INTERNATIONAL CONFERENCE ON ENGINEERING DESIGN, ICED15
27-30 JULY 2015, POLITECNICO DI MILANO, ITALY

USING BALANCE VARIABLES TO DESCRIBE SYSTEM INTERFACES AND ASSESS IN-PROGRESS DESIGNS

Salustri, Filippo Arnaldo (1); Rogers, Damian (2)
1: Ryerson University, Canada; 2: Ionada Incorporated, Canada

Abstract

To balance effectiveness and efficiency in design, a method is needed that can capture the range of expected behaviours (effectiveness) as well as their quality (efficiency) of artefacts existing in specific situations.

The authors introduce "balance variables" for this purpose. A balance variable (BV) is a pair of triplets, each of the form (minimum, nominal, maximum) that represent the range of values that a flow between two system interfaces can reasonably obtain. One triplet quantifies a flow from outside the object system and represents a requirement on the object system; the other triplet represents the flow handled by the object system and represents a measure of performance. BVs can capture cases of over- and under-design. A BV that is neither over- nor under-designed is said to be balanced. Measures of balance are normalized and non-dimensional, thereby permitting (a) balance measures for difference system flows to be reasonably compared, and (b) overall design balance measures to be trivially calculated. Using the example of an electric kettle, all these characteristics are shown.

Keywords: Design methods, Systems engineering (SE), Product modelling

Contact:
Prof. Filippo Arnaldo Salustri
Ryerson University
Mechanical and Industrial Engineering
Canada
salustri@ryerson.ca

Please cite this paper as:
Surnames, Initials: *Title of paper*. In: Proceedings of the 20th International Conference on Engineering Design (ICED15), Vol. nn: Title of Volume, Milan, Italy, 27.-30.07.2015

1 INTRODUCTION

A key purpose of design research is to facilitate the practice of good, successful design. Over time, this has led to the identification of many specific research questions. One of these questions is: How can one balance effectiveness (i.e., robustness, resilience, etc.) and efficiency during design so that the end product is (near) optimal with respect to both? Effectiveness treats the *range* of product responses, while efficiency treats the *quality* of those responses.

Key to research in this area is the development of methods to handle *ranges* of possible product responses (and not just user- or designer-preferred single values) to external forces, including but not limited to the physical forces typically handled in engineering. Existent work has focused on describing imprecision through probabilities, fuzzy logic, or interval methods – see (Wood and Antonsson, 1989; Antonsson and Otto, 1995; Habib and Ward, 1990). More recently, Bayesian logic has been used to model imprecision, such as reported in (Beer et al, 2013). In all the literature reviewed by the authors, the emphasis was on using developing engineering methods that conformed to existing mathematical structures, some of which are cognitively burdensome (e.g., fuzzy logic). We also noted that the methods we found in the literature make assumptions of significant background information being available; for instance, while fuzzy logic can be a useful tool, it is only as valid as the mapping between actual "real-life" values and their fuzzy equivalents – something which is not commonly known for the early stages of many types of design problems. The authors sought instead to begin with accepted norms from design practice and then seek the simplest (both cognitively and computationally) possible means to represent those norms with as much formal rigour as possible.

Managing ranges of values is especially important in early design, when relatively little is known about the design and the problem it ought to solve. Generally, one assumes quite large possible ranges of key performance values at the outset of a process, then progressively narrows the ranges as designing proceeds until those values are sufficiently "tight" to allow manufacture. Ceteris paribus, the tighter the values obtained, the more efficient (but less effective) the design; the looser the values obtained, the more effective (but less efficient) the design.

Furthermore, as *systems thinking* becomes more pervasive in design, it becomes more common to think of a product as a system hierarchy within a larger contextual system hierarchy, elements of which encapsulate structure and operation while interacting with each other by exchanging mass, energy, and information through interfaces. The systems abstraction is particularly useful to manage complexity as designs evolve and grow during designing. Given systems thinking, effectiveness and efficiency become characteristics of system interactions.

The goal of this paper is to propose a means to manage the ranges of values at systems interfaces to support better trade-offs between effectiveness and efficiency.

The authors did not follow any particular research method to develop the tools and methods described here. Rather, the work *emerged* naturally in the course of other research carried out by the authors and colleagues on design theory and methodology. A short history of that emergence is given below.

In 2009, the author and colleagues began to develop the notion of designing as *balancing* (Salustri et al, 2009). Based on the ideas of (Simon, 1969) and (Alexander, 1964), we proposed a framework wherein designing is an activity of balancing *forces* - physical, psychological, cultural, social, environmental, financial, etc. - in specific situations, that are not in balance with the needs and desires of the agents in that situation. We take "situation" in the sense of Gero's work on situated design (Gero and Kannengiesser, 2004). We noted that multi-valued variables could be used to describe a single force and quantify the extent of any imbalances (or "misfits" in the language of Alexander). Such quantification can help inform designers on the state of designs and on avenues to pursue to improve their designs.

Within this framework, we originally considered *mathematical intervals* (e.g., as in Habib and Ward, 1990) as a good way to represent the range of possible values a variable can obtain during design. However, Rogers (2013) determined that the use of intervals led to certain counterintuitive phenomena when trying to assess overall design fitness, and proposed an alternative formulation to manage value ranges. Rogers's approach is used in our current paper.

In 2010, Salustri identified the notion of a *balance variable* (Salustri, 2010) as a pair of ranges. There, he argued that there is *always* a form-context "ensemble" (per Alexander (1964)), even in situations that immediately precede the conception of highly innovative/disruptive designs, and that describing only a network of "misfits" (per Alexander) ignored important "fits;" that is, aspects of existent situations that are worth maintaining in new or improved designs. Adopting a systems perspective, balance variables can capture both fits and misfits in design situations, and represent them with ranges. This helps one ensure that solutions intended to re-balance misfits do not also unbalance fits, such as is done in Rogers (2013). Balance variables as presented in the current paper are based on Rogers's value range method.

In this paper, the authors will describe in greater detail what balance variables are and how they can be used to represent design information.

2 TERMINOLOGY AND BASIC CONCEPTS

The following terminology is adopted in this paper (terms in italics are those being defined):

Structure is the collection of inherent properties of a system (product); typically, these include shape, size, material constitution, thermal coefficients, electrical resistances, etc. *Behaviour* is the collection of responses (outputs) of the structure to stimuli (inputs) of mass, energy, and information arriving from outside the system. *Function* is the collection of roles that the behaviours of a system play in a larger system. This is largely consistent with Gero's *situated FBS* framework (Gero and Kannengiesser, 2004).

A *system* is a collection of interacting elements (each of which may be other systems) that provide recognizable functions and is crisply distinguished from its environment (Karnopp et al, 1990). Systems are separated by boundaries that are assumed to be crisp, but are probably more accurately thought of as layers. (The authors defer the treatment of system boundaries as layers of some thickness to future work.)

System boundaries are defined as "regions" in which discontinuities exist in structural properties (per the definitions above). Examples include differences in relative motion of two mated parts, the disruption of atomic forces existing in the gap between two parts no matter how tightly mated they may be, differences in thermal conductivity between a hot filament and the air around it, etc. These regions are typically spatial, but may be temporal, or even entirely abstract. A system's boundaries encapsulate its structure and operation. Other systems outside the one can only "perceive" the one's behaviour (i.e., its responses to stimuli). Boundaries therefore mark where system interfaces meet, across which mass, energy, and information can flow. For each interface of one system that shares a boundary with another system, there is a corresponding interface in the other system. That is, a single boundary denotes where two systems abut; at that boundary will exist pairs of interfaces, one pair for each flow of mass, energy, or information.

3 AN INTRODUCTORY EXAMPLE

Consider designing an electric kettle. The kettle exists as an as-yet undesigned, encapsulated system, in a context that contains various other systems, including users, power outlets, countertops, water, etc. There are boundaries between the kettle and the context systems that permit the exchange of mass, energy, and information such that the kettle provides some functions within the situation. For each flow of mass, energy, or information across a boundary, there is an interface in the kettle system and a corresponding interface in some context system. The "fitness" of a kettle is defined as the degree to which its interfaces *fit* with their correspondents in the context systems.

Let us consider only the electric voltage from the context into the kettle during use. Conventionally, this is done by plugging the kettle into an outlet. The transfer of electricity from the outlet (a context system) to the kettle requires an interface in the kettle and another in the outlet; these interfaces must "fit" one another for the kettle to function in that situation.

While the voltage at the outlet will have a nominal value expected during regular operation, it will also have a range of possible, though not necessarily preferable, values. These will include various spikes, surges, dropouts, and sags. If the kettle cannot somehow handle these values, it will damage itself, other nearby objects, and possibly harm users. Thus, the kettle must not only use power nominally but also handle the broader range of possible voltages through these interfaces. The goal of our approach is to provide assistance in recognizing and assessing when a design is, and is not, meeting the range of possible values available via interfaces from context systems, and to ensure that such information is captured as early as possible.

The electric voltage at the outlet is therefore described by a triplet including a minimum value, a maximum value, and a nominal value between them, which we write as `(min, nom, max)`. The range denoted by the triplet represents every reasonably possible value that the voltage can obtain. The electricity expected by the kettle will also be represented as another triplet of the same form.

Together, the two triplets describe the voltage from one interface to another, across a system boundary. The triplet on the context side of the boundary (which we will hereunder call the *out-system* side) represents a requirement on the kettle's design; it describes the range of circumstances that a well-designed kettle must accommodate for that situation. The triplet on the kettle side of the boundary (hereunder, the *in-system* side) represents a measure of the kettle's (expected) performance in the situation (with respect only to electricity use). We call such a pair of triplets a *balance variable*.

4 BALANCE VARIABLES

Based on the foregoing, a balance variable (BV) is a model of the range of possible values that each of two corresponding system interfaces can obtain across a given boundary, in a given situation. It is a pair of triplets of values; each triplet is of the form `(min, nom, max)`. One triplet is specified with respect to some (out-system) context system, and the other is specified with respect to the (in-system) product system. A graphical representation of one interface (or triplet) of a BV is shown in Figure 1.

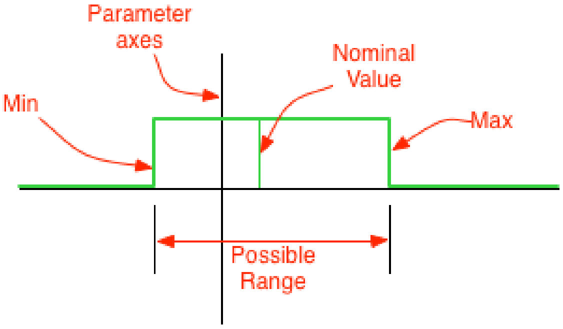

Figure 1. Graphical depiction of one side of a balance variable.

In the figure, the horizontal axis represents the value of the flow at the interface, and the vertical axis represents whether or not that system flow value is possible in the given situation. One may think of the figure as a graph of the probability distribution that the value of flow will obtain in a given case.

Figure 1 implies a uniform distribution of values, which will not happen in practice (e.g., spikes and lags of electric power do not happen as frequently as nominal power delivery). However, in the early stages of designing, it may well be that no information on actual probabilities is available; still, the designers do know the *limits of the range* of possible values. Therefore, for simplicity, we will only consider uniform distributions in this paper.

Figure 2 shows two interfaces for a single flow at a boundary between two systems. The lower side of the interface (in blue) models the in-system side; the upper side (in green) models the out-system side. (Associating the out-system with the upper interface is arbitrary and just an artefact of the visualization and has no special meaning.) The two interfaces together constitute one BV for a single flow across the boundary separating the system of interest from a context system.

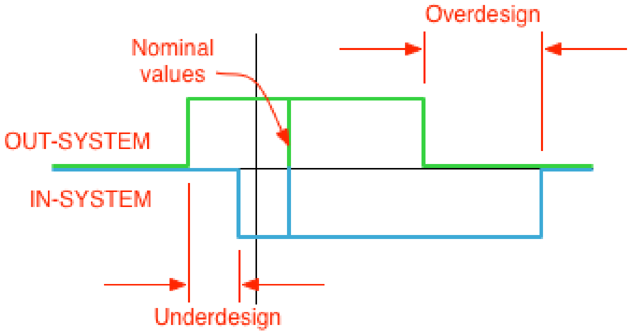

Figure 2. Graphical depiction of a whole balance variable.

The out-system side of a BV is typically beyond the designers' control. The designers of a kettle, for instance, cannot alter the power supplied from the wall outlet to accommodate their kettle design. Conversely, the specification of the in-system side of a BV is typically within the designers' control, since the design itself is (presumably) the designers' responsibility.

In Figure 2, the nominal values for both out-system and in-system interfaces correspond exactly. This signifies that the designers have set the ideal value for this parameter to correspond precisely to what is available from the context.

The minimum value for the out-system interface in Figure 2 exceeds that of the in-system interface. This means that there are possible parameter values that the context can obtain, but that the designed product will not handle. Therefore, the minimum performance of the system with respect to this parameter is *under-designed*.

Similarly, the maximum range value for the in-system interface in Figure 2 exceeds that of the out-system interface. This means there are possible parameter values that the product can handle, but that will never obtain in the context. This means that the maximum performance of the product with respect to this parameter is *over-designed*.

If the minimum, nominal, and maximum values of the two interfaces corresponded exactly, we would say that the design is *balanced* with its context, with respect to that one parameter. Any variations from this condition are called *imbalances*.

As an example, consider again the parameter of voltage across a boundary between a wall outlet and the plug of an electric kettle. Let us assume that Figure 2 is a fair representation of this situation. We can make the following statements about the design of the kettle, with respect to the particular power source(s) to which it is expected to be connected.

- The nominal value (e.g., 120 Volts) is balanced (i.e., perfectly aligned on both the kettle and the outlet sides of the boundary).
- We can expect under-voltage conditions to arise from the outlet that the kettle cannot handle.
- There are some over-voltage conditions that the kettle would be able to handle, but will never occur.
- The out-system range would be defined by the possible worst-case scenario for the Utilization Voltage typically found in houses: 87% - 106% of the nominal 120V for North America (Esser, 2011).
- The in-system side suggests that various design interventions are required, possibly implementing a variety of behaviours for each of several different sub-ranges, to manage the possible inputs from the context. For instance, the low-voltage condition in the kettle may trigger a behaviour that does not normally occur, and that behaviour may be provided by structural features (e.g., an auto-off feature or warning light), indicating to the user that the kettle might not work properly under the current conditions being supplied by the wall outlet.

In this particular example, then, we see that the kettle is both under-designed (for low-voltage conditions) and over-designed (for high-voltage conditions). If we had an actual kettle, we could attach values to the degree of over- and under-design and describe in a uniform way how "good" the kettle design is with respect to its handling of electric voltage. We would also have information useful to drive the design forward toward a more "balanced" state. These points are discussed in the following sections.

This approach supports the embedding of information about effectiveness (or robustness, per (Ullman, 2001; Taguchi et al, 2000)). The minimum and maximum values in a BV are useful to represent the range a parameter can obtain including sources of environmental variation (i.e., flows into a system from its context). In this way, one establishes at the outset of a design task the full range of a parameter, and when one then starts to optimize the design, one need not worry about making it hyper-efficient at the expense of resilience and robustness. That is, using this approach, we should be able to better optimize efficiency within the constraints of effectiveness - something we do not do well today. More generally, BVs allow separate and explicit monitoring and control over both effectiveness and efficiency, thus facilitating the development of better designs.

We note that we do not mean to imply that the over-design scenario in this example is *necessarily* bad. The over-design may actually a result of some system element over which the designers have no control. That is, the kettle over-design for over-voltage may result entirely from the use of certain components (like resistors, capacitors, transistors, etc.) that happen to be rated for 110-130V, for some unavoidable reason such as needing to lower unit costs. This leads us to consider the coupling of many BVs leading to the assessment of overall design performance. For instance, is over-design on one parameter acceptable to ensure balance on another parameter? This is discussed below.

5 ASSESSING BALANCE VARIABLES

BVs would be especially useful if there were some simple, systematic way to quickly assess how well balanced a given design is. This is possible using the method in Rogers (2013).

One BV describes a pair of interfaces. Something flows out of one system through one interface, into another system through the other interface. If one assumes a standard "stock and flow" perspective of systems, one may reasonably denote the flow out of a system as a negative value and the flow into a system as a positive value.

If we represent the range of an interface by a range $[X_{min}, X_{max}]$ (i.e., set aside the nominal value in BV triplets for now), we can use addition to determine the "difference" between the two interfaces through which a flow occurs at a system boundary. Given the two interfaces in a BV, O (output) and I (input), their sum is represented as $[-(O_{min} + I_{min}), O_{max} + I_{max}]$.

If the ranges of O and I are the same (but obviously of opposite sign per the convention above), then the sum is zero and we have a balanced BV. For the under-designed condition, the corresponding element (minimum or maximum) of the sum will be less than zero; for the over-designed condition, the sum will be greater than zero. These are both imbalances.

So far, we have considered only individual BVs. It would be useful to be able to evaluate the overall balance of a design described by a set of BVs. This is not possible using sums alone as described above, because the values of different BVs may have different units of measurement and significantly different ranges; comparisons made under these circumstances are meaningless (as would, for instance, saying that 200K is "more than" 20kg).

We can recover the ability to compare different BVs by normalizing the sum that represents the balance of a BV. We can normalize a BV by dividing the sum by some sensible quantity. Such a quantity must have the same units as the BV, thus making the normalized sum non-dimensional. If chosen appropriately, the value of this divisor would also normalize the magnitude of the sum. Finally, we should choose this normalizing divisor such that its value can be expected to remain constant through most, if not all, of the design process, so as not to artificially or unnecessarily change the assessment of a design's overall balance.

The authors first considered using the nominal value of the out-system interface of a BV as the normalizing factor. However, Rogers (2013) discovered that this led to counter-intuitive values of overall design assessment; i.e., numerically equal differences between out-system and in-system values resulted in numerically *unequal* measures of balance. Rogers's solution is to use the out-system value of the extremum being considered. That is, the measure of balance, b, of a BV is a triplet $[b_{min}, b_{nom}, b_{max}]$ such that:

$$b_{min} = -(O_{min} + I_{min})/|O_{min}|,$$

$$b_{nom} = |(O_{nom} + I_{nom}) / O_{nom}|,$$

$$b_{max} = (O_{max} + I_{max})/|O_{max}|.$$

The out-system interface is less likely to be under the control of the designer, and so can be taken as a contextual value outside the scope of changes that the designer may induce; that is, we contend that the out-system values are less likely to change during designing than the in-system values.

We note that the nominal balance is calculated using absolute value because one cannot tell *a priori* whether exceeding the out-system value is considered over-design or under-design.

Let us reconsider the example of the electric voltage in a domestic kettle.

- Let the out-system interface (the wall outlet) have a nominal value of 120V, a minimum of 104V, and a maximum of 127V (assuming the 87-106% range and all else in the house works properly); i.e., `[-104, -120, -127]V`. Values are negative because the voltage flows "out" of the context system.
- Let the in-system interface (the kettle) have a nominal value of 120V, a minimum of 110V, and a maximum of 130V; i.e., `[110, 120, 130]V`. Values are positive, because they represent flows "into" the kettle system. We note these are not necessarily the final values of the design, but rather the values that obtain at some point *during* the kettle's design.

- The sum of the ranges of the interfaces is `[-6, 0, 3]V`.
- Therefore the balance of this interface - i.e., of the kettle in its operating context - is `[-0.06, 0, 0.02]`. Notice the balance is dimensionless.

This means that the interfaces are balanced at the nominal values, but over-designed at the upper end and under-designed at the lower end.

Let us now consider a simple example with multiple BVs. Let us start by assuming two other BVs of the electric kettle are of interest: the heat energy produced by the kettle (as an engineering requirement corresponding to the customer requirement of heating time) and the volume of water able to be boiled at one time. Data for this example is derived from estimates based on diverse samples found on the Web and in kettles owned by the authors. The precise values are not important; rather, the point of the example is to demonstrate the kind of conclusions that can be drawn from the assessment of the BVs.

For the heat energy of the kettle, let us define the BV as `[-95, -280, -750]kJ` for the out-system and `[90, 300, 750]kJ` for the kettle (in-system). For the volume of water, let us define the BV as `[-0.25, -1, -2]L` for the out-system, and `[0.5, 1.2, 2.25]L` for the kettle (in-system). Using the equations above, we find that the balance of the heat energy BV is `[0.05, 0.075, 0]` and the balance of the volume BV is `[-1, 0.2, 0.125]`.

In summary, the balances of the three BVs are:

- Voltage: `[-0.06, 0, 0.02]`
- Heating energy: `[0.05, 0.075, 0]`
- Water volume: `[-1, 0.2, 0.125].`

Since the balance measures for the three BVs are non-dimensional, we can even add the balance values together to arrive at a measure of the overall balance of the whole kettle. If these three BVs were the only ones in the design (though clearly there would be others in a real case), the current kettle design's overall balance is `[-1.01, 0.275, 0.145]`. We might well argue that, overall, our kettle is significantly under-designed at this point in its development.

If we consider the individual balance calculations for the BVs in this example, we can see quickly where the imbalances are in the design. Because the balances are normalized and non-dimensional, we can also tell which imbalances are greatest and – presumably – more urgent based on their relative values. The worst imbalance is a comparatively significant under-design of the minimal water volume. That the nominal value of the water volume is also over-designed suggests some kind of misunderstanding of the scope of the kettle design. The voltage BV is the most balanced of the three. The heating energy aspect appears to be slightly over-designed generally. Overall, this analysis suggests that the most significant flaws in the design are around water volume and, *ceteris paribus*, these should probably be addressed first.

This example is meant to suggest that BV assessment can be a useful tool to drive designing tasks. While it is impossible to do in a general case such as this example, one can easily imagine that in a specific and fully described situation, project managers could use this information to help them allocate resources to the design work needed to address the imbalances.

The authors now present another small example, this time situated at the systems level of the design of a gas station in an urban setting. For brevity, we will only look at the flow rate of vehicles through the station. We assume a preliminary design already exists. Values are derived from exercises on gas station design used by Salustri in his classes and are loosely based on a particular station in downtown Toronto, and does not take into account any strongly atypical equipment failures (e.g., failure of more than one pump at a time). Again, the values themselves, though reasonable, are not important; what matters are the conclusions that can be drawn about the station design.

The station has four pumps. Vehicles arrive at the station at a rate ranging between 1 and 90 vph (vehicles per hour), depending on time of day, with a nominal rate of 23 vph. The rate of departure of vehicles from the station is between 1 and 75 vph (due to traffic congestion at the exit during peak hours) with a nominal rate of 23 vph. The rate at which the station can process vehicles is between 1 and 80 vph, with a nominal rate is 48 vph. The measures of balance are `[0, 1.09, -0.11]` at the entrance and `[0, 1.09, -0.07]` at the exit.

It is not surprising that the station is balanced at the minimum condition.

The nominal condition is strongly overdesigned at both the entrance and exit. This results from the fixed nature of the services (especially the number of pumps and the rate at which they dispense fuel). It also is strongly suggestive that there are other parameters in play that have the effect of increasing the nominal flow rate of vehicles for the overall benefit of all the stakeholders; e.g., ceteris paribus, more pumps mean more potential revenue and more satisfied customers. This is in fact a guideline for the use of BVs: a strongly overdesigned BV may indicate that some other BV has not been accounted for. This therefore should trigger a check by the designers for other system parameters that are not (yet) part of the model. In the case of the gas station as presented here, there are at least two such parameters: the speed at which the gas pumps dispenses fuel (customers are more likely to use the station the faster the pumps can fill their vehicle tanks), and the costs of multiple pump configurations versus the revenue they generate (there will be a point of diminishing returns from adding more and more pumps). BVs for these parameters are not currently complete and are the subject of future work. We hypothesize that if these BVs were available, we would find that the overall gas station was reasonably well balanced even though the nominal vehicle flow rate is overdesigned.

Finally, we find that the station is slightly underdesigned for the maximum conditions at both the entrance and exit, and that the underdesign is more pronounced at the entrance than at the exit. This means that neither can the station accept all the vehicles at peak conditions, nor can the streets accept all the vehicles at peak conditions (due to traffic congestion). The difference in the amount of underdesign arises from the differences in vehicle flow at the entrance and exit, compared to the maximum flow of the station itself.

One may address the underdesign at the entrance by adding more pumps, but this may adversely affect the finances of the station (the revenues of more pumps may not offset their costs); a full accounting of BVs for revenues versus costs would again be needed. However, increasing the number of pumps will also *worsen* the underdesign at the exit; that is, increased vehicle flow through the station will further congest the exit. One alternative in this case could be to look for a way to artificially congest the entrance during peak hours so that the flow rate into the station does not exceed the flow rate out of it (such as an illuminated sign at the entrance that indicates the station is full and attempts to wave off incoming drivers). While this may inconvenience some *potential* customers, it may reduce the overall inconvenience of *actual* customers, especially those “stuck” at the exit because of street congestion. To fully explore this option, we again see the need to add at least one BV to somehow capture customer satisfaction.

6 DISCUSSION AND FUTURE WORK

This paper has presented a technique to capture ranges of possible (though not necessarily preferable) values of design parameters throughout a design process, taking advantage of principles of systems thinking. The modelling construct is a *balance variable*, which captures the range of values on both sides of a boundary between two systems (one contextual, and the other of the artefact being designed).

The contribution of this work is to present a new tool and associated method to represent and quantify systems interface information for engineering design, with the aim of providing a way for designers to assess the quality of designs and determine best ways to improve designs, while imposing a minimum cognitive burden on them (so they can remain focussed on the actual design tasks).

BVs can be assessed in a non-dimensional, normalized way. This supports simple, straightforward checking for design balance, and quick identification of the existence and degree of over- and under-design conditions. Consequently, this kind of analysis can be useful in setting priorities and allocating resources to aspects of the design that need further work.

Using BVs allows one to establish that a design may be simultaneously overdesigned with respect to some parameters and underdesigned with respect to other parameters. This is not common; that is, we often hear of a design being *either* overdesigned or underdesigned, but not both. The authors believe our approach is a more nuanced one, enabling designers to understand in greater detail the nature of their designs as they are being designed.

While our approach does not actually ensure design effectiveness / robust design, it does support the embedding of information pertinent to robustness. The range of a BV can be used to set constraints derived from considerations of effectiveness – that is, a balanced design will manifest BVs to cover the ranges of corresponding BVs in the artefact’s environment. This can be established in the early stages of designing, thus helping to ensure that robustness is “built into” the design. Thereafter, the design can be optimized for efficiency with less concern of generating functionally brittle (ineffective) artefacts. We note that, in the examples shown in the previous sections, the out-system ranges of BVs could well be interpreted as the environmental variations that robust design is intended to accommodate.

The careful observer may have noted that the equations for balance b in section 5 above are problematic if one or more of the out-system BV values is zero. While it is unlikely in real situations that this condition will occur, the authors acknowledge that it is a concern; we are investigating various possible solutions.

Rogers (2013) has developed a variation of BVs as presented here, specifically intended to fit into a broader concept design framework based on an analogy with natural systems and evolution. In that work, BVs are conceptually the same – representations of ranges of required and expected performances, the combination of which supports a calculation of overall fitness of a design – but vary in implementation details. For instance, Rogers uses BVs to calculate only imbalances because those are where more design work is needed; whereas, in the work presented here, both balance and imbalance is captured. It is not clear to the authors at this time which approach is more practicable and efficient *in vivo*. Comparing the two approaches is a matter of future work.

It is easy to implement simple tools to support the assessment of BVs with readily available computer tools such as spreadsheets. The calculations are programmatically trivial. Conditional formatting can be used to automatically change the visualization of BVs to highlight them for designers. For instance, over-design conditions can be rendered in yellow and under-design conditions in red. One can also use built-in functions to include some sense of accuracy in the calculations. (For instance, an under-design imbalance of -0.0000001 may be well-treated as negligible in some situations.)

Ideally, a BV management tool would connect directly to other existing software for requirements management (to capture out-system BVs) and CAD software (to capture in-system BVs). This kind of tool could also be embedded in PLM software systems; indeed, this might well be the preferred implementation because most PLM software already provides the computational infrastructure to route data between computational elements / programs. While this would be more difficult than simple spreadsheets, it would be comparatively more powerful as an embedded part of a PLM system.

Salustri will be developing a simple spreadsheet based tool to support the assessment of BVs for educational purposes. It is hoped that the experiences gained in that exercise will lead to further insights into how to make the BV modelling method more useful in practice.

Currently, BVs are modelled with a uniform probability that any value in the range of a BV will obtain. While this is a reasonable simplification for early design (i.e., when information about probabilities is not likely available) it limits the application of the BV method in later design stages. The authors will also be investigating ways to model more realistic probabilities that certain values will obtain for a BV.

The boundaries of BV intervals, as shown in Figures 1 and 2, are crisp. However, there is a very good case to be made that in reality such edges are better modelled by transitional regions. This would turn a boundary into something akin to a *boundary layer*, in analogy to boundary layers in fluid flow. A graphical depiction of such boundary layers is shown in Figure 3. Although more realistic, whether such a refinement will result in a more useful model of design balance is an open question, at least insofar as the early stages of designing go. Nonetheless, there may well be situations in which an investment in transition modelling will result in benefits that make that investment worthwhile. Studying this is also a matter of future work.

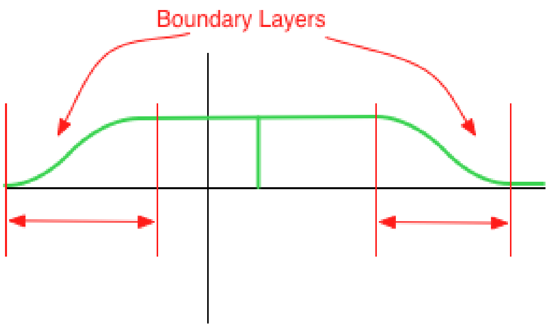

Figure 3. Boundary layer representation of one side of a balance variable.

Another possible improvement involves the use of Likert scales to measure BVs for which actual values are not available (this too can happen in early design). Likert scales may also allow a number of non-physical flows to be representable; this would expand the applicability of the BV method to areas outside conventional engineering design so as to include more subjective characteristics of product design and industrial design, etc.

Finally, BVs (or BV-like structures) may be useful representational elements of other methods and tools, such as Failure Mode and Effects Analysis (FMEA). In conventional FMEA, severity, likelihood, and detectability are measured on non-dimensional scales normalized to a range of 1 to 4. However, FMEA measures are single values, which can naively suggest a level of precision that is not actually present. It may well be that at least in some applications, using ranges of values, as is done with BVs for system interfaces, will yield more useful results.

The authors hope to investigate these and other applications in the future.

REFERENCES

Alexander, C. (1964). Notes on the Synthesis of Form. Harvard University Press.

Antonsson, E.K. & Otti, K.N. Imprecision in Engineering Design. (1995). ASME Journal of Mechanical Design, 117(B), 25-32.

Beer, M., Ferson, S., & Kreinovich, V. (2013). Imprecise probabilities in engineering analyses. Mechanical Systems and Signal Processing, 37(1), 4-29.

Esser, W. (2011). Export relevant information on supply distribution networks in North America. Technical Paper. Eaton Industries GmbH. http://www.moeller.net/binary/ver_techpapers/ver965en.pdf (accessed 4 December 2014).

Garcia, R. (2014). Roadway Design Manual. Texas Department of Transportation. URL: http://onlinemanuals.txdot.gov/txdotmanuals/rdw/index.htm.

Gero, J.S., & Kannengiesser, U. (2004). The situated function–behaviour–structure framework. Design studies, 25(4), 373-391.

Habib, W. & Ward, A.C. (1990). Proving the Labeled Interval Calculus for Inferences on Catalogs. Proc ASME Conf on Design Theory and Methodology. New York.

Karnopp, Dean C, Donald L Margolis, and Ronald C Rosenberg. "System dynamics: a unified approach." (1990).

Rogers, D. (2013). A framework for early design process stages based on an analogy to evolution. PhD Thesis, Ryerson University, Department of Mechanical and Industrial Engineering.

Salustri, F.A. (2010). Misfits, Balance, Requirements, and Systems: thoughts on Alexander’s Notes on the Synthesis of Form. Proc 2010 Conf of the Design Research Society, Montreal.

Salustri, F.A., Rogers, D., & Eng, N.L. (2009). Designing as Balance-Seeking Instead of Problem-Solving. Design Principles and Practices, 3(3), 343-356.

Simon, H.A. (1969). The Sciences of the Artificial. MIT Press.

Taguchi, G., Chowdhury, S., & Taguchi, S. (2000). Robust engineering. McGraw-Hill Professional. Chicago.

Ullman, D.G. (2001). Robust decision-making for engineering design. Journal of Engineering Design, 12(1), 3-13.

Wood, K.L. & Antonsson, E.K. (1989). Computations with Imprecise Parameters in Engineering Design: Background and Theory. ASME Journal of Mechanical Design, 111(4), 616-625.