

Modeling Redundancy: Quantitative and Qualitative Models

A. Mili, Lan Wu
College of Computer Science
New Jersey Institute of Technology
Newark NJ 07102-1982
mili@cis.njit.edu

M. Shereshevsky
Lane Department of EE and CS
West Virginia University
Morgantown WV 26506
m_shereshevsky@hotmail.com

F.T. Sheldon
U.S. DOE Oak Ridge National Lab
PO Box 2008, MS 6085, 1 Bethel Valley Rd
Oak Ridge TN 37831-6085
sheldonft@ornl.gov

J. Desharnais
Département d'Informatique et de Génie Logiciel
Université Laval
Québec QC G1K 7P4 Canada
Jules.Desharnais@ift.ulaval.ca

July 9, 2007

Abstract

Redundancy is a system property that generally refers to duplication of state information or system function. While redundancy is usually investigated in the context of fault tolerance, one can argue that it is in fact an intrinsic feature of a system that can be analyzed on its own without reference to fault tolerance. Redundancy may arise by design, generally to support fault tolerance, or as a natural byproduct of design, and is usually unexploited. In this paper, we tentatively explore observable forms of redundancy, as well as mathematical models that capture them.

Keywords

Redundancy, Quantifying Redundancy, Qualifying Redundancy, Error Detection, Error Recovery, Fault Tolerance, Fault Tolerant Design, Redundancy as a Feature of State Representation, Redundancy as a Feature of System Function.

1 Redundancy, an Evasive Concept

As a concept, redundancy is widely known and widely referenced; also, at an intuitive level, it is widely (though perhaps cursorily) understood. We have felt the need to model system redundancy in the midst of a research project that we conducted on analyzing a fault tolerance flight control system [2]. Subsequent work led to some tentative, preliminary, insights [9]. What strikes us most about this concept is a paradoxical combination of three premises, which are:

- Even though redundancy is a widely used and widely referenced concept, there appears to be little effort to model it in a generic manner that captures all observable forms.
- In addition to *artificially built-in* redundancy, such as modular redundancy, most systems have ample *natural* redundancy. This natural redundancy is seldom acknowledged and exploited, for example to build fault tolerance and improve system reliability.
- Although it is tantalizingly easy to understand at an intuitive level, redundancy has proven rather hard to model formally in a way that is general and meaningful.

In this paper, we attempt to catalog diverse manifestations of redundancy, then we explore means to model them. In Section 2 we present the many faces of redundancy, that we have observed and attempted to model; then we discuss some questions that we wish to address in the long. In Section 3 we introduce qualitative models of redundancy that we have tentatively explored, and briefly comment on the insights that these models afford us; for the sake of readability, we will keep the discussion informal, referring the interested reader to bibliographic references for technical details. In section 4 we explore a quantification of some aspects of redundancy, and outline issues that pertain to the quantification of other, more advanced concepts. In the conclusion, we summarize our main results and discuss future prospects.

2 Forms of Redundancy

While redundancy is widely recognized as an important system attribute, it has not been modeled and analyzed in a commensurate manner. Most bibliographic references appear to use this term in a literary sense, rather than a technical, well-defined, widely agreed upon, sense. In [10], we present a survey of various definitions and classifications of redundancy [1, 14, 16, 4]. From a modeling standpoint, we have found the following categorization to be useful for our purposes:

- *State Redundancy*. This arises when the representation of the system state allows a wider range of values than are needed to represent the set of reachable states. This is the traditional form of redundancy, that arises in parity-bit schemes, error correcting codes, modular redundancy schemes, etc.
- *Functional Redundancy*. This form arises when, for example, we compute the same function using three different algorithms, and we take a vote on the outputs. We do not distinguish, in functional redundancy, between whether the components that compute the same function are running concurrently, or in sequence.
- *Temporal Redundancy*. Consider the state defined by two variables: the altitude (Z) and the vertical speed (V_Z) of an aircraft. There is no redundancy between the values of Z and V_Z at a given time t (i.e. $V(t)$ and $Z_V(t)$ can take arbitrary values, for a given t), but there is redundancy between the values of these variables within small time intervals, e.g.

$$Z' = Z + V_Z \times dt.$$

- *Control Redundancy* (for lack of a better name). Control redundancy arises in control applications whenever we can achieve the same effect on the system by several distinct control settings. Our interest in redundancy stems from a project we conducted to analyze a fault tolerant flight control system [2]. This system is fault tolerant in the sense that it can keep flying the aircraft (under some restrictive conditions) even if some control fails. This stems from redundancy between the controls that the system operates: though the throttle, elevators, ailerons, flaps, and rudder have distinct functions, some may be used to make up for the loss of others.

This classification is neither complete nor orthogonal. It illustrates the diversity of forms of redundancy, and the interest in trying to model them and possibly unify them.

3 Qualitative Models of Redundancy

Whereas in Section 2 we cataloged forms of redundancy from an external standpoint, by characterizing their observable manifestations, we focus in this Section on models of redundancy, which attempt to provide a mathematical rationale for the observed manifestations.

3.1 Redundancy as a Feature of State Representation

In this section, we discuss a qualitative model of redundancy, which equates redundancy with functional attributes of the representation of the system state. In this view, we study redundancy as a representational issue, i.e. as a feature of the relation that maps states to their representations, which is the *state representation relation*. The simplest representation relations are those that are

- total (each state value has at least one representation),
- deterministic (each state value has at most one representation),
- injective (different states have different representations), and
- surjective (all representations represent valid states).

Not all representation functions satisfy these four properties—in practice hardly any satisfy all four, in fact.

- When a representation relation is not total, we observe a *partial representation* (for example not all integers can be represented in computer arithmetic).
- When a representation relation is not deterministic, we observe an *ambivalent representation*. Consider the representation of signed integers between -7 and +7 using a sign-magnitude format; zero has two representations, -0 and +0 [3].
- When a representation relation is not injective, we observe *loss of precision* (for example, real numbers in the neighborhood of a representable floating point value are all mapped to that value).
- When a representation relation is not surjective, we observe *redundancy* (for example, in a parity-bit representation of characters, not all bit patterns represent legitimate characters).

For the purposes of our discussions, we equate redundancy with non-surjectivity; for the sake of simplicity, we limit our discussion to representation relations that are deterministic, total, and injective—whence each state value has exactly one representation (by virtue of totality and determinacy) and different state values have different representations (by virtue of injectivity).

A simple example that illustrates why redundancy can naturally be equated with non-surjective representation relations is the parity bit representation of information in a computer. Out of eight bits in a byte, seven are used to represent different values, and one is used to represent the parity bit. This makes the representation non surjective since the range of the representation function has a cardinality of (at most) 2^7 while the set of representations has a cardinality of 2^8 . This non-surjectivity is crucial to error detection; one out of two representations is illegitimate, making it possible to detect some errors. Had the representation been surjective, an error would map a legitimate state into another legitimate (though incorrect) state, making it impossible to detect any error.

3.2 Redundancy as a Feature of System Function

Whereas the previous Section focuses on the redundancy of state representations, this Section focuses on the redundancy of system functions. Specifically, in this Section we will equate redundancy with fault tolerance capability, and explore how functional attributes of systems and their specifications introduce sources of redundancy that can be used to support fault tolerance. For the sake of readability, we resolve to keep the discussion non-technical, referring the reader to publications [9] for technical details.

3.2.1 Redundancy and Surjectivity

Building on earlier work [7], and on widely accepted fault tolerance ideas [5], we have proposed in [9] a hierarchy of correctness levels for intermediate states in computations:

- *Strict Correctness*, characterizing an error-free state.
- *Maskability*, characterizing a state which, while it may not be strictly correct, will spontaneously avoid failure, nevertheless.
- *Recoverability*, characterizing a state which, while it may not be maskable, does nevertheless contain all the necessary information that allows a recovery function to produce a maskable state.
- *Partial Recoverability*, characterizing a state which, while it may not be recoverable, does nevertheless contain some information about maskable states.
- *Non Recoverability*, which characterizes an erroneous state that contains no information that could be used for recovery.

In [10], we have established the link between function surjectivity and system redundancy by highlighting two relations: increasing levels of fault tolerance depend on in-

creasing levels of non-surjectivity; increasing levels of non-surjectivity depend on enlarging state spaces, by introducing state variables, which increase redundancy.

3.2.2 Redundancy and Injectivity: Past Functions

A function is said to be *injective* if it maps distinct arguments into distinct images. By and large, we refer to the equivalence relation ($P(x) = P(y)$) that a function P defines on its domain as the *nucleus* of P (represented, in relational notation, by $P\hat{P}$), and we refer to its equivalence classes as the *level sets* of P (terminology of Mills et al [6, 11, 12]). A function is injective if its nucleus is the identity relation (whence its level sets are singletons), and it grows increasingly less injective as its level sets grow larger and larger. The least injective function is a constant function, which maps all its arguments into the same image; the nucleus of such a function is the total relation (I), which has one equivalence class (the whole space).

To illustrate the relationship between redundancy and injectivity, we consider a program (or system) structured as a sequence of two components, say P (past) and F (future), as follows

$$P; K: F$$

We have shown, in [10], how maskability, recoverability, partial recoverability and un-recoverability of system states at label K (after execution of P , before execution of F) can be determined by the degree of injectivity of function P .

3.2.3 Redundancy and Injectivity: Future Functions

In the previous Section, we have observed how we can equate redundancy with the injectivity of past functions. It is possible to argue that we can also equate redundancy with the *non-injectivity* of future functions. The less injective the future function, the bigger its level sets (elements that have a common image), the more we can depart from a correct state without affecting the image by F . It is easy to see, intuitively, why redundancy increases when past functions grow more injective and future functions grow less injective:

- The injectivity of past functions ensures the preservation of important information.
- The non-injectivity of future functions reduces the amount of information that must be preserved to avoid failure.

3.3 Redundancy as a Feature of System Specifications

If we equate redundancy with fault tolerance, which we have throughout Section 3.2, we must recognize that the

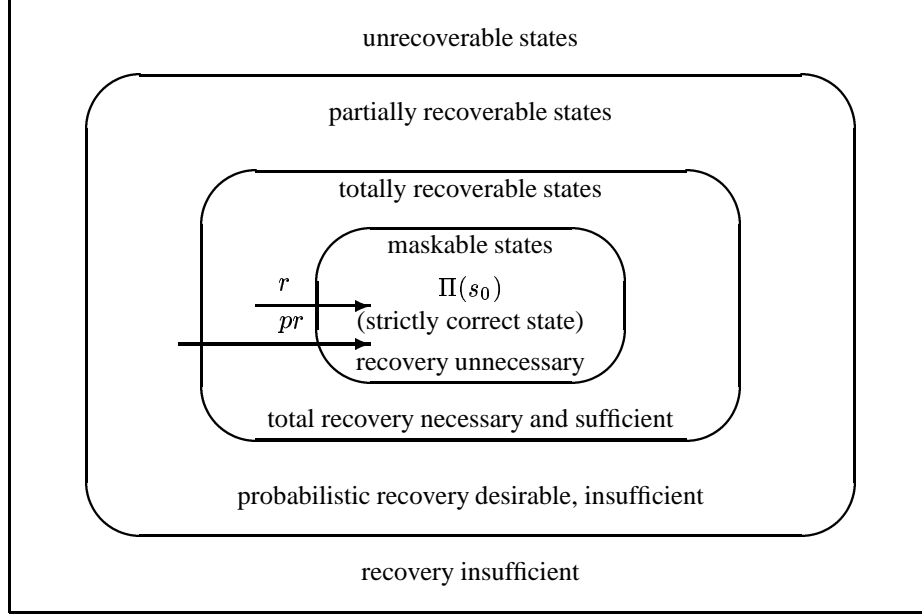


Figure 1: A Hierarchy of Correctness Levels.

non-determinacy of system specifications is an important source of (often unexploited) redundancy. In all the discussions we had so far, we have assumed that for a given input to the system, there exists a single correct output, and have equated recovery with the ability to retrieve this single correct output from current information. But there is usually a wide gap of determinacy between system functions and system specifications [13], for a variety of reasons (design decisions that characterize specific implementations but bear no relation to the system requirements; weak functional requirements, that require approximate values; state variables whose final value does not matter to the user; etc). Taking into account the non-determinacy of specifications would add several levels to the classification depicted in Figure 1; it would in fact double the number of categories, creating an orthogonal classification to that of Figure 1. We do not show the full classification obtained if we were to introduce non-deterministic specifications, but use Figure 2 to illustrate the impact of non-determinacy on maskability; in this Figure, we now have two levels of maskability, one for the system function (F-Maskable) and one for the system specification (S-maskable). Imagine a version of Figure 1 where each level of correctness is duplicated: one level pertains to the system function, and one pertains to the system specification.

What makes this redundancy model of particular interest to us is our hypothesis (yet to be confirmed) that this is an adequate model for control redundancy, introduced in Section 2. To illustrate this idea, we consider the *Flight Control Loop* depicted in Figure 3, taken from [8]. In this




Figure 2: Non Determinacy of Specifications: An additional Dimension of Redundancy

loop, the *Flight Control Software* (FCS) takes as input sensor readings (depicting relevant flight parameters), Pilot Commands (as Auto-Pilot settings) and possibly Navigation Signals (for navigating an approach to a landing), and produces as output actuator settings. The specification of FCS describes the relation that must be maintained between inputs and outputs. The control redundancy means that the same state of the aircraft can be achieved by more than one control setting; this is reflected by the non-determinacy of the specification of FCS, whereby the same input will be mapped onto a wide range of possible settings, corresponding to equivalent control combinations. This is obviously a tentative interpretation, which we envision to investigate and formalize further in this project.

4 A Quantitative Model: Measuring Excess Information

4.1 Space Redundancy

We focus our attention in this Section on state redundancy, and we try to quantify the amount of excess information of the state representation by means of a numeric function that matches the information carried by a state against the number of bits used to represent the state. We envision our function to satisfy the following identities: It takes its values in the set of non-negative real numbers; it takes value zero whenever the state carries no redundancy; it takes value 1 if a redundancy-free state is duplicated; generally, it takes value $N - 1$ whenever a redundancy-free state is duplicated N times. In [15], C. Shannon introduces a quantitative measure of state redundancy as follows:

”The ratio of the entropy of a source to the maximum value it could have while still restricted to the same symbols will be called its relative entropy. One minus the relative entropy is the redundancy”.

We consider a state s ranging over a state space S , we let $P(s)$ be the probability of occurrence of s , and we denote by $H(S)$ the entropy of S modulo probability P and by $M(S)$ the maximal entropy of space S (typically, $M(S) = \log(|S|)$). According to Shannon, the relative entropy of S is given by the following formula

$$\mathcal{E}(S) = \frac{H(S)}{M(S)}.$$

We use Shannon’s formula to obtain the *relative* redundancy of space S (though Shannon merely calls it redundancy)

$$\mathcal{R}(S) = 1 - \mathcal{E}(S).$$

As defined by this formula, relative redundancy varies between 0 and 1, and measures the quantity of excess information normalized to the maximum entropy. The redundancy function (\mathcal{D}) we are interested in ranges between 0 and infinity, and measures the amount of excess information normalized to the entropy of the space we are representing. We define it by the following formula,

$$\mathcal{D}(S) = \frac{M(S) - H(S)}{H(S)},$$

where the numerator represents excess information and the denominator normalizes the excess information to the actual information being represented. We find readily that the relative redundancy and the (absolute) redundancy satisfy the following equations:

$$\mathcal{D}(S) = \frac{1}{\mathcal{R}(S)} - 1,$$

$$\mathcal{R}(S) = \frac{1}{1 + \mathcal{D}(S)},$$

which are consistent with their respective interpretations.

In practice, we did encounter a difficulty, however: we were not sure how to interpret Shannon’s concept of maximal entropy, $M(S)$. Also, we found situations where all possible interpretations of Shannon’s definition lead to unsatisfactory results. Examples include:

- If we consider 8 states represented by taking binary codes 000 to 111 and duplicating them (to obtain: 000000, 001001, 010010, 011011, 100100, ... 111111), we are not sure what is the maximum entropy in this case: Is it $\log(8) = 3$ since we have only eight values to represent, or is it $\log(64) = 6$ since we can represent up to 64 values (if we view the symbol 100100, for example, as a sequence of six bits that can take arbitrary values)?
- In the case of Huffman prefix codes, where the length of the code varies according to the probability of occurrence of each symbol, we are not sure what maximal entropy means. There seems to be no way to reflect the property that we have a variable length code, since the maximal entropy is achieved by a uniform probability distribution, which yields a code of constant length.

To overcome this dilemma, we had to reinterpret/ redefine Shannon’s formula, by replacing the *maximal entropy* ($M(S)$) with the *weighted length* of the code, which we denote by $W(S)$. Whence we revisit the definition of redundancy and rewrite as:

$$\mathcal{D} = \frac{W(S)}{H(S)} - 1,$$




Figure 3: Outline of a Flight Control Loop.

where $W(S) = \sum_{s \in S} p(s) \times W(s)$, and $p(s)$ is the probability of occurrence of s , and $W(s)$ is the length of s . Intuitively, it makes sense that the formula of redundancy should mention the size of the representation somewhere, which it now does.

To illustrate this function, we consider some sample examples below; in the interest of brevity, we will not show details of our computations, but all we are doing is applying the formula above.

1. **No Redundancy.** If S contains 8 states that are equally likely to occur and are coded on 3 bits, then $\mathcal{D}(S) = 0$.
2. **Duplication.** If S contains 8 states that are equally likely to occur and are coded on 6 bits where the code of each state is duplicated, then $\mathcal{D}(S) = 1$. This value means that one hundred percent of the information needed to represent these states is added to the representation, which reflects the situation at hand.
3. **Error Correcting Code.** We consider a space of 8 values (of equal probability) and we represent it by four bits, say three bits of data and a parity bit. We find that the redundancy function is then equal to: 0.333, whose interpretation is obvious.
4. **Non Uniform Distribution.** If S contains 8 states that are not equally likely to occur and are coded on 3 bits, then $\mathcal{D}(S) > 0$. For example, if the probability distribution is:

$$(0.03, 0.05, 0.1, 0.12, 0.13, 0.15, 0.2, 0.22)$$

then the redundancy function yields the value: 0.069.

5. **Tame Distribution.** If the variance in the probability is more tame, the redundancy is much smaller. Hence

for the following probability distribution,

$$(0.09, 0.1, 0.1, 0.1, 0.125, 0.125, 0.14, 0.22)$$

we find the value: 0.0215.

6. **Huffman Coding.** If we consider the Huffman prefix code obtained from the first probability distribution, i.e.

$$(0.03, 0.05, 0.1, 0.12, 0.13, 0.15, 0.2, 0.22)$$

we find the redundancy as: 0.011, which is significantly less than the value found above for fixed size coding (item 4).

7. **Huffman Coding, Tame Distribution.** If we consider the Huffman prefix code obtained from the second probability distribution, i.e.

$$(0.09, 0.1, 0.1, 0.1, 0.125, 0.125, 0.14, 0.22)$$

we find the redundancy as: 0.012, which is significantly less than the value found above for fixed size coding (item 4) but more than the previous example (item 6), on account of having a tamer probability distribution.

8. **Prefix Code, Non Uniform Distribution.** We consider the non uniform probability distribution,

$$(0.03, 0.05, 0.1, 0.12, 0.13, 0.15, 0.2, 0.22)$$

but this time we use a prefix code that is not necessarily a Huffman code (some codes are longer than they have to be). In this case we find the redundancy value of 0.021, which is greater than the value found for the same probability distribution when we use Huffman coding (item 6).

9. **Non Prefix Code.** If we use the following distribution

$$(0.09, 0.1, 0.1, 0.1, 0.125, 0.125, 0.14, 0.22)$$

and build a non prefix code (by building the Huffman tree then placing one symbol on an internal node of the tree), we find the following value of redundancy: -0.081.

The results computed by this formula are consistent with our intuition, for the examples we have explored above. Although we do not consider these developments to be a proof, they illustrate the following premises: Redundancy takes a minimal value for Huffman coding. The redundancy of Huffman coding decreases as the distribution grows less uniform. Redundancy takes non negative values for all lossless codings. Redundancy takes negative values for codings that produce a loss of information.

4.2 Functional Redundancy: Non Surjectivity

Though we have, in our qualitative analysis, equated redundancy of a function with a variety of functional/ set theoretic attributes, we introduce, in this section, a quantitative measure that reflects the non-surjectivity of a function. We consider a function F defined from some input space X to some output space Y , and we let (by abuse of notation) X and Y be random variables that range over these spaces, respectively. We quantify the redundancy of function F by the formula

$$\mathcal{D}(F) = \frac{H(Y)}{H(F(X))} - 1,$$

where the entropy of Y is computed with respect to the uniform probability distribution on Y and the entropy of $F(X)$ is computed with respect to the probability distribution of X . The intuition behind this formula is that $H(Y)$ measures the entropy of the output space of F whereas $H(F(X))$ measures the entropy of the range of F ; their ratio measures how many times the output space has more entropy than is needed to represent the range of F ; subtracting 1 gives us the excess entropy. Note that when we duplicate, triplicate, or otherwise multiply a function (as is done traditionally to enhance error detection and error correction capability), we increase the size (whence, typically, the entropy) of the output space, while preserving the size and entropy of the range of F . Specifically, if we take a function F from X to Y and triplicate it, we obtain a function, say ϕ , which has the same input space as F , and the same domain as F . However, the output space of ϕ is Y^3 , and the range of ϕ is

$$\{ \langle y, y, y \rangle \mid y \in \text{rng}(F) \}.$$

This space is homomorphic to $\text{rng}(F)$.

To illustrate the significance of this formula, we consider the table of Figure 4. We let the reader contemplate in what sense the relations observed in this table reflect one's intuition about the properties one wants to see in a redundancy function. In this table, we denote by $B5$ (stands for 5-bit integers) be the set defined by:

$$B5 = [0..31].$$

5 Conclusion: Summary and Prospects

5.1 Summary

In this paper we have attempted to analyze system redundancy from a variety of angles, focusing first on external manifestations of redundancy, then on possible mathematical models that capture it. Our qualitative analysis has attempted to characterize redundancy by equating it with functional properties of state representations or system functions. As for our quantitative analysis, it is summarized by two formula: a formula for state redundancy and a formula for functional redundancy.

5.2 Prospects

We have offered no definite solutions in this paper, only tentative observations about the multiple forms, multiple models, and multiple observable manifestations of redundancy. Two measurable steps that we have achieved in this paper are the computation of state redundancy and functional redundancy. We are currently exploring the implications of these formulas. Examples of outstanding issues include: Are there necessary or sufficient conditions between quantities of (state or functional) redundancies and specific fault tolerance capabilities? Can we define redundancy in a manner that allows us to capture all the relations we have explored in a unified formula? How are the qualitative attributes of state representation and system function combined to produce redundancy? These matters are under investigation.

References

- [1] Anish Arora and Sandeep S. Kulkarni. Detectors and correctors: A theory of fault tolerance components. In *Proceedings, International Conference on Distributed Computing Systems*, pages 436–443, May 1998.
- [2] D. Del Gobbo and A. Mili. Re-engineering fault tolerant requirements: A case study in specifying

Name	Expression	Input Space	Output Space	Redundancy
F_1	X	B^5	B^5	0
F_2	$2 \times X$	B^5	B^5	0.25
F_3	$X \bmod 4$	B^5	B^5	1.5
F_4	$X \bmod 16$	B^5	B^5	0.25
F_5	$X \div 2$	B^5	B^5	0.25
G_1	$\langle F_1, F_1 \rangle$	B^5	B^{5^2}	1.0
G_2	$\langle F_2, F_2 \rangle$	B^5	B^{5^2}	1.5
G_3	$\langle F_3, F_3 \rangle$	B^5	B^{5^2}	4.0
G_4	$\langle F_4, F_4 \rangle$	B^5	B^{5^2}	1.5
G_5	$\langle F_1, F_2 \rangle$	B^5	B^{5^2}	1.5

Figure 4: Functional Redundancy

- fault tolerant flight control systems. In *Proceedings, Fifth IEEE International Symposium on Requirements Engineering*, pages 236–247, Royal York Hotel, Toronto, Canada, 2001.
- [3] E.C.R. Hehner. Quantifying redundancy. Private Correspondence, 2003.
- [4] Barry W Johnson. *Design and Analysis of Fault Tolerant Digital Systems*. Addison Wesley, 1989.
- [5] J.C. Laprie. Dependability —its attributes, impairments and means. In *Predictably Dependable Computing Systems*, pages 1–19. Springer Verlag, 1995.
- [6] R.C. Linger, H.D. Mills, and B.I. Witt. *Structured Programming*. Addison Wesley, 1979.
- [7] A. Mili. *An Introduction to Program Fault Tolerance: A Structured Programming Approach*. Prentice-Hall, Englewood Cliffs, NJ, 1990.
- [8] Ali Mili, Frederick Sheldon, Fatma Mili, and Jules Desharnais. Recoverability preservation: A measure of last resort. In *Proceedings, First Conference on the Principles of Software Engineering*, Buenos Aires, Argentina, 2004.
- [9] Ali Mili, Frederick T Sheldon, Fatma Mili, Mark Shereshevsky, and Jules Desharnais. Perspectives on redundancy: Applications to software certification. In *Proceedings, 38th Hawaii International Conference on System Sciences*, Hilton Waikoloa Village, Big Island, HI, 2005.
- [10] Ali Mili, Lan Wu, Frederick Sheldon, Mark Shereshevsky, and Jules Desharnais. Qualitative and quantitative models of redundancy. Technical report, New Jersey Institute of Technology, web.njit.edu/mili/red.pdf, 2005.
- [11] H.D. Mills, V.R. Basili, J.D. Gannon, and D.R. Hamlet. *Structured Programming: A Mathematical Approach*. Allyn and Bacon, Boston, Ma, 1986.
- [12] H.D. Mills, R.C. Linger, and A.R. Hevner. *Principles of Information Systems Analysis and Design*. Academic Press, 1985.
- [13] David L Parnas. Precise description and specification of software. In Daniel M. Hoffman and David M. Weiss, editors, *Software Fundamentals*, chapter 5. Addison Wesley, 2001.
- [14] Laura L Pullum. *Software Fault Tolerance Techniques and Implementation*. Artech House, Norwood, MA, 2001.
- [15] C. Shannon. A mathematical theory of communication. *Bell Syst. Tech. Journal*, 27:379–423, 623–656, 1948.
- [16] Martin L Shooman. *Reliability of Computer Systems and Networks: Fault Tolerance, Analysis and Design*. John Wiley and Sons, New York, NY, 2002.