Checking fairness of a shuffling algorithm numerically

Source: http://sci.tech−archive.net/Archive/sci.math/2007−10/msg05875.html

---

- *From*: Yaroslav Bulatov <yaroslavvb@xxxxxxxxx>

- *Date*: Wed, 31 Oct 2007 00:37:18 −0000

Suppose you have a shuffling algorithm on n cards. You would like to test whether k shuffles randomizes the deck sufficiently. "Randomize sufficiently" means that the variational distance between uniform distribution and actual distribution is less than epsilon. Or alternatively, maximum expected profit that a perfect gambler could make by betting $1 on certain card combinations against a fair house is epsilon.

An obvious difficulty is that there are n! possible permutations, which makes it infeasible to get a decent empirical distribution by counting # of times each permutation appears even for a deck of 52 cards. On other hand, the distribution might be smooth, so perhaps with smoothness assumptions one could use less samples.

The shuffling algorithm I'm looking at is Thorp shuffling where you separate deck into two halves, riffle them, then randomly flip every pair, ie 1234 is riffled to become 1324, then 13 or 24 could be flipped, so 1324,3124,1342,3142 are possible outcomes of one shuffle. It's interesting to get numerical estimates because tight theoretical estimates are not available for this kind of shuffling (O(log(n)^44) is the tightest known upper bound for shuffles needed to mix a deck of n cards)

So the question is −− how might one test for 1−100 cards (with small probability of error) whether distribution over permutations is within epsilon after k shuffles

.