

NOTICE WARNING CONCERNING COPYRIGHT RESTRICTIONS:
The copyright law of the United States (title 17, U.S. Code) governs the making of photocopies or other reproductions of copyrighted material. Any copying of this document without permission of its author may be prohibited by law.

CMU-CS-86-116

Measure Theory and Fair Arbiters

David L. Black
Carnegie-Mellon University
Pittsburgh, PA 15213

7 April 1986

Copyright © 1986 David L. Black

This research was partially supported by NSF Grant Number MCS-82-16706

Table of Contents

- 1. Introduction
- 2. Modeling Mutual Exclusion Elements
 - 2.1. Choice Sets
 - 2.2. Probabilistic Behavior
 - 2.3. Fairness and its Analysis
- 3. Measure Theory
 - 3.1. Introduction
 - 3.2. Product Measures
- 4. Fairness of Mutual Exclusion Elements
 - 4.1. The Basic Result
 - 4.2. Extension to Larger Elements
- 5. Fairness of Multiple Input Arbiters
 - 5.1. Modeling Arbiters
 - 5.2. Fairness Result
- 6. Conclusion
- References

1. Introduction

The existence of fair arbiters and formal specifications for them was a major topic of discussion at the Workshop. One of the many results discussed is that it is possible to create a fair arbiter by adding output delays to a mutual exclusion element [1]. This work builds on that result by investigating the basic fairness properties of mutual exclusion elements and combinations thereof. Rather than working with a particular mutual exclusion element, we abstract the behavior of a class of such elements using a choice set model and a probabilistic specification of the choice inherent in mutual exclusion. This allows us to capture the choice behavior of a mutual exclusion element in a probabilistic structure containing finite and infinite traces. To analyze such structures we employ techniques from the mathematical discipline of measure theory, and in particular the measure theoretic treatment of probability. The major result from this analysis is that mutual exclusion elements are fair under a strong probabilistic notion of fairness. This notion is similar to the standard notions used in [1]; its major advantage is that unlike the standard notions, it can be analyzed using probability and measure theory techniques. A complete introduction and explanation of measure theory is neither possible nor desirable in this context; the interested reader is urged to consult a standard text such as [3]. Similarly we assume a familiarity with the common notions of fairness discussed in [1].

2. Modeling Mutual Exclusion Elements

A mutual exclusion element enforces mutual exclusion among grants to concurrent competing requests for a shared resource. For the present we consider only two input mutual exclusion elements; generalizations to larger elements and arbiters containing multiple mutual exclusion elements will be discussed in a later section. To quantify the fairness of a mutual exclusion element, there are three facets of its behavior that must be modeled:

- Which requests are granted in what order?
- What are the possible choices for each grant?
- For those grants that require a choice between requests, how is that choice made?

We use a choice set model for the first two items and a probabilistic model for the final item; these models are discussed in the following sections. For a two-input mutual exclusion element we denote the competing requests by r_0 and r_1 , and their corresponding grants by g_0 and g_1 .

2.1. Choice Sets

For an individual grant by a mutual exclusion element, we define the *choice set* to be the set of outstanding requests that could be granted (including the request actually granted). To analyze fairness we consider sequences of such grants and corresponding choice sets. Since previous work has shown that infinite notions of fairness are necessary for delay-insensitive arbitration among independent processes [1], we restrict our attention to infinite sequences. Therefore our choice set model of an arbiter is the set of all infinite sequences of $\langle \text{grant}, \text{choice set} \rangle$. We call such a sequence an infinite choice sequence.

Because any two-input mutual exclusion element makes dichotomous choices between competing inputs, it must exhibit metastable behavior. One can intuitively understand metastable behavior as that of a man sitting on a fence who cannot decide on which side of the fence the grass is greener and therefore remains on the fence. The corresponding behavior of circuits (making no choice because both are equally attractive) is real and has been demonstrated for synchronizers, arbiters, and related circuits [2]. Just as the man will eventually get off the fence, a well-designed circuit will eventually make a choice; there is no upper bound on the time required to make this choice, but it is known to be exponentially distributed [4]. In the correspondence of our choice set model with any actual mutual exclusion element, we require that any non-singleton choice set correspond to metastable behavior.

2.2. Probabilistic Behavior

The final aspect of mutual exclusion element behavior that we need to model is the mechanism of making choices among elements of a choice set. For this work we adopt the following simplifying assumptions:

- In any sequence of choices, all choices are mutually independent.
- If a choice set contains both requests, then there are fixed probabilities for each request (r_0, r_1) to be granted, and these probabilities are identical for all such choice sets.

These assumptions formalize our requirement that two-element choice sets correspond to metastable behavior, and model metastable behavior as essentially invariant over choices by fixing the probabilities associated with each grant for a given mutual exclusion element. We denote these probabilities by p_0 and p_1 , and note that their sum must be 1. In practice one would expect these exit or resolving probabilities to have values close to 0.5 [5].

Minor variations of these probabilities do not change our results, but drastic variations (such as one of the probabilities becoming zero) invalidate the results. Chaney has observed short term

dependencies in exit probabilities among choice sets for some flip-flops [6]; in the absence of long-term dependencies we rely on the law of averages to remove any short-term effects from our analysis of infinite sequences of choices (i.e. our first assumption above is justified in part because we are interested only in infinite sequences of choices; it may be an unwarranted assumption for short finite sequences). Finally, Sproull has noted that although these assumptions are reasonable for well-designed mutual exclusion elements, there are pitfalls in designing MOS mutual exclusion elements that may produce elements that violate these assumptions. The major potential problem is that every wire in a MOS circuit has capacitance and therefore memory properties, including the wires used to cross-couple the inverting gates in a mutual exclusion element; the inadvertent introduction of memory in these wires can produce a mutual exclusion element that unfairly favors one of its requests [7]. This is one of the many areas in which appropriate care and diligence are required on the part of the circuit designer.

2.3. Fairness and its Analysis

To analyze the fairness of a mutual exclusion element we need consider only its choice behavior when there is a choice to be made; therefore from any choice set sequence we are interested in the unique subsequence consisting solely of choice sets containing both requests. Such a subsequence is completely specified by the sequence of grants because all the choice sets are identical. Furthermore, we are interested only in such subsequences which are infinite; if this subsequence is finite, then after some finite point in the original sequence every request is immediately granted (i.e. each choice set is of size 1, therefore that request is the granted request). This is certain to be fair under an infinite notion of fairness as it excludes starvation. A consequence of our modeling assumptions is that the likelihood of any such subsequence or collection thereof depends solely upon the probabilities p_0 and p_1 corresponding to the grant choices in the subsequence(s).

We now define fairness for both choice set sequences and mutual exclusion elements.

Definition 1: Weak Fairness for Sequences: A choice set sequence is fair iff any continually asserted request is eventually granted (i.e. the sequence is fair iff for any request appearing in any choice set, EITHER the corresponding grant appears in that or a subsequent sequence element, OR there is a subsequent choice set in which the request does not appear).

When applied to our two-element choice set subsequences this requires that every request eventually be granted in the subsequence; this is a stronger condition than requiring every request to be granted in the corresponding original sequences. The following definition employs this stronger condition.

Definition 2: Probabilistic Fairness for Mutual Exclusion Elements A mutual

exclusion element M is fair if the probability of unfair infinite two-element choice subsequences in the space of all infinite two-element choice subsequences is zero.

Weak fairness for a mutual exclusion element requires that every continually asserted request eventually be granted [1]; the above definition implies the similar condition that every continually asserted request be granted with probability 1. The difference between these two conditions is that the latter allows unfair behavior with probability zero. In a finite domain of possible behaviors, zero probability behaviors cannot occur; for an infinite domain the corresponding condition is that within any infinite sequence of randomly chosen behaviors from the domain, a behavior of probability zero may appear at most finitely many times. Although the difference between this condition and an event never occurring can be of great theoretical importance, we believe that it is relatively unimportant in practice; in both cases the probability of a request not being granted is zero, and this is an adequate notion of fairness for most purposes.

To establish our desired result on the fairness of mutual exclusion elements we need to consider the probability of unfair infinite two-element choice sequences (an infinite set) within the space of all infinite two-element choice sequences (another infinite set). Towards this end we now discuss measure theory, which provides a rich set of tools for this task.

3. Measure Theory

Measure theory is a branch of mathematics concerned with functions that produce the 'size' of sets; these set functions are called *measures*. A primitive example of a measure for intervals of the real number line is the function that produces the length of the interval. Measures are most interesting and most useful when the sets involved are infinite; this is precisely the situation that we are considering in analyzing the fairness of mutual exclusion elements. The particular class of measures we use are often called *probability measures* because they determine the measure of a set by the probability of the 'events' it contains. (The quoted terms are for the reader's intuition only; we do not intend to provide formal definitions.)

3.1. Introduction

This section presents a quick introduction to measure theory abstracted from [3]. Further details can be found there. This section may be skipped on first reading, as it is necessary only for the proof of the main result.

Rings and σ -rings are the mathematical objects on which measures are defined.

Definition 3: Rings and algebras of Sets: A *ring of sets* is a non empty class R of

sets which is closed under the formation of set unions and differences. Formally a non-empty class R is a ring iff

For all $E \in R$ and $F \in R$, $E \cup F \in R$ and $E - F \in R$

A σ -ring of sets is a ring of sets closed under the formation of countable unions. Formally a ring of sets S is a σ -ring iff

If $E_i \in S$ for $i = 1, 2, \dots$ then $\bigcup_{i=1}^{\infty} E_i \in S$

A(n) (σ) -algebra is a (σ) -ring that contains the union of all its elements. An equivalent definition of a(n) (σ) -algebra is that it is a (σ) -ring that is closed under set complementation of its individual elements.

This work relies primarily upon σ -algebras; the other rings and algebras are used to present the definitions and results of measure theory in full generality.

A measure is a function that 'measures' the 'size' of sets in a ring. We now formally define this concept:

Definition 4: Measure: An extended real valued function is a function whose range is $\mathfrak{R} \cup \{+\infty, -\infty\}$ where \mathfrak{R} denotes the real numbers.

A set function is a function whose domain is a class of sets (a ring or an algebra for example).

An extended real valued set function μ defined on a class E is *countably additive* if, for every disjoint sequence $\{E_n\}$ of sets in E whose union is also in E , we have

$$\mu\left(\bigcup_{n=1}^{\infty} E_n\right) = \sum_{n=1}^{\infty} \mu(E_n)$$

A *measure* is an extended real valued, non-negative, and countably additive set function μ , defined on a ring R such that $\mu(\emptyset) = 0$.

For an algebra A , let X be the union of all sets in the algebra. Then a measure μ is *totally finite* iff $\mu(X)$ is finite.

Probability measures are always totally finite because the set X includes all possible 'events', and the probability of at least one 'event' occurring is 1. (i.e. for a probability measure μ , not only is $\mu(X)$ finite, but in fact $\mu(X) = 1$).

Any measure on any ring satisfies the following property, called *countable subadditivity*: for every sequence $\{E_i\}$ of sets in ring R whose union is also in R the following is true:

$$\mu\left(\bigcup_{i=1}^{\infty} E_i\right) \leq \sum_{i=1}^{\infty} \mu(E_i)$$

The difference between this and countable additivity is that the sequence $\{E_i\}$ was of disjoint sets in the definition of countable additivity, whereas the sequence in this definition need not be disjoint; this

results in the equality of countable additivity being replaced by the inequality contained in the above definition.

Our desired result requires reasoning from probabilities of individual choices to probabilities of sequences of choices. In a probability domain, we would state that different universes of choice are involved; the corresponding objects in measure theory are called measure spaces.

Definition 5: Measure Spaces: A measurable space is a set X and a σ -ring \mathbf{S} of subsets of X such that $\cup \mathbf{S} = X$. A subset E of X is measurable iff E belongs to the σ -ring \mathbf{S} .

A measure space is a measurable space (X, \mathbf{S}) and a measure μ on \mathbf{S} . A measure space (X, \mathbf{S}, μ) is totally finite iff the measure μ is totally finite.

In most cases it causes no confusion to denote a measure space or measurable space by the same symbol as the underlying set X ; if it is necessary to call attention to the specific space involved, we use the notation (X, \mathbf{S}) for measurable spaces and (X, \mathbf{S}, μ) for measure spaces.

3.2. Product Measures

To prove our main result we define a measure space corresponding to a single choice, and then extend to sequences of choices by taking cartesian products. The basic result we need from measure theory is the existence of the product probability measures on cartesian products. To state this theorem we need the following additional definitions and results. All results from measure theory are stated here without proof; the proofs can be found in [3].

Theorem 6: If \mathbf{E} is any class of sets, then the class of rings containing \mathbf{E} has a unique minimal element. In addition the class of σ -rings containing \mathbf{E} also has a unique minimal element.

Definition 7: Generated Rings: In view of the preceding theorem, we define the ring generated by \mathbf{E} to be the unique minimal element of the class of rings containing \mathbf{E} and denote it by $\mathbf{R}(\mathbf{E})$. Similarly we define the σ -ring generated by \mathbf{E} to be the unique minimal element of the class of σ -rings containing \mathbf{E} and denote it by $\mathbf{S}(\mathbf{E})$.

For the following if $\{X_i\}$ is an infinite sequence of sets, then X denotes their cartesian product, $X = \prod_{i=1}^{\infty} X_i$. On this basis we can define the remaining concepts that are needed to state the main result of this section.

Definition 8: Measurability for Cartesian Products: For each set X_i , let \mathbf{S}_i be a Σ -algebra of subsets of X_i , and let μ_i be a measure on \mathbf{S}_i such that $\mu_i(X_i) = 1$.

A rectangle is a set of the form $\prod_{i=1}^{\infty} A_i$ where $A_i \subseteq X_i$ for all i and $A_i = X_i$ for all but a finite number of values of i . A measurable rectangle is a rectangle for which each A_i is a measurable subset of X_i .

A subset of $X = \prod_{i=1}^{\infty} X_i$ is *measurable* iff it belongs to the σ -ring (which is actually a σ -algebra) generated by the class of measurable rectangles. We denote this σ -algebra by $\mathbf{S} = \prod_{i=1}^{\infty} \mathbf{S}_i$.

The next definition defines useful terms and notations for dealing with products of measure spaces.

Definition 9: Cylinders: Let J be any subset of I , the positive integers. Then two points (in X), $x = (x_1, x_2, \dots)$ and $y = (y_1, y_2, \dots)$ agree on J if $x_j = y_j$ for every j in J . We write this as $x \equiv y(J)$.

A set E in X is called a *J-cylinder* if $x \equiv y(J)$ implies that x and y either both belong or both do not belong to E .

Finally we adopt the shorthand $X^{(n)} = \prod_{i=n+1}^{\infty} X_i$ for non-negative n . This allows us to write any J -cylinder E as $A \times X^{(n)}$ where $A \subseteq \prod_{i=1}^n X_i$ for some n .

The following theorem is the principle measure-theoretic result necessary to prove our main fairness result.

Theorem 10: Existence of Product Measures: If $\{X_i, \mathbf{S}_i, \mu_i\}$ is an infinite sequence of totally finite measure spaces with $\mu_i(X_i) = 1$, then there exists a unique measure μ on the σ -algebra $\mathbf{S} = \prod_{i=1}^{\infty} \mathbf{S}_i$ with the property that for every measurable set E of the form $A \times X^{(n)}$,

$$\mu(E) = (\mu_1 \times \dots \times \mu_n)(A).$$

The measure μ is called the *product* of the given measures μ_i , and the measure space (X, \mathbf{S}, μ) is called the *cartesian product* of the given measure spaces.

4. Fairness of Mutual Exclusion Elements

4.1. The Basic Result

Our model characterizes mutual exclusion elements by their resolving probabilities from the metastable state (p_0, p_1) . For real mutual exclusion elements one would expect these probabilities to be close to 0.5, but it is unreasonable to assume that they are both exactly 0.5 for all mutual exclusion elements; one would expect slight variations around 0.5. From a theoretical standpoint, the size of the interval around 0.5 that contains these probabilities is immaterial as long as zero is excluded. In other words, expanding that interval from a small area around 0.5 to the entire open unit interval between 0 and 1 changes neither the results nor their proofs, despite our expectation that the probabilities do not lie in most of the interval. Excluding 0 and 1 as probabilities requires that the metastable state exhibit an indecision among possible choices; admitting 0 and 1 allows the situation

in which the same request is almost always granted in any conflict, and thus the corresponding mutual exclusion element cannot possibly be fair. Our main result is the following:

Theorem 11: Any mutual exclusion element for which p_0 and p_1 are not zero is fair under our model and definitions.

Proof: According to our definition of fairness, it is sufficient to show that the probability of the set of unfair infinite two-element choice sequences is zero. For arbitrary positive probabilities p_0 and p_1 (such that $p_0 + p_1 = 1$) we proceed as follows:

Corresponding to a single choice we define the choice measure space $C = (C, \mathcal{T}, \mu_c)$ as:

- $C = \{g_0, g_1\}$
- $\mathcal{T} = \{\emptyset, \{g_0\}, \{g_1\}, \{g_0, g_1\}\}$
- $\mu_c(\emptyset) = 0, \mu_c(\{g_0\}) = p_0, \mu_c(\{g_1\}) = p_1, \mu_c(\{g_0, g_1\}) = p_0 + p_1$

C is a totally finite measure space such that $\mu_c(C) = 1$. C corresponds directly to a single choice from the two-element choice set.

For all positive integers i , let $P_i = C$. Now let (P, \mathcal{S}, μ) be the cartesian product of the P_i . From theorem 10, we conclude that the probability measure μ on P is well defined. P is therefore the probability measure space of infinite two-element choice sequences.

Lemma: Every point in P is measurable and has measure zero.

Proof: Consider an arbitrary one point subset of P , $\{x\}$. x is the infinite sequence (x_1, x_2, \dots) where each $x_i \in \{g_0, g_1\}$. For $n > 0$, let $J_n = \{1, 2, \dots, n\}$. Then (x_1, \dots, x_n) defines a J_n -cylinder E_n which contains $\{x\}$. More precisely we have

$$\{x\} = \bigcap_{n=1}^{\infty} E_n = (\bigcup_{n=1}^{\infty} E_n')'$$

where A' denotes the complement of the set A . \mathcal{S} is a σ -algebra, hence it is closed under complementation and countable union by definition. Since all of the E_n are in \mathcal{S} it follows from the equalities above that $\{x\} \in \mathcal{S}$, therefore $\{x\}$ is measurable by definition.

To find the measure of $\{x\}$, assume without loss of generality that $p_0 \geq p_1$. From theorem 10 and the definition of the measure space P , we know that the measure of E_n is

$\prod_{i=1}^n \mu_c(x_i)$. Under our assumption on p_0 and p_1 , this quantity is bounded from above by $(p_0)^n$. Since $\{x\} \subseteq E_n$ for all n we conclude that

$$\mu(\{x\}) \leq \liminf_{n \rightarrow \infty} (p_0)^n$$

Since $(p_0)^n$ is a geometric series with common ratio p_0 and $0 < p_0 < 1$, the lim inf and the limit are identical and equal to 0. As a measure μ must be non-negative, therefore $\mu(\{x\}) = 0$. □ Lemma.

Lemma: The set of unfair infinite two-element choice sequences is countable.

Proof: An unfair infinite two-element choice sequence must fail to grant some continuously asserted request. This requires that after some finite point the sequence becomes either g_0 repeating or g_1 repeating. Such sequences can be enumerated by enumerating their prefixes (assuming the last element of the prefix is the first element of the repeating portion) and casting out duplicates as part of the enumeration process.

□ Lemma.

Since the set of unfair infinite two-element choice sequences is countable, it is a countable disjoint union of the sets containing its individual elements. Each of these sets has measure zero by the first lemma, therefore the entire set of unfair infinite two-element choice sequences also has measure zero because any measure is countably additive. This proves the theorem. □

4.2. Extension to Larger Elements

Having proved our main result for two input mutual exclusion elements, we now extend it to larger mutual exclusion elements. The major new idea in this proof is the introduction of the excluded choice sets, a useful class of sets that includes the Cantor Set. We informally discuss our methodology before presenting the extended proof.

As the number of inputs increases, the complexity of our choice set model also increases. To avoid changing the spirit of our approach it is necessary to extract subsequences corresponding to each of the finitely many possible multi-element choice sets; each choice set corresponds to a different measure space and product space, so the proof that the unfair sequences have measure zero must now be performed once for each multi-element choice set. For choice sets containing exactly two elements the proof has been given above, but for choice sets containing more than two elements, the set of unfair sequences becomes uncountable requiring a slightly more sophisticated approach.

For choice sets containing three or more elements the main theorem above is still true, but its proof involves the infinite sequence analog of the Cantor Set and related sets; we refer to these as excluded choice sets, and note that they are needed to prove the larger choice set versions of the lemmas in the original proof.

Definition 12: Excluded Choice Sets: For a choice set C and an excluded set $E \subseteq C$, the n^{th} excluded choice set, $\chi_E^n(C)$ for $n \geq 0$, is the set of infinite sequences of elements from C such that no element of E occurs in the subsequence starting at the n^{th} position. (The 1st excluded set forbids the occurrence of elements of E , the 3rd allows them to occur only in the first two positions, etc.).

One of these sets corresponds to the mathematical oddity known as the Cantor Set. The Cantor Set is constructed by taking the unit interval from zero to one including endpoints and deleting its open

middle third (the interval from $1/3$ to $2/3$ without endpoints), then deleting the open middle third from the two intervals thus formed and so on ad infinitum. The Cantor Set is an uncountable nowhere dense set of measure zero which contains all of its limit points; as such it is a common counterexample in mathematics. In our framework the Cantor set corresponds to $\chi_{\{1\}}^1(\{0,1,2\})$ where each infinite sequence is considered as the ternary expansion of a number in the unit interval between 0 and 1 inclusive. To our knowledge, this is the first use of the Cantor Set in computer science.

For an n -input mutual exclusion element, we denote the requests, grants and respective probabilities by $\{r_0, \dots, r_{n-1}\}$, $\{g_0, \dots, g_{n-1}\}$, and $\{p_0, \dots, p_{n-1}\}$. We can now state and prove the generalized theorem.

Theorem 13: Any n -input mutual exclusion element for which all of $\{p_0, \dots, p_{n-1}\}$ are non-zero is fair under our model and definitions.

Proof: Based on our above discussion, it suffices to show that for any choice set containing more than one element, the probability of unfair infinite choice sequences from all the elements is zero. Let n be the size of an arbitrary such set, and without loss of generality assume that it contains $\{g_0, \dots, g_{n-1}\}$. For arbitrary positive probabilities

$\{p_0, \dots, p_{n-1}\}$ such that $\sum_{i=0}^{n-1} p_i = 1$ we proceed as follows:

Corresponding to a single choice we define the choice measure space $C = (C, \mathcal{T}, \mu_c)$ as:

- $C = \{g_0, \dots, g_{n-1}\}$
- $\mathcal{T} = \mathcal{P}(C)$, the powerset of C .
- $\mu_c(A) = \sum_{g_i \in A} p_i$ for $A \in \mathcal{T}$.

C is a totally finite measure space such that $\mu_c(C) = 1$. C corresponds directly to a single choice from the n -element choice set.

For all positive integers i , let $P_i = C$. Now let (P, \mathcal{S}, μ) be the cartesian product of the P_i . Theorem 10 implies that the probability measure μ on P is well defined. P is therefore the probability measure space of infinite n -element choice sequences.

Lemma: The set of unfair infinite n -element choice sequences is the countable non-disjoint union of the of $\chi_E^m(C)$ where E ranges over all singleton sets consisting of exactly one element from C and m ranges over the positive integers.

Proof: An unfair infinite n -element choice sequence must fail to grant some continuously asserted request, say r_i . Such a sequence is in $\chi_{\{g_i\}}^m(C)$ where m is the sequence position at which starvation begins.

Conversely, any element of $\chi_{\{g_i\}}^m(C)$ for m positive must fail to grant the request r_i after position m , and is therefore unfair. □ Lemma.

Lemma:

For $g_i \in \{g_0, \dots, g_{n-1}\}$ and m positive, $\chi_{\{g_i\}}^m(C)$ is measurable in \mathcal{P} and has measure zero.

Proof: For $k > 0$, let $(J_k = \{m, m+1, \dots, m+k-1\})$. Let E_k be the measurable J_k -cylinder defined by excluding g_i from the positions in J_k . E_k contains the desired excluded choice set. More precisely we have

$$\chi_{\{g_i\}}^m(C) = \bigcap_{k=1}^{\infty} E_k = (\bigcup_{k=1}^{\infty} E_k')'$$

where A' denotes the complement of the set A . \mathcal{S} is a σ -algebra, hence it is closed under complementation and countable union by definition. Since all of the E_k are in \mathcal{S} it follows

from the equalities above that $\chi_{\{g_i\}}^m(C) \in \mathcal{S}$, therefore $\chi_{\{g_i\}}^m(C)$ is measurable by definition.

To find its measure, let p_j be the smallest probability in $\{p_0, \dots, p_{n-1}\}$. From theorem 10 and the definition of the measure space \mathcal{P} , we know that the measure of E_k is $\prod_{i=1}^k \mu_c(C - \{g_i\})$. This quantity is bounded from above by $(1-p_j)^k$. Since $\chi_{\{g_i\}}^m(C) \subseteq E_k$ for all k we conclude that

$$\mu(\chi_{\{g_i\}}^m(C)) \leq \liminf_{k \rightarrow \infty} (1-p_j)^k$$

Since $(1-p_j)^k$ is a geometric series with common ratio $1-p_j$ and $0 < 1-p_j < 1$, the \liminf and the limit are identical and equal to 0. As a measure μ must be non-negative, therefore

$$\mu(\chi_{\{g_i\}}^m(C)) = 0.$$

□ Lemma.

From the above two lemmas we know that the set of unfair infinite n -element choice sequences is a countable non-disjoint union of sets of measure 0. Since μ is both non-negative and countably subadditive, the set of unfair infinite n -element choice sequences has measure 0. This proves the theorem. □

This completes the proof of our main fairness theorem for mutual exclusion elements with three or more inputs. We now consider the implications of this result for (multi-input) arbiters implemented with more than one mutual exclusion element.

5. Fairness of Multiple Input Arbiters

This section builds upon our fairness result for mutual exclusion elements to establish a general fairness result for arbiters containing one or more mutual exclusion elements. As with our approach to mutual exclusion elements, we first define an abstract model of such arbiters and then prove our main result within the framework of that model. Throughout this section we are most interested in asynchronous arbiters, i.e. those arbiters that impose no timing constraints on the presentation of the requests among which they arbitrate, and continue to function in the absence of requests from one or more connected processes. Our model and result are general enough to be applicable outside this domain, but they do not apply to prioritized or daisy chain arbiters in which one request has absolute priority over another. Such arbiters must favor a request of higher priority over one of lower priority; this can result in starvation of the lower priority request.

5.1. Modeling Arbiters

Common design and implementation techniques for multiple input arbiters use multiple mutual exclusion elements arranged in a hierarchical or ring formation. A request from one process using the arbiter must then be granted by one or more of these elements before actually being granted by the arbiter. We assume for this work that the arbiter has been correctly designed to implement mutual exclusion and prevent internal deadlock.

The two main methodologies for designing multi-input arbiters using multiple (usually two-input) mutual exclusion elements are the hierarchical and ring methodologies.

- A hierarchical arbiter uses multiple levels of mutual exclusion elements and surrounding circuitry to reduce the number of requests to two which are finally arbitrated by a top-level mutual exclusion element. Thus, a particular request entering the arbiter must obtain a fixed sequence of grants culminating in the grant from the top-level element to receive a grant from the arbiter.
- A ring arbiter arranges mutual exclusion elements and surrounding circuitry in a ring around which a single token or privilege travels. Possession of the token or privilege gives a ring element the right to issue a grant to its attached request; absence of the token requires a request to obtain it, or a wait for it. The mutual exclusion element is needed in each ring element to arbitrate between holding the token to grant a pending request at that element and passing the token along the ring so that some other pending request may be granted. Thus a request must obtain grants from all the mutual exclusion elements required to pass the token around the ring to the corresponding element, and finally from that element itself. This sequence of required grants depends upon the token position when the request is made; since the ring has a fixed number of elements, there are finitely many possible positions and therefore finitely many such sequences are possible.

Based on the above, our model of arbiters is that for each process connected to the arbiter there are one or more finite sequences of mutual exclusion elements from which grants must be obtained in order to cause the arbiter to grant a request from that process. For any particular request from a process, exactly one of these sequences is appropriate; this sequence is determined by the state of the arbiter (token position in the case of a ring) when the request is made.

5.2. Fairness Result

The following definition of fairness is the probabilistic analog of the standard notion of weak fairness. We again emphasize that these two notions are very similar; a fairness result using this definition all but establishes weak fairness for arbiters.

Definition 14: Probabilistic Weak Fairness for Arbiters: An arbiter is probabilistically weakly fair if any continually asserted request is eventually granted with probability 1.

Using our model and this definition, we can prove the following result:

Theorem 15: All arbiters represented by our model are fair in the weak probabilistic sense.

Proof: Let r be an arbitrary request and M_1, \dots, M_n the corresponding sequence of mutual exclusion elements from which grants must be obtained to grant r . Our result on mutual exclusion element fairness implies that M_1 will eventually grant r with probability 1, thus allowing it to be presented to M_2 , which eventually grants it with probability 1, etc. Finally r is presented to M_n which eventually grants it with probability 1, thus causing the arbiter to grant r . Multiplying these probabilities produces the result that r is eventually granted with probability 1 as in the definition of probabilistic weak fairness. Since r was arbitrary this establishes the theorem. □

We believe that this theorem essentially settles the weak fairness question for arbiters; essentially any non-prioritized asynchronous arbiter that has or will be designed falls under our model, and is therefore fair under our definition of fairness. This definition is close enough to weak fairness to essentially settle the weak fairness question for arbiters.

6. Conclusion

In this paper we have considered the fairness of mutual exclusion elements, the most important building block for any arbiter. A probabilistic choice set model has been introduced to capture the choice behavior of such elements. Using this model on infinite sequences we have defined a probabilistic notion of fairness, and shown that mutual exclusion elements are fair in general, provided that a simple assumption about their probabilistic behavior is satisfied. (Any well-designed mutual exclusion element does satisfy the assumption.) We have also extended this result to

establish the fairness of a wide class of arbiters including virtually all known non-prioritized multi-input designs. This essentially settles the weak fairness question for non-prioritized arbiters; in general such arbiters are fair in a sense that is very close to the standard notion of weak fairness.

References

- [1] Black, D. L.
On the Existence of Fair Delay-Insensitive Arbiters: Trace Theory and its Limitations.
Distributed Computing , 1986.
To Appear, draft available as CMU Technical Report CMU-CS-85-173.
- [2] Chaney, T. J.
Anomalous Behavior of Synchronizer and Arbiter Circuits.
IEEE Transactions on Computers C-22:421-422, 1973.
- [3] Halmos, P. R.
Measure Theory.
D. Van Nostrand, 1950.
- [4] Marino, L. R.
General Theory of Metastable Operation.
IEEE Transactions on Computers C-30(2):107-115, February, 1981.
- [5] Rosenberger, F. U.
Personal Communication.
- [6] Sproull, R., I. Sutherland, and C. Molnar.
Seminar on Self-Timed Systems.
Course Notes.
Carnegie-Mellon University, Spring 1985 Includes some unpublished work of Tom Chaney.
- [7] Sproull, R.
Personal Communication.