Mereotopology for Product Modeling: A New Framework for Product Modeling Based on Logic

Filippo A. Salustri
Ryerson University

Mereotopology for product modelling

A new framework for product modeling based on logic

Filippo A. Salustri, Ph.D., P.Eng.

Ryerson University, Toronto, Canada
salustri@ryerson.ca

Abstract: Mereotopology is the branch of logic that treats the qualitative formalisation of parthood and connection relations between entities. Although it has apparently not yet found use in spatial reasoning for designed product modelling, the author proposes that it is well suited to the task. This paper introduces mereotopology and discusses some of the principles guiding the development of *design mereotopology* (DMT), a logic being developed by the author for product modelling. Typical MT logics cannot be applied directly to engineering applications because they assume a “common sense” approach, whereas design requires a more strict “engineering sense.” DMT can provide a framework for improved understanding of product modelling knowledge and will lead to better computer-based aids to manipulate that knowledge.

Key words: logic, mereotopology, product model, formal method

1. Introduction

Mereotopology (MT) is the branch of logic dealing with the qualitative formalisation of two fundamental relationships between entities: parthood (i.e. one entity being part of another) and connection. As such, MT seems to have obvious application to the modelling of engineered products. However, while the author has found some evidence of work in formal spatial reasoning for design – such as Clementini et al (1997) and Dym et al (1995), and the various papers of Borgo and his colleagues, such as Borgo et al (1996) – the author has found nothing about MT in the current design literature.

Systems of logic can provide new insights into the nature of design knowledge by providing a framework in which to express facts about products in formal ways. They can also reveal the fundamental principles underlying a discipline or field of investigation. The qualitative aspect of logic is also important. Since logic does not require quantitative values such as actual dimensions and masses, it is very well suited to the early stages of design processes and product modelling, where little if any quantitative product information is available. The author has therefore embarked on a research project to study MT and to develop it into a form that could facilitate product representation and reasoning in the upstream stages of design.

MT approaches are substantively different from conventional topologies based on set theory. Most importantly, MT captures formally the actual perceived characteristics of "real" entities in the world; this is something for which so-called point-set topologies were never intended. Thus, MT is grounded in a sense of reality that is more pertinent to design and product modelling than the alternatives.

Furthermore, MT theories such as RCC (Randell et al, 1992) can have tractable subsets that are amenable to computerisation. This opens the possibility of developing new computer-based tools to aid practising designers.

This paper provides a brief overview of general MT, specifically addressing its potential use in product modeling. The author's current research in MT is also presented the goal of which is the development of a specific MT theory for product models and computer-based applications. This theory, called *design mereotopology* (DMT) will be a formal theory of reasoning that interprets the principles of MT from the point of view of what the author calls *engineering sense* (as distinct from *common sense*). Some of the governing principles of DMT are discussed, and the core of the theory itself is presented. Since DMT is still under development, no complete solution is available at this time. Some possible applications of DMT are suggested to show that the theory would have implications for design practice and computer-based applications.

Most of the work in MT is relatively recent, and has been driven by developments in artificial intelligence, robotics, cognitive science, and geomatics and medical imaging. Current trends in AI emphasise systems that can reason with *common sense* knowledge about the world (that is, knowledge typical of the general population); this work has contributed to the development of autonomous vehicles and robots, economic theories, etc. In the area of knowledge-based systems, MT is contributing to the resolution of inconsistencies arising from a lack of distinction between *properties* and *parts* of objects. In geomatics and medical imaging, MT is helping to define systems of reasoning with the complex and abundant data collected with modern imaging and sensing hardware, e.g. Galton (1996). Common sense knowledge is an important aspect of all these research areas.

Common sense is, however, usually insufficient in engineering environments. A different perspective is needed – what the author calls *engineering sense*. Engineers require a perspective that is more structured, more based on scientifically acceptable views of reality, and less tolerant of contradiction and inconsistency, than does the "common" person. They do not care about emotional states or belief systems of a non-technical nature.

Still, this is not to say that we require the most strictly verifiable perspective, in the sense of the natural sciences. For the most part, product engineers do not (yet) care about quarks, dark matter, physical space described by more than four dimensions, or events that happened billions of years ago. Nonetheless, engineers require a far stricter perspective than does the average lay person. They seek a level of consistency consistent with the classical sciences, and scaled within a few orders of magnitude of the human body. But within these limits, design engineers require a strict level of rigour. Design engineers are also interested in maintaining a physical sense that takes into account the macroscopic behaviour of materials and systems that may seem counter-intuitive or bizarre to lay persons (e.g. non-newtonian fluids). This engineering sense, then, sits between the rather naïve *common sense* and the very strict sense of the natural and applied scientists.

Engineering sense impacts on the requirements of logical systems for design engineering. Since most existing logics are based on a common sense perspective, they are likely not completely consistent with engineering sense. Therefore, to develop a proper MT for design, it is very likely that some reinterpretation of “standard” logical notions will be required.

Development of logics for design engineering must not be seen as an attempt to automate design. It is certainly unclear that such a goal would be attainable, let alone desirable. The development of logics serve three purposes in design engineering. First, logic promotes structured thinking. The real worth of logic, as exemplified by the scientific method and by mathematics, is that they provide a framework for thinking about complex problems. These frameworks are particularly important as problem complexity increases. Design is creative, to be sure; but an idea that cannot be pursued to its *logical* conclusion will not be implemented. Also, in today’s highly competitive economy, structured thinking can shorten time-to-market, improve product quality, and lower product cost.

Second, logic helps prune the search space of possible designs by eliminating designs that are logically inconsistent. Logic helps organise knowledge and information about design, which in turn can make evident logical flaws that would otherwise remain hidden in the vast amounts of data generated during a design process. Logic can temper creativity and direct it towards, reasonable, attainable goals.

Third, logic can be used to construct improved computer-based tools. The speed and repeatability of computers combined with advanced reasoning capabilities embodied in systems of logic can lead to new classes of knowledge-based systems. The net gain is that humans now have more time to focus on aspects of product development and design for which the human mind is best – perhaps ideally – suited.

2. Background

The author has been working on product model formalisation for several years. In that work, axiomatic set theory was used as the logical foundation of the *Axiomatic Information Model for Design*, AIM-D (Salustri, 1996). Although some interesting results were uncovered, structural problems were found that arose directly from the nature of set theory itself and its interpretation for design engineering. While trying to address these problems, the author came across mereotopology and therein found a new tool that the author believes is more compatible with design engineering and engineering sense.

A MT theory is any theory combining mereology (a logic of part-whole relations) and topology (a logic of connections between parts). Although mereology has existed for almost a century (Lesniewski 1927-1930, 1982), it is only in the past 15 years that it has found popularity in combination with topology, especially in the field of spatial reasoning. The obvious pertinence of spatial reasoning and product design makes it clear that MT can contribute to the advancement of product modelling and computer-based tools. Historically, mereological and MT theories were overlooked in favour of set-theoretic ones, but recent work (Smith, 1996) indicates that MT has the potential to resolve several open issues in set theory.

Perhaps most importantly for product modelling and computer-based tools, *region-based* theories of MT – such as Eschenbach (1994) and Randell et al (1992) – suggest that MT can be used to represent entities that exist in other spaces besides the usual physical one. A *region* is a portion of a space, typically the portion occupied by some entity, material (e.g. a physical part) or otherwise (e.g. a hole). The overarching goal of any MT theory is to describe the nature of regions and the entities that occupy them, and the interrelations between regions. Regions themselves, however, need not be primitive entities within the theories. There are many MT theories that do not take regions as primitive entities – such as Casati and Varzi (1997) and Smith (1997) – but rather build them up from more primitive entities.

The author, however, believes that regions in a MT for design should be primitive because they are generally accepted primitive items in engineering sense. Design engineers think in terms of *real* objects that are made of a material and that occupy space. The spaces occupied by these objects are regions. All other entities – boundaries, surfaces, points, etc. – are conceptual entities that result from our observations of the interactions between objects (and regions); for example, a boundary is taken as the entity that separates two regions. Eschenbach and Heydrich (1995) show how theories for different domains can be constructed simply by changing how a *region* is defined. Indeed, although the MT theories that have been reported in the literature have nearly always defined regions as physical, spatial ones, the author believes MT can also be applied successfully to non-physical domains that are relevant to product modelling (one example is a "space" of product function). Thus, one properly developed theory can be reused many times in many design domains by swapping out one characterisation of *region* for another, just as is suggested in Eschenbach and Heydrich (1995). Examples of this are given later in this paper.

3. A design mereotopology

In this section, the core of *design mereotopology* (DMT) is presented. The development proceeds by identifying the primitive relations that underlie DMT, characterising them logically (the *terminological* component of the theory), and imposing further *ontological* axioms to restrict the theory to represent only those entities of interest to the product modelling domain. It is assumed that full 1st-order logic is available. The development presented here is based upon the

Closed Region Calculus (CRC) (Eschenbach, 1999) and the theory presented in Smith (1996).

3.1. The Core of DMT

Parthood is represented by *xPy* and is read *"x is part of y"* and represents parthood in the most general sense. That is, if there is any reasonable way to consider one entity as part of another, the *P* must apply. Examples of this include components of assemblies, moments in a timeline, items in a batch, surfaces of a volume, sub-functions of functions, regions of space, and so on. To maintain generality, entities are considered parts of themselves. Other useful relations can then be defined, including proper parthood (*PP*), overlap (*O*), complement ('), binary mereological summation (+), and general mereological summation (*S*).

xPPy	$:= xPy \wedge \neg yPx$	(D1)
xOy	$:= \exists z\ (zPx \wedge zPy)$	(D2)
x = y'	$:= \forall z\ (xOz \Leftrightarrow \neg zPy)$	(D3)
x = y+z	$:= \forall v\ (xOv \Leftrightarrow yOv \vee zOv)$	(D4)
x = Sy[ψ(y)]	$:= \forall z\ (xOz \Leftrightarrow \exists y\ (\psi(y) \wedge yOz))$	(D5)

In D5, ψ stands for a predicate that is used for choosing entities (e.g. "all entities in a product made of steel"). The theory remains a 1st-order logic because existential quantification ($\exists$) and quantifiers embedded in any given ψ are not needed. Furthermore, summation bears a superficial resemblance to set-wise union, but it is in fact very different. The union of two sets includes only the contents of the sets. But the mereological summation of the entities includes their parts, the parts of their parts, and so on. This may not seem sensible, but is in fact a truer depiction of reality. When automobiles are put on a truck – for delivery, say – all the automobiles parts, parts of parts, etc. are on the truck as well. This goes without saying in engineering sense, but must be spelled out within a formal system. The fact that MT theories can capture these "obvious facts" of engineering sense while set theoretic approaches cannot indicates strongly the relative potential of MT compared to set theory as a useful design research tool.

An important feature of *P* is that every entity is a part of itself. This may seem counter-intuitive with respect to engineering sense, but it serves two very important purposes. First, it generalises the sense of parthood to be as all-inclusive as possible. One would naturally expect fundamental relations to be universal (or nearly so); so selecting fundamental relations in a logic to be as broadly applicable as possible is considered "good style". Second, it greatly simplifies the development of the theory in general. A *proper part* relation that is conventionally thought of in engineering sense can be defined easily from *P*. Any MT theory can be rewritten assuming *proper* parthood as the primitive relation without loss of soundness, but with some loss of simplicity.

Parthood in DMT is characterised by the following axioms.

$\forall xyz$	$[xPy \wedge yPz \Rightarrow xPz]$	(A1)
$\forall xy$	$[xPy \wedge yPx \Rightarrow x=y]$	(A2)
$\forall x$	$[xPx]$	(A3)
$\exists y$	$[\psi(y)] \Rightarrow \exists x\ (x = Sy[\psi(y)])$	(A4)
$\forall xy$	$[xPy \Leftrightarrow \forall z\ (zOx \Rightarrow zOy)]$	(A5)

Axioms A1-A3 ensure that *P* is transitive, anti-symmetric, and reflexive. Axiom A4 ensures that the mereological sum of entities that satisfy ψ exists. Axiom A5 asserts *extensionality*, a basic property of any logical theory. In this case, any thing that is a part of another shares parts with the other too; put another way, in combination with A2, A5 says that things that have identically the same parts are in fact identical themselves.

Extensionality is a distinguishing feature of different approaches to MT. Some authors have based extension on purely mereological grounds (e.g. Eschenbach 1999), while others such as Vieu (1993) have added extensionality based solely on topology, and still others theories, such as RCC (Randell et al., 1992) defined extension with respect to both mereology and topology.

Mereological extension says that two entities are identical if the have exactly the same parts. Since MT theories are independent of time (i.e. can be thought of as describing things at an instant only), then there is nothing wrong with mereological extension in an engineering sense. Topological extension says that two entities are identical if they are connected to exactly the same things. This is problematic for two reasons. First, topological extension leads to a prohibition of *atomic* entities, which the author contends are essential for product modelling logics (the importance of atoms will be discussed below). Second, it is possible to model many engineered products in ways that are useful to designers, yet topologically indistinguishable, which can lead to false equivalencies of product models. Thus, in DMT we choose in favour of only mereological extensionality, and prohibit topological extensionality.

The second fundamental relation in DMT is the topological one. The connection relation is written *xCy*, and is read *"…x is connected to y…."* It is intransitive, reflexive, and symmetrical. This primitive covers and form of geometric, physical, or other form of connection or contact, permanent or temporary, independent of the criterion by which a particular sort of connection is defined. Again, various distinctions (e.g. two welded parts versus two parts that are merely in contact) must be made eventually, but the underlying universal sense of connection is defined terminologically at the outset.

Many useful relations can be derived from *C*, including contact (*EC*), tangential and non-tangential proper parts (*TPP* and *NTPP*), self-connection (*SC*), and enclosure (*E*). Self-connection is particularly important for product modelling in that it distinguishes between single-region entities (e.g. physical parts and assemblies) and entities whose elements are not so connected (e.g. the Earth, Moon, and Sun; the stator and the rotor of a motor; a television set and its remote control unit).

xECy	$:= xCy \wedge \neg xOy$	(D6)
xTPPy	$:= xPPy \wedge \exists z\ [zECx \wedge zECy]$	(D7)
xNTPPy	$:= xPPy \wedge \neg\exists z\ [zECx \wedge zECy]$	(D8)
SC(x)	$:= \neg\exists yz\ [x = y+z \wedge \neg(yCz)]$	(D9)
xEy	$:= \forall z\ [zCy \Rightarrow zCx]$	(D10)

Topological structure is defined by the following axioms. Connection is reflexive (A6) and implies a mereological sharing of parts (A7). A8 propagates connectivity from entities to their parts: connected entities must have some connected parts. Similarly, A9 propagates connectivity from parts to their wholes: something connected to a part of a thing is also connected to the thing itself. Finally, A10 ensures that an entity can be uniquely identified by its external connections (this prevents entities from occupying the same space at the same time).

$\forall xy$	$[xCy \Rightarrow yCx]$	(A6)
$\forall xy$	$[xCy \Rightarrow xOy]$	(A7)
$\forall xy$	$[xCy \Rightarrow \exists z\ [zPx \Rightarrow zCy]]$	(A8)
$\forall xyz$	$[zCx \vee zCy \Rightarrow zC(x+y)]$	(A9)
$\forall xy$	$[x=y \Leftrightarrow \forall z\ [zECy \Leftrightarrow zECx]]$	(A10)

It is tempting to equate topological enclosure (D10) and mereological parthood; the two concepts appear quite similar. However, such an equivalence would mean that (a) every region encloses its parts, and (b) every enclosed entity is part of its enclosing entity. While the first condition is perfectly reasonable in engineering sense, the second one is not. An entity may enclose other things besides its parts: for spatial regions, an automobile may enclose its passengers, but its passengers are not among its parts by any definition of parthood consistent with spatial regions. Therefore, only condition (a), above, holds in the form of an implication. Some existing MT theories assume only condition (a), others assume only (b), and still others assume the complete equivalence of enclosure and parthood. Based on the simple argument presented here, however, the author contends that all MT theories supporting either condition (b) in any way must be discounted as foundations for DMT.

The preceding discussion has set the fundamental bounds on the primitive relations in DMT: *P* for parthood, and *C* for connection. The interrelations between the two primitives have also been established.

From these 10 axioms, a number of very reasonable theorems about spatial entities can be proved. Theorem T1 shows that an entity that is not part of another entity has a part that shares nothing with the other. Also, things that share parts are connected (T2), and connected things are parts of things that are connected (T3). These theorems have been proved in Eschenbach (1999). It is important that logical systems be able to prove such “obvious” statements. By doing so, the system proves its worth as a representation of knowledge, and clearly demonstrates that many “obvious” properties of the world can be deduced logically from simple primitive relations.

$\forall xy \quad [\neg xPy \Rightarrow \exists z\, [zPx \wedge \neg yOz]]$ (T1)

$\forall xy \quad [xOy \Rightarrow xCy]$ (T2)

$\forall yz \quad [zCy \Rightarrow \forall x\, [yPx \Rightarrow zCx]]$ (T3)

3.2. Extending the Core System

The modelling functionality of the core system can be extended by defining other mereotopological relations. Two such examples are presented here to show the potential for enhancement of the core system. Other extensions are possible and are being studied by the author. For brevity, axiomatisations for these extensions are omitted; they will be included in a subsequent publication.

It is possible to develop axiomatic representations of two relations, *everywhere in* (*EI*) and *somewhere in* (*SI*) that are valid with the DMT framework. If it is true that a predicate applies everywhere in an entity, then it must also apply somewhere within it; that is, *EI* and *SI* can be both true. The limiting conditions of an entity being homogeneous (*HO*) (i.e. a predicate applying either no where or everywhere within it), or heterogeneous (*HE*) (i.e. a predicate applying somewhere but not everywhere within it) can be easily defined. The underlying work on these was originally presented in Eschenbach (1999). For example, *yEIx* is read *"…y holds everywhere within x…."* Some provable theorems include the following. T4 states that if ψ holds nowhere in an entity, then it holds nowhere in each of its parts. T5 states that if ψ is true everywhere in an entity, then ψ is true somewhere in each entity that shares parts with it.

$\forall x \quad \neg\psi SIx \Rightarrow \forall y\, [yPx \Rightarrow \neg\psi SIy]$ (T4)

$\forall x \quad \psi EIx \Rightarrow \forall y\, [xOy \Rightarrow \psi SIy]$ (T5)

Theorems such as T4 and T5 are important in abstract (e.g. systems) models of products, where it is difficult to identify clearly where one component starts and another ends, especially at the early stages of design processes. Knowing that there exist formal systems that allow such intuitions to be derived suggest that there do exist formal foundations for these intuitions. Hopefully, a detailed study of these formal systems will lead to new insights about the processes that require those intuitions in order to reach a successful conclusion. Also, of course, new computer-based designers' aids may be developed that are able to compute these otherwise intuitive insights, because of the formal systems from which they can be derived.

Another extension is a relation MAX-C (for *maximal component*, a single self-connected entity exhibiting certain characteristics, and being as large as possible with respect to the underlying space. For example, It is possible to derive the intuition that anything connected to a maximal component cannot exhibit the same characteristics as the maximal component. This relation can be used to explore the characteristics of other components implied by the known characteristics of a given component.

3.3. Open Issues

3.3.1. Atoms

As noted above, *atoms* are vital entities in product modelling. In the language of MT, an atom is an entity that has no parts but themselves; that is, they have no mereological structure. Given the fundamental split between mereology and topology, even in combined MT theories, care must be taken to distinguish clearly between mereological and topological notions of *atomic* entities. Atoms are mereological. Their topological counterparts, *points*, can exist separately. A point is an entity with no *topological* structure. Similarly, an atom has no mereological structure, but may have a topological structure. While this may seem counter-intuitive, it is quite reasonable. We commonly refer to entities as *atomic* to indicate the most primitive entities of interest, even if those entities can be extended in space (i.e. they are not points). Examples are a mechanical part when viewed from the perspective of design for assembly, and an assembly that is viewed as a single entity from a systems perspective. To maintain this distinction, it seems that both atoms and points are required in DMT. Eschenbach has provided two examples of MT theories that allow atomic entities (Eschenbach, 1994; Eschenbach, 1999). This is possible because, unlike other MT theories, there is no restriction that every entity must have an interior part. While this seems like a perfectly reasonable assumption, the implications of this for product modelling must be studied further.

In conventional topology, an atom can be any element in a set-theoretic sense. This is similar to the mereological perspective: a topological atom can be any mereological entity. The differences between mereotopological theories and topologies based on set theory arise from the differences between mereology and set theory. For example, the *axiom of foundation* in set theory induces a hierarchy of aggregates (sets that contain sets and so on); there is no such hierarchy in mereology. This has major implications for the mapping of structures in the theory to entities in the theory's domain of application.

3.3.2. Boundaries

Boundaries are entities of dubious status in MT. Some researchers include them in the types of entities to be treated, while others choose to exclude them completely, and still others have found ways of deriving boundaries from the existence of other entities (with some restrictions). Boundaries are difficult to treat because logicians cannot agree on whether they are "real" entities or only cognitive artefacts of human perception and reasoning.

Whatever the epistemological status of boundaries, there is little doubt that they are of fundamental importance in engineering and so must be included in one manner or another in DMT. The intuitive sense of boundary that we want to capture in DMT is of the limit of some property or properties at which point the values of the property or properties change. That is, a boundary is the divider between two entities. A boundary cannot exist around a thing without some other thing being on the other side of the boundary. Boundaries are thus emergent properties of connected entities, not intrinsic properties of single entities. It would make sense, then, to seek a formalism that defines a boundary as a sort of entity shared by the connected and thus mutually bounding entities. If the boundary is shared, then the connected parts should overlap (refer to D2). However, while the two connected things that share a boundary are full entities, the boundary itself is in some way less. In spatial reasoning, boundaries are generally of lower dimension than the entities they bound (i.e. volumes are bound by surfaces, surfaces by curves, and curves by points); the only entities are the volumes and the boundaries are not counted as entities.

Now consider the relation of *contact* (D6), which is defined on two entities that are topologically connected but do not share any parts. So if boundaries are treated as parts then they cannot be shared by entities in contact; and if they are treated as non-parts, then they may be shared by connected entities, but cannot be covered by DMT in its current form.

There are three possible solutions to this conundrum. First, some MT theories derive boundaries by considering special relations between conventional entities (e.g. volumes), such as that in Smith (1996). Second, we can follow the approach of Eschenbach (1994), wherein boundaries are defined but where an extra primitive to distinguish mereological entities ("regions") from topological entities (points, lines, etc.) is introduced. Thirdly, we can ignore the boundaries themselves in favour of identifying the other entities that bound a given one; this solution is based on the use of the MAX-C relation (see below) and the fact that boundaries only occur where entities connect. It is unclear at this time which, if any, of these solutions is best for DMT. The author is currently studying them all.

3.3.3. Granularity

As designs are developed, different product models are generated at different levels of detail. Each level of detail requires a different granularity of the entities described by the model. As granularity increases, items that have no mereological structure – atoms – may come to have parts; similarly, as granularity decreases, entities with parts may become atoms. In order to support such changes of granularity, some restrictions must be placed on what can and cannot change from atom to non-atom and vice versa. Eschenbach (1994) has proposed one way of embedding this into a MT theory, but her solution is based on the existence of a predicate that can distinguish between mereological entities and topological ones. It is not evident to the current author at this time that such a predicate is valid within DMT as defined here. This matter must be studied further.

This completes a brief introduction to DMT and some of its governing principles. Further details will appear in another paper, currently in preparation.

4. Possible application areas

Obviously, DMT is not of direct use to practising designers. However, given a formal theory such as DMT, it is possible to envision a variety of tools based on it. The author's goal here is not to explain in detail how DMT would be applied, but rather to give an overview of the potential benefits DMT could provide in various practical cases.

4.1. Computer-Based Design Tools

Perhaps the most important potential development is that of computer-based reasoning and representation tools to aid designers in building product models. Since formal systems are usually amenable to software implementation, the author sees DMT as the basis of many different computer-based design tools. The author believes that the best near-term application in this regard is in the development of knowledge-based systems (KBSs) using *description logics*, such as that in Brachman et al (1991). One substantial problem with all such KBSs of which the author is aware is that the distinction between properties of entities and parts of entities is not well formalised. Typically, parts will appear as values of slots that are named *has* or *has-parts* (as distinct from other characteristics or properties, usually named for the property itself – such as *colour*). This superficial distinction between parts and properties leads to models that do not reflect the intention of their creators and thus do not communicate design information accurately. (This problem is typically dealt with by placing extra constraints and error checks in the reasoning engine, which makes the system more complex and error-prone.) It also leads to increased computational complexity in the system itself (because the system cannot partition parts from properties). If a system such as DMT were to be embedded into a KBS, a formal distinction between parts and properties could be made, thus likely improving the robustness and computational efficiency of the system.

4.2. Geometric and Configuration Modelling

All MT theories target the representation of complex physical objects; thus, DMT could correspondingly be used for geometric and configuration modelling, especially in the upstream stages of design processes, when quantitative information is rarely available. This is a straightforward application of DMT, since all MT theories take as their basis a domain of physical space and the objects within it.

More research is needed, however, to specify DMT to the domain of physical space. A physical part (of an assembly, for example) is a restricted kind of mereological part. The characteristics that distinguish physical parts (e.g. that it was manufactured, that it is intended to serve a function, that it corresponds only within certain tolerances to its own specification, etc.) are neither mereological or topological. A useful and specific theory for physical parts and assemblies will therefore have DMT embedded within it. Some other theories of this sort exist, such as Borgo et al (1996). They will be studied in detail as part of this aspect of DMT's development.

4.3. Material Knowledge-Bases

One important special case is that of *mass nouns* that name a material without naming an object (e.g. *steel* or *lubricant*). Technically, mass nouns are strictly conceptual structures that do not exist in reality (i.e. there is no *steel* per se, only actual pieces of the stuff having dimension and volume). They are, however, also fundamental to engineering and designing, and so should be treated somehow. Lacking volume, mass nouns must therefore be treated separately from physical objects. This can be achieved by considering a different, non-physical space. To treat mass nouns, we consider an abstract space whose axes measure material properties (e.g. density, tensile strength, electrical resistance, and chemical volatility) as well as other characteristics of significance in design settings, such as cost and availability.

We can then envision a space populated by regions, each of which represents one kind of material. Specific materials would be represented by small volumes (in order to account for statistical deviation from standard values for those characteristics). In turn, these would be *parts* of larger regions representing classes of materials (e.g. *steel*). Regions of overlap indicate regions where different materials have similar properties, and regions that abut one another indicate materials that together cover a wider range of characteristics. A void in this space indicates a combination of characteristics for which no materials exist – and is suggestive of areas of research for materials scientists. Indeed, such a space can be looked upon as a "periodic table" of materials, where a material's location relates directly to its properties – and provides much the same kind of visualisation capability as the periodic table of the elements. The problem of material selection now becomes equivalent to a geometric problem involving searching a *space*, a process that can be reasoned about in a semi-automatic way.

One can even imagine a computer-based tool that depicts the space graphically by mapping characteristics to physical dimensions, and allowing designers to navigate a materials space, guiding a semi-automated reasoning engine in order to conduct materials selection tasks.

4.4. Function Modelling

Another area where DMT could be applied is in modelling product *function* rather than structure. Function modelling and analysis is acknowledged as an essential aspect of the upstream stages of product development because early design decisions have the strongest effect on product quality and cost.

Many approaches to function modelling are reported in the literature – such as Yang and Salustri (1999) and Qian and Gero (1996); no clear advantage has been identified for any of these approaches so far. Generally, the goal of all these formalisms is to establish a set of primitive and independent functions from which all others are derived. This approach is consistent with DMT: each independent primitive function type would be represented as an axis in a space. The resulting abstract space would be described with DMT. A product would now be represented as the region in the function-space that contains all the functions the product provides. Each physical part of the product would fill a sub-region corresponding to the functions provided by the part.

Various inferences can be drawn from the relations that occur between parts and products in the function-space. Parts whose function-regions lie partly outside that of the product contribute unnecessary functionality; such parts could be redesigned to remove functionality without affecting the product itself. Regions of product function for which there are no overlapping part function regions suggest functionality that is *emergent* from the combination of parts in a given way. One might then consider the notion of "functional efficiency" as the ratio of total product "volume" in the function-space to the volume of all the product's parts in that space. The larger the ratio, the more emergent functions are exhibited by the product per part function. Again, being able to reason in this abstract function-space facilitates reasoning about products.

4.5. System Modelling

A system is defined as a set of interacting components that provides a definite set of functions, and as being crisply distinct from its operating environment (Karnopp et al, 1990). Systems are abstractions of physical assemblies, the differences being (1) system components need not interact through physical contact alone, and (2) system components need not have any physical manifestation at all. Because of this similarity, it is reasonable to expect DMT to represent systems as well as it can potentially represent physical objects.

Of particular note in this case is the broad applicability of systems theory (including mechanical, electrical, software, biological, organisational, and other areas). Because of this, one may also propose that DMT – or at least some variation of it – could form a foundational logic for *all* such disciplines. It could also be a mechanism to support knowledge transfer between disciplines by providing a common knowledge representation language.

5. Conclusion

Clearly, there remains a substantial amount of work to do on DMT. Some open issues that the author is currently pursuing are discussed briefly in Section 3.3: atoms, boundaries, and granularity. Furthermore, it would be useful to compare DMT to other forms of MT, for the sake of identifying problems in DMT that have not yet been detected as well as seeking features of other theories that could be incorporated into DMT.

As has been indicated above, it can be the case that counter-intuitive notions make engineering sense on closer examination, such as the existence of regions that are closed (as opposed to regions that are bounded). Engineering sense suggests that closed entities are required to appropriately represent physical product models, but the details of the best form of closure to use is not obvious. For other kinds of models (e.g. function-space models) open regions may serve important purposes.

This paper has introduced the beginnings of a *design mereotopology*, a logical theory suitable for application to product modelling and spatial reasoning of designed products. The theory is still in its infancy, but there are indications that once mature, it will be able to set a qualitative logical foundation for product modelling. Future work by the author will treat various outstanding issues in the logic as well as exploring further the mechanisms to develop practicable systems that embed DMT in knowledge-based environments. The author hopes that the development of DMT will further our understanding of the design engineering endeavour and improve our collective ability to develop innovative, high-quality, and low-cost products.

Acknowledgements

The author graciously acknowledges the financial support of the National Sciences and Engineering Research Council of Canada.

References

Borgo, S., Guarino, N., and Masolo, C. (1996) A pointless theory of space based on strong connection and congruence. *Principles of Knowledge Representation and Reasoning* eds L.C. Aiello and J. Doyle, pp. 220-229. Morgan Kaufmann, Boston.

Brachman, R.J., McGuinness, D.L., Patel-Schneider, P.F., and Resnick, L.A. (1991) Living with CLASSIC: when and how to use a KL-ONE-like language. *Principles of Semantic Networks: Explorations in the Representation of Knowledge* ed J. Sowa, pp. 401-456. Morgan Kaufmann, San Mateo.

Casati, R. and Varzi, A.C. (1997) Spatial entities. *Spatial and Temporal Reasoning* ed O. Stock, pp. 73-96. Kluwer, Dordrecht.

Clementini, E., Di Felice, P., and Hernandez, D. (1997) Qualitative representation of positional information. *Artificial Intelligence*, 95, 317-356.

Dym, C.L., Summers, M.D., Demel, C.T., and Wong, C.S. (1995) DEEP: A knowledge-based (expert) system for electric plat design. *Computing Systems in Engineering*, 6(6), 497-509.

Eschenbach, C. (1994) A mereotopological definition of point. *Topological Foundations of Cognitive Science, Papers from the Workshop at the FISI-CS,* Buffalo, NY, pp. 63-80.

Eschenbach, C. (1999) A predication calculus for qualitative spatial representation. *Spatial Information Theory: Proc. COSIT 99*, Heidelberg, pp. 157-172.

Eschenbach, C. and Heydrich, W. (1995) Classical mereology and restricted domains. *Intl J Human-Computer Studies*, 43, 723-740.

Galton, A. (1996) The mereotopology of discrete space. *Spatial Information Theory: Proc. COSIT 99*, Heidelberg, pp. 251-266.

Karnopp, D.C., Margolis, D.L., and Rosenberg, R.C. (1990) *Systems Dynamics: a Unified Approach*. Wiley & Sons, New York.

Qian, L. and Gero, S. (1996) Function-behavior-structure paths and their role in analogy-based design. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 10, 289-312.

Lesniewski, S. (1927-1930) O podstawach matematyki. *Przeglad Filozoficzny*, 30, 164-206; 31, 261-291; 32, 60-101; 33, 75-105, 142-170.

Lesniewski, S. (1982) On the foundations of mathematics. *Topoi*, 2, 7-52 (abridged English translation of 'O podstawach matematyki').

Randell, D.A., Cui, Z., and Cohn, A.G. (1992) A spatial logic based on regions and connection, *in* B. Nebel et al (eds), *Principles of Knowledge Representation and Reasoning* eds B. Nebel et al, pp. 165-176. Cambridge, Mass.

Salustri, F.A. (1996) A formal theory for knowledge-based product model representation. *Knowledge-Intensive CAD II: Proc. IFIP WG5.2 Workshop* eds S. Finger, T. Tomiyama, and M. Matyla, pp. 59-78. Chapman & Hall, London.

Smith, B. (1996) Mereotopology: a theory of parts and boundaries. *Data and Knowledge Engineering*, 20, 287-303.

Smith, B. (1997) Boundaries: an essay in mereotopology. *The Philosophy of Roderick Chisholm* ed L. Hahn, pp. 534-561. LaSalle: Open Court.

Vieu, L. (1993) A logical framework for reasoning about space. *Spatial Information Theory: Proc. COSIT 93*, Elba, pp. 25-35.

Yang, B. and Salustri, F.A. (1999) Function modeling based on interactions of mass, energy and information. *Proc 12th Florida Artificial Intelligence Research Symposium*, Sanibel Island, USA. pp. 384-388.