

Trading reliability targets within a supply chain using Shapley's value

John Quigley, Lesley Walls*

Department of Management Science, University of Strathclyde, Glasgow, G1 1QE Scotland, UK

Accepted 22 September 2006

Available online 4 December 2006

Abstract

The development of complex systems involves a multi-tier supply chain, with each organisation allocated a reliability target for their sub-system or component part apportioned from system requirements. Agreements about targets are made early in the system lifecycle when considerable uncertainty exists about the design detail and potential failure modes. Hence resources required to achieve reliability are unpredictable. Some types of contracts provide incentives for organisations to negotiate targets so that system reliability requirements are met, but at minimum cost to the supply chain. This paper proposes a mechanism for deriving a fair price for trading reliability targets between suppliers using information gained about potential failure modes through development and the costs of activities required to generate such information. The approach is based upon Shapley's value and is illustrated through examples for a particular reliability growth model, and associated empirical cost model, developed for problems motivated by the aerospace industry. The paper aims to demonstrate the feasibility of the method and discuss how it could be extended to other reliability allocation models. © 2006 Elsevier Ltd. All rights reserved.

Keywords: Reliability allocation; Reliability targets; Supply chain; Reliability growth model; Pareto optimality; Shapley value

1. Introduction

The design, development, manufacture, delivery and service of complex systems are fulfilled by many organisations within a supply chain. Aerospace is one industry where prime contractors and system integrators lead co-ordination of the value chain where, according to [1], approximately 70% of the final value of the aircraft is outsourced. A five-tier model represents the aerospace supply chain [2] with the first tier corresponding to the prime contractors. The second to fourth tiers involve large and medium sized manufacturers with responsibility for subsystems, major components and structural parts. The fifth tier comprises small companies and sub-contractors specialising in materials, manufacturing, processes, software and services.

The aerospace industry is interesting because its supply chain contains dedicated subsystem suppliers serving a relatively narrow market. In relation to other manufacturing industries, aerospace production volume is low;

product lifecycles and lead times between projects are long; investment in research, development and unit manufacturing costs are high; largely due to certification. The relatively small number of sub-system suppliers at upper levels of the chain combined with the highly specialised small-to-medium sized enterprises within the lower levels of the chain, means prime contractors are highly dependent on developing longterm relationships with suppliers and this has implications for the management of reliability in the light of emergent changes in organisational structures and processes [3].

Currently, the prime contractor translates customer reliability requirements into appropriate targets for the system then allocates these targets across sub-systems. This is repeated down the supply chain. Reliability targets may be expressed as, for example, failure free operating times, failure rate or mean time between specified events, such as failure, replacement or overhaul. To assure targets are met, each organisation implements a programme of reliability activities that aim to uncover potential failure modes and/or to prove its product design. The activities may include, for example, failure mode and effects analysis, fault tree analysis, physics of failure, reliability enhancement testing

*Corresponding author. Tel. +44 141 548 3616; fax: +44 141 552 6686.
E-mail address: lesley.walls@strath.ac.uk (L. Walls).

etc. It is not unreasonable to expect the cost of a reliability programme to depend upon the target and the current level of knowledge about the inherent system reliability. For example, the higher the target levels and lower the knowledge about product heritage then the greater investment needed in the reliability programme.

There is no evidence that organisations within a supply chain work together rationally to re-negotiate targets, in the light of information feedback from observations obtained through implementation of reliability activities. In part this may be due to lack of contractual incentives. For example, contracts that differentiate payment for product delivery and operational support provide no incentive for reliability improvement since profit can be generated through spares provisioning. In contrast, service level contracts [2] do incentivise the supplier to minimise operational failure rate since non-availability represents a cost. Further, such contracts imply the customer does not have to micro-manage the reliability programme, or administer repair, and suppliers must be given freedom to modify systems to improve reliability.

This paper describes a method to support trading of reliability targets across a supply chain by minimising the cost of the combined suppliers reliability programmes in the light of information generated from implementation of activities in each, under the constraint that the system reliability target is to be met. To achieve this requires a means of trading reliability targets between organisations. We assume this is feasible during long project gestations for complex systems because different sub-systems and parts will intentionally evolve over different horizons, hence the opportunity to review information and costs at multiple decision points is available. Section 2 positions the proposed approach in relation to the existing literature in reliability allocation, hence establishing the gap in knowledge that this paper seeks to fill. A criterion for measuring the value of exchanging reliability targets within a supply chain based upon game theoretic arguments is proposed and Section 3 shows how the model can be developed into a framework for trade within the supply chain by examining a mechanism for deriving a fair price for transferring reliability targets between suppliers. Section 4 presents an example to illustrate use of the method. Conclusions and suggestions for further work are presented in Section 5.

2. Summary review of reliability allocation literature

Reliability requirements are regarded by [4] as *‘the overall needs of the user and define the objectives for the product’*, while [5] believes they *‘define the problems that the designer must overcome’*. Both are valid although they assume differing perspectives: the customer and the technology respectively. Guidance about the precise specification of requirements given in [5] proposes that they should: define failure in relation to the system function; provide full descriptions of the environments in

which the system will be stored, transported, operated and maintained; and state failure modes, and effects, which are critical and therefore must have a very low probability of occurrence which can be subject to verification.

Reliability targets will be generated from requirements and are typically higher than those specified to compensate for uncertainties and shortfalls in measurement. Reliability targets contribute to the portfolio of constraints, which includes functionality, cost and time, bounding the design solution space. Reliability allocation is the process of apportioning the system target to lower indenture levels in relation to the design architecture. This permits sub-contractors to work to relevant targets for sub-systems and parts; permits evaluation of achieving requirements; and allows designers to understand relationships between parts, sub-system and system reliability as well as the relationship between reliability and other system characteristics.

Methods for allocation range from apportioning reliability equally across sub-systems to assign targets based on importance criteria, feasibility or historical data [6]. All make similar assumptions. For example each sub-system has a constant failure rate and fails independently of others; hence each sub-system is assigned a percentage of the system failure rate. The AGREE method [7] assumes a series system comprising identical sub-systems and assigns each equal failure rates. The ARINC method [8] also assumes a series system but computes weights in proportion to sub-system failure rate where the latter are estimated from historical data. The feasibility of objectives method [9] weights sub-system failure rates as a function of factors, such as complexity of sub-system design, where such factors may be assessed using subjective expert judgement.

Regardless of the approach taken, reliability allocation provides an input to a decision support model which can help engineers determine an optimal system design where reliability is a constraint and choices are made about component selection, configuration and level of redundancy. Numerous optimisation algorithms exist. Ref. [6] gives examples where technical, or economic measures are optimised subject to specified constraints. A comprehensive review of optimisation in system reliability is given in [10]. Typically, one decision maker is assumed and models are constructed to determine an optimal design as a one-off decision, an exception being [11] who explore Pareto optimal sets for system reliability optimisation when uncertainty exists within the component reliabilities. Research effort has tended towards efficient and effective algorithms. While mathematical programming approaches are useful, they provide little insight into the manner in which reliability targets are sequentially traded off against costs and negotiated across the system.

The extent to which cost is captured in allocation models varies. Typically continuous parametric cost functions are assumed. For example, [12,13] model hardware reliability costs for system task allocation, assuming

linear cost functions. Ref. [14] proposes three criteria for cost functions used in reliability optimisation: it should be positive definite; non-decreasing; and increase rapidly as reliability approaches 1. They develop a generic cost model as an exponential function of the inverse of the system reliability. However, as argued in [15], continuous costs functions are not always appropriate as decisions are discrete and alternatives are finite. This will be particularly true when design choices are few. Hence there is a need for further consideration of discrete cost functions.

This review has reported models to support engineers allocate reliability across the system but no models have been identified for the stated management problem of allocation and trade of reliability targets within the supply chain. Hence the development of a method to support trading, given that the system design uncertainty will change through development, should contribute to the literature. This paper also addresses some specific shortcomings of the existing literature by developing, for example, discrete cost functions empirically, which has relevance to other resource allocation models.

3. Method for trading reliability targets

The proposed method to facilitate transfer of reliability targets between organisations in a supply chain is based upon game theoretic principles and is first developed at a conceptual level for the simplest case of a set of two suppliers before extension to a more general solution. The underlying assumptions are stated in Section 3.1. Section 3.2 illustrates the basic principles and discusses the characteristics of optimal exchange of targets. Section 3.3 presents the general solution.

3.1. Assumptions

The following assumptions are required in order to develop the method. The assumptions have been informed by both our understanding of industry practice as well as elements of game theory.

1. Customer reliability requirements are stated and fixed for the system.
2. Contractual obligations between suppliers and the customer are agreed prior to the start of the reliability programme.
3. The supply chain is defined by the organisations that provide sub-systems and parts with measurable reliability targets for integration into the system (the suppliers) as well as the system integrator (the prime supplier).
4. There exist contractual incentives for the supply chain to meet reliability targets at minimum cost.
5. The target reliability of the prime supplier is a function of the target reliabilities of all sub-suppliers in the chain and the reliability of the system integration.
6. There are multiple decision makers; one within each organisation within the supply chain.
7. There are multiple decision points corresponding to time points within the system project plan at which reliability is assessed and can be re-negotiated.
8. Target reliabilities can be traded off between suppliers such that the system reliability target is met at the minimum cost to the supply chain.
9. Cost is measured as the investment in resources for the activities conducted within the reliability programmes of the suppliers.
10. Each organisation uses a model that explicates the relationship between the reliability of the product design and the reliability activities to which it can be exposed.

Most of the assumptions are self-explanatory. For example, assumptions 1 through 9 are dependent upon the industry context hence their validity can only be assessed for a given problem. Assumption 10 requires that each organisation within the supply chain should be using a model to estimate reliability through development and to select the optimal portfolio of activities in formulating their initial reliability programme. Ref. [16] proposes such a reliability growth model and show how Bayesian updating can be performed for an individual organisation. This model is implemented for the examples presented in this paper and its salient properties will be explained as required.

3.2. Edgeworth box and Pareto optimality for a two-supplier chain

A simple example is used to illustrate the basic principles of the trading model. Consider a simple chain consisting of two suppliers, each contributing a component part. In addition to the aforementioned assumptions in Section 3.1, assume each part has a constant failure rate and the system failure rate is the sum of the two part failure rates since the system requires the parts to operate in series.

First let us consider one supplier. Conceptually, the expected cost of their reliability programme will be negatively related to the level of the target failure rate. Note that failure rate is used simply as one example of a reliability indicator for illustration. The approach is not constrained to a failure rate, although the diagrams require adaptation to reflect the appropriate direction of the relationship. For example, the cost would be positively related to probability of successful operation for a specified duration.

The trade-off between cost and revenue from performance will translate to a profit for the system. If profit is the only criterion to be satisfied, then the supplier will be indifferent to the trade-offs between programme costs and failure rate for fixed profit which is represented by points on a profit curve. Fig. 1 illustrates these so-called indifference curves and shows that as profit levels increase,

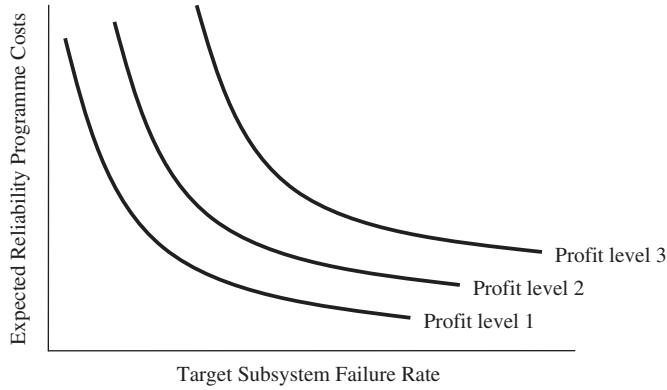


Fig. 1. Indifference curves capturing trade-off between expected cost and reliability for one supplier at a different system profit level.

the curves shift to the right. For illustration the indifference curve is drawn as a smooth continuous function, although in practice, the curve will comprise a series of discrete points [15] but will have the characteristic of being monotonically non-increasing.

Each supplier will have an indifference curve for each fixed profit level. For the case of two suppliers an Edgeworth box [17] can be constructed, as shown in Fig. 2. This represents the superposition of the indifference curves of the two suppliers from their respective origins in the bottom left hand vertex (supplier A) and top right hand vertex (supplier B), for fixed profit levels, either A_{P1} or A_{P2} for supplier A or B_{P1} or B_{P2} for supplier B. The vertical axis of the box corresponds to the total expected reliability programme costs for the system, which is the sum of the costs for the two suppliers, while the system failure rate is plotted on the horizontal axis and is the sum of the failure rates of the component parts.

Fig. 2(a) illustrates a target failure rate and associated cost for supplier A denoted by (μ_{A1}, c_{A1}) and for supplier B denoted by (μ_{B1}, c_{B1}) . Measuring these from their respective origins so that $(\mu_{A1}, c_{A1}) = (\mu_{B1}, c_{B1})$, the box can be partitioned according to the contribution of the two suppliers to the system failure rate and costs. Any point in the box represents an allocation of costs and failure rate, whereby the combined costs and failure rate remain constant. Note not all points may be feasible given the discrete nature of cost as a function of reliability. The current allocation of reliability targets is sub-optimal, since moving to the right along indifference curve B_{P1} , for example, would result in an improved allocation. This is because at the current allocation of reliability targets, a marginal increase in the failure rate for A, say μ_{A1S} , compensated by a decrease in the failure for B, to μ_{B1S} , will result in a reduction in total costs since $c_{A1S} + c_{B1S} < c_{A1} + c_{B1}$. See, for example, Fig. 2(b). The vertical difference between the two indifference curves at any particular allocation of the failure rate represents the reduction in combined cost. The shaded area between these curves represents improved allocations because at least one

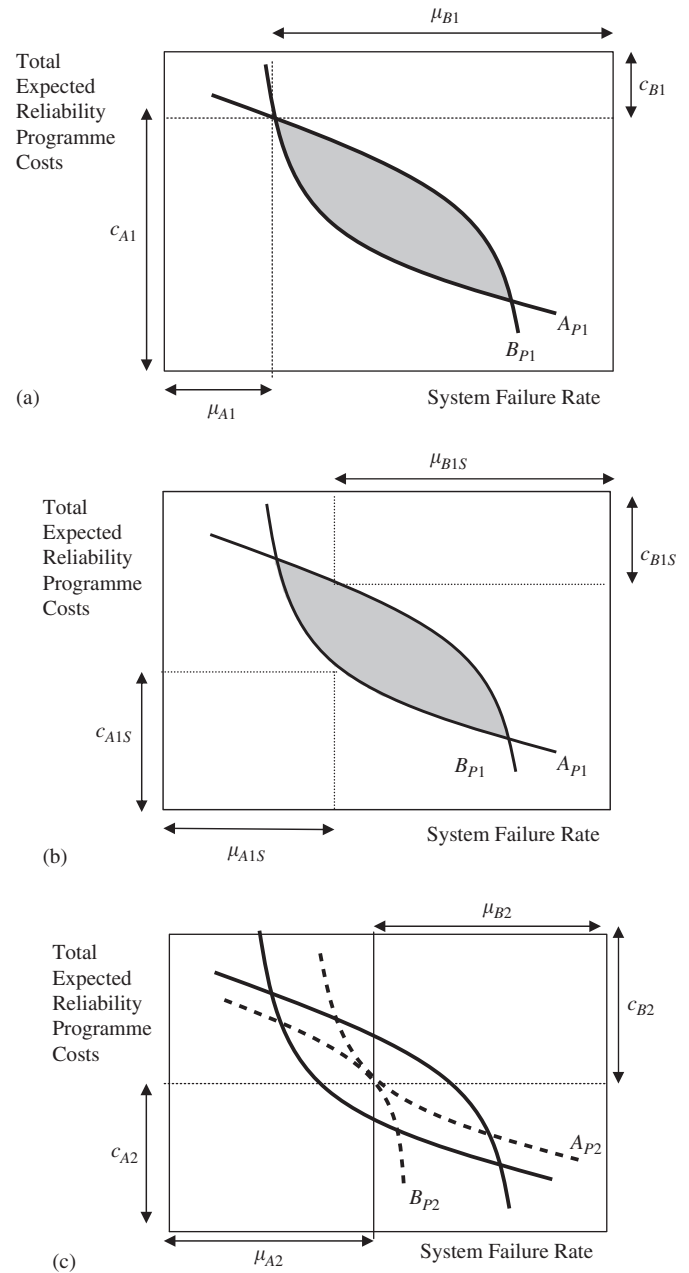


Fig. 2. (a) Edgeworth box for suppliers A and B. (b) Edgeworth box shift in failure rates and total savings. (c) Optimal allocation between suppliers A and B.

of the suppliers would be on a higher profit curve without changing the combined costs and failure rates. However, there is no guarantee that a feasible allocation exists within the shaded area.

Fig. 2(c) represents another possible allocation of the failure rate. Again this leads to cost savings, which if divided between supplier A and B can be conceptualised as each having indifference curves with a greater profit level, denoted by A_{P2} and B_{P2} , respectively. This allocation is optimal in a sense, since the characteristic of both indifference curves are tangent to one another, representing equivalence in the marginal rates of substitution between costs and failure rate at the point $(\mu_{A2}, c_{A2}) = (\mu_{B2}, c_{B2})$.

While all allocations within the shaded region area in Fig. 2(a) are improvements over the initial allocation, they are not all optimal. Multiple optimal solutions may exist for this re-allocation problem and can be found using mathematical programming assuming data are available to construct a total-cost model.

Sharing benefits from cost reduction between suppliers requires their cooperation and therefore depends on bargaining between both parties. The concept of Pareto optimality characterises efficient allocations within Edgeworth boxes, such as those presented in Fig. 2, and hence informs bargaining. An allocation is considered Pareto optimal if there does not exist an alternative allocation where one supplier is better off without making the other supplier worse off. A detailed discussion of Pareto optimality is given in [17], while [11] develops Pareto optimal solutions for optimising system reliability when uncertainties exist in component reliabilities.

3.3. Shapley value to support trading

A characteristic of a Pareto optimal solution is that the marginal rate of substitution for fixed profit level (i.e. the derivative of the expected reliability programme cost function with respect to target reliability) is the same for each supplier. Each problem can contain multiple allocations with this characteristic; we seek a method to determine a unique solution by generalising the concepts introduced in Section 3.2 and developing a criterion to support the allocation of benefits derived from the co-operation between those suppliers who have agreed to form a coalition using the Shapley value. This provides a measure of the contribution of each supplier towards a co-operative reliability strategy.

The Shapley value has its roots in co-operative game theory where it is reputedly the most studied and widely used single-valued solution concept [18]. Applications to valuation cover a range of areas. Recent cases include environmental pollution cost allocation [19]; production decisions [20]; identification of drivers for customer satisfaction [21]; transportation problems [22]; optimal cartel size [23]; allocation of electricity transmission costs [24]; and insurance pricing [25]. There are no reported applications in reliability analysis.

Shapley [26] derived a formula for evaluating the contribution a player makes to the value of a cartel within a co-operative game. The formula is based on three axioms. Here we express these axioms in terms of our problem. Hence organisations are the players, coalitions within the supply chain represent the cartel and value corresponds to the savings made through reallocation of reliability targets that are to be shared between organisations which form coalitions to trade targets:

1. The savings attributed to the contribution of an organisation depends upon whether they belong to the coalition or not, and does not depend on the order in which the organisation joined the coalition;

2. The sum of the savings attributed to the individual organisations should equal the savings made within the coalition, with organisations making no contribution to the coalition being assigned zero value;
3. There is no expected gain or loss for an organisation in delaying trade at any given decision point.

For the case of a two supplier chain, only one coalition is possible. For the general case of a supply chain comprising more than two organisations, many possible different coalitions could be formed ranging from one coalition comprising all suppliers to many coalitions containing a few suppliers. We assume that the number of organisations joining a coalition will be specified a priori and is denoted by $|N|$. The marginal savings each organisation generates to a coalition depends upon the size of that coalition. Let $|Z|$ denote the set of organisations in a coalition prior to organisation i joining. The saving arising from the inclusion of the i th organisation in a coalition of size $|N|$ has been shown by [26] to be given by the Shapley value:

$$\phi_i = \sum_{Z \subseteq N-i} \frac{|Z|!(|N| - |Z| - 1)!}{|N|!} [v(Z \cup \{i\}) - v(Z)], \quad (1)$$

where $v(Z \cup \{i\})$ is the savings of the coalition of Z organisations and organisation i , $v(Z)$ is the savings of the coalition of Z organisations; $|Z|$ and $|N|$ represent the number of elements in the sets Z and N , respectively.

Thus Shapley value is essentially a weighted average of the marginal savings organisation i makes to a coalition, averaged over all possible points of entry into the coalition. That is, first, second and so on. The weights measure the probability that a coalition of size $|Z|$ is formed prior to supplier i joining, assuming a purely random distribution of joining sequences. A detailed discussion is given in [27].

4. Example application of trading method

The proposed method has not been applied in practice; hence we develop an illustrative example that builds upon our experience of reliability modelling during the development of aerospace products [28].

Consider a system with four sub-systems, each supplied by different organisations. All four sub-systems must work for the system to operate and each fails independently of the others. The target reliability for the system is 0.73 where reliability is measured as the probability of successful operation at 100 time units. Targets allocated to suppliers 1 through 4 are 0.95, 0.95, 0.95 and 0.85, respectively. While these targets are negotiable, the system target is fixed.

Sections 4.1–4.3 explain the implementation of the proposed approach. In Section 4.1 we show how the process is initiated by estimating the sub-system reliabilities and selecting the activities within the programmes for each individual supplier. In Section 4.2 we explore whether there are less costly allocations of reliability between suppliers

that will meet system targets. If so, we identify the optimal allocations and use Shapley's value to share the cost savings between those suppliers who agree to trade as shown in Section 4.3.

4.1. Model for reliability and allocation of activities in individual organisations

The model proposed by [16] is used to explicate the relationship between the reliability of the system, or sub-system, design and the reliability activities to which it can be exposed. The model assumes the design has enumerable failure modes that will result in failure in operation if not corrected. The purpose of the reliability activities is to detect these failure modes or confirm whether or not they exist. Each activity may provide information on many failure modes, although not necessarily all.

Ref. [16] shows that the reliability of a system at time t exposed to a portfolio of activities can be written as

$$R_S(t) = \prod_i \prod_j \left[1 - (1 - R_{ij}(t)) \lambda_{ij} \prod_k (1 - p_{ijk})^{\delta_{ik}} \right], \quad (2)$$

where $R_S(t)$ is the reliability function at time t , i the i th supplier, j the j th potential failure mode; $R_{ij}(t)$ the reliability function of the i th supplier; λ_{ij} the subjective probability j th potential failure mode of i th supplier that exists; k the k th activity; p_{ijk} the subjective probability of the j th potential failure mode of i th supplier realized by activity k , and δ_{ik} the decision variable representing selection (1) or non-selection (0) of activity k by the i th supplier.

Since the exact number of failure modes and the actual effect of an activity upon a failure mode are not known, probabilities measuring the likelihood of each of these events can be elicited through structured expert judgement. Historical failure data, for example from heritage systems or generic databases, can be used to provide estimates of the time of realisation of failure modes in operation.

The target reliability can be stated in terms of the reliability function. Typically, the objective is to minimise costs, while satisfying the required target reliability and completion of the programme within specified calendar project duration. Assuming all activities will be conducted sequentially, Ref. [16] has shown that solving the following integer programming problem will give an optimal portfolio of activities:

$$\text{Min} \sum_k \sum_i \delta_{ik} C_{ik},$$

subject to:

$$R_S(t|\underline{\delta}) = \prod_i \prod_j \left[1 - (1 - R_{ij}(t)) \lambda_{ij} \prod_k (1 - p_{ijk})^{\delta_{ik}} \right] \geq R_0, \quad (3)$$

$$\sum_k \sum_i x_{ik} \delta_{ik} \leq X_0,$$

where R_0 is the target reliability, C_{ik} the cost of k th activity for the i th supplier; X_0 the specified project duration; x_{ik} the duration of k th activity for the i th supplier and $\underline{\delta}$ the matrix of decision variables.

4.2. Instantiation of reliability and cost models with data

Tables 1 and 2 summarise data from each of the four suppliers, $i = 1, 2, 3, 4$.

For the corresponding sub-system, Table 1 lists the potential failure modes and the probabilities that the j th failure mode for the i th supplier will be realised in operation (λ_{ij}), assuming the current design remains unchanged. The sets of failure modes for each sub-system are assumed mutually exclusive. This is not unreasonable given one organisation will be required to take responsibility for addressing a particular issue.

The time to realise a potential failure mode can be modelled by a probability distribution. In Table 1, we express this as the empirical survival probability for the j th failure mode of the i th sub-system at 100 time units for the particular failure mode, $R_{ij}(100)$. There are no constraints on the expression of the probability distribution and the method supports both parametric and non-parametric approaches.

Each organisation selects relevant reliability activities. These sets of activities may be overlapping between organisations, although for this example we have assumed they are unique to an organisation. The subjective probabilities that the k th activity may remove the failure mode j th failure mode of the i th sub-system, p_{ijk} , are given for all combinations in Table 1.

The cost of the k th each activity for the i th supplier, C_{ik} , (in standardised monetary units) can be estimated from, for example, person effort for implementation and are given for each activity for each supplier in Table 2.

Table 3 presents the sub-system reliability estimates and the expected cost of meeting targets through selected activities based upon implementation of the model described in Section 4.1 for each independent organisation. This shows that activities costing 360 units require to be conducted by 4 suppliers in order to meet the system target and represents the default situation. The total costs obtained through pooling all 4 suppliers together to meet system reliability requirements regardless of individual reliability targets is 200. We propose using Shapley's Value to divide the 120 savings amongst the 4 suppliers. This is discussed in detail in the next section.

4.3. Computing the Shapley value to support trading between coalitions

Assessing the contribution made by a particular supplier requires computation of their marginal contribution on

Table 1
Input data on failure modes, activities and probabilities for each supplier

Failure mode	Probability failure mode realised in operation	Survival probability at 100 time units	Probability failure mode j for supplier i is realised through activity k							
j	λ_{ij}	R_{ij} (100)	p_{ijk}							
(a) Supplier 1 ($i = 1, j = 1, \dots, 5, k = 1, \dots, 6$)										
			1	2	3	4	5	6		
1	0.1	0.9999	0.9	0.1	0	0	0	0.1		
2	0.5	0.9990	0.7	0.9	0	0.1	0.3	0.05		
3	0.2	0.9050	0.8	0.1	0	0	0.2	0.02		
4	0.1	0.9050	0.5	0.01	0.5	0	0.3	0.1		
5	0.6	0.9050	0.5	0	0	0.4	0.2	0.06		
(b) Supplier 2 ($i = 1, j = 1, 2, k = 1, 2$)										
1	0.5	0.3680	0.9	0.1						
2	0.3	0.9050	0.05	0.8						
(c) Supplier 3 ($i = 3, j = 1, 2, k = 1, 2$)										
			1	2						
1	0.3	0.9900	0.5	0.15						
2	0.2	0.9900	0.6	0.75						
(d) Supplier 4 ($i = 4, j = 1, \dots, 3, k = 1, \dots, 8$)										
			1	2	3	4	5	6	7	8
1	0.9	0.3680	0.1	0.05	0.2	0.3	0.5	0	0.75	0
2	0.8	0.9050	0.3	0.06	0.05	0.05	0.003	0	0	0.01
3	0.75	0.9050	0.1	0.01	0.04	0.01	0	0.05	0	0.05

Table 2
Input data on costs of activities for each supplier

Supplier (i)	Cost of activity (C_{ik})							
	1	2	3	4	5	6	7	8
1	100	10	5	25	30	50		
2	100	100						
3	10	20						
4	10	10	20	30	20	21	45	10

Table 3
Reliability estimates, targets and expected activity costs with no coalitions

Supplier (i)	Target reliability (R_0)	Estimated design Reliability $\hat{R}(100)$	Optimal set activities	Expected activity costs
1	0.95	0.92	4, 5	55
2	0.95	0.66	1, 2	200
3	0.95	0.995	None	0
4	0.85	0.37	1, 2, 5, 7	105
System	0.73	0.22		360

joining the coalition. Intuitively, the marginal contribution, relative to the other suppliers, will decrease the later a supplier joins the coalition. Therefore we consider the marginal benefit of all possible joining sequences, so that

the share of one supplier is not influenced by the sequence in which organisations join the coalition. The results of the analysis using the methods developed in Section 3 are summarised in Table 4.

Table 4 shows the 12 possible coalitions from the 4 suppliers. For example, say suppliers 1 and 2 already form a coalition, $Z = \{1, 2\}$, then they may be joined by supplier 3 giving the set $Z \cup \{i\} = \{1, 2, 3\}$. In summary, the rows in Table 4 correspond to the possible coalitions of none, 2, 3 or 4 organisations.

For each coalition, the total cost of the programme that will meet system reliability targets at minimum cost is computed using Eq. (3) given in Section 4.1. The total cost for the default case of no coalition is 360 as before. For convenience, we do not carry any information in Table 4 about the activities to be undertaken, although such a list can be produced for all suppliers under all coalition combinations following the example of Table 3.

The marginal contributions of each supplier involved in a coalition can be obtained by computing the value of the supplier to the coalition. For example, the value of supplier 1 to the coalition $\{1, 2\}$ is 50 monetary units corresponding to the saving 360–310 monetary units. Similarly supplier 2 has a marginal contribution of 50 monetary units to the coalition $\{1, 2\}$. As another example, consider a coalition of 3 suppliers. If a fourth supplier joins the coalition then the marginal contributions are: 30 monetary units (230–200) for supplier 1; 95 monetary units (295–200) for supplier 2; 65 monetary units (265–200) for supplier 3; 5 monetary units (205–200) for supplier 4.

Table 4
Share of savings to suppliers forming coalition computed using Shapley value

Supplier coalition $Z \cup \{i\}$	Total cost $\sum_{vi} \sum_{vk} C_{ik}$	Marginal contributions of each supplier [$v(Z \cup \{i\}) - v(Z)$]				Weight $\frac{ Z !(N - Z - 1)!}{ N !}$
		$i = 1$	$i = 2$	$i = 3$	$i = 4$	
None	360	0	0	0	0	0.25
1, 2	310	50	50			0.083
1, 3	305	55		55		0.083
1, 4	350	10			10	0.083
2, 3	260		100	100		0.083
2, 4	280		80		80	0.083
3, 4	330			30	30	0.083
1, 2, 3	205	55	100	105		0.083
1, 3, 4	295	35		55	10	0.083
1, 2, 4	265	15	85		45	0.083
2, 3, 4	230		100	50	30	0.083
1, 2, 3, 4	200	30	95	65	5	0.25
Expected, share ϕ_i		25.83	66.67	49.17	18.33	

The weights allocated to each coalition measure the probability that a coalition of size Z is formed prior to the i th supplier joining. For example, the number of organisations is $|N| = 4$. For the case of a coalition of all suppliers, $|Z| = 3$ represents the set of all suppliers less the i th organisation. Hence the weight is computed by $3!(4-3-1)!/4! = 0.25$. That is, randomly ordering the four suppliers implies a 25% chance that the i th organisation will join last. Similarly when there is no coalition formed, the weight is 0.25. For all other cases, the weight equals $1/12$ (0.083). This is a third of the weight allocated to the cases of no coalition and all suppliers forming a coalition and represents the three ways in which the i th supplier can join coalitions of two or three suppliers.

The Shapley value is calculated using Eq. (4) as a weighted average of the marginal savings organisation i makes to a coalition, averaged over all possible coalitions of different sizes within the supply chain. Therefore the expected share of the savings generated through the formation of the optimal coalition of all suppliers are found to be 26, 67, 49 and 18 monetary units for suppliers 1 to 4, respectively.

The results in Table 4 show that the most cost-effective plan for the coalition between all 4 suppliers reduces the cost of the combined reliability programmes from 360 to 200, saving 160 monetary units. This requires suppliers 1 and 3 to do no reliability development, supplier 2 to perform activity 1 and supplier 4 to perform 4 activities (1–4). The savings will be shared according to the Shapley values. Intuitively supplier 2 obtains the greatest share because the marginal contributions are, on average, highest as shown in Table 4; hence coalitions involving supplier 2 will generate the greatest savings.

Fig. 3 shows graphical analysis of the relationship between suppliers 1 and 4 to illustrate some of the workings underpinning the results within Table 4. Fig. 3(a) shows points representing the set of activities

where cost is minimised for a given target reliability for supplier 1. Fig. 3(b) shows the equivalent results for supplier 4. If suppliers 1 and 4 form a coalition then the possible combinations of the sets of activities for the two organisations are shown by the grey points in Fig. 3(c). The joint optimal set of activities for a given target reliability are denoted by the black points which bound the cluster of points as we might expect. For example, the target reliability for a sub-system comprising the parts of suppliers 1 and 4 is 0.8075. Analysis of the data in the plot indicates that the minimum cost achievable by a coalition of suppliers 1 and 4 is 150 monetary units. This constitutes a saving of 10 monetary units over the cost associated to these suppliers if they do not trade.

5. Conclusions and further work

This paper has explored a mechanism for deriving a fair price for optimal allocation of reliability targets between suppliers by sharing information gained about potential failure modes through development and the costs of activities required to generate such information. The approach draws upon ideas proposed in the co-operative game theory literature. An illustrative example demonstrates the feasibility of the approach under an initial set of assumptions.

It is feasible to relax some assumptions made in the examples presented. In the first example for the 2 supplier chain, we assumed constant failure rates to introduce basic concepts. In the illustrative example we relaxed this assumption and considered a target expressed as a reliability function at a specified time. More generally, the model can be used with both parametric and non-parametric distributions.

Similarly the examples consider a system that is configured as a series of sub-systems provided by suppliers. The use of the Shapley value is not restricted to this

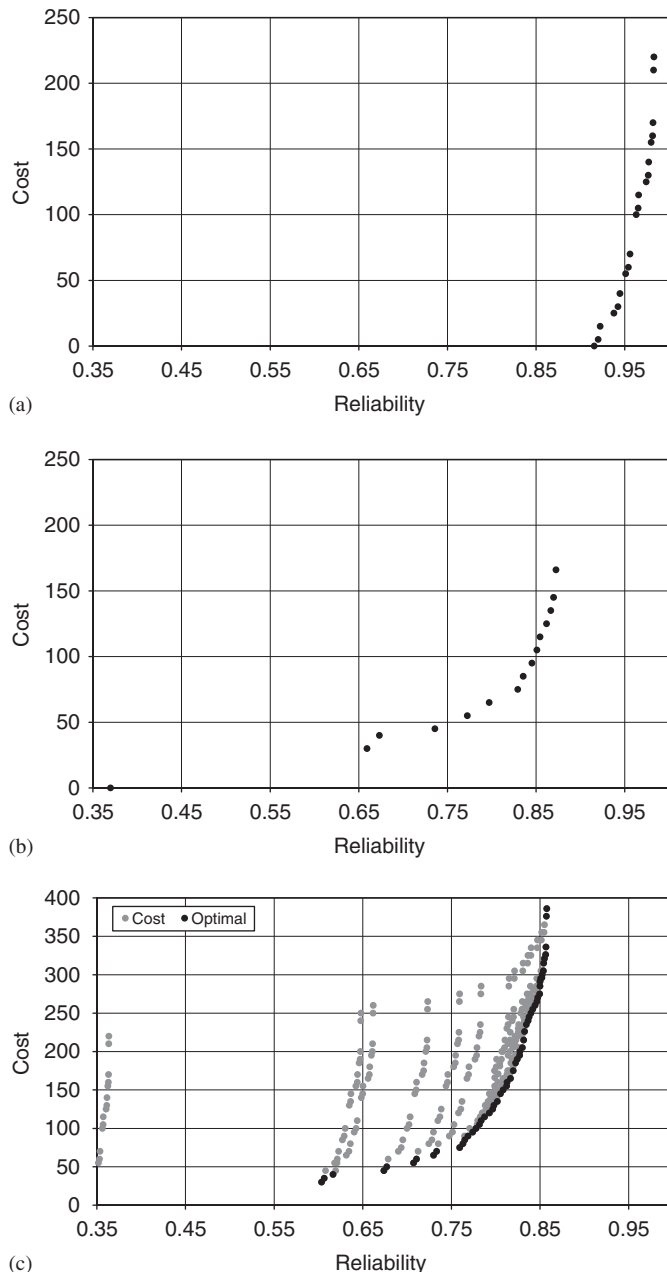


Fig. 3. (a) Cost of activity sets for target reliability for supplier 1, (b) costs of activity sets for target reliability for supplier 4 and (c) all possible (grey) and optimal (black) sets of activities for coalition between suppliers 1 and 4.

situation because it is used in conjunction with a model that explicates the relationship between the reliability of the system design and the reliability activities to which it can be exposed. In the examples within this paper, the model used did assume a series structure. This is not a requirement of the proposed trading mechanism as it could be developed for use with more complex models.

The examples considered two simple situations with 2 or 4 organisations in the supply chain. Hence $|N|$ was taken to equal 2 or 4, respectively. There is no need to assume that a coalition will be formed by all organisations. Indeed practically this would not appear appropriate. The method

can be applied where only subsets of suppliers agree to form such a coalition. The only requirement of the approach is that the number of suppliers should be agreed prior to evaluation of the Shapley value.

There are limitations of the proposed approach and these require exploration through further research. For example, we provide a means of allocating a fair share of the benefits of co-operation, assuming all organisations share information honestly.

The motivation for this framework was generated through action research with the UK aerospace industry [17]. A further cycle of such research would allow ideas to be extended and validated practically. For example, service agreements may be expressed as availabilities requiring a means of combining reliability and supportability targets across suppliers. This would require operational costs, as well as development costs, to be modelled through the trading mechanism.

Given the narrowness of the aerospace industry, supplier behaviour may be influenced by the knowledge that they will co-operate with others in future. Hence there is a need to examine the optimal strategic behaviour of organisations over different horizons and explore the consequences of games that might be played. For example, many political factors may affect collaboration, hence different behavioural scenarios should be examined.

The idea of trading reliability is an interesting issue. This paper has sought to present a rational mechanism for supporting such trade and sharing the financial benefits. The approach is based on established theory and has been motivated by practical problems. The approach would be equally valid for intra as well as inter-organisations.

Acknowledgements

The authors would like to thank the industrial partners involved in the UK DTI/aerospace industry funded research project, reliability enhancement methodology and modelling (REMM), for discussions that have motivated ideas for this paper.

References

- [1] Williams T, Maull R, Ellis B. Demand chain management theory: constraints and development from global aerospace supply webs. *J Oper Manage* 2002;20:691–706.
- [2] <http://www.aecma.org>
- [3] SBAC. The UK aerospace industry: economic impact and competitiveness. London: Society of British Aerospace Companies; 2000.
- [4] Clausing D. Total quality development—a step-by-step guide to world class concurrent engineering. New York: ASME Press; 1994.
- [5] O'Connor PDT. Practical reliability engineering. Great Britain: Wiley; 2001.
- [6] Blishchke WR, Murthy DNP. Reliability: modeling, prediction and optimization. Great Britain: Wiley; 2000.
- [7] Advisory Group of Reliability of Electronic Equipment (AGREE). Reliability of military electronic equipment. Office of the Assistant Secretary of Defense Research and Engineering, Washington DC, 1957.

- [8] Von Alven WH. Reliability engineering. Englewood Cliff, NJ: Prentice-Hall; 1964.
- [9] Engineering design handbook: design for reliability. AMCP 706-196, ADA 027370, 1976.
- [10] Kuo W, Prasad VR. An annotated overview of system reliability optimization. *IEEE Trans Reliab* 2000;49(2):176–87.
- [11] Coit DW, Jin T, Wattanapongsakorn S. *IEEE Trans Reliab* 2004;53(3):369–80.
- [12] Hsieh C, Hsieh Y. Reliability and cost optimization in distributed computing systems. *Comput Oper Res* 2003;30:1103–19.
- [13] Hsieh C. Optimal task allocation and hardware redundancy policies in distributed computing systems. *Eur J Oper Res* 2003;147:430–47.
- [14] Elegbede AOC, Chu C. *IEEE Trans Reliab* 2003;52:106–11.
- [15] Guikema S, Pate-Cornell ME. Component choice for managing risk in engineered systems with generalized risk/cost functions. *Reliab Eng Syst Saf* 2002;78:227–38.
- [16] Johnson W, Quigley J, Walls L. Optimal allocation of reliability tasks to mitigate faults during system development. *IMA J Math Appl Business Ind* 2006; 157–69.
- [17] Varian H. Microeconomic analysis. New York: Norton; 1992.
- [18] Perez-Castrillo D, Wettstein D. Bidding for the surplus: a non-cooperative approach to the Shapley value. *J Econ Theory* 2001;100: 274–94.
- [19] Petrosjan L, Zaccour G. Time-consistent Shapley value allocation of pollutions cost reduction. *J Econ Dyn Control* 2003;27:381–98.
- [20] Mutuswami S, Winter E. Efficient mechanisms for multiple public goods. *J Public Econ* 2004;88:629–44.
- [21] Conklin M, Powaga K, Lipovetsky S. Customer satisfaction analysis: identification of key drivers. *Eur J Oper Res* 2004;154:819–27.
- [22] Sanchez-Soriano J. The pairwise egalitarian solution. *Eur J Oper Res* 2003;150:220–31.
- [23] Rothschild R. On the use of a modified Shapley value to determine the optimal size of a cartel. *J Econ Behav Organ* 2001;45: 37–47.
- [24] Kattuman PA, Green RJ, Bialek JW. Allocating electricity transmission costs through tracing: a game-theoretic rationale. *Oper Res Lett* 2004;32:14–120.
- [25] Tsanakas A, Barnett C. Risk capital allocation and cooperative pricing of insurance liabilities. *Insur: Math Econ* 2003;33:239–54.
- [26] Shapley LS. Contributions to the theory of games. Princeton University Press; 1953.
- [27] Myerson R. Game theory: analysis of conflict. Harvard University Press; 1991.
- [28] Walls L, Quigley J, Marshall J. Modelling to support reliability decisions in product development with applications in the UK aerospace industry. *IEEE Eng Manage* 2006;53:263–74.