A Risk Hypothesis and Risk Measures for Throughput Capacity in Systems

James Bradley
Department of Computer Science
University of Calgary,
Calgary, Alberta, Canada

Abstract A basic risk hypothesis, expressed as a risk equation, for system throughput capacity (I), and governing all non-growth, non-evolving, agent-directed systems, is proposed and derived. The equation relates throughput capacity, resources and risk relative to the system environment, for efficient environments. The risk equation may be combined with, and thus enhances, a resource-sharing equation relating throughput capacity, resources and the time required to execute complex coordinated sharing procedures, an equation derived in an earlier paper.

The *basic risk equation* shows how expected I increases [decreases] linearly with positive [negative] risk of loss of I in efficient environments. The conventional *standard deviation risk measure with respect to the mean,* from financial systems, may be used. A proposed, new, usually equivalent measure, called *the mean-expected loss risk measure with respect to the hazard-free case*, is shown to be more appropriate for systems in general. The concept of an *efficient system environment* is also proposed.

All quantities used in the equation are precisely defined and their units specified. The equation reduces to a numerical expression, and can be subjected to experimental test. The equation clarifies and quantifies basic principles, enabling designers and operators of systems to reason correctly about system risk.

Key words *resources, risk, risk measure, system, throughput capacity.*

Introduction

In recent years, complex systems are increasingly having to cope with risk in their operating environment. Some systems are exposed to risk whose source is intrinsic, for example risk of deadlock [9, 11, 12], or risk of inter-process interference [13, 18], in computer operating systems [21]. Other systems deal with risk whose source is external to the system, such as an on-board aircraft database system using a terrian database to manage the risk of colliding with a mountain, or merely the risk of loss of a database's integrity [6].

Such systems are all human-agent directed systems. Unfortunately, the responsible agents may not always be able to reason clearly about the relationship between system throughput capacity, resources, resource sharing procedures [2,3], environment risk, and risk-reducing resources and procedures [2]. As a result, serious mistakes in system design and system operation can occur, and particularly with systems whose throughput is human, not infrequently throughput loss means loss of human life.

In this paper we present a hypothesis, expressed as a risk equation, that relates system throughput capacity to resources and risk, and which governs all non-growth, non-evolving, agent-directed systems. Agent-directed systems are systems that are designed, constructed and operated by human agents; thus we exclude naturally occurring systems such as meteorological systems and living organisms. Non-growth, non-evolving systems are systems where the output is not fed back to alter, usually grow, the resources of the system; we therefore exclude, for example, all living organisms, whose output is used to both grow and reproduce. This still leaves a very large class of systems to which the proposed equation applies. The equation applies as well to a chemical plant producing a polymer, as it does to a transportation system moving people or

freight from input locations to output locations, or to a computer system converting input data to useful information, or retrieving data from a file or database [5]. The research leading to the risk hypothesis was inspired by phenomena in complex computer systems, such as operating systems [21] and database systems [5], and it is expected that it is in the computer systems arena that the equations will throw the most light. Nevertheless the risk equation expresses a unity of principle in a great diversity of system phenomena.

Although the risk equation can be used for precise numerical analysis, it is important to grasp that its main virtue is that it clarifies and quantifies basic system principles. It thus enables clarity of reasoning about specific costs and benefits of different system approaches and different methods of system operation, and so enables avoidance of serious error.

In this paper every attempt has been made to ensure that concepts are clearly defined, and that symbols appearing in the risk equation denote quantities that are measurable and expressible in clearly defined units.

1.0 Overview of the risk hypothesis and risk equation

A system is considered to be any entity that converts inputs to outputs, and which may be composed of subsystems, connected either in series or in parallel, each of which is also a system. We also consider a system to be an entity that functions under the direction of one or more agents, usually human. An agent employs essential resources R, often portable, in an environment E, enabling the system to convert a set of inputs N, informational or material or both, to a set of outputs U per unit time, so that U is the system throughput. A system environment will be defined later.

If the maximum value for U is I, then I is said to be the throughput capacity of the system, measured as the units of output per unit time when the system is operating at maximum capacity. In general, if the system is operating at a fraction f of throughput capacity I, then the throughput is U = fI.

In some cases there will be more than one type of throughput product. In such a case, we can declare throughput capacity of one of them to be the main throughput product capacity I, in units of product per time period. We can then treat each other product throughput capacity as a byproduct capacity that is a function of I, in most cases a linear function. Accordingly, in this paper, we deal only with I, considered to be either the only or the main throughput product capacity.

1.1 Summary of equations

A reference summary for the equation relating throughput capacity I to resources and risk is given below. First, however there is a summary of another system equation, the resource sharing equation, to be published elsewhere [2, 3] for system throughput capacity I, when subsystem threads share a common resource of the system [1, 19]. The risk equation is built on this resource equation, but only on that elementary subset of the resource equation (I = KR) for the case when there is no resource sharing. The resource-sharing equation can be used to evaluate system performance in a wide variety of contexts [14].

1. The resource sharing equation:

$$I = KR(1 - T_S/T)[1 + sF_T(T_S)] \quad (1A)$$

$$I = KR \quad (1Aa) \text{ no resource sharing}$$

Independent variables under the control of the agent are R and T_S. K is a constant and R measures resources available to the system, alterable only in valid units (basic harmonic units) of the type comprising R. T is a constant and measures the time for which I is computed. T_S has units of time per basic resource unit, and measures the execution time, and thus resource-sharing complexity, of a *normally complex, coordinated, resource-sharing procedure* for sharing resources within R. $F_T(T_S)$ is a growth function that has value 0 when T_S is zero and increases at an decreasing rate with increasing T_S to saturate at 1.0. The constant s is the *available sharing-enhancement potential*, where sR is the effective increase in R due to execution of a coordinated resource sharing procedure sufficiently comprehensive to saturate $F_T(T_S)$ at 1.0. T_S/T is a measure of the resource capacity diverted from normal operations to carrying out the resource sharing procedure, and is thus a measure of the negative impact of the sharing procedure on throughput capacity I. When there is no resource sharing, T_S and $F_T(T_S)$ are zero, and the expression reduces to I = KR. There is a value for T_S at which I is maximized, found by solving $dI/dT_S = 0$.

2. The basic risk equation

$$I = R[K + (b_{pb} - 1)r(E)] = RK + Rb_{pb}r(E) - Rr(E) = R[K + b_p r(E)] \quad (1Ba)$$ /* for positive risk */

$$I = R[K + (b_{nb} - 1)r(E)] = RK + Rb_{nb}r(E) - Rr(E) = R[K - b_n r(E)] \quad (1Bb)$$ /* for negative risk */

Independent variables under the control of the agent are R and r(E). I is now a mean or expected throughput capacity. r(E) is risk per unit R; the risk r(E) is risk of loss of throughput capacity, and is a function of the environment E relative to the system; the risk can be varied for constant R, by varying the environment E relative to the system; b_{pb}, b_p, b_{nb}, and b_n are constants. Positive risk is risk it can pay to take on average; negative risk is risk it can not pay to take on average. Allowed risk measures are statistical measures, for example standard deviation (SD) risk with respect to the mean I, or a new proposed measure, more suitable for systems in general, called mean expected loss (MEL) risk with respect to the hazard-free I. Positive and negative risk, using statistical risk measures, cannot occur together. The environments allowable to the system must be *efficient environments*, as explained later, otherwise b_{pb}, b_p, b_{nb}, and b_n are not constants. $Rb_{pb}r(E)$ is the gross extra throughput capacity generated by exposure to the risk if the hazard risked does not occur; Rr(E) is the average loss in throughput capacity due to the hazard actually occuring. The difference is the net extra throughput capacity added (if positive risk), on average, by exposure to the risk.

1.2 Axiomatic foundation

In deriving the risk equation we employ an axiomatic proposition: Where there is no coordinated sharing of resources by a system among different threads, then for valid changes to R

$$I = KR \quad (1Aa)$$

must hold true. This equation is implicit in the resource sharing equation (1A) above. Note that we cannot prove this assertion from any more fundamental proposition. It seems to be implicit in the nature of things and so we take it as axiomatic. Space does not permit a discussion of this proposition here; a thorough discussion examining its implications more closely can be found in [2, 3].

2.0 Risk and Resources

Risk of loss of anything of value is normally run to secure some gain in a value, and is therefore of primary interest to humans. Consequently, the concept of risk has been thoroughly studied in the financial industry. A major result is that we have a sound statistical measure of risk, at least in a financial context [7, 8], that is, *the standard deviation risk measure with respect to the mean*, primarily due to Markowitz [17]. Risk is, however, known to be a slippery concept, and interested readers unfamiliar with the statistical concept of risk may need to become familiar with some of the literature on the subject. When systems in general are considered, we find that a new measure, somewhat different from the conventional standard deviation risk measure used in finance, but still often equivalent to it, is very useful, as will be developed shortly,

One key to understanding risk is grasping the distinction between exposure to *the certainty* of future loss and exposure to merely *the possibility* of future loss. In both cases there is exposure to future loss. But *only exposure to the possibility of future loss is a source of risk*. In some cases there is exposure to future loss consisting of both exposure to the certainty of future loss and to the possibility of an additional future loss. For example, if exposure to a future loss of 20 where a future loss of 16 is certain and a future loss of 4 is additionally possible, there is only risk of loss of 4. Thus there has to be a variability aspect to risk, since risk has to do with only a possibility of future loss: in some future periods the loss will occur fully, in others partly and in others not at all. Failure to grasp these distinctions has caused many fruitless arguments and debates about correct measures of risk

In this paper we are primarily concerned with risk of loss of system throughput capacity I, the system being considered to be financial only if throughput capacity I is in actual currency entities, such as dollars, and not in other physical units valued in currency units. But first we need to consider risk measures, both the conventional standard deviation measure from financial systems, and a new proposed mean-expected-loss measure for systems in general.

2.1 Risk measures: SD-risk and MEL-risk

In general, risk of loss of throughput capacity has two components, namely the probability of a hazard occurring and the size of the loss in throughput, with respect to some standard level, should the hazard occur. However, in a system situation, where there is exposure to possible loss with respect to some standard level of throughput capacity, there will often be exposure to possible gains in addition, depending on the standard level used. An accurate risk measure must therefore combine these different aspects of risk, but must not include any measure of certain future loss.

Suppose the system is exposed to unpredictable losses and gains in throughput capacity, that the statistics of these fluctuations are constant (or stationary [4], in statistical terminology), and that over a long enough period of time to be representative of these statistics, the mean, and thus expected, throughput capacity is I_m, and in n fully representative time periods the actual capacity values are:

$I_m - L_1$, $I_m - L_2$, ... $I_m - L_i$, $I_m + G_1$, $I_m + G_2$, ... or $I_m + G_j$

where L_1, L_2, ... are deviations downward (losses) from the average I_m, and G_1, G_2, ... are deviations upward (gains) from I_m, with $n = i + j$, so that

$$(L_2 + L_2 + ...L_i) = (G_1 + G_2 + ...G_j)$$

The same throughput capacity deviations can then be expected to occur in the future in unpredictable order, all equally likely.

[Note that the above quantities should be interpreted as follows: Suppose n = 10, i = 6 and j = 4, and the losses L_1 .. L_i are 30, 20, 20 10, 10, 10, and the gains G_1 .. G_j are 40, 20, 20,

and 20. Then imagine an urn containing 10 balls: 6 red balls each with one of the losses marked, and 4 green balls each with one of the gains marked, and thus the distribution of future (and past) losses and gains; to simulate what will happen in the next time period, select a ball randomly from the urn, the number on the ball giving the amount; to simulate for the subsequent time period the ball removed must first be replaced, i.e. we must use selection with replacement. The sum of the numbers on green balls divided by 10 is the average or expected gain, equal to the red ball sum divided by 10, the expected or average loss. Nevertheless, suppose the period involved was 1 week; then in some weeks there would be a gain, in others a loss, and in few weeks a severe loss or a large gain. Thus the observed result is merely a sequence of unpredictable losses and gains per time period, with respect to the mean, that is, I will fluctuate from one period to the next. In terms of probabilities, the risk is due to probabilities 0.1. 0.2, 0.3 of losses 30, 20 and 10 respectively, and probabilities 0.1, 0.3 of gains 40 and 20, all with repeat to the mean.]

The expected or average I_m actually rarely occurs if at all. In reality all we have is the unpredictable sequence of losses (L_1, or L_2, ...) and gains (G_1, or G_2, ...) with respect to an average or expected throughput capacity I_m in a given time unit, and it is such losses and gains with respect to expected throughput I_m that must be used in the measure of risk of loss of throughput.

For a meaningful measure of risk there are now two choices, the traditional standard deviation measure, and a new measure that in many cases is more suitable for systems in general.

Choice 1. Take the standard deviation of the deviations (L_1, L_2, ... G_1, G_2, ...) from the mean throughput capacity I_m, as *a standard deviation measure of possible loss with respect to (or down from) the mean I_m*, to give the *Standard Deviation (SD) risk* measure.

If we use twice the standard deviation we have an even stronger risk measure, the *2-Standard Deviations (2-SD)* risk measure.

***Interpretation*:** A SD-risk of s means that in the next time unit, there is a 50% chance or possibility of a loss down from the expected I_m, and a 34.1% chance of a loss between 0 and s down from the expected I_m, and a 15.9% chance of a loss > s. In addition there a 47.7% chance of a loss between 0 and 2s with respect to the mean I_m. This implies, that there is a 13.6% chance of a loss between s and 2s, and a 2.3% chance of a loss >2s, both losses with reference to the mean throughput capacity I_m. In specifying an SD risk, *we must both specify the standard deviation of the variations in throughput and specify with respect to what standard level of throughput*. A 2-SD-risk of 2s means that in the next time unit, there is a 50% chance of a loss with respect to the expected I_m, and 47.7% chance of a loss between 0 and 2s and a 2.3% chance of a loss > 2s.

The percentages used are from a normal distribution function table, and assume that losses and gains in each time unit are distributed normally. The numbers needed are different if the distribution departs from normal.

[The SD-risk measure is the one widely used in finance, particularly for stock and bond portfolio management, for which it is both correct and adequate [8, 15, 17], since stock and bond prices follow close to a random walk, which gives rise to a near-normal distribution of price changes [4]. Notice that where there is exposure to future loss, where the future loss includes a certain loss and a possible loss, the SD-risk measure selects out only the possible loss, that is, the true risk. For example, suppose a system where ideally I = 400 if there were no future loss exposure, but where actually the system has exposure to a future loss in I whose mean is 100 and whose standard deviation is 14, where the least loss is always greater than 70. That means a certain loss of 70 plus a loss whose mean is 30 that can be as small as 0 and as large as about 60, with a standard deviation of 14. That means a certain loss of 70 plus a standard deviation of 14

about the mean of the loss variations of 30, that is, a certain loss of 70 -- plus an SD-risk of 14 with respect to a mean of 300.]

To deal with the problems and possibilities in arbitrary systems, where the distribution of gains and losses is not anything like normal, an additional and complementary risk measure is very useful. This is the MEL-risk measure defined below. [The reader who is expert in financial risk analysis using the SD-risk measure may be want to immediately dismiss this additional risk measure as nothing but an intellectual crutch; the author asks such readers to suspend judgement until after studying the measure. It is particularly useful where the variations in throughput depart significantly from the normal distribution, and makes practical sense, even in a financial context.] The author therefore proposes:

Choice 2. Suppose that for a system exposed to risk, there is at least one hazard-free time period, in which, by chance, the hazard risked does not occur, and where the gain with respect to the mean throughput capacity I_m is G_b in this hazard free time period, and where a gain exceeding G_b is thus not possible (but a gain under G_b is very possible). Thus, in the best case scenario, the total hazard-free throughput capacity is $I_m + G_b$. Then all other throughput capacities $I_m - L_1$, $I_m - L_2$, ... , $I_m + G_1$, $I_m + G_2$, …, each in a time period where the hazard does occur in varying degrees of intensity, may be considered as exhibiting losses, or loss deviations, $G_b + L_1$, …$G_b - G_1$ … down from, or with respect to, the value of I in the hazard-free time period. We may use the mean of these loss deviations (down) from I for the hazard-free time period as a measure of the risk, that is, a measure of expected losses in the future with respect to the throughput capacity for a hazard-free time period, that is, the *Mean Expected Loss (MEL) with respect to, or down from, the throughput capacity in any hazard-free time period,* or *MEL-risk.*

Note that in specifying a MEL-risk, *we must both specify the mean deviation, and specify with respect to what level.*

Interpretation. An MEL-risk of L means that the average loss, with respect to the (best) value for I in a time period where the hazard does not occur, is exactly L. However, there are two extreme possibilities with regard to what is to be considered as I for a hazard-free time period.

(a) Natural, or explicit, hazard-free case In this case, we can find a naturally-occurring, best-case, hazard-free throughput capacity $I_m + G_b$ that cannot be exceeded for the value of R. This hazard free throughput capacity will thus occur in a time period when all goes well and no hazard occurs. Such a time period is certain to occur sometime. Thus, over a long period of time, there will be a distribution of n gains and losses about the mean for the time period; and in at least one of the n time periods there will occur a gain deviation G_b (with respect to the mean) up to the hazard-free case. However, no gain deviation exceeding G_b will ever occur -- the concept of a hazard free throughput capacity level implies (1) that no variation in throughput capacity can occur leading to a throughput capacity above the hazard-free level level, and (2) that the only variations in throughput capacity that can occur must lead to a throughput capacity at or below that hazard-free level. This ensures that *all possible variation is included in, and certain loss is excluded from, the MEL-risk measure*

(b) Artificial or implicit hazard-free case In this case, the values in each time period fluctuate about the mean I_m, and distribution of the per-period deviations from the mean follows some reasonably bell-shaped distribution, where large but usually improbable gain deviations from mean I_m do sometimes occur, and where no explicit hazard-free throughput capacity can be determined. In such a case, we may define an *artificial hazard-free case* for throughput capacity $I_m + G_b$, by defining an imaginary hazard-free time period where the

gain G_b is 2 standard deviations up from the mean. For this case, we then define the MEL-risk as the mean expected loss with respect to $I_m + G_b$ for this imaginary hazard-free time period, with the throughput capacity in each time period being considered as exhibiting a loss with respect to the hazard-free $I_m + G_b$, except for the rare time period with a throughput capacity value lying beyond 2 standard deviations above the mean, which is taken as a negative loss (a gain) with respect to the hazard-free I.

In both cases MEL-risk can therefore be quite simply viewed as the hazard-free deviation, either natural or artificial, up from the mean, but also equal to the average loss to be expected in the future with respect to, or down from, throughput capacity I for the hazard-free time period (real or artificial).

2.2 Comparison of SD-risk and MEL-risk measures

When there is no natural hazard-free case, and the deviations from the mean follow a bell shaped distribution, a common situation in finance and many physical systems, the SD-risk measure and the MEL-risk measure are equivalent, since MEL-risk is exactly twice SD-risk or equal to 2-SD-risk, although they each are with respect to different standard levels.

If there is an actual hazard-free throughput capacity, with only deviations down from the hazard-free I, the distribution of deviations about the mean will tend to be skewed on the left (or truncated on the right), since upward fluctuations are blocked by the hazard-free I that cannot be exceeded, and yet very large, if rare, downward deviations from the mean of I can occur. In such a situation there seems to be no simple equivalence between the SD and MEL-risk measures. But the MEL-risk measure has an obvious advantage here, since it is precisely equal to the average loss that can be expected with respect to the hazard-free or best-case situation, and since the SD-risk measure would now be applied to a skewed distribution, something for which it is really not designed, and means very little. This situation can be expected to occur frequently in systems in general, but rarely in financial systems involving stocks and bonds (but it does occur in insurance related systems); hence the need for the MEL-risk measure.

Notice that where there is exposure to future loss, where the future loss includes a certain loss and a possible loss, then the MEL-risk measure also selects out only the possible loss, that is, the true risk. Suppose again a system with ideally I = 400 if no future loss exposure, but which is actually exposed to a future loss in I whose mean is 100 and whose standard deviation is 14 where the least loss is always greater than about 70. Once more this means a certain loss of 70, plus a loss whose mean is 30 that can be as small as about 0 and as large as about 60, with a standard deviation of 14. Applying the MEL-risk measure, we have an artificial hazard-free throughput capacity 2 standard deviations up from the mean, that is at I = 328. Thus there is a certain loss of 70, or an almost certain loss of 72, plus a mean loss of 28, the MEL-risk, down from the hazard-free level of 328, giving a mean throughput capacity I_m of 328 – 28 = 300. Occasionally a fluctuation up from the mean of 300 may reach I = 330, just slightly above the supposedly (but artificial) hazard-free best case of 328, but this may be simply considered as a "loss" of –2.

If, on the other hand, there is only exposure to possible loss, and never to certain loss, there is no need for an artificial hazard-free throughput capacity. In the example above, if for I = 400 with no loss exposure, there is exposure only to a possible loss whose mean is 30, this means the average loss down from the true hazard free I of 400 is 30, the MEL-risk, for a mean I_m = 370.

The true advantage of the MEL-risk measure is that it continually forces the user to think in terms of the distinction between certain loss and possible loss (so arguably it *is* merely an intellectual crutch), whereas the SD-risk does not. As a result, although SD-risk is absolutely

correct as a risk measure whre there is a normal (or near-normal distribution), in the author's opinion it can lead to both confusion in thinking about risk and to obscuring some of the fundamental attributes of risk with aspects of systems in general.

2.3 System environments, non-limiting resources and risk

As pointed out in an earlier paper, if resources R of a system can be divided into R_1 and R_2 for subsystems in series, then R_2 can be operating at full capacity and R_1 not at full capacity, so that R_1 is the non-limiting resource (for the non-limiting subsystem). If R_1 is non limiting then throughput capacity I for the total system obeys $I = K_1R_2$ for increases in R_2 up to the point where R_1 has started to operate at full capacity. However, if R_1 is so far away from full capacity that the normal range of changes upwards in R_2 will never cause R_1 to operate at full capacity, then we can neglect R_1 in figuring the effect of changes in system resources on throughput capacity I. In such a case *we can regard very non-limiting resources R_1 as part of the system environment.*

Every system operates somewhere, and it is always possible to consider the whole Earth as comprising the system. However, most of the Earth, often even the building housing, or geographic area containing, the system, will behave as a non-limiting serial subsystem, whose resources can consequently be neglected from the system. Thus *we can define the system environment as a collection resources that are non-limiting, very far from full capacity, but necessary for operation of the system*. Thus what is considered part of the system and what is part of the environment will be somewhat arbitrary. As an example, consider a computer system doing infrequent information retrieval from a small <1.0 Mbyte file on a 3 Gigabyte hard disk. Because the disk is so far from being used at full capacity it would be legitimate to regard it as part of the system environment, although most computer systems specialists would probably include it as part of the system. As another example, if a double-track railroad is currently being used infrequently by only a few trains each with exclusive control of the track, then the track is so far from being used at full capacity that it would be legitimate to consider it as part of the environment.

For a given system, risk depends on the system relative to the environment, which we denote by r(E) in this paper. If we move the system from one environment to another, the risk may change, or if we keep it in the same environment but alter the layout of the system without changing each subsystem's functionality, or if we change the subsystem functionality, risk may also change. Thus if we have different risks $r(E_1)$ and $r(E_2)$ we may either have two different environments E_1 and E_2, or we may have the same environment but a different environment relative to the system. In the following discussion *an environment E_n always means a specific environment relative to the system.* Also, Equations 1Ba and 1Bb is valid only for systems operating in *efficient environments*, a concept to be defined later.

3. The influence of risk on throughput capacity

We now derive the basic risk equation (1B), and show that, for a risk measure r(E) -- the risk of a system in an efficient environment E relative to the system, the expected throughput capacity I is given by:

$$I = R[K + (b_{pb} - 1)r(E)] = RK + Rb_{pb}r(E) - Rr(E) \quad (1Ba)$$
$$= R(K + b_p r(E))$$

where the independent variables are R and r(E), and where the risk is positive, that is, it is risk it can pay to take (on average), and where, it is crucial to realize, risk measure r(E) is a measure of possible throughput capacity loss, per time period, per (basic harmonic resource) unit R, that is,

measured as a fraction of R, and not as a fraction of throughput I. The risk r(E) is SD-risk or MEL-risk per unit R.

We also show that for negative risk:

$$I = R[K + (b_{nb} - 1)r(E)] = RK + Rb_{nb}r(E) - Rr(E) \quad (1Bb)$$
$$= R(K - b_n r(E))$$

Expression (1Ba) is a statement of the following rule: throughput capacity for a given system resource level R increases linearly with R (for valid R changes, as specified in Equation 1A) and also *linearly* with increases in the risk measure r(E) of distinct system environments relative to the system, provided the risk is positive and the environments are efficient. If the risk is negative, throughput capacity will decrease linearly with increasing r(E) (expression (1Bb)). If the environments are not efficient, the equation's constants b_{pb} , b_p, b_{nb}, and b_n will no longer be constant.

3.1 Efficient environments and linear relationship between throughput and risk

Consider now two environments E_1 and E_2 relative to the system. Suppose E_1 is a risk-free environment in which the system has an unvarying throughput capacity I = KR per time unit for a system with resources R, in accordance with expression (1A) with no sharing procedure. E_2 is the same as E_1 except that in E_2 the system is in a positive risk environment.

Suppose gross throughput capacity in E_2 is KR + G per time unit, in each of one (or more) time periods in which the risk in E_2 is run but where it just happens (by good luck) that the hazard does not occur. Thus KR + G can be viewed as the hazard-free throughput capacity in the presence of risk but where the hazard does not occur. However, when a risk is run repeatedly, throughput capacity losses must occur over time. If the average throughput capacity loss per time unit, due to the hazard occurring, is L_r, then the net throughput capacity from running the risk in reality will on average be KR + G - L_r per time unit.

But expression KR + G - L_r must also give the expected or average throughput capacity I, so that we can take L_r as the MEL-risk with respect to the hazard-free capacity of I = KR + G. In general there are now two possibilities for this risk: it can be r*isk which it can pay to run repeatedly (positive risk)*, where G > L_r, or it can be r*isk it cannot pay to run repeatedly (negative risk)*, where G < L_r.

Assume now for E_2 that G-L_r is positive, so that E_2 exposes the system to a risk it can pay to run repeatedly. Now recall that environment E_1 was risk free, with I = KR at full capacity for resources R applied. It is clear that there will be an increase in throughput capacity by shifting R from an environment E_1 with no risk and throughput capacity KR, to environment E_2 with average throughput capacity KR + G - L_r, that differs only in E_2 having a risk it can pay to run repeatedly.

But out of E_1 and E_2 we can construct an arbitrary number of synthetic environments, each with risk it can pay to run repeatedly, intermediate between the zero risk in E_1 and the MEL-risk L_r in E_2. We can do this in actual practice by taking a (valid) fraction of R and applying it to E_1 and the remainder of R to E_2. For example, suppose such a synthetic environment E_S when the fraction is 0.5. In E_S, for the resources R/2 operating risk free, the throughput capacity I will be KR/2, and for the remaining resources in risky operations, it will be (KR + G - L_r)/2 for a total of KR + G/2 - L_r/2, on average. Thus in E_S the MEL-risk, with resources R applied, will be L_r/2, and the increased throughput capacity in excess of KR, due to running the risk, will be (G - L_r)/2,

on average. We can repeat this with any fraction, so that it is clear that for such synthetic environments, expected throughput capacity I will increase linearly with MEL-risk.

We can even have a synthetic environment E_S where the average throughput loss, and thus MEL-risk, exceeds L_r, if we include the case where additional resources are borrowed and applied to E_2; for example, if we borrow an additional R resources (at the cost KR of the risk-free throughput from the borrowed R), and apply them to E_2, the throughput capacity I is now (2KR + 2(G - L_r) - KR) for the system, or KR + 2G - 2L_r, on average, so that the extra throughput capacity for agent's resources R is 2(G - L_r), on average, with the MEL-risk being 2L_r, consistent with throughput capacity increasing linearly with risk. Thus in general, average or expected throughput capacity I is given by:

$$I = KR + nG - nL_r$$

where nL_r is the MEL-risk, R and n are independent agent-controlled variables, and n >= 0 and may vary with the synthetic environment chosen for the system by the agent.

Now, the above expression should be pondered over, for it is absolutely correct, regardless of the distribution of losses over time, provided the following condition holds: There must exist, among all the hazard-occurring loss-generating time periods, a time period where no hazard occurs; in this loss free time period, we must have an extra throughput capacity nG. In addition, we must have an average loss nL_r, per time period, averaged over all the time periods, and measured with respect to the throughput capacity in the best-case hazard-free time period, where throughput capacity is KR + nG.

Now suppose that for a given system, out of the set of all non-synthetic naturally occurring environments in which the agent is free to operate the system, that is, accessible environments, we select an environment for which G/L_r is the highest, and let us call that environment the *reference environment* E_e. We call G/L_r the *risk efficiency coefficient* for the reference environment. Using that environment and the risk free environment we can now construct any number of accessible synthetic environments for which the gross extra gain per time period is nG , for an average loss per time period nL_r, where n is > 0, so that in every one of these synthetic environments the gross extra throughput capacity per unit of average loss, or risk efficiency coefficient, is the highest and the same as in the reference environment. Thus there will exist a set of environments, made up of synthetic environments *and natural environments*, in which the system could operate, and in each of which, for that system, the risk efficiency coefficient is the same as that of the reference environment E_e. We call this set of environments the *efficient set* based on a specific E_e, and call each of its members an *efficient environment.* An efficient environment can thus be either natural or synthetic.

3.2 Derivation of the risk equation

From this it follows that average throughput capacity I must increase linearly with the environment MEL-risk nL_r, for all efficient environments in which the system could run, and not all merely synthetic environments, where it is positive risk (meaning $nG > nL_r$). Accordingly :

$$I = KR + nG - nL_r$$

also holds for efficient environments. [Notice that if the agent chooses, irrationally, a range of environments in which to operate the system, some of which are inefficient, the equation will not hold; instead we will have $I = KR + mG - nL_r$ where n/m can vary from one environment to another.]

Now it is clear that if we have a valid increase in R to Rf, where f is a positive real value, for example, 2.0, it is clear that this is equivalent to adding a parallel system with resources R(f-1). Hence throughput capacity I will increase to If, and throughput capacity for the risk free environment will be KRf, and gross extra throughput capacity due to the presence of risk will be nGf, and the risk will be nfL_r. From this consideration it is clear that we may legitimately be concerned with gross extra throughput capacity per (basic harmonic resource) unit R, namely ng, where g = G/R, and MEL-risk per unit R, namely nl_r where $l_r = L_r/R$, so that we can rewrite the above expression as:

$$\begin{aligned} I &= KR + ngR - nl_rR \\ &= R(K + ng - nl_r) \\ &= R(K + ng - r(E)) \end{aligned}$$

if we define MEL-risk per unit R as $r(E) = nl_r$. Since, for positive risk, ng > r(E), we can rewrite ng as $(g/l_r)r(E)$, where g/l_r is greater than 1.0. Hence, for average throughput capacity I:

$$\begin{aligned} I &= R(K + (g/l_r)r(E) - r(E)) \\ &= R(K + b_{pb}\, r(E) - r(E)) \qquad (2b) \end{aligned}$$

Hence $I = R(K + (b_{pb} - 1)r(E)) \qquad (1Ba)$

or $I = R(K + b_p r(E)) \qquad (2c)$

where b_{pb} is the gross gain per unit of risk, or *risk efficiency coefficient*, for any member of the set of efficient environments, and is a constant that measures the extra gross throughput obtained for the best case of no hazard actually occurring, that is, the constant G/L_r. In practice it can be expected that while L_r is smaller than G for risks it can pay to take, it is not much smaller, so that b_{pb} will be only a little greater than 1.0, and b_p much less than 1.

Expression (1Ba) is the basic risk equation for positive risk, where r(E) is the risk of an efficient environment relative to the system. If some of the environments in which the agent (irrationally) operates the system are not efficient, equation (1Ba) still holds, but b_{pb} will no longer hold constant as the system is shifted from one inefficient environment, and therefore one risk, to another, each with a different risk efficiency coefficient $b_{pb} = nG/nL_r$. It obviously behooves the agent to discover the set of efficient environments for the system and select from that set. Since agents can be assumed to be rational and risk averse, they are not likely to run the system in an accessible risky natural environment for a smaller gross extra throughput capacity nG than could be obtained from a synthetic environment with the same risk nL_r constructed using the reference environment E_e, that is, they are likely to run the system only in efficient environments. This principle might be called the *risk equivalence principle*. It is self-evident, we believe, because its converse makes no sense. In the literature there is no sign of any research having been done into efficient system environments, a neglect it could clearly pay to remedy. For what it is worth, the author suspects, based on anecdotal evidence, that efficient system environments are highly orderly.

[K is assumed constant over long periods of time, although the expression allows for K varying (slowly in the long run) independently of the benefit of running the risk r(E), that is, independently of net addition to mean throughput capacity per unit R as measured by $b_p r(E)$. Such independence happens in financial systems where K corresponds to the risk free (per unit) interest rate, which does change over the long run, independently of the benefits of risk taking [10]. Expression (1Ba) assumes such independence of K and b_p for systems in general. If for some system it can be shown that the benefits of risk-taking have a long-run linear variation with

K, as might sometimes happen, then we could write $b_p = a_pK$, giving $I = KR(1+a_p)r(E)) = KR[1 + (a_{pb} - a_{pc})r(E)]$, which is a possible variant of the risk equation.]

The quantities $Rb_pr(E)$ in (2c) and $KR(b_{pb}-1)r(E)$ in (1Ba) are each expressions for the average net extra throughput capacity achieved by taking the risk r(E) of the environment, measured as MEL-risk per unit R. Where SD-risk is preferred, so that typically SD-risk is 0.5 times MEL-risk, one can convert from MEL-risk by inserting the corresponding SD-risk measure r(E) into the equations of risk and adjusting the constants b_{pb} and b_p.

In (1Ba) KR is the unvarying throughput capacity I for a risk-free environment, and which therefore induces no fluctuations in I due to hazards. Particularly when using MEL-risk, but ultimately also with SD-risk, $Rb_{pb}r(E)$ corresponds to the gross extra throughput capacity nG in a time period where the risk is present but the hazard does not occur, and $Rr(E)$ corresponds to the average throughput loss nL_r over all time periods due to the hazard occurring.

Financial risk equation

In the finance arena, SD-risk r(E) is used; the equivalent of equation (2c) is also used [16, 20], but equation (1Ba) is unknown. With financial systems, using:

$$I = R(K + b_pr(E)) \qquad (2c)$$

the system resources R become the principal sum invested, and K becomes the risk-free per-unit interest rate obtainable from (risk-free) Treasury bills, so that $b_pr(E)$ is the extra per-unit return gained by exposure to risk, that is, in percentage terms if K is 10%, $b_pr(E)$ might be 3%, for a total return of 13%. If the principal R is invested in common stocks, that would put the system in one risky environment with one r(E) value, if R is invested in bonds that would put the system in another risky environment with a different (smaller) r(E) value, and if R is in Treasury bills that would put the system in a risk-free environment; if we distribute R over stocks, bonds and Treasury bills that puts the system in a synthetic environment with a further r(E) value. b_p varies from decade to decade, but is of the order of 0.3; this means, crudely, that a lot of fluctuation in return has to be endured to get an small increase in mean return, i.e. to get an extra 3.0 percentage points of return (of R) on average, $b_pr(E)$ must be 0.03 so that r(E) must be 0.1, so that about 10% (of R) standard deviation fluctuations in overall return must be endured. Similarly to get an extra 6.0 percentage points on average, 20% standard deviation fluctuations must be endured, and so on [20]. For financial systems, the SD-risk r(E) is the standard deviation in annual return per unit time per unit R, thus, the standard deviation risk is expressed as a fraction of R, and not as a fraction of the mean or expected return itself.

If investors are never irrational, then for the same return in two different environments the risk should be the same, otherwise the environment with the lower risk would be preferred and securities price levels will adjust. If one environment, say the common stock environment S, has a fluctuation level and thus SD-risk r(S) that is q times the SD-risk r(B) of another environment, say the bond environment B, so that r(S) = qr(B), then the extra return s in S above the risk free return K must be q times the extra return b in B above K, or s = qb, and securities prices will adjust to make it so. Typically $q > 1.0$, since stocks are riskier than bonds. This s = qb must hold, since when principal sum R is in stocks (say the S&P500 index), it is possible to construct a synthetic environment (T) with R consisting of a principal sum qR in bonds with a sum (q-1)R borrowed at the Treasury bill rate K, for which the risk r(T) is qr(B), the same as for the stock portfolio. The stock and synthetic bond-based portfolios must now each give the same return s = qb, otherwise the environment S or T with the lower risk will be preferred and securities prices will adjust to make the returns identical. It is this principle of equivalent risks giving the same return that was used originally to derive the relation (2c), which is probably the fundamental expression of finance [20]. This method was not used earlier in this paper, and

indeed cannot be used, to derive the basic risk equations (1C) and (2c) for systems in general, since this method is based on the assumption that markets will adjust their price levels to force all financial environments, stock, bonds, treasury bills, and combinations of these, to form a single set of efficient environments for investment. In this paper the basic risk equations were derived for all non-growth non-evolving agent-directed systems only from considerations of the fundamental nature of risk.

3.3 Risk equation for negative risk

It should also be clear from a similar analysis of a set of efficient environments based on a reference environment for which the ratio nG/nL_r is maximal but less than 1, that is, containing only risks it can not pay to run, that we must also have

$$I = R[(K + (b_{nb} - 1)r(E)] \qquad (1Bb)$$
$$= R(K - b_n r(E))$$

This is the basic risk equation where r(E) is a measure of negative risk for an efficient environment, or risk it can not pay to run (that is, $b_{nb} < 1$ or b_n is positive). Thus b_{pb}, for positive risk, is greater than 1, and b_{nb}, for negative risk, is fractionally less than 1.

3.4 A sample application of the risk equation

A numerical example of the use of the positive risk equation will give the reader a better understanding of its implications. Suppose a system with R = 1000, with K = 0.075 or 7.5%, so that in a risk free environment, I is 75 x-units per week. Suppose we place the system in an efficient environment E with a MEL-risk r(E) of 0.02 x-units per week per unit R, and this placement raises mean throughput capacity from 75 to 80 x-units per week. In that case

$$I = R[0.075 + 1.25\, r(E) - r(E)] = R[\, 0.075 + 0.25 * 0.02]$$
$$= 1000[\, 0.075 + 0.005] = 75 + 5 = 80 \text{ x-units per week}$$

The best case throughput capacity is:

$$R(0.075 + 1.25r(E)) = 1000[\, 0.075 + 0.025] = 75 + 25 = 100 \text{ x-units per week}$$

from which level there are downward throughput capacity fluctuations, with a mean of 20, to mean I = 80 on average, with a minimum downward fluctuation of zero, and a maximum probably of 40, to a minimum throughput capacity of 60. If we shift the system to another efficient environment with double the MEL-risk r(E) of 0.04 x-units per week per unit R, the mean of I will be raised to 85, and the hazard free case to I =125, for a mean fluctuation of 40 down from125 to a mean I = 85. If instead the system is placed in an environment with an SD-risk of 0.02 x-units per week per unit R, throughput capacity rises from 75 to a mean of 80 as before, but with the standard deviation of fluctuations about the mean of 80 being 20 x-units per week, instead of the maximum fluctuation about the mean of 80 being about 20 if the risk of 0.02 were MEL-risk.

3.5 Risk combining and destructive interference

Readers are cautioned about naively using the risk equation to combine parallel systems with risks from different efficient environments. Suppose a composite system C made up of two parallel identical systems in two different efficient environments, with the same risk per unit R, and the same gross extra throughput per unit R, except that the risks are not well, or are negatively, correlated. Suppose throughput capacities are I_1 and I_2 such that total average system throughput capacity for C is $I = I_1 + I_2$, where

$$I_1 = R(\, K + dr(E_1)\,) \quad \text{and} \quad I_2 = R(\, K + dr(E_2))$$

Although the mean throughput capacities are additive the risks are not, since throughput capacity fluctuations in opposite directions will cancel, and thus destructively interfere. Thus for the composite system we must have

$$I = 2RK + dRr(E_1) \sim(+)\sim dRr(E_2)$$

where the tilde+ notation indicates that risk addition to give a composite risk is based on the underlying statistics, allowing for destructive interference, where:

$$Rr(E_1) \sim(+)\sim Rr(E_2) = 2Rx + 2Rr(E)$$

and where r(E) is risk per unit R for the composite (and necessarily efficient) environment E_{12} derived from E_1 and E_2, and Rx is *a positive certain loss, with respect to the hazard free throughput capacity, due to destructive interference between the capacity fluctuations of the two constituent systems.* Hence for the composite system we must have:

$$I = 2R(K + b_p r(E))$$

for any environment E, where E_{12} can be a valid value for E, where $b_p = d(r(E_1) + r(E_2))/2r(E_{12})$, that is, a value b_p that will be larger than d.

A numerical example is instructive. Suppose the two systems obey

$$I_1 = R(0.075 + 0.25r(E_1)) = R(0.075 + (1.25 - 1)r(E_1))$$

and $I_2 = R(0.075 + 0.25r(E_2)) = R(0.075 + (1.25 - 1)r(E_2))$

Suppose R = 500, and $r(E_1) = r(E_2) = 0.02$, giving mean throughput capacity of 40, and a total risk of 10, with best case of 50, for each system, and throughput capacity of 37.5 for each system in a risk free environment, so that the extra throughput capacity of 2.5 is the benefit of running the risk in each system.

If we combine the systems, the mean throughput capacity will be 80 no matter what, but if the total risks destructively interfere, so that instead of adding to 20, suppose they add to only 2, so that when $Rr(E_1)$ and $Rr(E_2)$ are each 10, $2Rr(E_{12})$ is 2. The composite system will therefore be described by:

$$I = 2R(0.075 + 2.5r(E)) = 2R(0.075 + (3.5 - 1)r(E))$$

And for 2R = 1000, and $r(E) = r(E_{12}) = 0.002$, the best case throughput capacity is 82, the mean is 80, and the total risk is 2. [This is not what would be the case if there were no destructive interference, but instead perfectly correlated risks, allowing simple addition of the risks, so that the best case is 100, the mean is 80, the total risk is 20, and b_p remains unchanged (same as d) at 0.25.] For this combined case with the destructive interference:

$$b_p = 2.5 = d(r(E_1) + r(E_2))/2r(E_{12}) = 10d = 10*0.25$$

As a result of the above analysis, it is clear that if a given system is exposed to a future loss consisting of two distinct risks r_1 and r_2, equal in magnitude, measured as MEL-risks, the resultant future loss exposure consists of a certain loss x(c) plus a risk or possible loss r_3, where

$$r_1 \sim(+)\sim r_2 = x(c) + r_3$$

where x(c) is zero only when the correlation coefficient c for the two underlying risks is +1. The certain loss factor x(c) will increase with decreasing c, with r_3 decreasing to zero as c approaches –1, the point at which all possible loss is converted entirely to certain loss and risk is zero! This is the principle of risk reduction by means of investment diversification in finance, although the fact that it merely involves converting possible loss to certain loss is rarely discussed [16].

3.6 Risk as a function of system environment relative to the system

The basic risk equation states that R may be altered independently of the risk r(E) (per unit system resource) of the efficient environment relative to the system. However, it should be recalled that, with equation (1A), R can be validly altered only by decreasing or increasing R in valid (basic harmonic resource) units of the existing resource type of R, so that if we n-fold increase the value of R we n-fold replicate the system. Only with such an alteration of R will r(E) not change and expression (2A) hold. If an addition of another type of resources is made we may

bifurcate the system into two systems, the original system and a new one consisting of the new type of resources, and although the environment of the two systems may be identical, the *environment relative to that new system* may well be different, so that the risk per unit R may be different for the new system. In addition the environment relative to the new system may not be efficient either; it would have to be analyzed to find out.

A simple physical example should illustrate these points where risk is involved. Suppose a trucking system with resources R that ships disk drives from A to B over a gravel road designed to take 10 ton trucks at the heaviest (so that the risk of a truck being delayed because of the wheels sinking into the highway is zero), and suppose that R consists of 4 10-ton trucks each carrying 20 disk drives from A to B per day, so that I is 80 disk drives per day, and the basic harmonic resource unit is 1 10-ton truck. If we increase R by 6 10-ton trucks the risk of the efficient environment relative to the system remains zero, and throughput capacity I climbs by 6*20 to 200 disk drives per day. If instead we had added 2 20-ton trucks each capable of carrying 60 drives per day, throughput capacity I would also increase to 200 drives per day, except that the increase in R is not valid, and worse, the 2 20-ton trucks on a road designed to carry 10-ton trucks will be exposed to risk of delay due to getting stuck. Thus there are two systems, a risk-free system where R consists of 4 10-ton trucks, for which the risk of the efficient environment relative to the system is zero, and a system where R consists of 2 20-ton trucks, with the same efficient environment (the road), but for which the environment *relative to the system* is risky, and which may or may not be efficient. In this practical context, "not efficient" would mean that there is an alternate road for the heavy trucks with less risk for the same extra throughput capacity benefit.

3.7 Combination of risk with a resource-sharing procedure.

Sometimes a system exposed to risk in an efficient environment will also have an active resource-sharing procedure [2, 3] at the same time to better utilize the available resources. For this case we must combine the efficiency equation (1A) and the basic risk (1B) giving, for the case of just positive risk:

$$I = R[K + (b_{pb} - 1)r(E)](1-T_S/T)[1+ sF_T(T_S)]$$

or

$$I = R(1-T_S/T)[1+ sF_T(T_S)][K + (b_{pb} - 1)r(E)]$$

This indicates that system resources R, in use in the risky efficient environment before application of the sharing procedure, have effectively been increased to $R(1-T_S/T)[1+ sF_T(T_S)]$ by the use of the coordinated resource sharing procedure, independently of the prior existing risk per unit R in the efficient environment. Note however, that total risk is increased to $R(1-T_S/T)[1+ sF_T(T_S)]r(E)$, so that the absolute size of losses when a hazard does occur will increase proportionally. A simple example would be a computer system with an unshared cpu executing 20 jobs per hour but exposed to the risk of corrupt input data that can cut throughput capacity to 10. If we alter the system by increasing the level of sharing, so that capacity is 100 jobs per hour, the risk is now of having capacity cut to 50, for the same level of risk r(E) per unit R of corrupt input data.

4. Future Research

Since the risk equation and resource sharing equation are valid only for agent-directed, non-growth, non-evolving systems, an obvious area for future research is development of similar equations for agent-directed systems that allow for system growth and evolution. There is every

reason to believe that such a set of equations could be developed. An even more challenging project, however, would be to develop a set of equations for naturally occurring, growing and evolving systems such as biological systems. The obvious problem here is the absence of any directing agent other than nature, whose impressive organizing capabilities with regard to biological systems remain to be fully explained.

5. Concluding Remarks

A basic risk equation, valid for all non-growth, non-feedback agent-directed systems, but especially for computer and information systems, has been presented and justified. This basic equation may be combined with a resource-sharing equation, developed in a separate paper [2, 3], enabling expressions for a very wide variety of system situations (1F).

The resource sharing equation (1A) states (a) that throughput capacity increases linearly with valid increases in resources, and (b) that for a given throughput capacity, as system resources are decreased, system operating complexity must increase and vice versa. The basic risk equation (1B) states that for a given system resource level, mean throughput capacity increases linearly with the risk in an efficient environment relative to the system. Many obvious examples of the use of these expression are to be found in multiprogramming operating systems and file and database systems.

The risk equation holds only for specific classes of environments, namely efficient environments. The risk equation is new, except that there is a version of the risk equation that is widely used in financial systems, and which was derived by Sharpe and others in the 1960s on the basis of the behavior of financial markets and their participants; however the basic risk equation presented here was derived from the basic nature of risk, and applies to all agent-directed no-growth, non-evolving systems in efficient environments.

In the equation the constants and variable parameters can be reduced to numbers and measurable quantities, so that the equation is subject to experimental verification. Nevertheless, the reader who has taken care to fully grasp the equation will no doubt see that it fits with his or her own experience of practical functioning systems, especially real-time systems. However, the major benefit of the equation, together with the resource sharing equation, to system designers and operators, is that they promote and simplify clear thinking and accurate reasoning about complex system situations and possibilities that have hitherto been in shrouded in nebulous obscurity and complexity.

Although beyond the scope of this paper, three additional variations of the risk equation are possible, each one for one of three practical methods for reducing or eliminating a quick without significantly impairing the benefits of running the risk. These equations will be presented in later papers, but information is currently available in a research report [2].

References

1. Anderson, T.E., Lazowska, E. D., and Levy, H. M., The performance implications of thread management alternatives for shared-memory multiprocessors, IEEE Trans. On Computers, 38(12), 1989, pp1631-1644.

2. Bradley, J. "Five equations relating throughput capacity to systems resources and risk for all agent directed non-growth system. University of Calgary, Department of Computer Science, Research Report No. 1999-642-05 (WWW-accessible)

3. Bradley J. A resource-sharing hypothesis for coordinated sharing of system resources, submitted for publication, 2001.

4. Beaumont, G. P., "Probability and Random Variables", Ellis Horwood Limited, 1986.

4. Bunemann, O. P., and Clemons, E. K., Efficiently monitoring relational databases, ACM TODS, 4(3), 1979.

5. Date, C. J., Introduction to Data Base Systems, Addison Wesley, 7th Edition, 2000.

6. Date, C.J, A contribution to the study of database integrity, In C.J. Date, "Relational Database Writings", pp 1985-1989, Addison-Wesley, Reading, Mass., 1990.

7. Elton, J. E. and Gruber, M. J., "Finance as a Dynamic Process" Prentice-Hall, Englewood Cliff, NJ, 1975.

8. Fama, E.F., Risk, return and equilibrium, some clarifying comments, Journal of Finance 23(1), 1968, 29-40.

9. Habermann, A. N., Prevention of system deadlocks, Com. of ACM, 12(7), 1969, pp373-377.

10. Hennessy, J. L., and Patterson, D. A., "Computer Architecture: A Quantitative Approach", Morg an Kaufmann Publishers, Palo Alto, California, 1990

11. Holt, R.C. Some deadlock properties of computer systems Computing Surveys, 4(5), 1972, pp 179-176

12. Isloor, S. S., and Marsland, T.A. The deadlock problem: An overview, Computer 13(9), 1080, pp58-78

13. Lamport, L., The mutual exclusion problem, Journal of the ACM, 33(2), 1986, pp 313-348.

14. Laowska, E.D, Zahorjan, J., Graham, G.S., and Sevcik, K. C., "Quantitative System Performance," Prentice-Hall, Englewood Cliffs, NJ, 1984.

15. Lintner, J., The valuation of risk assets and selection of risky investments in capital budgets, Review of Economics and Statistics, 47(1), 1965, 13-37.

16. Lintner, J., Security prices, risk and maximal gains from diversification, J. of Finance 20(4), 1965, 587-615, 13-37

17. Markowitz, H.M. Portfolio selection, Journal of Finance, 1(1) 1952, 77-91.

18. Raynal, M., "Algorithms for mutual exclusion," MIT Press, Cambridge, MA, 1986

19. Sauer, C. H., and Chandy, K. M., "Computer System Performance Modeling", Prentice-Hall, Englewood-Cliffs, N.J., 1981.

20. Sharpe, W. F., Capital asset prices: a theory of market equilibrium under conditions of risk, Journal of Finance 19(3), 1964, pp 425-42.

21. Silberschatz, A, Peterson, J. L., "Operating Systems Concepts.", 5th Ed, Addison Wesley, Reading, Mass., 1998