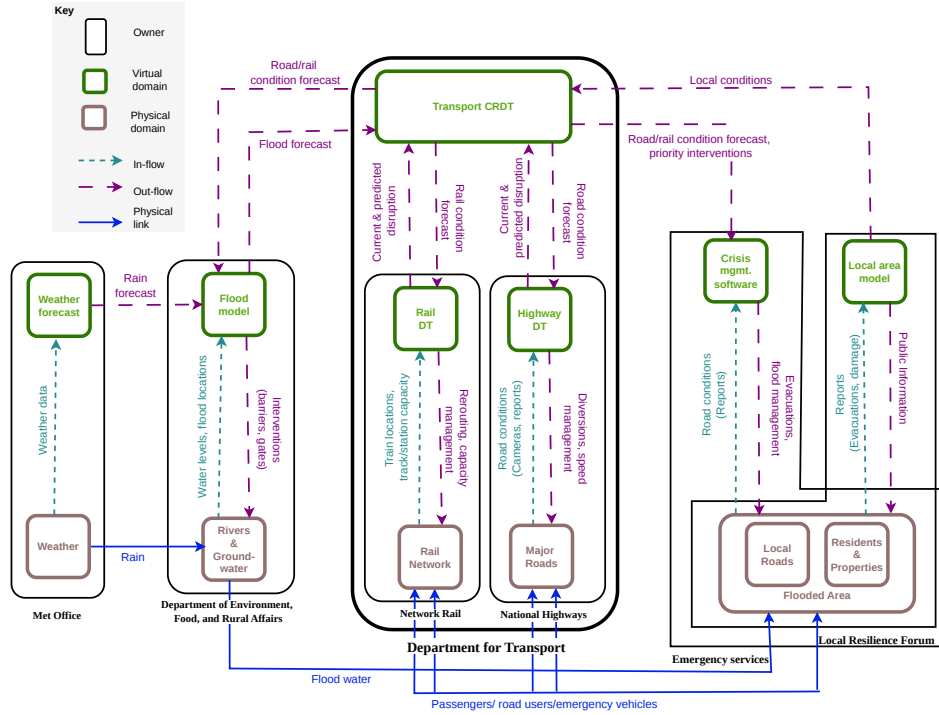


Fig. 5: Domain relationship diagram describing one possible architecture for a crisis resilience digital twin (CRDT) architecture for transport resilience, with the Department for Transport as the viewpoint owner.

- Identify gaps in data collection and modelling availability;
- Enable technical design of digital twin software and hardware; and
- Support inter-party cooperation, such as data-sharing agreements.

4.2 Case Study 2: Electrification of the Port of Dover

The Port of Dover plays a critical role in the United Kingdom's transport and trade infrastructure. It is the busiest international ferry port in the country, handling approximately 2 million freight vehicles annually and facilitates £144 billion in trade each year. This level of activity concentrated in a relatively compact operational footprint presents unique challenges and opportunities for decarbonisation. As targets for emissions reduction intensify, the Port of Dover is a valuable case study for exploring how high-throughput, high-dependency, transport nodes can transition towards more sustainable, low-carbon operations without compromising their critical logistical function. An overview of the landscape and surrounding operations at the Port of Dover can be viewed within Fig. 6 highlighting key road routes which Heavy Goods Vehicles (HGV) take, surrounding approximate energy infrastructure locations, port operations and maritime routes between Dover and Calais.

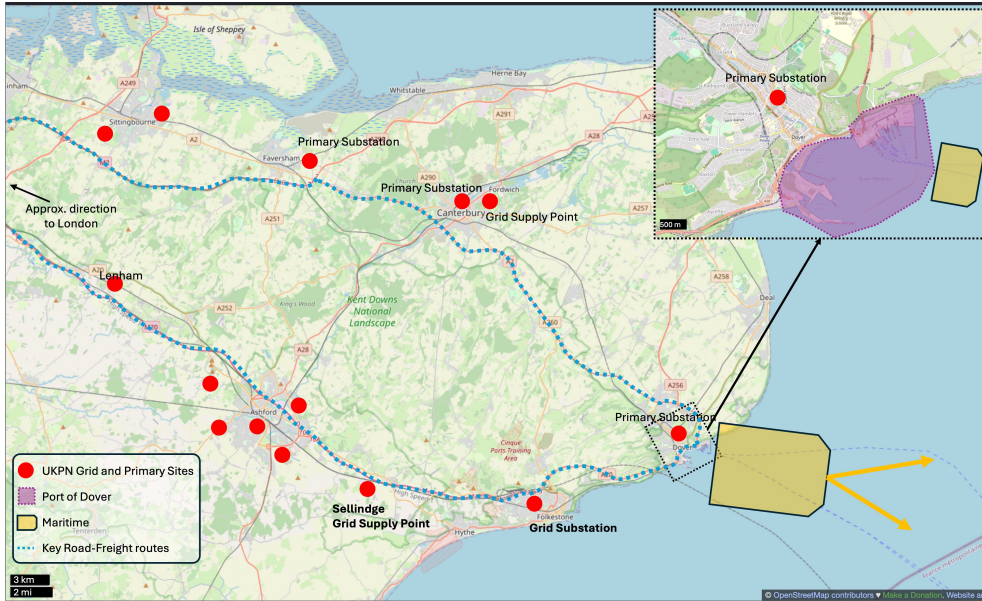


Fig. 6: Map of Dover utilising open access UK Power Networks data relating to substation positions and regions of ownership presented in Fig. 7.

As ports look to explore decarbonisation strategies, digital twins offer powerful capabilities for simulating, testing, and optimising low-carbon interventions before they are physically deployed. Their effective application in environments like the Port of Dover is far from straightforward: the operational landscape of the port is shaped by a diverse and interdependent set of stakeholders, each with its own data sources, priorities, and governance frameworks. This makes it a key area of research for the TransiT national research hub for transport decarbonisation digital twins (TransiT 2023).

The DRD is shown in Fig. 7. We choose TransiT as the viewpoint owners for this DRD. TransiT represents a EPSRC funded project relating to digital twinning for transport decarbonisation. In this context, TransiT owns the port decarbonisation strategy along with the Port of Dover (PoD) TransiT are also the sole owners of the digital twin to keep that strategy current. Along with PoD, the strategy relies on inputs from the UK energy model owned by the National Energy System Operator (NESO), as well as the DfT transport decarbonisation model. It can be seen how TransiT acts as a broker between these three owners. Despite having no direct feedback to any physical domain, TransiT's digital twin does in fact operate as one twin through the networked connections with other owners.

We now add more detail to the DRD by considering the physical domains which interact with the port. In terms of energy infrastructure, we care primarily about gas and electricity, which impact port operations through power availability, cost and carbon intensity. These are owned by the UK Power Networks (UKPN) and SGN,

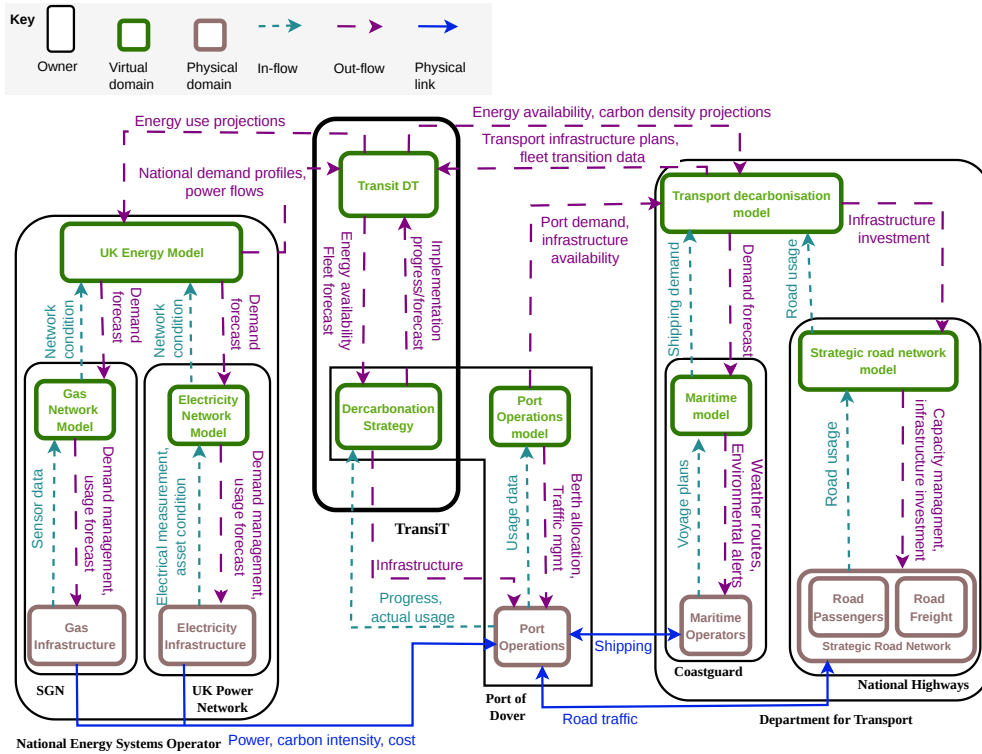


Fig. 7: Domain relationship diagram describing a digital twin architecture designed to meet the challenge within the Port of Dover for a Decarbonisation use case. Note: SGN was previously known as Scotia Gas Network Limited.

respectively and are part of a hierarchical relationship with the national energy network, owned by NESO. Additionally, the port receives both maritime and road traffic, which are in a hierarchical relationship with the national transport network and owned by the Coastguard and National Highways respectively. It is useful to represent the strategic road network as a shared domain containing both passengers and freight, as the practical requirements of both differ.

4.3 Comparison and evaluation

Contrasting Fig. 5 and Fig. 7 allows us to make some observations about the differences, and similarities, between the two use cases. While both projects fall under the broad category of digital twins for transport infrastructure, there is only a partial overlap between the two ecosystems. This, combined with different relationships between the owners, and timescales involved in both projects, leads to the specific DRDs being notably distinct in detail.

The overlap between both projects lies in the strategic road network, both as a physical domain and the virtual domain of its representation. In both cases these are

represented as being owned by National Highways, in a hierarchical federation with the Department of Transport. However, the specific representation of the physical domain differs in each diagram, with Fig. 5 using the term ‘Major Roads’ to contrast with ‘Local Roads’ elsewhere in the diagram, while in Fig. 7, ‘Strategic Road Network’ is used, and the domain subdivided into passengers and freight. This ability to represent the same system in different ways, based on the role it plays in the specific CPES under consideration, demonstrates the flexibility of the DRD. Equally, that this comparison can be made in the first place shows that the minimal structure provided supports meaningful conversation beyond that afforded by unstructured methods.

This flexibility is further evidenced by considering the timescales involved in the flows in each diagram. The first use case is tactical in nature, focused on hour- to day-scale issues, while the second is more strategic, focusing on the year- to decade-scale. This distinction is only evident through familiarity with the specific use cases. As the specific definition of ‘right-time’ data could be expected to be an important point of negotiation between project stakeholders, the inability to represent this on the diagram is arguably a weakness of the format. Similar issues exist for other quantities which might be considered as part of design for CPES, such as sequencing or flow intensity. We believe that the need to reduce cognitive load on those interpreting the diagram outweighs the potential benefits from representing these quantities in some way, especially given the lack of commonly understood natural mappings (in contrast with e.g. the use of arrows to indicate directionality). However, we consider this a matter for open research, along with other questions of interpretability (as discussed in Section 5.1).

Issues of interpretability presented a large degree of difficulty in the construction of both diagrams, especially given the restriction imposed by Principle 1, that the DRD be interpretable by non-experts. Both balancing the amount of text for each element, and the question of which elements to include or omit were a cause for much iteration. In part this can be ascribed to the nature of creating design objects in a vacuum as a demonstration aid for an unknown audience, rather than as part of a collaborative exercise with other stakeholders. As indicated by Principle 2, and further discussed in Section 5.1, the DRD is inherently considered to be an object to be iterated on and negotiated over. However, some consideration should be given to the ways in which software tools can support the behaviour of designers, as discussed by Eng et al. (2017). While being able to produce a static diagram provides some practical value – both in the ability to create pen and paper versions in collaborative design exercises, and in the practical requirement to include static figures in design specifications and documents – the design of bespoke software to support interpretability and iteration is a key issue.

One final point of interest can be observed by looking at the viewpoint owners for both use cases (Department for Transport and TransiT, respectively). Considered in isolation, it could be argued that neither is the owner of a digital twin in the strict sense (of the kind shown in Fig. 1a). Both viewpoint owners only have direct ownership over virtual domains (in contrast to the National Highways in both figures, who are represented as owning a strict digital twin). Nevertheless, the situating of both viewpoint owners in their respective ecosystems makes clear the ways in which

these domains are components of a digital twin. The evidence for this is the closure of a loop between the virtual and physical domains, even though the path passes through domains with other owners.⁶ While not the primary purpose of the DRD, enabling this form of analysis of digital twins in CPES is an interesting benefit.

5 Discussion and Future Directions

The task of creating digital twins of complex cyber-physical ecosystems presents many challenges beyond simply understanding the ecosystems. Here we discuss a few of these challenges, how they relate to the DRD, and the future work required to meet them.

5.1 The DRD as a boundary object

As stated in Section 3.2, the primary aim of constructing the DRD is for it to serve as a boundary object in the design process. Specifically, the proposed uses design scenarios are to request information; negotiate clarity and constraints; facilitate handover; and enable joint design. General methods for each of these are well documented in the design literature (Eckert et al. 2005; Knight et al. 2025) and there is no obvious reason for them not to apply to design for CPES. However, we present some potential examples of how the DRD might be used for each of these design scenarios, to help contextualise the method.

Conceptually, we can consider each of these scenarios to be conversations between two broad classes of stakeholder – system designers and system experts. Designers are those involved in actualisation of the digital twin(s), at any point in the design process from planning, through development, to operation (Tadeja et al. 2026). They might be involved holistically (e.g. as customer or project manager) or in a specific niche (e.g. software engineer). Experts, on the other hand, are the actors with knowledge about some part of the ecosystem. These labels are descriptive and impermanent: the same actor might be an expert at one stage and a designer at another. For example, an engineer might lend their expertise regarding data availability during the ‘analysis’ stage, and then later be design the data pipeline to another part of the ecosystem in the ‘develop’ stage.

Through this lens, requesting information can be seen as an interaction between designers and experts, where the designer seeks to elicit knowledge from the expert regarding their area of the ecosystem. Here the DRD can serve two purposes. Firstly, as a holistic vision of the ecosystem, the DRD can be used to clarify the scope of the design project and the nature of the expert’s niche within it. Secondly, the DRD might aid in reducing information distortion by the highlighting knowledge gaps (the ‘unknown unknowns’) between expert and designer.

Negotiating clarity and constraints happens between designers at an eco-system level as they seek to both bound the totality of the ecosystem under consideration as well as the limits of each domain. Again, the nature of the DRD as a holistic vision can lend itself to these discussions. The requirement to separate the ecosystem into

⁶As a counterexample, consider the Met Office in Fig. 5. There is no way to draw a closed loop starting from the virtual domain of the weather forecast. This matches our intuition – there is no mechanism for the weather forecast to affect the weather itself – and indicates the forecast’s nature as a (high value and immensely detailed) digital shadow.

constituent subsystems attributed to owners intuitively leads to negotiation between designers as to who those owners are and how the constituent systems might interface. An alternative approach in cases with multiple high interest/high influence stakeholders might be to construct multiple DRDs of the same ecosystem with each of those stakeholders as the viewpoint owner. If an equivalence can be drawn between the overlapping domains within each representation, then this can be taken as an indication that the stakeholders share a vision. A lack of correspondence would indicate that further negotiation is needed, with the mismatching areas serving to indicate where to focus these efforts.

In order to support these constraint negotiations, there needs to be some way to compare different DRD representations of overlapping CPES. This can be achieved by recording explicit correspondences between domains, owners, and flows – for example equivalence, refinement, or containment (Nuseibeh et al. 1994). A lightweight comparison method could make use of a shared glossary plus scope notes. An alternative, more formal approach would be to treat overlaps between elements of each DRD as an “ontological overlap” problem and documenting agreed mappings before comparison (Spanoudakis and Finkelstein 1997; Shvaiko and Euzenat 2013). With correspondences in place, a structured comparison would highlight missing or extra domains/flows, disagreements over ownership boundaries, and mismatches in flow direction or type.

Finally, joint design and handover are both interactions between designers at the subsystem/system level. The former concerns information exchange between designers in different domains, while the latter concerns information exchange in the same domain but at different stages of the design process. In both cases, the DRD can serve to provide clarity about what information needs to be exchanged with which parties. Additionally, construction of the DRD might support the design process at a meta level through early negotiation of the means and frequency of communication between design teams operating in different domains. Although not its main purpose, in addition to identifying the live flows required for the design, the DRD might be used to alleviate the risks of miscommunication during the design process caused by ‘missing information provision’ (Eckert et al. 2005) due to power structure or confidentiality concerns, as well as information distortion due to hierarchical communication paths or intermediary expertise.

In each case, while the ways of working proposed in this section are based on existing design theory, their use in this context requires evidencing. Given the large scale and immense variability of CPES, this presents a not inconsiderable amount of future work.

5.2 Fragmented ownership

Whilst digital twins are being developed across these domains, a clear gap relates to fragmented ownership where it remains uncertain who should act as the custodian of a system of systems approach. This fragmentation is often compounded by ‘data hoarding,’ where organisations retain valuable data in silos due to commercial sensitivity, security concerns, or lack of trust. Future directions may include the design of multi-party governance models, inspired by federated data governance approaches

and Catapult-led initiatives, to enable controlled data sharing, establish clear custodianship, and move from isolated ownership to collaborative stewardship. The inclusion of owners in the DRD means it is well-suited to supporting this work, providing a consistent way to identify and interrogate the fragmentation of the ecosystem.

As discussed in Section 5.3, ontologies are a crucial issue here. In transportation, ontologies such as the Transportation Ontology (TRANSIT) and SUMO (KPIonto) can model various aspects of mobility systems, infrastructure, and modes of travel (Katsumi n.d.; Benvenuti et al. 2017). In the energy sector, ontologies such as SAREF (Smart Applications REFerence ontology) and SEAS (Smart Energy Aware Systems) support the modelling of energy usage, production, and devices (Lefrançois 2017). Climate-related ontologies, such as SWEET (Semantic Web for Earth and Environmental Terminology) provide standardized terms for environmental processes and climate events (McGibbney et al. 2022). As cities become smarter and more connected, integrating these ontologies becomes crucial to understanding the interplay between transportation, energy consumption, and climate impacts, ultimately enabling more informed decision-making and planning.

5.3 Communication and Data Flow in Federated Digital Twins

A specific challenge of designing digital twins for complex cyber-physical ecosystems is enabling data sharing between virtual domains, especially those with separate owners. While the DRD aids in the identification of where data needs to be shared (through allocating in- and out-flows), in order to progress we need a machine-readable and semantically coherent framework for data sharing.

Ontologies provide such a framework by enabling translation of abstract organisational interdependencies into explicit specifications of entities, roles, data flows, and governance (Fan and Flood 2018). An ontology allows the relationships between organisations to be expressed in terms of classes, properties, and constraints that formalise key aspects of responsibility and data stewardship. For example, physical domains can be defined as classes with associated data sources, while virtual domains are linked through properties that specify provenance, access rights, and update mechanisms. Figure 8 shows a potential ontology based on the use case described in Section 4.2. Nodes with yellow circles define conceptual classes, nodes with purple diamonds specify data instances of that class, and arrows define the relationships between them, as shown in the legend on the right. For example, in the bottom left, the concept *Organisation* has data instance *CityTransit* and is linked to the concept *Domain* by the green relation *manages domain*.

Ownership can be expressed as higher-level constructs that define which agents—human or organisational—hold authority over particular datasets or processes. These can be created and updated independently of a use case. For example, the UK’s Data Protection Acts (DPA 2018) outline roles and responsibilities for entities that are always required which can be encoded into an ontology. The knowledge in an ontology can then be used with the *instance* of a particular DRD to create a formalised representation. This formalised representation can complement the DRD’s qualitative information descriptions with *quantitative* information descriptions. From Figure 8, yellow marked boxes show quantitative terms, such as (*measurement*) *unit*,

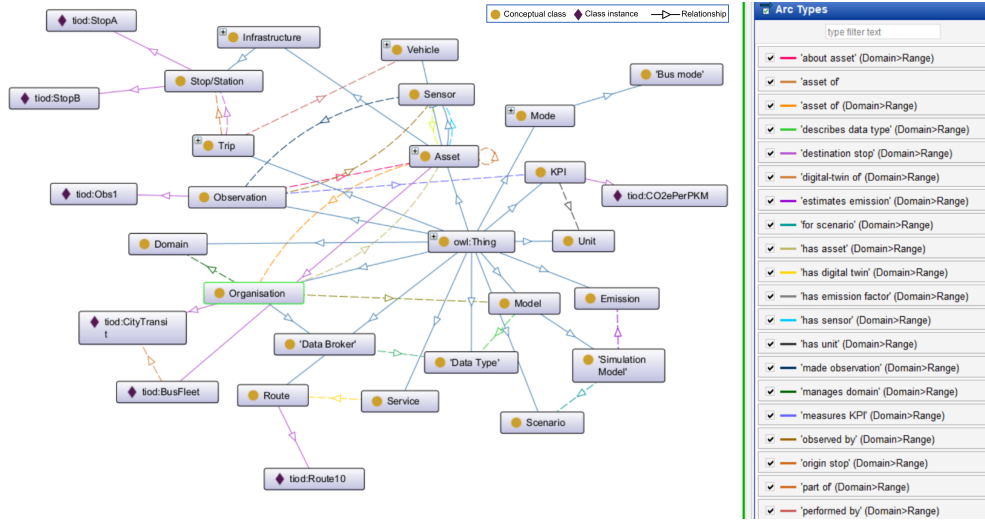


Fig. 8: Example ontology based on Use Case 2. **Note:** 'tiod' is a placeholder name space, with no particular meaning. OWL refers to the Web Ontology Language (Semantic Web Standards 2012).

KPI, emission factor. This can be extended to other uses cases, providing flexibility and adaptability for new scenarios. Also, purple marked boxes show specific instances of data aligned to the defined concepts, also known as knowledge graphs (Ehrlinger and Wöb 2016), enabling explainable inference on specific contexts. While not omitted here to improve legibility, cardinality can also be added to capture the mappings between terms, such as one UK energy model to many supply networks (Figure 7), and timestamps can be added to describe cause-effect or temporal sequencing relationships (Sawesi et al. 2022). In this way, the ontology functions as a semantic bridge between conceptual visualisation and technical design.

These ontological mappings enable several critical features in supporting a federated digital twin. First, they provide a shared vocabulary across organisational boundaries, ensuring that the meaning of digital twin components is preserved even in a distributed environment. Second, they allow for the derivation of data exchange protocols by encoding not only what information must be shared, but also the constraints under which it can be accessed or transformed. Third, they facilitate modularity by encapsulating organisational domains as semi-autonomous sub or *unit-twins* (Ding et al. 2025), which can interoperate through defined ontological interfaces. Finally, the ontology makes governance explicit by embedding roles, responsibilities, and accountability mechanisms into the very structure of the digital twin.

Through this translation, the graphical depiction of organisational relationships is elevated from an informal systems view to a formalised, computable specification (for example, by using a knowledge graph). The resulting design supports a decentralised architecture, in which each organisation retains control over its respective data and

models, while still contributing to a coherent, distributed digital twin that reflects the ecosystem as a whole.

5.4 Trust, Security, Uncertainty & Resilience

Linking digital twins across critical infrastructure domains raises significant challenges of trust, security, and resilience. Each twin carries varying levels of sensitivity, from proprietary port logistics to real-time energy system data, creating concerns over data sovereignty and controlled access. Interconnections between domains introduce new cyber-attack surfaces and amplify the risks of cascading failures across the system-of-systems. Ensuring resilience requires robust cybersecurity mechanisms and strategies for graceful degradation which enable dependent models to continue to function meaningfully if one twin becomes unavailable or compromised. Research is needed into secure, federated architectures that balance openness with protection, enabling trust among diverse stakeholders without compromising system reliability (Mitchell et al. 2025a).

Managing uncertainty in digital twins is essential. Uncertainty can arise at multiple levels, including: inconsistencies between the physical system and the virtual model; specific components of an interconnected digital twin becoming unavailable; and, at the system design level, uncertainty whether the relationships are properly captured (and evolved if required). It might be desirable in future to extend the DRD to explicitly represent these qualities. Similarly, the nature of trust between owners is worth considering (e.g. Governments might be more trusting of the inputs from other Government departments than third party operators). One method for provoking discussion around these issues is the removal of links from the DRD – viewing the diagram with missing links is also likely to spark discussion between designers to interrogate the further consequences and determine if a redesign is necessary.

The DRD is about expressing a shared notion of what a design does, or should, look like, e.g. an intermediary object (Boujut and Blanco 2003), and so is a method of starting discussions around uncertainty in the design of the system itself. However, the intent is to use the information captured through this process to support more structured/formal representations later in the design process. In the context of uncertainty and confidence, we might consider the value obtained through converting the qualitative assessments in the DRD into quantitative values, for example as part of a formal, graph-based representation of the ecosystem. Being able to quantitatively evaluate these metrics would enable a more rigorous assessment of the ecosystem, including assessing trust of individual domains and links; identification of critical weak points; and the quantification of the extent to which uncertainty in any individual component perturbs across the rest of the ecosystem.

Large, complex adaptive systems pose unique challenges and involve varying levels of operational risk. For instance, determining the urgency of repairs and understanding their impact on the wider system are critical—especially when these systems support essential services. In this context, probabilistic model checking is a valuable tool for assessing impact, prioritising and scheduling repairs after component failures, and forecasting maintenance costs (Calder and Sevegnani 2019). Ensuring the reliability and resilience of interconnected digital twins is also a high priority. Similar to their

application in cyber-physical architectures for multi-robot fleets—where they formally verify system behaviour under uncertain and dynamic offshore environments—these approaches could be extended to the energy–transport–maritime ecosystem (Mitchell et al. 2025b). By employing model checking, temporal logic, or other formal verification techniques, it becomes possible to rigorously monitor how system-of-systems twins behave under stress, detect emergent risks such as cascading failures, and guarantee survivability properties even when individual subsystems degrade.

In all cases, the DRD provides a potential tool to support this research. The framework inherently supports the comparison between system-level and ecosystem-level approaches, providing a way to describe emergent behaviours between the two. The representation of owners on the diagram also supports discussion regarding the responsibility for the security of individual systems and the ecosystem as a whole and the trust requirements across data flows. Future work could expand this approach to the subsystem/component level, as well as investigate the best representation for security and resilience features on the diagram.

6 Conclusion

Effective design for complex cyberphysical ecosystems requires a high level, accessible language in order for us to effectively engage with the domain experts and stakeholders who are already involved in those systems. In this paper we presented the domain relationship diagram (DRD) as a concept mapping framework for knowledge synthesis early in the design process. We represent the ecosystem as consisting of a number of ‘domains’ (either physical or virtual), connected either by physical links or data flows. We assigned ‘owners’ to these domains, representing the parties responsible for discreet parts of the ecosystem. In this way, we aim to map responsibilities, describe existing and desired linkages, identify gaps, enable technical design, and support inter-party co-operation. As a concept mapping tool, the DRD balances the flexibility and low accessibility of less structured approaches, while still providing some level of structure and consistency to support the development of formal models later in the design process.

We discussed six archetypical relationships (centralised, shared physical domain, interacting physical domain, horizontal federation, hierarchical federation and brokered federation) described by the DRD, representing classes of architectural relationship between components. These relationships enable discussion around the different design philosophies which might be applied to the cyber-physical ecosystem, by representing different ways to connect the same components. We then described how to create the DRD in general terms, including the four principles of the interpretability, iteration, viewpoint, and domain relationships.

We then demonstrated our approach to creating the DRD in two different contexts, based on existing projects designing digital twins for complex cyber-physical ecosystems. The first context was based on recent research supported by the Department for Transport in the UK and focused on understanding the use of digital twins for crisis resilience in the event of flooding, while the second focused on ongoing work decarbonising the Port of Dover. Between them, these high impact use cases represent

projects that cross multiple fields (including transport, climate, energy, and logistics) and multiple timescales (hour-day scale for the former; year-decade for the latter). This reflects the versatility of the DRD.

Finally we discussed some broader challenges in CPES design which are captured by the DRD, including fragmented ownership and communication & data flows. We discussed the ways in which ontologies play a key role in these challenges by offering a coherent framework for data sharing between domains. We described how we can map the DRD to an ontology, providing a conceptual bridge between the vision offered by the DRD and technical design. We also describe how mapping the connections between domains using the DRD might enable better resilience and security, by supporting the analysis of cascading risks, the assessment of different architectures for security, and the providing a framework to move between sub-system, system and ecosystem level methods.

Acknowledgments

The authors gratefully acknowledge the support and guidance provided by DTNet+, DTBOC the UK Government Department of Transport and the Port of Dover. We would also like to acknowledge the contributions of individuals and organisations to the DTBOC project, including D. McIlhatton, A. Fenton, E. Koutsouroupa, and the many sandpit participants who volunteered their time and expertise.

Funding declaration

This work was supported in part by EPSRC under Grants EP/Z533221/1 and EP/Y016289.

Disclosure

The authors report there are no competing interests to declare.

References

- (2018) Data protection act 2018. UK public General Acts, <https://www.legislation.gov.uk/ukpga/2018/12/contents/enacted> (Accessed 16 September 2025)
- (2023) A general contractor’s perspective on construction digital twin: Implementation, impacts and challenges. *Buildings (Basel)* 13(4):978–
- Aheleroff S, Xu X, Zhong RY, et al (2021) Digital twin as a service (dtaas) in industry 4.0: An architecture reference model. *Advanced Engineering Informatics* 47:101225. <https://doi.org/10.1016/j.aei.2020.101225>
- Almasan P, Ferriol-Galmés M, Paillisse J, et al (2022) Network digital twin: Context, enabling technologies, and opportunities. *IEEE Communications Magazine* 60(11):22–27. <https://doi.org/10.1109/MCOM.001.2200012>

- Archibald B, Calder M, Sevegnani M (2025) Practical modelling with bigraphs. *Form Asp Comput* 37(3). <https://doi.org/10.1145/3721142>, URL <https://doi.org/10.1145/3721142>
- Atalay M, Murat U, Oksuz B, et al (2022) Digital twins in manufacturing: systematic literature review for physical-digital layer categorization and future research directions. *International journal of computer integrated manufacturing* 35(7):679–705
- Benvenuti F, Diamantini C, Potena D, et al (2017) An ontology-based framework to support performance monitoring in public transport systems. *Transportation Research Part C: Emerging Technologies* 81:188–208. <https://doi.org/10.1016/j.trc.2017.06.001>
- Blanche J, Mitchell D, Gupta R, et al (2020) Asset integrity monitoring of wind turbine blades with non-destructive radar sensing. In: 2020 11th IEEE Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), pp 0498–0504, <https://doi.org/10.1109/IEMCON51383.2020.9284941>
- Bolton W (2002) *Control systems*, 1st edn. Newnes, Oxford
- Boujut JF, Blanco E (2003) Intermediary objects as a means to foster co-operation in engineering design. *Computer supported cooperative work* 12(2):205–219. <https://doi.org/10.1023/A:1023980212097>
- British Standards Institution (2025) Bs iso 23247-5 automation systems and integration - digital twin framework for manufacturing. part 5: Digital thread for digital twin. [online] <https://bsol.bsigroup.com/Bibliographic/BibliographicInfoData/0000000030480650>
- Bujorianu ML (2023) Cyber-physical ecosystems: Modelling and verification. In: *International Conference on Engineering of Computer-Based Systems*, Springer, pp 215–230, https://doi.org/10.1007/978-3-031-49252-5_17
- Cabinet Office (2023) *The Roles of Lead Government Departments, Devolved Administrations and Other Public Bodies (HTML)*. URL <https://www.gov.uk/government/publications/list-of-lead-government-departments-responsibilities-for-planning-response-and-recovery-from-emergencies/the-roles-of-lead-government-departments-devolved-administrations-and-other-public-bodies-html>
- Cabinet Office (2025) *National risk register 2025 edition*. Tech. rep., GOV.UK, URL https://assets.publishing.service.gov.uk/media/67b5f85732b2aab18314bbe4/National_Risk_Register_2025.pdf, accessed: 2026-01-23
- Calder M, Sevegnani M (2019) Stochastic model checking for predicting component failures and service availability. *IEEE Transactions on Dependable and Secure Computing* 16(1):174–187. <https://doi.org/10.1109/TDSC.2017.2650901>

- Davies M (2011) Concept mapping, mind mapping and argument mapping: what are the differences and do they matter? *Higher education* 62(3):279–301. <https://doi.org/10.1007/s10734-010-9387-6>
- Defence Science and Technology Laboratory (2025) Guidance digital twin (official). [online] <https://www.gov.uk/government/publications/digital-twin-definition/digital-twin-official>, accessed 20/04/2026
- Department for Business and Trade (2024) What a digital twin is and how you can contribute. [online] <https://www.gov.uk/government/publications/what-a-digital-twin-is-and-how-you-can-contribute/what-a-digital-twin-is-and-how-you-can-contribute>, accessed 06/06/2025
- Department for Transport (2024) Integrated network management digital twin economic benefits analysis project. [online] <https://assets.publishing.service.gov.uk/media/6705442ae84ae1fd8592effa/integrated-network-management-digital-twins-economic-benefits-analysis-executive-brief.pdf>, accessed 06/06/2025
- Ding S, Flynn D, Harvey P (2025) Automatic generation of digital twins for network testing. In: 2025 IEEE 45th International Conference on Distributed Computing Systems Workshops (ICDCSW), pp 653–658, <https://doi.org/10.1109/ICDCSW63273.2025.00122>
- DTBOC (2025) [unpublished data]
- Eckert CM, Maier AM, McMahon C (2005) Design Process Improvement A review of current practice, Springer, chap 9. https://doi.org/10.1007/978-1-84628-061-0_10
- Ehrig H, Ehrig K, Prange U, et al (2006) Fundamentals of Algebraic Graph Transformation. Monographs in Theoretical Computer Science, Springer
- Ehrlinger L, Wöß W (2016) Towards a definition of knowledge graphs. SEMANTiCS (Posters, Demos, SuCCESS) 48(1-4):2
- Eng N, Aurisicchio M, Bracewell R (2017) Mapping software augments engineering design thinking. *Journal of Mechanical Design* 139(5):051103. <https://doi.org/10.1115/1.4036129>
- Fan L, Flood MD (2018) An ontology of ownership and control relations for bank holding companies. In: Proceedings of the Fourth International Workshop on Data Science for Macro-Modeling with Financial and Economic Datasets. Association for Computing Machinery, New York, NY, USA, DSMM'18, <https://doi.org/10.1145/3220547.3220551>
- Giese H, Rumpe B, Schätz B, et al (2012) Science and engineering of cyber-physical systems (dagstuhl seminar 11441). *Dagstuhl Reports* 1(11):1–22. <https://doi.org/10.4230/DagRep.1.11.1>

- Grieves M, Vickers J (2017) *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*, Springer International Publishing, Cham, pp 85–113. https://doi.org/10.1007/978-3-319-38756-7_4, URL https://doi.org/10.1007/978-3-319-38756-7_4
- Jensen JC, Chang DH, Lee EA (2011) A model-based design methodology for cyber-physical systems. In: 2011 7th International Wireless Communications and Mobile Computing Conference, pp 1666–1671, <https://doi.org/10.1109/IWCMC.2011.5982785>
- Katsumi M (n.d.) Transportation system ontology. <http://ontology.eil.utoronto.ca/i-city/TransportationSystem/1.2/>, revision: 1.2
- Khaitan SK, McCalley JD (2015) Design techniques and applications of cyberphysical systems: A survey. *IEEE Systems Journal* 9(2):350–365. <https://doi.org/10.1109/JSYST.2014.2322503>
- Knight A, White LM, Rask J, et al (2025) Design logic visualizations as boundary objects in design. In: *Proceedings of the Design Society*, vol 5. Cambridge University Press, Cambridge, pp 1185–1194, <https://doi.org/10.1017/pds.2025.10132>
- Lefrançois M (2017) Planned ETSI SAREF Extensions based on the W3C&OGC SOSA/SSN-compatible SEAS Ontology Patterns. In: *Proceedings of Workshop on Semantic Interoperability and Standardization in the IoT, SIS-IoT*,
- McGibbney LJ, Whitehead B, Rueda-Velásquez CA, et al (2022) Semantic Web for Earth and Environmental Terminology (SWEET). URL <http://sweetontology.net>
- Mitchell D, Blanche J, Zaki O, et al (2021) Symbiotic system of systems design for safe and resilient autonomous robotics in offshore wind farms. *IEEE Access* 9:141421–141452. <https://doi.org/10.1109/ACCESS.2021.3117727>
- Mitchell D, Baniqued PDE, Kivrak H, et al (2025a) Cyber-physical-human systems for mobile robots in energy asset management: Current practices and future opportunities. *Energy and AI* 21:100570. <https://doi.org/10.1016/j.egyai.2025.100570>
- Mitchell D, Gu Y, Blanche J, et al (2025b) *A Cyber Physical Architecture for Symbiotic Multi-robot Fleet Management*, Springer Nature Singapore, Singapore, pp 379–421. https://doi.org/10.1007/978-981-96-2332-7_17
- Mori K, Chen Y (2014) *Concept-oriented research and development in information technology*, 1st edn. Wiley series in systems engineering and management, Wiley, Hoboken, New Jersey
- Newman J (2001) Failure-space: A systems engineering look at 50 space system failures. *Acta astronautica* 48(5):517–527. [https://doi.org/10.1016/S0094-5765\(01\)000](https://doi.org/10.1016/S0094-5765(01)000)

- Nuseibeh B, Kramer J, Finkelstein A (1994) A framework for expressing the relationships between multiple views in requirements specification. *IEEE Trans Software Eng* 20(10):760–773. <https://doi.org/10.1109/32.328995>
- Papyshev G, Yarime M (2021) Exploring city digital twins as policy tools: A task-based approach to generating synthetic data on urban mobility. *Data & Policy* 3:e16. <https://doi.org/10.1017/dap.2021.17>
- Phillips D (2018) *Python 3 Object Oriented Programming [Third Edition]*. Packt Publishing
- Podvalny SL, Vasiljev EM (2021) Digital twin for smart electricity distribution networks. *IOP conference series Materials Science and Engineering* 1035(1):12047–
- Port of Dover (2024) Funding awarded for development of marine digital twin of dover harbour. [online] <https://www.portofdover.com/news/developmentofmarinedigitaltwinofdoverharbour/>, accessed 29/08/2025
- Rajhans A, Bhave A, Ruchkin I, et al (2014) Supporting heterogeneity in cyber-physical systems architectures. *IEEE Transactions on Automatic Control* 59(12):3178–3193. <https://doi.org/10.1109/TAC.2014.2351672>
- Rosas SR (2012) The utility of concept mapping for actualizing participatory research. *Cuadernos Hispanoamericanos De Psicologia* 12(2):7–24
- Rosen R, von Wichert G, Lo G, et al (2015) About the importance of autonomy and digital twins for the future of manufacturing. *IFAC-PapersOnLine* 48(3):567–572. <https://doi.org/10.1016/j.ifacol.2015.06.141>
- Sawesi S, Rashrash M, Dammann O (2022) The representation of causality and causation with ontologies: A systematic literature review. *Online journal of public health informatics* 14(1):e4
- Semantic Web Standards (2012) Web ontology language (OWL). <https://www.w3.org/OWL/>, accessed 15/05/2026
- Shvaiko P, Euzenat J (2013) Ontology matching: State of the art and future challenges. *IEEE Trans Knowl Data Eng* 25(1):158–176. <https://doi.org/10.1109/TKDE.2011.253>
- Spanoudakis G, Finkelstein A (1997) Reconciling requirements: A method for managing interference, inconsistency and conflict. *Ann Softw Eng* 3:433–457. <https://doi.org/10.1023/A:1018998207415>
- Subrahmanian E, Monarch I, Konda S, et al (2003) Boundary objects and prototypes at the interfaces of engineering design. *Computer supported cooperative work*

12(2):185–203. <https://doi.org/10.1023/A:1023976111188>

- Tadeja S, Jarrett J, Kipouros T (2026) Mapping of the digital twin design process landscape: the results of a domain-expert visual qualitative study. *Research in Engineering Design* 37(1):7. <https://doi.org/10.1007/s00163-025-00468-8>
- Taentzer G (2000) AGG: A graph transformation environment for modeling and validation of software. In: *Applications of Graph Transformations with Industrial Relevance*. Springer, pp 446–453
- Tang W, Mitchell D, Blanche J, et al (2021) Machine learning analysis of non-destructive evaluation data from radar inspection of wind turbine blades. In: *2021 IEEE International Conference on Sensing, Diagnostics, Prognostics, and Control (SDPC)*, pp 122–128, <https://doi.org/10.1109/SDPC52933.2021.9563264>
- Tao F, Zhang H, Liu A, et al (2019) Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics* 15(4):2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
- Tellioglu H (2009) Model-based collaborative design in engineering. *Cooperative Design, Visualization, and Engineering*, 6th International Conference, CDVE pp 85–92. https://doi.org/10.1007/978-3-642-04265-2_12
- Thelen A, Zhang X, Fink O, et al (2022) A comprehensive review of digital twin — part 1: modeling and twinning enabling technologies. *Structural and multidisciplinary optimization* 65(12)
- Tipuric M (2025) DRD Templates. [online] <https://github.com/MTipuric/Domain-Relationship-Diagrams>
- Tipuric M, Fenton A, Wagg D, et al (2025) Digital twinning for crisis response in transport-based scenarios: Research outcomes. Report, URL https://eprints.whiterose.ac.uk/id/oai_id/oai:eprints.whiterose.ac.uk:229620
- TransiT (2023) Who we are. <https://transit.ac.uk/who-we-are/> (Accessed 22 September 2025)
- Wagg DJ, Worden K, Barthorpe RJ, et al (2020) Digital twins: State-of-the-art and future directions for modeling and simulation in engineering dynamics applications. *ASCE-ASME journal of risk and uncertainty in engineering systems, Part B Mechanical engineering* 6(3). <https://doi.org/10.1115/1.4046739>
- Wang Y, Su Z, Guo S, et al (2023) A survey on digital twins: Architecture, enabling technologies, security and privacy, and future prospects. *IEEE Internet of Things Journal* 10(17):14965–14987. <https://doi.org/10.1109/JIOT.2023.3263909>

Willems J (2007) The behavioral approach to open and interconnected systems. <https://doi.org/10.1109/MCS.2007.906923>