

SIGNPOSTING : MODELLING UNCERTAINTY IN DESIGN PROCESSES

B D O'Donovan, P J Clarkson and C M Eckert

Abstract

In engineering design, there is often considerable uncertainty in knowledge about the product, for example, the lift generated by a wing might be unknown until wind tunnel tests have been performed. Equally, at the start of a design project, there is considerable uncertainty about the direction the process will take, which tools will be necessary, which tasks will be successful, which aspects of the design requirements will be difficult to meet. This process uncertainty causes difficulty in planning engineering design projects, and in managing these projects during their execution. Existing process modelling tools represent processes on a high level, as well as only a subset of the possible variations in a design project, neglecting such factors as non essential information transfer, use of alternative tools or methods, and the effect of assigning different engineers to different tasks. This paper presents the most recent iteration of the Signposting system, which has been developed in Cambridge to provide a framework for engineering design process modelling. Such a framework must meet a number of requirements which are discussed, and the performance of the Signposting framework against these metrics is assessed. The practicalities of building models in industry using this framework are discussed, and an number of model building protocols are introduced.

Keywords: process modelling, engineering process, risk management

1 Introduction

The challenges facing industry today are no longer primarily technical or technological; as the complexity of engineering products has increased, the issues surrounding the organisation of the design process are becoming increasingly vital. It is no longer sufficient to design a better product than your competitors, companies now have to bring a product to market faster and more cost effectively than their rivals [1]. A number of approaches are available to meet these new targets, including the application of new design tools, and the closer integration of existing tools. The modelling of design processes also has a role to play, providing new insight into and understanding of existing processes, and suggesting means by which they might be improved.

Modeling initiatives are already underway in a number of industries, most notably the aerospace industry, where the size and complexity of the design effort make the understanding that models can deliver particularly valuable. These initiatives are using a number of existing modelling frameworks, including IDEF0 [2], Design Structure Matrices [3], Critical Path analysis and GDM [4]. Each of these tools has a number of strengths and weaknesses, but none of them are well suited to modelling the uncertainty and unpredictability that are characteristic of the design process. There is, we believe, a need for a new tool to assist industry in modeling their processes.

2 A specification for design process modelling

The creation of a new design process modeling tool can itself be viewed as a design process. And as such, it is important to begin by defining the requirements of and constraints on such a tool.

The modelling technique described here is not intended as a general theory, or prescriptive methodology of design, such as has been presented by other authors [5]. Instead we seek to provide a framework for describing design processes, with the intent that models for specific industries or classes of design, for example a model describing the jet engine design process or automotive design, could be constructed later based on this framework. This fundamental decision about the scope of our work has been based on a number of factors. It has been observed by others that the general descriptions of design, those which might be applied directly without contextualisation for a specific industry, have broadly reached a state of convergence [6]. Furthermore, while these models are valuable in teaching engineers in universities about the abstract structure of design, they provide relatively little concrete guidance to engineers in industry, who are concerned more with the details of the specific design task that faces them.

A framework-based approach, in which a core of central ideas can be customized and contextualized to a specific task offers a good compromise between the generic and the specific. Inevitably, such a framework will embody assumptions about an underlying theory of design, but this will takes the form of the range of elements that might be described under the framework, without imposing any constraints or rules on how the elements might interact, or indeed whether they should exist at all in a particular design process.

2.1 The audience for a model

To begin to develop a specification, we should first consider the potential audience for a design process model. There are many stakeholders in the design process, the engineers, their managers, the directors of engineering companies, the company’s shareholders, the customers of these industries, and even the academic researchers carrying out case studies in industry. To each of these groups, the design process will appear different, and their corresponding concerns and interests in the process will likewise differ.

Table 1. Audiences for design process modelling

Participant	View of design process	Involvement / interest
Engineer / designer	Detailed, limited scope	<ul style="list-style-type: none"> • Execution of design work
Operational Management	Wide scope, limited detail	<ul style="list-style-type: none"> • Optimisation of design process • Deployment of resources • Scheduling of tasks • Choice of design tools / methods • Risk management • Crisis response
Strategic Management	Multiple projects	<ul style="list-style-type: none"> • Evaluation / choice of potential projects • Process cost / timing • Investment in training / new tools
Academia	Dependent on research	<ul style="list-style-type: none"> • Forming / testing hypotheses • Communicating research to industry

The audience for which a model is designed will determine both the aspects of the design process which are of interest, and also the intervention in the design process which we may want the tool to support. Clearly, given the range of interests and views of the design process, no single model can accommodate all the possible audiences. However, a single framework might allow a design process model to be customized for a given audience, just as it could allow customization for a specific industry. The core of the process model, the task breakdown, remains constant for all audiences, but the attached information and the detail represented will vary. Building a model to address the concerns of a new audience within a process will consist of a mixture of appropriation of elements from previous models of the process (providing a hook for integration), and the capture of new, audience-specific information.

The primary audience for this work will be the operational management and strategic management of engineering companies, as it is at this level that much of the existing modelling effort is concentrated, but the framework is also compatible with the specific interests of non-managing engineers and academia.

2.2 Requirements for a model

A generic framework for modelling the development process must be appropriate to a wide audience and range of applications (planning, process management, designer guidance, etc.) within the process as discussed, but there are other underlying requirements that must also be met:-

- Scalability – Design processes in industry vary between the simple, e.g. design of a fastener, and the complex, e.g. design of communications satellites or passenger jets. A general framework must be compatible with both scales of project, and useful at each. This creates a need for hierarchical groupability of model elements.
- Practicality – For the model to become widely used it must be possible for industry to build models of processes without direct support from researchers.
- Maintainability – As design processes change dynamically, it is important that a model can be updated by the user with a minimum amount of effort. Local changes to a process should require only local changes to a model. It is not uncommon for academic methods introduced to industry to be dropped due to the cost of maintenance.
- Robustness and trustworthiness – The quality of data available on past design processes is limited, both in the quality of the information available, and also in quantity. In design process modelling, unlike other fields of simulation, a model may need to be constructed based on only one or two similar precedent projects (or no precedents) rather than on a representative sample of the total population. Because of the limited quality of the information, the model must be robust enough that any guidance derived from it can be trusted.

2.3 Defining an approach

A major differentiator between the available design process models is the degree of emphasis placed on representation and interaction. An analogy would be that a drawing of an engineering component is primarily representative, whereas a parametric 3D CAD description of the same component has a more interactive nature. Both CAD model and drawing contain geometric information, but a distinction can be made based on the completeness of information, and also the possibility of variation – the drawing is a static view, while a parametric CAD description defines not only the component *as is* but also *as it could be*. A

similar metric could be applied to a design process model – it should not only describe the process as it is, but also capture the potential processes that could be. On this metric, many existing modelling techniques have limited potential, for example the input into an IDEF task could be modified, but the effect of the change is not suggested by the model. We can identify a number of process-related factors that can cause a new process to vary from the precedent process that provided the template for a model.

- Varying task sequencing and scheduling
- Variation in duration of task
- Variation in task outcome; degrees of success, failure modes
- Choice of alternative tasks
- Effects from different resourcing of, and information input into tasks

Existing methods support a subset of these factors, but no previous methodology offers the potential to model all of the above issues. In neglecting some factors, existing models ignore potential process outcomes, reducing the accuracy of predictions that can be made, and constraining the solution when the models are used to optimise the process.

A key feature of engineering design processes is the uncertainty in the outcome of the design activity. The final design depends on the ideas that have occurred to the engineers, which in turn might depend on individual knowledge, understanding, preferences, chance events providing insight or inspiration and so forth. As such, any model that treats the design process as purely deterministic constrains itself to a narrow range of design, in which it has been possible to reduce design to a mathematical problem to solve. A stochastic model accounts better for the impossibility of capturing all of the cause-effect relationships involved in the design process – a number of factors can be subsumed under a probability distribution of outcomes.

The decision to focus in this work on task-based modelling rather than product-based modelling is based on a number of factors. Because the product is abstracted in task-based modelling, a wider range of potential products can be represented – as fewer assumptions about product architecture and the relationships between elements of the design need to be made, a model based on a historic design process may be more applicable to a future design. Furthermore, many tasks in the design process are not directly related to the product being designed, e.g. market research, production of reports and documentation – these activities would not be captured in product-based modelling, so process risk assessments based on a product model may be incomplete.

3 The evolution of the signposting model

Signposting is the name of the design process modelling approach that has been developed over the past six years in Cambridge. The current system builds on the work carried out in three previous research projects.

The original implementation of signposting was created by J. Hamilton [7] in response to attempts to capture the rotor design process at Westland for the purpose of training newly recruited engineers. At this point the model was simply intended to indicate a ‘next step’ in the process, given the current state.

In later work, J. Jarrett [8] developed a tool for conceptual design in jet engines, marrying the signposting task-selection system to a range of design tools used in industry. This application is powerful in that the signposting system actually creates a design, rather than simply describing the process by which a design might be created, but the investment of time in creating such a model is significantly greater.

The work of A.F.Melo [9] offered signposting to a different audience, engineering management. Rather than analysing the only ‘next step’ to be performed in a design process, a Markov chain model was used to look at the sequence of all steps necessary to complete a design from a given start point. These different routes to completing a design could be compared, and the route with the lowest cost selected and used as the basis for a process plan.

4 Elements of the model

There are three core elements to the signposting model, tasks, states and ‘signposting parameters’. Most of the complexity and depth of modelling that signposting offers can be obtained from these basic elements, although some extensions such as numerical or decision-based design require additions to the basic model. To illustrate the model we will take the example of a hypothetical toaster-design process.

4.1 Signposting Parameters

Signposting parameters, henceforth referred to simply as parameters, are distinct from the conventional definition of a parameter in design. In Signposting a parameter is any element of the design process that might be considered an input to or an output from a task. This includes parameters in the conventional sense, such as geometry or electronic component values, but also includes less conventional elements, such as manpower, or learning about the design. This is a significant departure from earlier versions of signposting, in which ‘parameter’ was used primarily in the conventional sense.

Table 1. Classes of Signposting Parameter

Parameter Class	Qualifier	Example
Concept and manufacturing description	Maturity / level of definition	Toaster geometry Sketch → Working CAD → Final shape
Behaviour / performance	Fit to requirement	Toaster safety (fit) Unsafe → Marginally safe → Safe
	Accuracy of knowledge	Toaster safety (accuracy) Unknown → Predicted → Tested
Resources	Number available / quality	Product designer Available ↔ not available
Requirements	Maturity / level of definition	Customer aesthetic requirements Unknown → Surveyed → Provisionary → Frozen
Learning / understanding	Quality / quantity	Knowledge of safety issue No knowledge → Aware → Solution known

In Signposting, each parameter has both an identifier and a qualifier. The identifier is simply the name of the element, such as “toaster geometry” or “toaster safety”. The exact meaning of the qualifier varies with the parameter type, but generally is an abstract representation of the

quality or maturity of the parameter. The actual value of the parameter is not used in Signposting.

For example, if we consider the case of a parameter representing the toaster length, at some point in the process it could have a value of 200mm, but this length is specific to a single toaster design, and what is more important is the role of that information in the overall design process. So, in Signposting what we would need to know is whether that 200mm was an initial estimate, a working value, or the final manufactured dimension. The estimate, working value, and final value, correspond to three distinct levels of the parameter's qualifier. In earlier iterations of the Signposting framework, three levels of qualifier were used, high, medium and low. More generally, a parameter can have any number of qualifier levels, as many as are needed. To reflect this, in the current implementations of the model, numerical qualifier values are used, with zero indicating the absence of a parameter, and the positive integers representing progressively higher levels of maturity and quality.

The parameter qualifier is not simply a reflection of the form of an element of data (sketch, CAD, prototype), but may also reflect the work done and tests performed on, and increasing designer confidence of a single data form. Hence, a CAD model might have a higher associated parameter qualifier value after a finite element analysis has been performed on it, even if the actual CAD data hasn't changed.

4.2 States

The complete set of all parameters and their associated qualifier values at any instant in time defines a process *state*. This embodies the net result of the work that has been completed, and in conjunction with the set of tasks involved in the design process, implies what further work will be required to complete the design.

Two special states can be defined, a start state and an end state. These define the limits of the design process, not in terms of the actual design values, but in terms of the context of the information. So, while the layout of the control circuit of the finished toaster design is unknown at the start of the process, we do know that by the end of the process the design will specify component values, manufacturing instructions, etc, which can be represented by a high parameter qualifier value. Similarly we can use a starting state to define in terms of parameter qualifiers what is known or previously completed at the start of a process.

4.3 Tasks

A typical Signposting task is shown in figure 1. This is composed of a number of elements

One or more *mappings* – A mapping is the transfer from a valid input state for the task, to the potential output states that could occur. Each output state has a probability of occurring, and the sum of all probabilities for the output states attached to each input state is one.

One or more *input states* – Each input state expresses the minimum levels of the qualifier for each parameter necessary for the task to be executed.

One or more *output states* – These define the new state of the parameters and qualifiers after the task is completed. The output state either defines a new value for the parameter qualifier, or a 'no change' value if the parameter is not involved

Attached information – Each mapping, or the overall task can have secondary information attached to it, for use by specific analysis tools. This can include, but is not limited to task cost, duration, pointers to documentation or named personnel involved in the task.

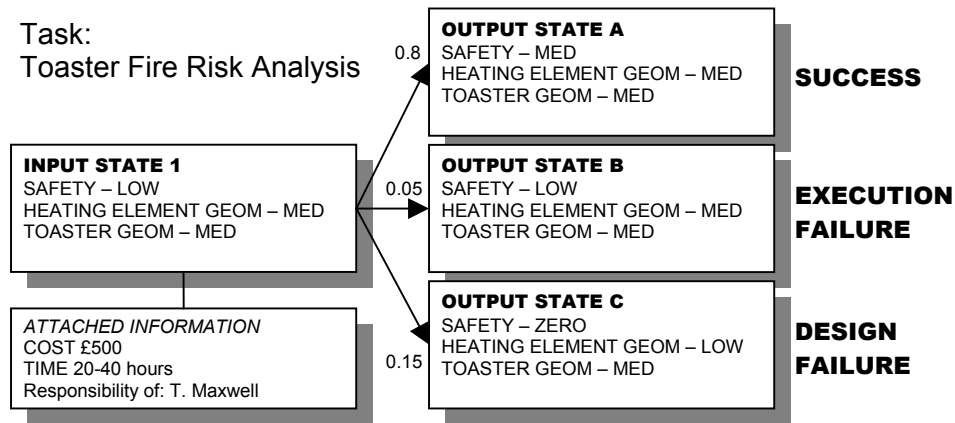


Figure 1. A typical signposting task

In Signposting, the effect of different resourcing inputs to a task may be modelled. This could include multiple engineers being assigned to the task, or the distinction between a novice and an experienced engineer. Each significant resourcing level can be expressed as an alternative input state for a task, with differences in time, cost, probability of success and even the types of outcomes possible (figure 2).

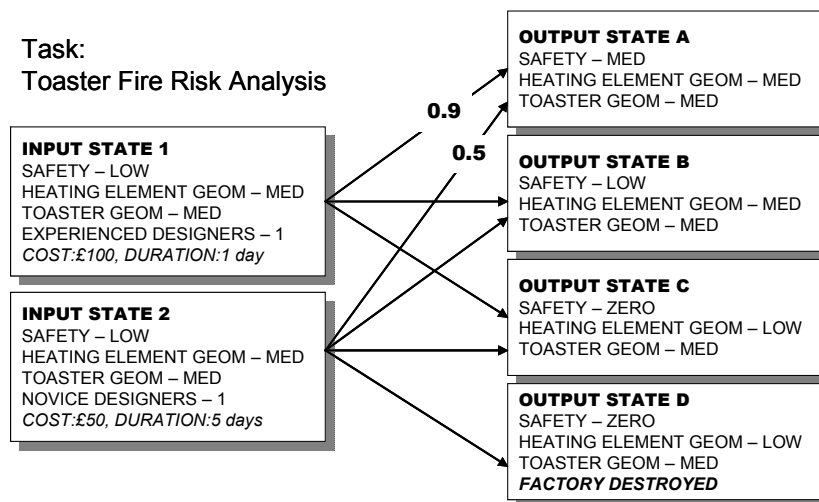


Figure 2. Modelling resourcing effects

Previous versions of Signposting did not model learning during the process, with the result that the distribution of outcomes had a characteristic exponential decay curve, with the shortest process duration the most likely, rather than there being a central average with a smaller number of faster or slower processes. This was not a good match for observed process in industry, where design iterations are normally used to learn about possible solutions and to converge towards a final design. This can be modelled in tasks where learning is significant by generating increases in learning parameters in the case of design failure, which are fed back into the task, reducing time or cost, or increasing the likelihood of success (figure 3).

5 Model Building

While signposting models do not require the level of detail necessary for artificial intelligence based design systems, not requiring capture of numerical relationships internal to the design,

or design rationale, the model is nonetheless complex, and building a specific model is a significant undertaking.

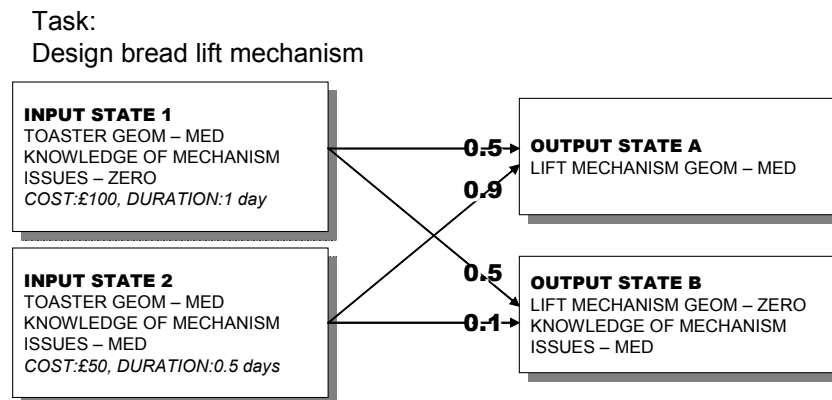


Figure 3. Modelling task learning

Although a researcher might become familiar with building Signposting models, for the approach to be widely accepted in its industry, it should be possible to create models with little previous experience. This requires the formalisation of the model building protocols that we have used in industry. Because the flexible nature of Signposting doesn't force model builders to incorporate all possible aspects of the design process into their model, a checklist is useful to ensure that the full range of potential factors are considered:

- What is a useful task breakdown?
- What is the minimum set of inputs for the task to be possible?
- What additional information / higher quality inputs will improve the execution of the task?
- If the task needs to be repeated, will learning from previous execution affect the outcome?
- Who can potentially carry out this task, and are there differences in ability between them?
- What are the possible failure modes, their likelihoods, and their outcomes?

Signposting models can be constructed either piecewise, eliciting individual tasks from the process participants in interviews; or as a group exercise with all, or a representative sample of the designers and managers.

In theory, the Signposting model software can assemble an entire process from its component tasks. However, this is complicated by a number of factors. One issue is the potential for disagreement over terminology, with different names being assigned to the same piece of information by different participants in the process. This problem has been observed in large scale model building exercises in industry [10]. A number of approaches can help with this problem; a common list of information terms could be assembled in a group meeting at the beginning of the model-building exercise; participants could be asked to whom the information comes from and who they send the information to. Another issue is that participants involved in the process at the point of execution might not be able to provide insights into the possible consequences of someone else carrying out their task. Finally, the model builder will need to be careful to maintain a common level of detail across all of the task decompositions, ensuring that the model is not compromised by a single interviewee providing a finer breakdown than his or her colleagues.

6 Applications of the signposting model

The signposting model, supported by the analysis tools created to serve it, can be used in a wide variety of ways:

- Process Capture – Recording existing processes for communication, increasing understanding, or training of new designers;
- Process Evaluation – Estimating cost or duration for a proposed design project;
- Planning – Optimising design processes, responding to unexpected events during a project;
- Design execution – either providing guidance to designers, or integrated with tools.

7 The future of the signposting system

The current belief is that no single model can serve all aspects of design process modelling, but there are compelling reasons for both industry and academia why convergence on a single or small number of models is appealing. For industry, a single model reduces costs of acquisition and ownership – each new model requires an investment in staff training, model building and possibly software purchase. For academia, the incompatibility between case studies built around different modelling frameworks reduces the potential for knowledge transfer between researchers. This is a concern in the field of design research where data collection is time consuming, and variation between projects, companies and industries can be large.

It is difficult for a single model to serve every interest well for a number of reasons; a monolithic model requires a larger time investment in model building; it is difficult to anticipate every requirement that every possible audience might have; assumptions made in the model may preclude some applications

To address this, allowing Signposting to act as a single model serving a wide range of applications, three layers of flexibility exist:

- Model flexibility – As different models are built using the framework, model-builders are free to invest more or less time in modelling specific aspects of their process. The industry independent nature of the framework means that application in the automotive industry is not constrained by the need to serve the electronics or software industries
- Framework flexibility – The core elements are common to all versions of the model, but model builders are free to include or ignore other elements of the modelling framework to suit the desired application. For example, a company interested in improving communication in its design processes might choose to tag tasks with the name of the person responsible, while someone interested in evaluating the strategic value of a project might attach cost and duration information to the tasks. This ensures that time is not wasted in building parts of the model that are unnecessary. Furthermore, the flexibility remains after a model has been built, so new interests can be accommodated by extending a model, rather than starting from scratch.
- Analysis flexibility – A range of tools have been developed for analysis of the model, each providing different insights intended for different participants in the design process.

A UML-based data structure for a central model repository has been developed. Between the user and the database, and between the database and the analysis modules, interpretation layers will format and process the model data. For example, when the user calls an analysis routine, the interpretation layer will check the data content of the model against the requirements of the specific analysis tool and prompt the user to extend the model if necessary.

Acknowledgements

The authors wish to thank the UK Engineering and Physical Sciences Research Council (grant GR/R64100/01) for funding this research.

References

- [1] DTI, “Successful Product Development”, Department of Trade and Industry, UK, 1992.
- [2] Godwin A.N., Gleeson J.W., et al., “An assessment of the IDEF notations as descriptive tools”, Information Systems Vol.14, No. 1, 1989, pp.13-28.
- [3] Browning T.R., “Applying the Design Structure Matrix to System Decomposition and Integration Problems: A Review and New Directions”, IEEE Transactions on Engineering Management, Vol.48, No. 3, 2001, pp.292-306.
- [4] van Langen P.H.G., “The anatomy of design : Foundations, Models and Applications”, SIKS Dissertation Series, Amsterdam, 2002.
- [5] Pahl G. and Beitz W., “Engineering Design” Springer, London, 1996.
- [6] Maffin D., Alderman N., Braiden P., Hills B. and Thwaites A., “Company Classification: A New Perspective on Modelling the Engineering Design and Product Development Process”, Journal of Engineering Design, Vol. 6, No. 4, 1995, pp.275-289.
- [7] Hamilton J.R., Clarkson P.J., et al., “The modelling of design knowledge for computer supported aerospace design”, Proceedings of ICED ‘97, Tampere, 1997.
- [8] Jarrett J.P., Clarkson P.J. et al. (2002). “Accelerating turbomachinery design”, ASME International Gas Turbine Institute TURBO EXPO 2002, Amsterdam, Netherlands.
- [9] Melo A.F. and Clarkson P.J., “Design process planning using a state-action model”, Proceedings of ICED ‘01, Glasgow, 2001.
- [10] Guivarch A.D., “Resolving timing issues using DSMs: Application at Ford Motor Company”, Proceedings of 2002 MIT DSM Workshop, Boston, 2002.

Brendan D. O’Donovan
Engineering Design Centre
University of Cambridge
Trumpington Street
Cambridge CB2 1PZ
United Kingdom
Tel. Int +44 1223 332 673
Fax. Int +44 1223 332 662
E-mail: bdo20@eng.cam.ac.uk
URL: <http://www-edc.eng.cam.ac.uk/people/bdo20.html>