[3] A. V. Oppenheim, R. W. Schafer, and T. G. Stockham, Jr., "Nonlinear filtering of multiplied and convolved signals," *Proc. IEEE*, vol. 56, pp. 1264–1291, Aug. 1968.
[4] N. A. Peppers and J. S. Ostrem, "Determination of wave slopes from photographs of the ocean surface: A new approach," *Appl. Opt.*, vol. 17, pp. 3450–3458, Nov. 1978.
[5] M. Born and E. Wolf, *Principles of Optics*, 4th ed. New York: Pergamon, 1970.
[6] G. I. Pokrowski, "Über die helligkeitsverteilung am himmel," *Phys. Z.*, vol. 30, pp. 697–700, 1929.
[7] R. G. Hopkinson, "Measurements of sky luminance distribution at Stockholm," *J. Opt. Soc. Amer.*, vol. 44, pp. 455–459, June 1954.

On Systems with Redundancy and their Inherent Bounds on Achievable Reliability

PETER W. BECKER, SENIOR MEMBER, IEEE

***Abstract*— A system which is supposed to bring some commodity (e.g., electric power, gas, oil, telephone messages) from one point P to another point T, is considered. To avoid total disruption of the services, several parallel connections between P and T may be built, i.e., the designer tries to enhance the system's reliability through use of redundancy. A method, the so-called θ-transformation, is described by which the highest and lowest system reliability achievable may be precisely determined for a redundant configuration. As a by-product we become able to pinpoint the statistical relationships which give rise to the highest and the lowest system reliabilities. By way of a numerical example it is shown that *applying redundancy to a system does not necessarily enhance the system's reliability*; we have thus disproven one of the most cherished beliefs of the reliability community.**

I. Introduction

In a developed society, distribution systems bring commodities like electric power, oil, gas, telephone messages, etc., from one point in space to another. To avoid total disruption of the services, critical sections are quite often duplicated or triplicated, i.e., the designer tries to enhance the system's reliability through the use of redundancy. The increase in systems' reliability obtained through use of redundancy is, however, hard to assess. Assessment of this increase, if any, is the topic of this correspondence. To make the presentation simple let us consider a power system as an example; the assumption results in no loss of generality.

The total loss of power, a blackout, is a major predicament to a modern society. Consider for instance a power system (PS) connecting a power plant P with a town T. If, say, the overhead wires break during a snow storm, the consequences may be appreciable before a team of repair men succeeds in restoring the function of the system. One way of ameliorating the situation would be to connect P and T by two or more different transmission systems (TS's) following different routes; then it is hoped that the same event which causes one TS to fail may spare the other TS's. This concept is known in reliability theory as redundancy. The increase in PS reliability when redundancy is used does, however, depend heavily on the statistical relationship among the states of the redundant TS's (and might be zero as shown in Section V). A method, the so-called θ-transformation, is described here by which the *highest and the lowest PS reliability may be determined for a system configuration using redundant TS's having specified reliabilities*. As a by-product, we become able to pinpoint the statistical relationships (among the states of the TS's) which give rise to the highest and the lowest PS's reliabilities. The computational procedures are similar when the highest PS reliability and the lowest PS reliabilitiy are determined; consequently only the computation of the lowest PS reliability will be dealt with by way of a numerical example in Section IV.

Manuscript received August 11, 1980; revised February 10, 1981.
The author is with the Electronics Laboratory, Technical University of Denmark, Building 344, DK-2800 Lyngby, Denmark.

II. A Lattice Illustrating the Power System's Reliability

A. The Lattice

Assume that P and T are connected by n redundant TS's called $x_1, x_2, \cdots, x_n$ each of which follows a different route through the countryside; *we will then model our PS as a parallel connection of n TS's*.

Let it also be assumed that TS number j, x_j can be in only one of its n_j possible states (which exclude each other and together exhaust all possibilities). Then the set $\bar{x}$ of n TS's, $\bar{x} = (x_1, \cdots, x_n)$, will be in one of ${}_pn$ states

$$ {}_pn = n_1 \times n_2 \times \cdots \times n_j \times \cdots \times n_n. \tag{1}$$

Clearly the state of $\bar{x}$ may be described as a lattice point in an n-dimensional lattice (an example with $n = 3$ is given below). Some of the points illustrate "good" $\bar{x}$-states where the power functions satisfactorily. The remaining lattice points illustrate "bad" $\bar{x}$-states where the PS functions unsatisfactorily. The ${}_pn$-states of the PS exclude each other and together they exhaust all possibilities.

B. f, the n-Variate Probability Density

The PS can be in each of its ${}_pn$-states with some (usually unknown) probability. The set of ${}_pn$ discrete probabilities can be regarded as an n-variate probability density called f; f is not any specific density but just the name of the typical density. The sum of the ${}_pn$-probabilities is unity. f has n marginal probability densities. They are called ${}_1f = {}_1f(x_1), \cdots, {}_jf = {}_jf(x_j), \cdots, {}_nf = {}_nf(x_n)$. ${}_jf$ consists of n_j discrete probabilities having a sum of unity; the set of n_j probabilities shows how likely each of the n_j possible states of transmission system number j are (while disregarding the actual states of the other $(n-1)$ TS's). The product of the n marginals is called f_p. f_p is a special f-density as it is the multivariate density which results when $x_1, \cdots, x_n$ all are statistically independent variables, i.e., the state of one TS being unrelated to the states of the other $(n-1)$ TS's, a situation which is unrealistic in most cases. f_p has the n marginals ${}_jf$, $j = 1, \cdots, n$. The important point here is that whereas we may be fairly ignorant about f, we can make educated guesses about the n marginals ${}_1f, \cdots, {}_nf$. Using the product of the marginals f_p and the θ-transformation we could then in principle (given an immense amount of computer time) compute all possible f concomitant with the n marginals. We will, however, concentrate on finding f with interesting properties, e.g., in Section IV we determine the very f called f_l which gives the exact minimum of a certain probability.

III. The θ-Transformation, a New Factotum of Analysis

A. The Unchanged Marginal Probability Densities

Consider two lattice points: $P_1 = (d_1, \cdots, d_j, \cdots, d_n)$, and $P_2 = (d_1 + e_1, \cdots, d_j + e_j, \cdots, d_n + e_n)$; the two values of x are observed with probabilities θ_1 and θ_2. Let θ be a quantity between zero and the smaller of θ_1 and θ_2. The θ-transformation consists of moving θ units of probability mass from P_1 to a lattice point P_3, while simultaneously moving θ units of probability mass from P_2 to a lattice point P_4; the locations of P_3 and P_4 are determined

0018-9472/81/0500-0387$00.75

as follows. The n coordinates for P_3 are obtained by using some of the P_1 coordinate-values and some of the P_2 coordinate-values; x_1 is d_1 or (d_1+e_1), x_2 is d_2 or $(d_2+e_2),\cdots,x_j$ is d_j or (d_j+e_j). The above mentioned coordinate-values not used for P_3 are used for P_4. As an illustration, consider the case $n=3$; P_3 could be (d_1, d_2, d_3+e_3) and P_4 would then be (d_1+e_1, d_2+e_2, d_3). *The θ-transformation leaves the n marginals unchanged because the probability of a lattice point having $x_j=d_j,{}_jf(d_j)$, or having $x_j=d_j+e_j,{}_jf(d_j+e_j)$, $j=1,\cdots,n$, is unchanged.*

B. A Theorem

We are concerned with joint densities f which consist of ${}_pn$ (1) nonnegative discrete probabilities all of which are assumed to be rational numbers (i.e., they are all multiples of some small quantity q); such joint densities will be called "admissible." The set n marginals is said to be "admissible" because all marginals consist of nonnegative discrete probabilities. f_p, the product of the n members of an admissible set of marginals, clearly is an admissible joint density. An admissible joint density always has an admissible set of n marginals. An "admissible sequence of joint densities" is a sequence of admissible joint densities, ${}^1f$, ${}^2f$, ${}^3f$, etc., each of which is obtained from its predecessor by *one* θ-transformation. When both the values of ${}_pn$ and the small quantity of probability mass q are specified, the number of possible admissible joint densities obviously is bounded. The theorem can now be stated; for a proof see [1]. Other applications of the θ-transformation have been described elsewhere in the literature [4] and [5].

Theorem: Let f be some discrete admissible joint density with the marginals $({}_1f,\cdots,{}_nf)$; then there exists at least one admissible sequence of joint densities which begins with f_p and ends with f, and which is a sequence of finite length.

The importance of the theorem lies in the fact that it ensures us against the following unpleasant possibility. One could imagine that in order to reach an optimum "admissible joint density" by hill climbing from f_p, it might be necessary to pass through nonadmissible joint densities; the theorem tells us that this is not so.

The theorem does not tell the hill climbing designer how to find the multivariate density which has some particular property in largest measure; the theorem only states that the interesting multivariate density can be generated from f_p by a finite number of applications of the θ-transformation, each of which, changes one admissible density to another admissible density.

C. Three Practical Questions

When using the θ-transformation the designer is faced with three questions.

1) How should the lattice points P_1 and P_2 be chosen? In the simple example, Section IV, the points are chosen by inspection so that the desired f-density f_l is obtained after only four θ-transformations. When ${}_pn$ is not too large *all possible (P_1, P_2)-combinations, C may be listed*. Different (P_1, P_2)-values are then selected for θ-transformations using some pseudorandom routine which ensures that each of the C-combinations is tried from time to time. When ${}_pn$ is so large that it becomes unpractical to list all possible (P_1, P_2)-values, a new set of $2n$ coordinates d_j and (d_j+e_j), $j=1,2,\cdots,n$, is selected before each θ-transformation using a random number routine.

2) How should the value of θ be selected? Let the four probability masses at points P_1, P_2, P_3, and P_4 immediately before θ-transformation number k be ${}_k\theta_1$, ${}_k\theta_2$, ${}_k\theta_3$ and ${}_k\theta_4$. With the simple example in Section IV [4] we are in a typical linear programming situation; θ should consequently be the smaller of ${}_k\theta_1$ and ${}_k\theta_2$ if it turns out to be beneficial to move θ-probability mass from P_1 and P_2 to P_3 and P_4; if the effect turns out to be detrimental instead the smaller of ${}_k\theta_3$ and ${}_k\theta_4$ should be added to P_1 and P_2 and subtracted from P_3 and P_4. When we are not in a linear programming situation e.g., [5] only small quantities of probability mass should at first be moved at each θ-transformation; if the effect of the transformation is beneficial the quantity could be increased.

3) What determines that the optimum has been reached and the algorithm is to be terminated? In the simple case where all C possible (P_1, P_2)-combinations can be listed, the answer is simple: an optimum has been reached when none of the C possible θ-transformations result in an improvement. When it is unpractical to list all possible (P_1, P_2)-combinations, the algorithm is terminated when no improvement has been recorded during the last L θ-transformations. L is a number chosen in advance by the designer. The problem is basically one of terminating a hill climbing procedure.

IV. An Illustrative Example

Let it be assumed that the town T obtains its power solely from the power plant P. Three TS's connect P with T, one called $A1$, runs through the valley A, the second $B2$ through valley B, and the third $C3$ through valley C (i.e., $n=3$, the lattice is three-dimensional). At the end of the winter—during which repairs are impossible—the first TS is in one of two states: state a where it functions, and state $\bar{a}$ where it has failed. Likewise the second and the third TS can each be in one of two states b and $\bar{b}$, c and $\bar{c}$, i.e., $n_1=n_2=n_3=2$, and ${}_pn=8$. From past experience we feel reasonably sure that the three marginal probability densities have the following values:

$$ {}_1f(a)=0.9, \qquad {}_1f(\bar{a})=0.1; $$
$$ {}_2f(b)=0.8, \qquad {}_2f(\bar{b})=0.2; $$
$$ {}_3f(c)=0.7, \qquad {}_3f(\bar{c})=0.3. $$

To supply T with power at least two of the three transmission systems must function, i.e., we have four "good" states (a,b,c); $(\bar{a},b,c)$; $(a,\bar{b},c)$; and $(a,b,\bar{c})$. The question now presents itself: given the three marginals *what is then the probability of uninterrupted power supply to T during the winter—assuming the most adverse statistical relationships*. The answer is obtained by first generating ${}_pf={}_1f\cdot{}_2f\cdot{}_3f$. We then perform some θ-transformations with the objective of minimizing R_T,

$$ R_T=f(a,b,c)+f(\bar{a},b,c)+f(a,\bar{b},c)+f(a,b,\bar{c}) $$

for $f={}_pf$, $R_T=0.504+0.056+0.126+0.216=0.902$. As shown in Fig. 1 and Table I, the absolute minimum (not some lower bound), $\min\{R_T\}=0.7$ can be obtained after as little as four θ-transformations, i.e., with the specified marginals; T is supplied with power uninterruptedly with a probability of at least 70 percent. Because the dimensionality is so low $n=3$, the results from Fig. 1 are readily illustrated by a Venn diagram (see Figs. 2 and 3); the diagrams show how $f(abc)=0.7$, $f(ab\bar{c})=0.2$, and $f(\bar{a}b\bar{c})=0.1$ will add up to unity when the five other Boolian functions have probability zero.

In the example illustrated in Fig. 1 we have shown how the θ-transformation may be used to pinpoint the "least favorable" failure pattern achievable by parallel redundancy of the three TS's. At this point the reader may rightfully object that the worst-case situation is unnecessarily pessimistic. It is for instance not realistic to assume that the joint probabilities of the three states $\bar{a}bc$, $a\bar{b}c$, and $ab\bar{c}$ all be zero, i.e., that the very situations which should justify our use of parallel redundancy are impossible. Therefore, when θ-transformations are used to determine worst-case system reliabilities, the designer should abstain from transformations which result in failure patterns which clearly are unrealistic. This may be readily achieved by *putting upper and lower bounds on the joint probabilities* associated with the ${}_pn$ lattice points.

TABLE I
FOUR θ-TRANSFORMATIONS WILL CHANGE $_pf$ TO AN f-DENSITY FOR WHICH R_T IS MINIMIZED

P_1	P_3	P_2	P_4		$(f(P_1);$	$f(P_3);$	$f(P_2);$	$f(P_4))$	=>	$(f(P_1);$	$f(P_3);$	$f(P_2);$	$f(P_4))$
(a,b,c),	$(a,b,\bar{c})$,	$(\bar{a},b,\bar{c})$,	$(\bar{a},b,c)$	:	(0.504;	0.216;	0.024;	0.056)	=>	(0.560;	0.160;	0.080;	0.000)
$(a,\bar{b},c)$,	$(a,\bar{b},\bar{c})$,	$(\bar{a},\bar{b},\bar{c})$,	$(\bar{a},\bar{b},c)$	:	(0.126;	0.054;	0.006;	0.014)	=>	(0.140;	0.040;	0.020;	0.000)
$(a,b,\bar{c})$,	$(\bar{a},b,\bar{c})$,	$(\bar{a},\bar{b},\bar{c})$,	$(a,\bar{b},\bar{c})$	:	(0.160;	0.080;	0.020;	0.040)	=>	(0.140;	0.100;	0.000;	0.060)
(a,b,c),	$(a,b,\bar{c})$,	$(a,\bar{b},\bar{c})$,	$(a,\bar{b},c)$	:	(0.560;	0.140;	0.060;	0.140)	=>	(0.700;	0.000;	0.200;	0.000)

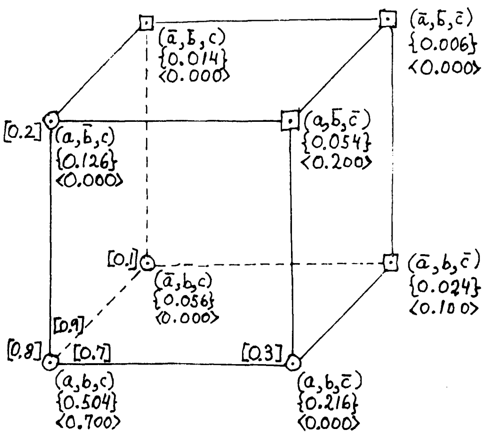

Fig. 1. Three-dimensional lattice mentioned in example, Section IV. The figures in $(\cdot)$, $[\cdot]$, $\{\cdot\}$, and $\langle\cdot\rangle$ are coordinates of eight corners, six marginal density values, eight values of f_p, and the eight f-values, respectively, obtained after four θ-transformations which yield $\min\{R_T\} = 0.7$. Four "good" states are indicated by circles, four "bad" states by squares. $R_T = 0.7$ is obtained only when five joint densities are zero—an unnecessarily pessimistic situation. When joint densities are bounded, more realistic R_T-values will be obtained. We notice that $f(\bar{a}, b, c) = f(a, \bar{b}, c) = f(\bar{a}, \bar{b}, c) = 0$; consequently *the worst case reliability is not enhanced by adding* $C1$ because when $C1$ functions, both $A1$ and $B2$ already function, and under our assumptions, this is sufficient to bring power from P to T.

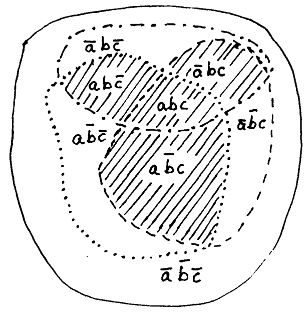

Fig. 2. Venn diagram illustrating the eight possible states of the PS: abc, $\bar{a}bc$, etc. Points enclosed by the $\cdots$ curve illustrate the a-state. Points enclosed by the $-\cdot-$ curve illustrate the b-state. Points enclosed by the --- curve illustrate the c-state. Cross-hatched area illustrates "good" states where PS functions satisfactorily. Venn diagrams are only practical for $n \leq 4$ [4].

The highest possible value of R_T, $\max\{R_T\}$, may likewise be determined by modifying $_pf$ through a series of θ-transformations.

V. THE INCREASE IN RELIABILITY DUE TO REDUNDANCY

The following question now presents itself. What is the probability r_T, of an uninterrupted power supply to T during the winter if only the two TS's $A1$ and $B2$ are available and both should function? Or to pose the question differently: how much is gained by adding $C3$? As before, we assume that

$$_1f(a) = 0.9 \qquad _1f(\bar{a}) = 0.1 \qquad _2f(b) = 0.8 \qquad _2f(\bar{b}) = 0.2.$$

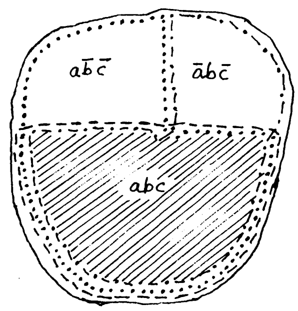

Fig. 3. Venn diagram illustrating worst possible statistical relationships where $f(abc) = 0.7$, $f(a\bar{b}\bar{c}) = 0.2$, and $f(\bar{a}b\bar{c}) = 0.1$ and R_T is minimized. Cross-hatched area illustrates $\min\{R_T\} = f(abc) = 0.7$. Notice that five of the areas from Fig. 2 have vanished. By trial and error we verify the minimum value of R_T, i.e., no θ-transformation can bring R_T-value below 0.7.

Clearly $r_T = f(a, b)$. Assuming statistical independence, we find that

$$f(a, b) = 0.72 \qquad f(a, \bar{b}) = 0.18$$
$$f(\bar{a}, b) = 0.08 \qquad f(\bar{a}, \bar{b}) = 0.02.$$

By adding $\theta = 0.02$ to $f(a, \bar{b}) = 0.18$ and $f(\bar{a}, b) = 0.08$ while subtracting 0.02 from $f(a, b) = 0.72$ and $f(\bar{a}, \bar{b}) = 0.02$, we obtain the smallest possible value of $f(a, b) = \min\{r_T\} = 0.7$. The reader will recall from Section IV that also $\min\{R_T\} = 0.7$. In other words assuming the most adverse statistical relationships among the states of the TS's, *the probability of uninterrupted power supply to T is* 70 *percent as well when we use only* $A1$ *and* $B2$ *as when we add* $C3$ *to* $A1$ *and* $B2$! *With this numerical example we have refuted the widely held belief that use of redundancy always enhances reliability.* (The reason why the reliability does not increase when $C1$ is added is of course that the numbers were so chosen that $\min\{f(a, b, c)\} = \min\{f(a, b)\}$.) In practical cases where the statistical relationships are less adverse we will of course gain by adding $C3$; e.g., assuming statistical independence $r_T = 0.72$ and $R_t = 0.902$ (as calculated in Section IV). If the PS consists of only $A1$ and $C3$, the minimum reliability $\min\{f(a, c)\}$ is 60 percent; if the PS consists of only $B2$ and $C3$, the minimum reliability $\min\{f(b, c)\}$ is 50 percent. In the two cases the minimum reliability is raised to $\min\{R_T\} = 70$ percent by adding, respectively, $B2$ or $A1$; i.e., in these two cases the minimum reliability was enhanced by adding a third TS.

VI. CONCLUSION

Quite often a designer is faced with a problem the solution of which requires knowledge of some n-variable probability density function f. We are concerned with the situation where the designer has access only to the n marginals $_1f, _2f, \cdots, _nf$ rather than f. To solve this problem the designer must answer the fundamental question: what n-variable densities are concomitant with a set of n marginals? This problem can be solved only with the help of a new factotum of analysis, the θ-transformation. In this paper the θ-transformation is used to compute the minimum reliability for a redundant system.

By way of a numerical example we show that applying redundancy to a system does not necessarily increase its reliability.

The more traditional methods of computing reliability figures have been discussed elsewhere in the literature [3], [7], [8], [9] and [10].

VII. Acknowledgment

The author wishes to thank Dr. John Aasted for many helpful discussions.

References

[1] P. W. Becker, "An Introduction to the Design of Pattern Recognition Devices." New York: Springer, 1972, addendum.

[2] ——, "Recognition of Patterns," third revised and enlarged edition. New York: Springer, 1978.

[3] P. W. Becker and F. Jensen, "Design of Systems and Circuits for Maximum Reliability or Maximum Production Yield." New York: McGraw-Hill, 1977.

[4] P. W. Becker, "The highest and lowest reliability achievable with redundancy," *IEEE Trans. Rel.*, vol. R-36, no. 3, pp. 209–213, Aug. 1977.

[5] ——, "What multivariate probability densities are possible with a specified set of marginals?" in *Proc. Second Int. Joint Conf. Pattern Recognition*, 1974, IEEE Catalog No. 74CHO885-4C, pp. 368–371.

[6] J. Venn, "The Logic of Chance," originally published in 1888, 4th edition. New York: Chelsea, 1962, pp. 116–117.

[7] A. E. Green and A. J. Bourne, "Reliability Technology." London: Wiley-Interscience, 1972.

[8] C. O. Smith, "Introduction to Reliability in Design." New York: McGraw-Hill, 1976.

[9] D. B. Brown, "Systems Analysis and Design for Safety." Englewood Cliffs, NJ: Prentice-Hall, 1976.

[10] K. C. Kapur and L. R. Lamberson, "Reliability in Engineering Design." New York: Wiley, 1977.