INTERNATIONAL CONFERENCE ON ENGINEERING DESIGN, ICED'07
28 - 31 AUGUST 2007, CITE DES SCIENCES ET DE L'INDUSTRIE, PARIS, FRANCE

PREDICTING CHANGE PROPAGATION ON DIFFERENT LEVELS OF GRANULARITY: AN ALGORITHMIC VIEW

O. O. Ariyo[1], R. Keller[1], C. M. Eckert[1], and P. J. Clarkson[1]
[1]Cambridge Engineering Design Centre

ABSTRACT
When using connectivity models to assess the potential impacts of component changes on other parts of a product, plausible inferences are readily assessed when such products are represented at the appropriate level of granularity to support specific queries. In this paper, we describe the development of a prediction algorithm, which enables coherent computations of the likelihoods of change propagating on several levels of detail of product description given component level change probabilities. The results show that a multilevel approach to change prediction supports an increased range of design queries beyond that achievable with a single level model. Such change prediction capability is useful when carrying out a comprehensive change impact assessment.

Keywords: Product modelling, Product overview, Change management

1 INTRODUCTION

In order for manufacturers of complex products such as diesel engines or gas turbines to be competitive, existing designs are often modified to meet new requirements. When adopting such a strategy, it is considered good practise to assess the implications of change proposals on other parts of a product prior to its implementation. A probabilistic approach to identifying components for assessment was put forward by Clarkson et al. [1]. It involves estimating the likelihood that a component '$\mathbf{a_i}$' of a product '**P**' will affect on another component '$\mathbf{a_k}$' directly as well as through a combination of effects on other components within a product as shown in Figure 1a. This assists design engineers in drawing attention to components which may be directly or indirectly affected by carrying out a design change. The practicality of this approach to addressing design queries depends on finding a suitable decomposition of the product in question from which relevant inferences can be made.

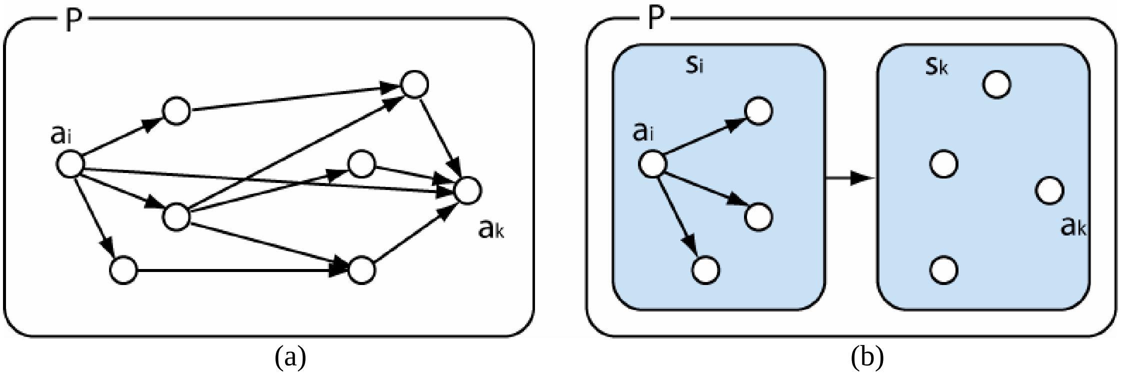

Figure 1: Representations of (a) component connectivity (b) system connectivity in a product

In order to evaluate effects of change on products which consists of several thousand components, it is important that the product is decomposed into understandable and manageable representations of interactions such that design queries can be supported at different levels of abstraction [2]. For example, consider the interactions between the *Starter System* and the *Ignition System* of a typical diesel engine. The *Starter System* consists of components such as a *Starter Aid* and a *Starter Motor* while the *Ignition System* consists of components such as the *Fuel Injectors*. The properties of the *Starter system* may change as a consequence of change to a component such as the *Starter Motor*. This change in the system properties may have significant implications on its interactions with the *Ignition System*. However, it is difficult to find a suitable product decomposition from which all the necessary inferences can be drawn. A system level decomposition does not support a query about effects of changing a component like the *Starter Motor*. Similarly, a component level decomposition may not reflect the dependency between the *Starter System* and the *Ignition System*. As such, in addition to assessing the likelihoods of a component '$\mathbf{a_i}$' affecting '$\mathbf{a_k}$', it would also be advantageous to estimate likelihood of systems '$\mathbf{s_i}$' affecting '$\mathbf{s_k}$' as shown in Figure 1b as there is currently no way to compute system-to-system likelihood values.

In this paper, we describe an algorithm which supports assessing the implications of change proposals on various levels of granularity. The modelling approach was developed as an extension to the Change Prediction Method (CPM), a modelling technique used to predict change propagation in complex design developed by Clarkson et al. [1]. The information presented in every possible view is coherent with the information stored in the original model.

2 HIERARCHIES

Describing complex products at different levels of granularity can be carried out using hierarchical structures. Yet while this mode of structuring has been successfully deployed in various fields, it can be a problematic concept [3]. In order that hierarchical structures are adapted into product connectivity models, it is important to identify its fundamental characteristics and analyse the implications of such properties on product representation.

2.1 Hierarchical structures

Hierarchy formation is a process of abstracting and classifying based on relations. The structure of a hierarchy basically represents abstract units, which have ranked relations between them. Each unit corresponds to nodes or levels of a hierarchy and they may be related to sub-units, based on some form of ranked relationship. For example the *Starter Aid* can be classified as a sub-unit of the *Starter System*.

The properties of abstracted units are dependent on factors such as the particular abstraction principle. A generalisation/specialisation type abstraction principle may be a useful way to describe varying configurations of a family of engines, but it will not support description of connectivity in products. To attain such a goal, it is important that the structure is composed of an interaction between parts and wholes in a product. These are types of relation, which sometime lead to hierarchy formation between parts that make up an arbitrary whole. Described in many texts as part-whole or whole-part relations, these types of relations are very common in design. Examples of part-whole relations include functional decompositions, component assemblies and organisation structures

The resulting structure of a hierarchy depends on the properties of the entity that are being abstracted and the relationships between entities on the nodes of the hierarchy. Simon [4] explains that complex systems are 'nearly decomposable' in that they can only be partially decomposed. While some interactions between component pairs are strong, there are interactions with other systems, which may be weak, but not negligible. Attempts to describe a product's structure hierarchically depend on the strength of interactions between the components. At some threshold of complication, it becomes difficult to represent entities which describe relations in complex systems as trees [5]. To this end, breaking down product description to fit into tree structures may require that a strict decomposition procedure is followed.

2.2 Need for hierarchies

The need for hierarchies can on the one hand be expressed as a cognitive argument. It has been shown that humans are limited in memory capacity [6], hence some information is too complex to be perceived on the lowest level of granularity. For example, being able to memorise information regarding up to 10,000 components that are part of a helicopter is impossible for humans. Hierarchical structures are a way of grouping information into manageable chunks. Also, some problems need to be represented in a particular level of abstraction. For example, it may be difficult to address issues of how the vibration of an engine affects a chassis, if such a product were as described at the level of its piston rings.

On the other hand, hierarchies allow a specific view on the model or product, which can be tailored to the stakeholders' needs. For example, a project manager, responsible for scheduling and costing, requires information on a much lower level of abstraction than a technical designer working on a particular aspect of a product. Furnas' fisheye views [7] are one way of modifying a hierarchical structure to allow points-of-interests to be highlighted.

2.3 Problems with hierarchies

While hierarchical structures are necessary and useful, they are also problematic. Forming a suitable hierarchical structure is difficult. Also, there is no one single context for structuring a hierarchy. For example, consider decomposing an engine into parts in which components such as sensors belong to a particular assembly. The reality may be that it is more practical to categorise such a component into functional clusters, perhaps dedicated only to sensors, despite each sensor not necessarily being in the same physical assembly. Neat clusters which fall strictly within the rules of a particular decomposition context are not always practical.

3 BACKGROUND

The CPM approach to assessing likelihoods of change impacting on other components is carried out in a series of steps [1]. The first of these is to identify linkages between component pairs within the product. Each component is equivalent to a node in a network of interacting components (see Figure 1a). The second step is to allocate the likelihoods of changes propagating between linked pairs. These values are based on expert opinions on the nature of interactions between components. Based on these direct likelihood estimates, the probability of a change to one component affecting another component within the product can be computed. Figure 2a shows an example of a hair dryer. The likelihood that a change to the power supply will affect the motor is calculated by creating a potential propagation tree as shown in Figure 2b. The sum of the total probabilities for each trial in this tree is the likelihood that change may propagate. This resulting likelihood estimate can be computed using equation 1, where "ρ" represents the probability value for each trail "***i***" on the tree.

Initiating Component

Affected Component

Dependency	a	b	c	d	e	f
a power supply	-	x	x			
b motor	X	-		X	x	X
c heating unit	x		-	x	x	x
d fan	X	x	x	-		x
e control system			x	x	-	x
f casing	X	x	x	x	x	-

(a)

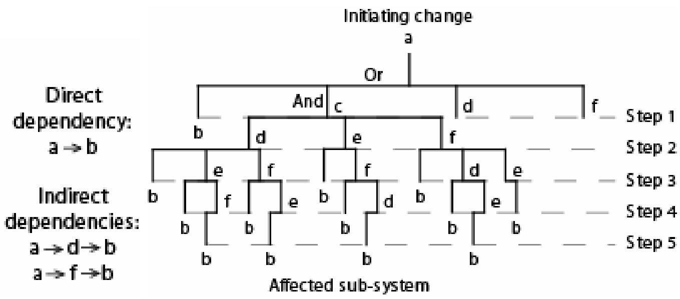

(b)

Figure 2: (a) Connectivity model of a hairdryer (b) Potential propagation tree between its power supply and motor [1]

$$likelihood(a\ to\ b) = 1 - \prod_i \left[1 - likelihood\left(\rho_i\right)\right] \qquad (1)$$

One possible approach to assessing likelihoods of change propagating at various levels of granularity involves repeating the CPM approach for different product decompositions. As such, the likelihood of change propagating between a pair of systems such as the *Starter System* and the *Ignition System* discussed in Section 1 can be obtained following the same steps used for component assessments. The main problem with this approach is that it is difficult to ensure that likelihood estimates are consistent across all levels of granularity. The reason is that the initial likelihoods are obtained using heuristic which are prone to bias and lead to inconsistent likelihood approximations. An example of such bias in probabilistic judgement is the representativeness bias i.e. when probabilities are derived based on how similar events are to each other. For example, the effect a change to a *Starter Motor* has on *Fuel Injectors* may be perceived similar to the effects a change to the *Starter System* has on an *Ignition System* and as a result, the frequency at which other components within the system initiate change may not be properly accounted for when estimating likelihoods of a system change on other systems [8, 9]. This drawback in this approach is in addition to the extensive effort required to build a model for each possible set of queries.

The algorithm described in this paper enables consistent likelihood estimation on various levels of granularity within a product, but it only does so in a bottom-up manner. This is because it is more practical to infer higher level interaction from lower level ones without any further expert input, than in a top-down situation. In addition to estimates of likelihoods that change may propagate from component-to-component which can be derived using the original CPM approach, the system level descriptions enable the possibility of assessing the likelihoods of changes propagating from (1) system-to-components, (2) components-to-systems as well as from (3) systems-to-systems. This significantly increases the range of queries that can be supported with a single model and can be valuable when assessing the effect of change in complex products.

4 ESTIMATING PROPAGATION LIKELIHOODS ON DIFFERENT LEVELS OF GRANULARITY

In order that change propagation likelihood can be assessed at different levels of abstraction, it is important to account for implications of the part-whole relationship that exists between components and systems. To this end we distinguish between intra-system connectives and inter-system connectivity in products. This distinction is particularly important when estimating both system-to-components as well as component-to-system likelihoods because each system is in itself a set of interacting components.

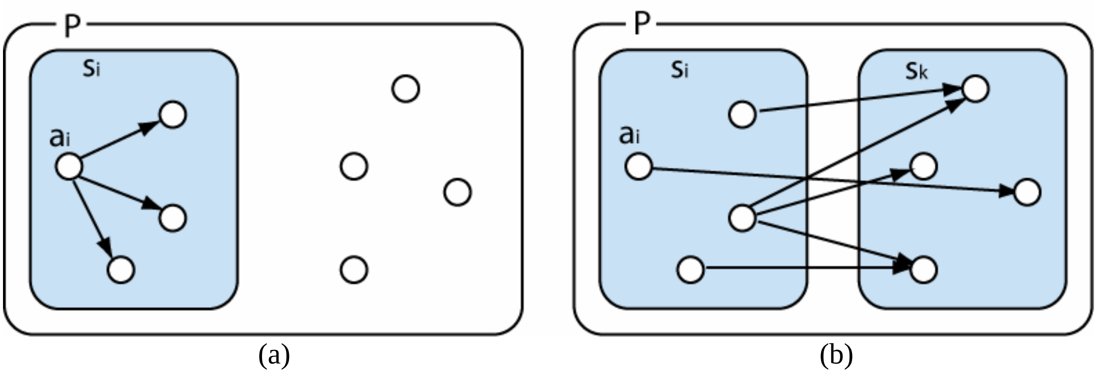

Figure 3: (a) Intra-system connectivity (b) inter-system connectivity

- Intra-system connectivity: this term is used to describe linkages between components within the same system as illustrated in Figure 3a. It is important to distinguish this type of relationship from an inter-system relationship to avoid introducing errors into the risk estimation process. For example, consider an assessment of system-to-component propagation likelihoods between a *Fuel System* and a *Fuel Pump*. A change to the *Fuel Pump* automatically implies a change to the *Fuel System*. However the converse is not necessarily true. This interaction between the system and the component should be taken into consideration when assessing likelihood estimates.
- Inter-system connectivity: The inter-system connectivity refers to the linkages between components of two separate systems as shown in Figure 3b. This type of interaction implies that when there is a linkage between components of two separate systems, then both system nodes are interacting. This understanding is core to ensuring consistency across various levels of granularity.

In order to assess the likelihood of change propagating on various levels of granularity within a product, the system-to-components, components-to-systems and systems-to-systems likelihoods can each assessed independently. This approach was taken so that the considerations affecting each type of relationship can be accounted for individually. These considerations are in addition to the factor of whether or not there is an intra-system or inter-system dependency between components and systems.

The diagram in Figure 4 represents interactions between components in two hypothetical systems **A** and **B**. The initiating component $\mathbf{a_i}$ in **A** may cause change to propagate to an intra-system component $\mathbf{a_k}$ that lies at the 'border' of system **A** which is in turn connected to a final component $\mathbf{b_j}$ in the second system **B**.

The combined component-to-component likelihood on the intra-system level is computed by applying the standard CPM algorithm as described by Clarkson et al. [1] and Section 3 only for the components within the system **A**. The link to the component in the other system **B** is the direct change likelihood of $\mathbf{a_k}$ affecting $\mathbf{b_j}$. These two terms can be combined to obtain a likelihood that one component $\mathbf{a_i}$ in system **A** affects a component $\mathbf{b_j}$ in system **B** via a particular 'border' component $\mathbf{a_k}$ in **A** (Equation 2). The likelihood that the initiating component $\mathbf{a_i}$ then affects a component $\mathbf{b_j}$ in the other system (inter-system connectivity) via all possible intra-system 'border' components is given in Equation 3, which is simply the summation of the probabilities of the paths via all possible 'border' components. This component-to-component inter-system likelihood will be used in the following sections to explain how system-to-components, components-to-systems and systems-to-systems likelihoods are calculated.

$$likelihood(a_i \text{ to } a_k \text{ to } b_j) = combined\_likelihood(a_i \text{ to } a_k) direct\_likelihood(a_k \text{ to } b_j) \quad (2)$$

$$likelihood(a_i \text{ to } b_j) = 1 - \prod_{k=1}^{n} \left[1 - likelihood\left(a_i \text{ to } a_k \text{ to } b_j\right)\right] \quad (3)$$

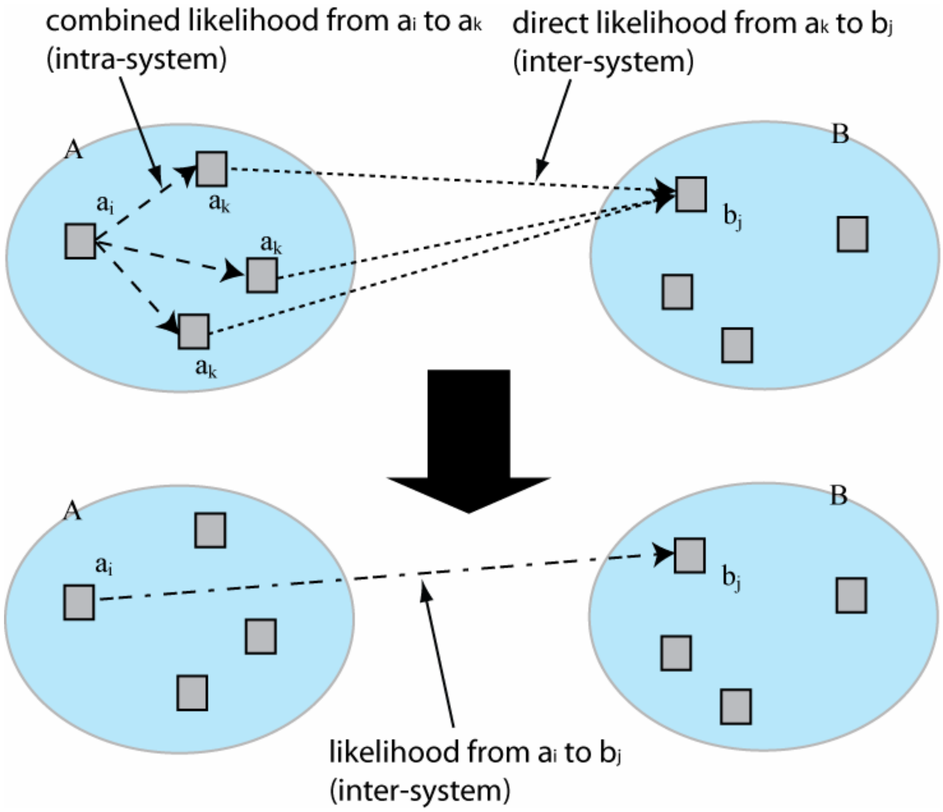

Figure 4: Likelihood estimation using node link representation

4.1 Component-to-system likelihood estimation

Following the notation shown in Figure 4, the likelihood that a change propagates from a component to a system is obtained simply by adding the probabilities that a change to a component $\mathbf{a_i}$ in system **A** will propagate to any of the components in system **B** (Equation 4). The likelihood of a single component $\mathbf{a_i}$ in **A** affecting a component $\mathbf{b_j}$ in **B** is equivalent to the inter-system likelihood via all 'border' components that was established in Equation 3.

$$likelihood(a_i \text{ to } B) = 1 - \prod_{j=1}^{n} \left[1 - likelihood\left(a_i \text{ to } b_j\right)\right] \quad (4)$$

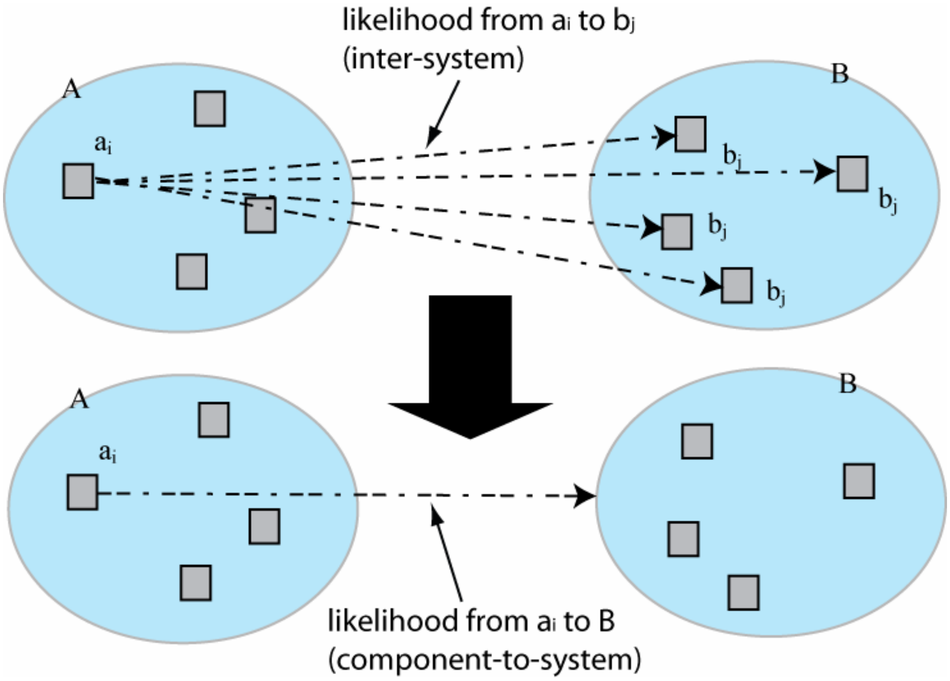

Figure 5: Estimating likelihood of change propagating from a component to a system

4.2 System-to-component likelihood estimation

Assessing system-to-component likelihoods is theoretically more complicated than component-to-system likelihood estimation as the change to a system does not necessarily mean a change to all of its components. As a result, in addition to the probability that a component within a system may affect another component in a product, it is important to account for the manner in which a change within a system could be initiated.

The probability that a change to a system will affect a component outside of the system is dependent on the chances of one or more components changing within the system. However, allocating probabilistic values to possibilities of initiating changes within a system introduces a chance of a non-event. In other words, there is a mathematical theoretical probability that a system may change without any of its components actually changing. Such a scenario would not conform to reality, since the essence of a system is brought about by the existence of interacting components. The approach used to compute both component-to-component and component-to-system change propagation likelihoods could result in such non-event. In order to avoid being caught in such paradox, the algorithm assumes that any one component can cause change to propagate in a system as illustrated in Figure 6 and Equation 5.

The system-to-component likelihood is derived using the numerical average of each component $\mathbf{a_i}$ in **A** initiating a change to a component $\mathbf{b_j}$ in **B**. This likelihood was given in Equation 3.

$$likelihood(A\ to\ b_j) = \frac{1}{n}\sum_{i=1}^{n} likelihood\left(a_i\ to\ b_j\right) \tag{5}$$

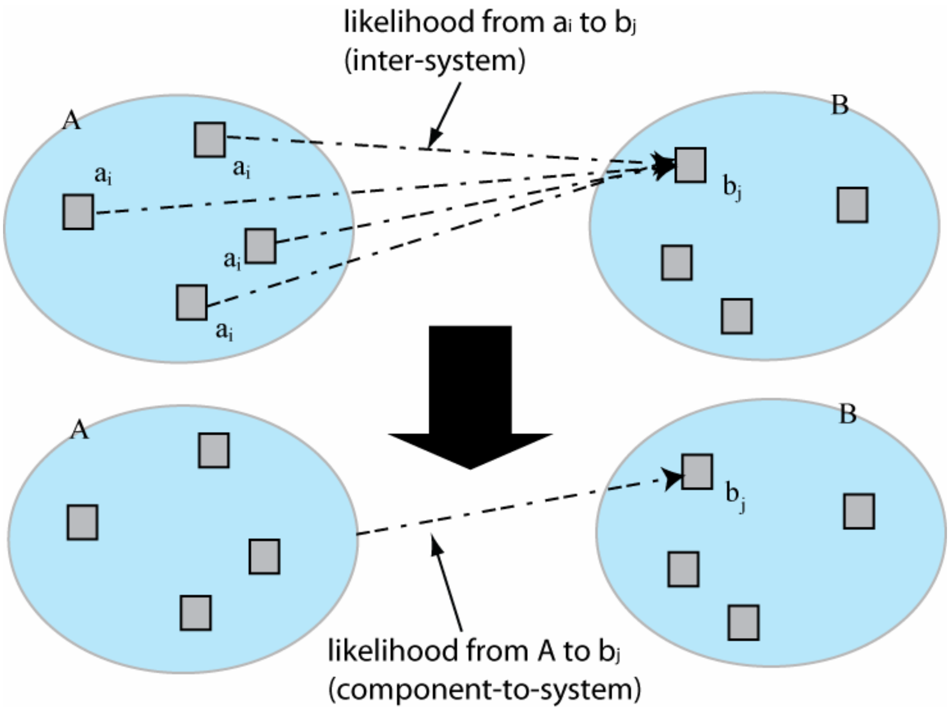

Figure 6: Estimating likelihood of change propagating from a system to a component

4.3 System-to-system risk assessment

The system-to-system likelihood estimation follows a similar approach to the system-to-component likelihood derivation. It is essentially the average of all component-to-system likelihood values which were established in Equation 4 (see Equation 6).

$$likelihood(A\ to\ B) = \frac{1}{n}\sum_{i=1}^{n} likelihood\left(a_i\ to\ B\right) \tag{6}$$

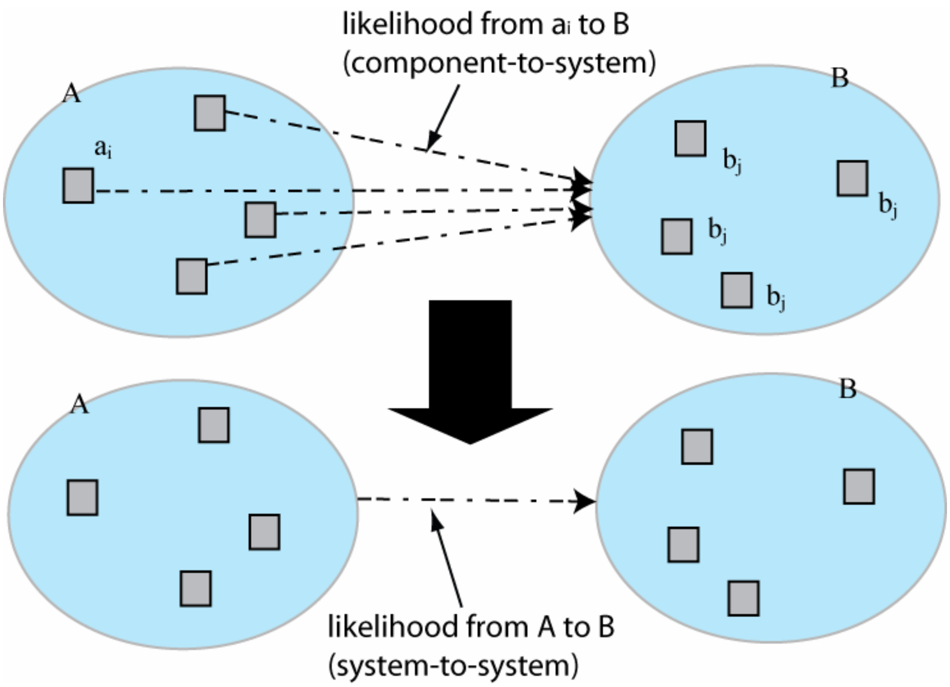

Figure 7: Estimating likelihood of change propagating between two systems

4.4 Single component systems

There may arise a situation where a system has a single component. This may arise as a result of a component not belonging to any particular system per se. In practical terms, such a system is just a duplication of the lower component likelihood and risk estimates.

5 PREDICTING CHANGES AT DIFFERENT LEVELS OF GRANULARITY

The approach to estimating likelihood of change propagating described in this paper was used to assess the effects of change within a diesel engine. We adopted a connectivity model built in a leading UK automobile company [10]. At the level of granularity chosen for the model, the diesel engine consists of 41 components. The model contained expert estimates of the likelihood that change propagates between component pairs. These components were grouped into systems in order to create another level of granularity to assess change effects. The process is summarised in the diagram in Figure 8.

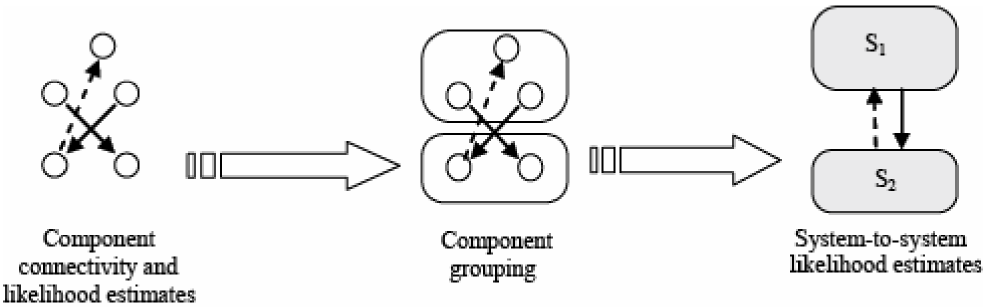

Figure 8: Estimating change likelihood at different levels of granularity

5.1 Product decomposition

For this study, only one additional level of granularity was developed. However it is important to note the method can be repeated for descriptions at several levels whilst following the bottom-up likelihood estimation described in Section 4. Attaining a suitable high level product description from the existing model is not an entirely straightforward process and does require some negotiation. In this study, the product was decomposed following a similar approach to hierarchical product decomposition used in [11].It was important that the product was decomposed to levels which were useful for supporting queries during change impact assessments. As a result, care was taken to consider the "flatness" of the product decomposition. The relation between system and component levels can easily result in a structure that is too narrow or too broad. For example, a high level description such as "engine" will not be practical for the intended purpose as it constitutes all 41 components. As such, any product breakdown which takes the form of a broad and shallow tree structure in principle, defies the aim of structuring change at various granularity levels. Similarly, an excessively narrow structure may also be inadequate for the intended purpose.

Generally, components can be grouped in various ways depending on the context. Each particular context of grouping may also be referred to as a "perspective" [12]. A set of components can be grouped into assemblies, modules, systems and so on. The choice of a specific perspective is dependent on the relation an observer is representing. There is a strong interdependence between the purpose and the perspective. For the purpose of our study, components were grouped into systems. All 41 components were grouped into 10 systems. In order to distinguish between intra-system and inter-system connectivity, only the main function of the component was considered when assigning components to systems.

5.2 Change impact assessment

Once a high-level decomposition for the diesel engine was agreed upon, likelihoods of change propagating from component-to-component, component-to-system, system-to-component and system-to-system were estimated based on values from the original model. The outputs were recorded using a Design Structure Matrix (DSM) since it is a compact way of structuring pair wise data. In order that the estimated values were easily accessible, information visualisation software techniques described in [7] were used to restructure the matrix. This enabled quick assessments of system-to-component, component-to-system and system-to-system likelihood estimates.

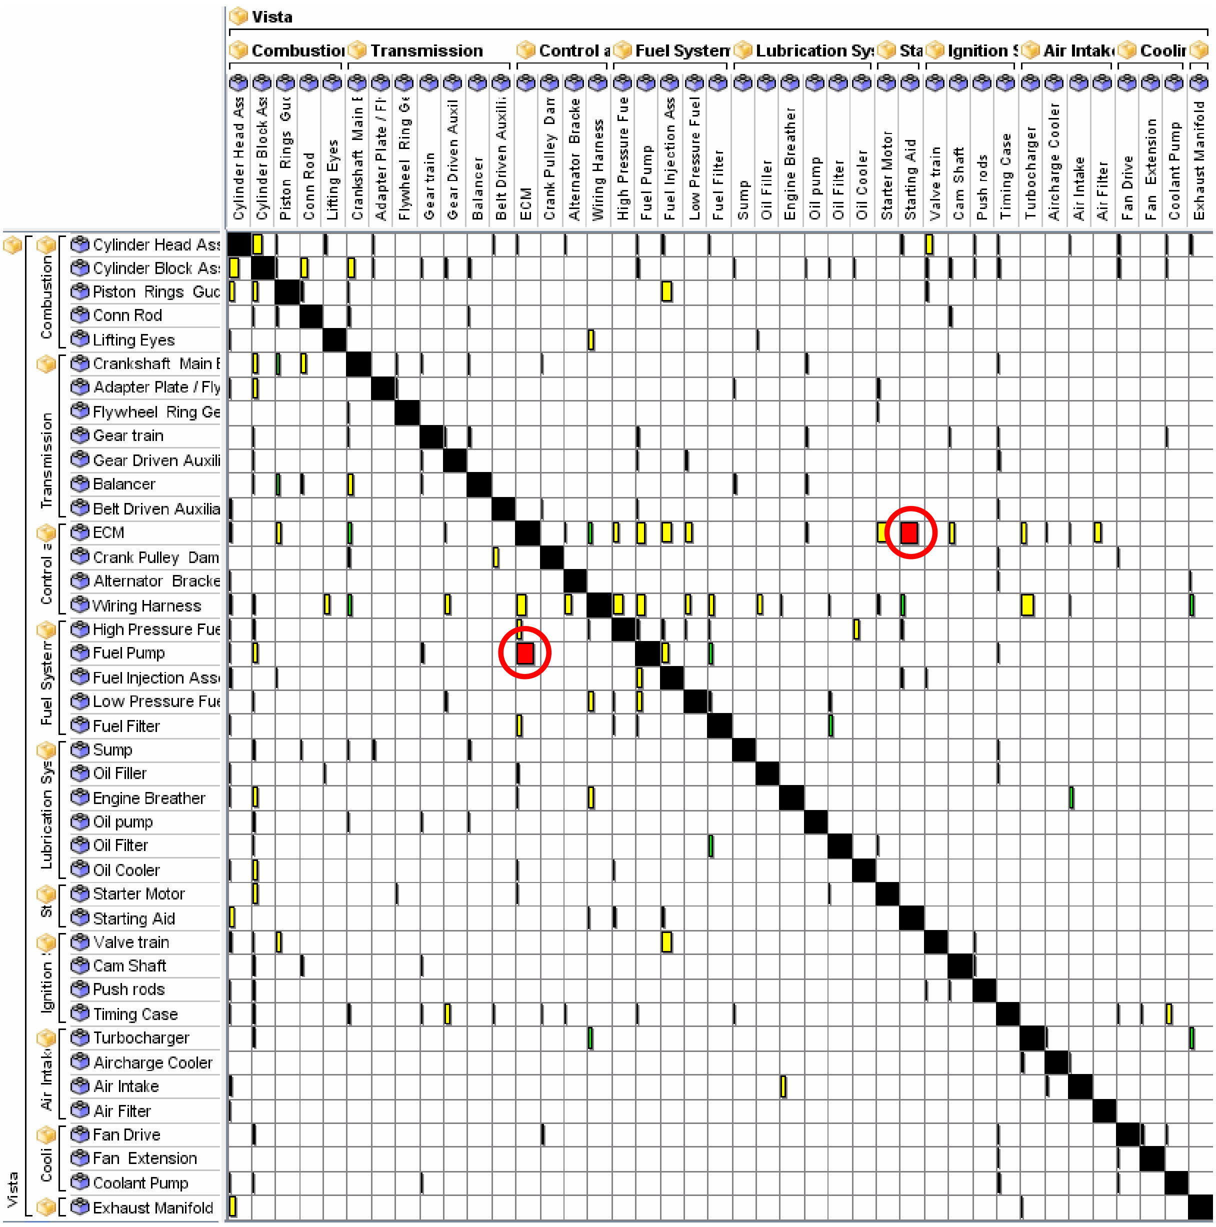

Figure 9: Component level model of the Diesel Engine

See Figure 9 for the original component-to-component likelihood matrix. In this diagram, the likelihoods of changes propagating between components are colour coded. High-likelihoods are shown in red, medium ones are represented with yellow, and low-likelihood are highlighted as green. One can see that this model has a high level of detail on a component level, but it is difficult to assess the effects of changes on a system level.

In Figure 10, the resulting system-to-system likelihoods are given based on the computations described in Section 4. The component-to-system allocation used for this system likelihood assessment is shown in Figure 9 by the entries on the axes. The system that is most likely causing other components to change is the *Combustion System* (highlighted column), while the *Control and Electricals System* is most likely to be affected by other system changes (highlighted row). The highest likelihood found in the system-system matrix is from the *Starter System* to the *Control and Electricals System*, which is mirrored by a high component-to-component interaction of the *Starting Aid* affecting *ECM* (Figure 9). This shows consistency in the results.

Further confirmation of this consistency can be seen in examples of other high-likelihood interactions on the component level that are also reflected on the system matrix. For example, consider the high component interaction from the *ECM* to the *Fuel Pump*. While at the component level, there is a high-likelihood of change propagating between these two components, a medium-likelihood value is estimated on the system level i.e. *Control and Electricals System* to the *Fuel System*. The reason for the reduced likelihood is the majority of the other components within the *Control and Electricals System* have a very low likelihood of causing change to propagate to the *Fuel System* e.g. *Wiring Harness* to *Fuel Filter*. As such the system level description represents the likelihood values attributed to changes propagating between both systems.

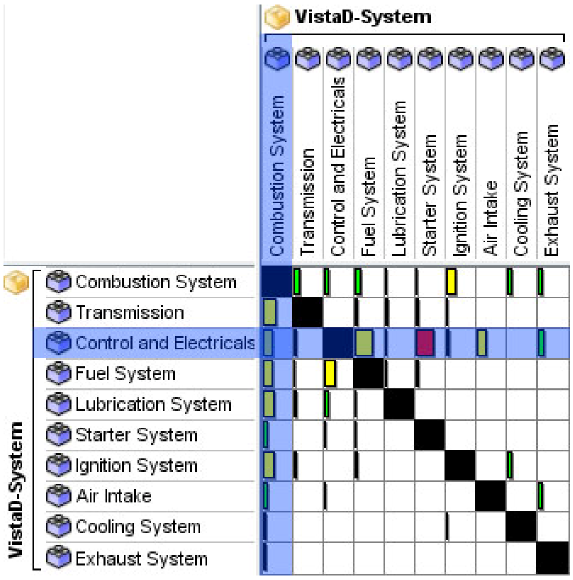

Figure 10: System level model of the diesel engine

5.3 Validation

In the previous section, it was shown that high-risk component connections are mirrored by high-risk system connections. However, the likelihood values obtained through the computations described in Section 4 still have to be validated.

One way to validate these values is to use values computed by the CPM algorithm that was already used to compute intra-system relations in Section 4 on the whole product and compare them to results obtained from this method. Given the component-to-component likelihood values and using simulations, it is possible to compute combined system-to-system likelihood values. These simulations essentially go through all propagation paths and count how often components that are part of a particular system are affected by an initiating change. Additionally, using the algorithmic approach from Section 4 and subsequently the deterministic CPM algorithm on the system-to-system likelihood values also gives combined system-to-system likelihood values. Both approaches should result in similar values for the combined system-to-system likelihood (Figure 11a).

For the diesel engine model, these analyses were carried out and the results can be seen in Figure 11b. The x-axis of this scatter plot represents the combined likelihood values obtained from the simulation CPM approach and the values on the y-axis correspond to the use of the algorithm described in this paper in combination with the CPM algorithm on a system level. One can see that there is a high correlation between the two variables validating the results obtained for the system-to-system likelihood. The differences in the values result from the fact that simulations only produce approximate results.

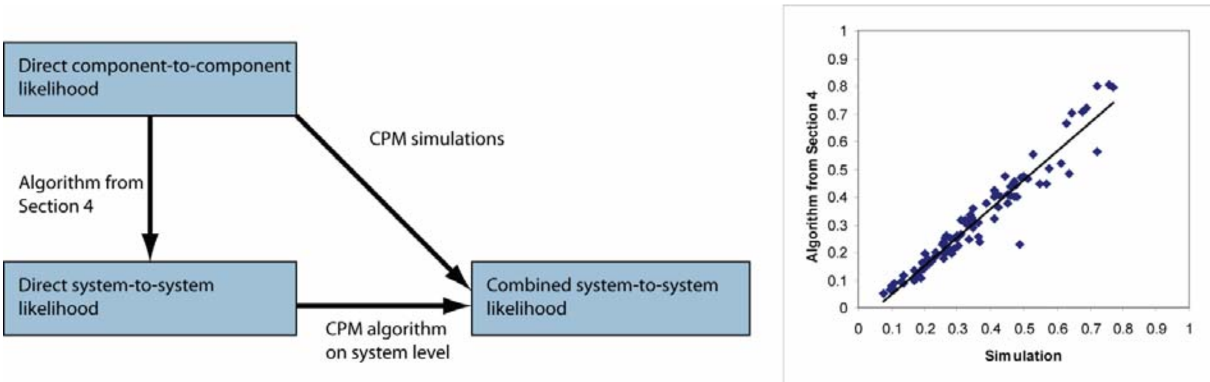

(a) (b)

Figure 11: (a) Validation via the use of the CPM algorithm and (b) comparing combined system-to-system likelihood values

6 SUMMARY AND CONCLUSIONS

Models of product connectivity play an important role in reasoning about engineering change. To this end, we have provided an approach for modelling connectivity on different levels of granularity. Hierarchical product decompositions (without any change propagation likelihood estimates) have been shown to be useful for supporting practical design queries sometimes by simply facilitating communications between design teams [13]. However, the techniques for propagation likelihood assessment discussed in this paper take the level of support a step further to include considerations across hierarchy levels.

Generally, likelihood estimation at various levels of granularity is prone to all sorts of bias arising from human judgement and experience [14]. By simply allocating components into groups, the algorithm described in this paper enables consistent likelihood estimation across several levels of a hierarchy. This approach not only reduces the effort required for building hierarchical models, it also reduces the chance of error being introduced into the model once the initial component DSM has been created.

Being able to assess multilevel likelihood values can give valuable insights into the product structure. This paper introduced an algorithmic approach to compute such multilevel likelihood values for the prediction of change propagation based on a bottom-up approach. It was shown how component-to-component change probabilities can be translated into coherent system-to-system, system-to-component and component-to-system relations allowing the assessment of change likelihood probabilities on all levels of abstraction.

Although this paper focused on the application of such an algorithmic approach in change propagation likelihoods, the approach can be applied in many other disciplines, which have underlying probabilistic networks such as epidemiology. The algorithm presented in this paper would allow the computation of the likelihood that one particular group of people would infect another group given that individual infection likelihood values are known.

REFERENCES

[1] Clarkson P.J., Simons C.S. and Eckert C.M. Predicting Change Propagation in Complex Design. *Journal of Mechanical Design*, 2004, 126(5), 788-797.

[2] Zeitz C.M. Some Concrete Advantages of Abstraction: How Experts' Representations Facilitate Reasoning. In Feltovich, P.J., Ford, K.M. and Hoffman, R.R., eds. *Expertise in Context: Human and Machine.*, pp. 43-65 1997 (AAAI Press, California).

[3] Grene M. Hierarchy: one word, how many concepts? In Whyte, L.L., Wilson, A.G. and Wilson, D., eds. *Hierarchical Structures*, pp. 56-58 1969 (American Elsevier, New York).

[4] Simon H.A. *The Sciences of the Artificial*, 1981 (MIT Press, Cambridge, MA).

[5] Shlichta P.J. Overlap in Hierarchical Structures. In Whyte, L.L., Wilson, A.G. and Wilson, D., eds. *Hierarchical Structures*, pp. 138-143 1969 (American Elsevier, New York).

[6] Miller G.A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 1956, 63, 81-97.

[7] Furnas G.W. Generalized fisheye views. In *Human Factors in Computing Systems Conference (CHI '86)*, Boston, Massachusetts, USA, 1986, pp.16-23.

[8] Skjong R. and Wentworth B.H. Expert Judgment and Risk Perception. In *ISOPE-2001: Eleventh (2001) International Offshore and Polar Engineering Conference,* Stavenger, Norway, 2001.

[9] Kahneman D., Slovic P. and Tversky A. *Judgment Under Uncertainty: Heuristics and Biases,* 1982 (Cambridge Univeristy Press, Cambridge).

[10] Jarratt T.A.W., Eckert C.M. and Clarkson P.J. Development of a Product Model to Support Engineering Change Management. In *in Tools and Methods of Competitive Engineering*, Lausanne, Switzerland, 2004.

[11] Jorgensen K.A. Product Modeling on Multiple Abstraction Levels. *Mass Customization: Challenges and Solutions*, pp. 63-84 2006 (Springer, New York).

[12] Bækdal L.K. and Kristensen B.B. Perpectives and Complex Aggregates. In *in 6th International Conference on Object-Oriented Information Systems,* London, UK, 2000, (Springer).

[13] Lockledge J.C. and Salustri F.A. Design Communication Using A Variation of the Design Structure Matrix. In *13th International Conference on Engineering Design,* Glasgow, 2001, (Professional Engineering Publishing).

[14] Kahneman D. and Tversky A. On the Psychology of Prediction. *Psychological Review*, 1973, 80(4), 237-251.

Contact: O.O. Ariyo
University of Cambridge
Engineering Design Centre
Trumpington Street
Cambridge CB2 1PZ
United Kingdom
Phone: +44 1223 332758
Fax: +44 1223 766963
e-mail: ooa24@cam.ac.uk