Lectures on Stochastic Matching: Guises and Applications

E. G. Coffman, Jr.
AT&T Bell Laboratories
Murray Hill, NJ 07974 USA

Extended Abstract

Since the pioneering work on multi-dimensional packing by Karp, Luby, and Marchetti-Spaccamela [KLM] in 1984, stochastic matching problems in two and three dimensions have surfaced in the analysis of a surprising number of algorithms and systems, with applications in operations research, electrical engineering, and computer science. Thus, the related theory has provided an important analysis tool, one that, unfortunately, has yet to be applied by more than a handful of researchers. Of the several variants of stochastic matching that we will discuss, the following two have played fundamental roles. Let n plus points and n minus points (i.e., points labeled with $+$'s and $-$'s) be chosen independently and uniformly at random in the unit square. Let M_n denote a matching of the plus points to minus points and let (P^-, P^+) denote a pair of matched points (or an edge) in M_n.

Euclidean Matching: Let $d(P^-, P^+)$ denote the Euclidean distance between P^-, P^+, and let M_n be a matching that minimizes the total distance between matched points, $D_n = \sum_{(P^-,P^+)\in M_n} d(P^-, P^+)$. Find $E[D_n]$.

Up-Right (or Ordered) Matching: Consider only up-right matchings M_n such that $(P^-, P^+) \in M_n$ only if P^+ is up and to the right of P^-, i.e., if $P^- = (x_1, y_1)$, $P^+ = (x_2, y_2)$, then $x_1 \le x_2$, $y_1 \le y_2$. Let U_n be the number of unmatched points in an up-right matching of maximum cardinality. Find $E[U_n]$.

As stated, these problems have proven to be quite difficult, a property that is inherited in particular cases from the problems in which stochastic matching finds application. However, tight asymptotic bounds have been derived. Specifically, for Euclidean matching, Ajtai, Komlos, and Tusnády [AKT] proved that $E[D_n]$ grows like $\sqrt{n \log n}$ for large n; to be more precise, we can express this in the standard $\Theta(\cdot)$ notation,

$$E[D_n] = \Theta(\sqrt{n \log n}) \ .$$

For up-right matching, Shor [Sho], Leighton and Shor [LS] and Rhee and Talagrand [RT] proved that

$$E[U_n] = \Theta(\sqrt{n}\log^{3/4} n) \ . \tag{$*$}$$

Subsequent, simpler proofs were given by Coffman and Shor [CS] and Talagrand [Tal]; see Chapter 3 in [CLu] for a general treatment. Asymptotic bounds on tail probabilities are also available; indeed, it is often the case that such bounds provide the desired estimates of expected values. The result in $(*)$ has also been generalized to the ordered matching problem in $d \geq 3$ dimensions. (If (P^-, P^+) is in a d-dimensional ordered matching, then in every dimension the coordinate of P^- must be at most that of P^+.) As shown in [KLM], one obtains $E[U_n] = \Theta(n^{(d-1)/d})$, $d \geq 3$.

As is common with asymptotic results of this sort, the absence of information about multiplicative constants is compensated by a greater generality of the results. Obvious examples of robustness are rescalings: n can be replaced by αn for any given $\alpha > 0$ and, in the planar case, the unit square can be replaced by an $a \times b$ rectangle for any given $a, b, > 0$; such changes modify only the hidden multiplicative constant. For less obvious examples, we note that the following changes to problem instances also leave $(*)$ unchanged:

(i) the numbers of plus and minus points are independently Poisson distributed with mean n; or the total number of points is Poisson distributed with mean n, and each point is equally likely to be a plus or minus independently of the others;

(ii) the points are restricted to lie on an $n \times n$ lattice, or the problem is similarly discretized in just one of the dimensions; or

(iii) n plus points are fixed at the vertices of a $\sqrt{n} \times \sqrt{n}$ lattice and n minus points are chosen at random as before.

These lectures will take results of the above type as background theory (proofs of such results would require at least four lectures by themselves). Instead, we concentrate on techniques that exploit these results in analyzing a variety of mathematical models of engineering problems. The techniques require only a modest background in applied probability, e.g., certain basic probability inequalities, elementary results from the theory of random walks, central limit theorems, and Chernoff bounds.

As illustrated in the problems below (e.g., see Problem 1), stochastic matching often arises quite naturally in the analysis; in such cases, some form of matching can usually be recognized in the sample functions of the underlying stochastic process. In other cases, however, as in Problem 2 below, the matching problem may be well disguised; it can be seen only after a substantial problem reduction has been made. Even after a stochastic matching problem has been identified, a nontrivial analysis may still be required to account for new variations in problem instances. For example, in Problem 3 below, an instance of up-right matching arises, but the horizontal components of the minus points do not have a uniform distribution, except in a certain asymptotic sense.

The lectures will expand at length on the above matters, as they emerge in the analysis of the problems listed below. If time permits, we shall also discuss a matching problem in processor-ring communications.

Problem 1. A Selection-Replacement Process on the Circle [CGS].

Given n points on a circle, a selection-replacement operation removes one of the points and replaces it by another. To select the removed point, an extra point P, uniformly distributed, arrives at random and starts moving counterclockwise around the circle; P removes the first point it encounters. A new random point, uniformly distributed, then replaces the removed point. The quantity of interest is $d = d(n)$, the distance that the searching point P must travel to select a point. In particular, consider the mean of d in the stationary version of the selection-replacement process. Sample functions of the process can be represented by $+$ and $-$ points on a cylinder representing the product space of the circle and a time axis; $+$'s denote the selected points on the circle and $-$'s indicate the points P that remove selected points. Up-right matchings of $-$'s to $+$'s pair off points P with the points they remove. The expected horizontal component of the distance between matched points gives $E[d] = \Theta\left(\frac{\log^{3/2} n}{n}\right)$, as shown in [CGS].

In a computer application, the circle represents a track on a disk memory, P is a read-write head, the n points mark the beginnings of n files and d determines the time wasted as the head moves from the end of the last file processed to the beginning of the next. The number n is a parameter of the service rule (the next service goes to one of the n customers waiting the longest).

Problem 2. First Fit Bin Packing with Discrete Item Sizes [CJSW].

A list L of n items is to be packed into a sequence of unit capacity bins. The first-fit (FF) rule packs each successive item into the first bin of the sequence that has room for it. We discuss an average-case analysis of FF in the discrete uniform model: the item sizes are drawn independently and uniformly at random from the set $\{1/k, \ldots, (k-1)/k\}$, for some $k > 1$. Let $FF(L)$ denote the wasted space in the FF packing of L, i.e., the total space still available in the occupied bins. It is proved in [CJSW] that $E[FF(L)] = O(\sqrt{nk})$, i.e., there exists a constant $c > 0$ such that $E[FF(L)] \le c\sqrt{nk}$ for all n, k sufficiently large. In the proof of this result, items in a problem instance are represented by points in 2 dimensions, one denoting the item's index in the list L and the other denoting its "folded" size, i.e., an item of size s is plotted as a $-$ with coordinate s if $s \le 1/2$ and as a $+$ with coordinate $1-s$ if $s > 1/2$. (Note that a $+$ item can fit in a bin with a $-$ item only if the size coordinate of the $+$ is larger than that of the $-$.) A certain class of up-right matchings of problem instances corresponds to the matchings produced by a modified FF rule that packs at most two items in a bin and never uses fewer bins than FF. The bound on $E[FF(L)]$ is then obtained from a bound on the expected number of unmatched points in such up-right matchings.

Problem 3. Dynamic Storage Allocation [CLe].

A computer storage device is represented by a sequence of adjacent cells with sizes s_i nondecreasing in $i = 1, 2, \ldots$. Files arrive and depart by Poisson processes, with each item placed at its time of arrival into a smallest empty cell large enough for the item. The known, average number of files in storage in the stationary regime is denoted by n, and assumed to be an integer. One assumes a given distribution of file sizes such that s_i, $1 \le i \le n$, can be chosen so that a file size is equally likely to fall in any of the intervals $[0, s_1], [s_1, s_2], \ldots, [s_{n-1}, s_n]$; for all $i > n$, one chooses $s_i = s_n$, a largest file size. Trajectories of the stationary storage process are represented in the two dimensions of time t and cell index i; a $-$ point at (i, t) denotes a new arrival at time t with a size in $(s_{i-1}, s_i]$, $s_0 = 0$. At plus point at (i, t) denotes a departure from cell i at time t. Up-right matchings pair arriving files ($-$'s) to the departures ($+$'s) creating the empty cells in which the files are placed. Of interest is the expected number of interior unused cells, i.e., the number W_n of empty cells among the first m, where cell m is the highest indexed occupied cell. An analysis of the up-right matchings gives $E[W_n] = \Theta(\sqrt{n}\log^{3/4} n)$, as shown in [CLe].

Problem 4. Probabilistic Analysis of a Vehicle Routing Problem [BCSS].

Consider n points distributed uniformly at random in some rectangular region. The points represent customers with demands for some commodity supplied by a depot, which is represented by an additional point with a given location in the region. An unlimited number of equal capacity vehicles are available at the depot for delivery of customer demands. Vehicles are to be routed to the n customers, with each customer being visited by at most one vehicle, and all vehicles making a round-trip tour, so that all demands are satisfied and the total route length L_n of all vehicles is minimized. The demands are i.i.d. uniform draws from $[0, 1]$, each giving the required fraction of a vehicle's capacity.

In a third (vertical) dimension plot customer demands directly above the corresponding customer locations in the (horizontal) plane, using the folding convention of Problem 2 to decide whether a point is labeled with a $+$ or $-$. Now construct a (three-dimensional) *upward matching* of $+$'s and $-$'s that minimizes the number of unmatched points; the only requirement that must be met by matched $+$ and $-$ points is that the $+$ must be above the $-$. This yields an obvious heuristic routing algorithm. Vehicles serve at most two customers in a round-trip tour; matched customers are served by the same vehicle, and unmatched customers are served by vehicles that serve no other customer. It is proved in [BCSS] that $E[L_n] = nE[d] + \Theta(n^{2/3})$, where $E[d]$ is the average distance between the customers and the depot. Moreover, it is shown that this is a best possible asymptotic result in the sense that the expected total route length under an optimal routing algorithm is also equal to $nE[d] + \Theta(n^{2/3})$.

References

[AKT] Ajtai, M., Komlós, J., and Tusnády, G., "On Optimal Matchings," *Combinatorica*, **4** (1984), 259–264.

[BCSS] Bramel, J., Coffman, E. G., Jr., Shor, P. W., and Simchi-Levi, D., "Probabilistic Analysis of the Capacitated Vehicle Routing Problem with Unsplit Demands," **40** (1992), 1095–1106.

[CGS] Coffman, E. G., Jr., Gilbert, E. N., and Shor, P. W., "A Selection-Replacement Process on the Circle," *Ann. Appl. Prob.*, **3** (1993), 802–818.

[CJSW] Coffman, E. G., Jr., Johnson, D. S., Shor, P. W., and Weber, R. R., "Bin Packing with Discrete Item Sizes, Part III: Tight Bounds on First Fit," AT&T Bell Laboratories, Murray Hill, NJ 07974 (to appear).

[CLe] Coffman, E. G., Jr. and Leighton, F. T., "A Provably Efficient Algorithm for Dynamic Storage Allocation," *J. Comp. Sys. Sci.*, **38** (1989), 2–35.

[CLu] Coffman, E. G., Jr. and Lueker, G. S., *Probabilistic Analysis of Packing and Partitioning Algorithms*, Wiley-Interscience, New York, 1991.

[CS] Coffman, E. G., Jr. and Shor, P. W., "A Simple Proof of the $O(\sqrt{n} \log^{3/4} n)$ Up-Right Matching Bound," *SIAM J. Disc. Math.*, Feb., 1991, 48–57.

[KLM] Karp, R. M., Luby, M., and Marchetti-Spaccamela, A., "A Probabilistic Analysis of Multidimensional Bin Packing Problems," in *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, (1984), 289–298.

[LS] Leighton, T. and Shor, P. "Tight Bounds for Minimax Grid Matching with Applications to the Average Case Analysis of Algorithms, *Combinatorica*, **9** (1989), 161–187.

[RT] Rhee, W. T. and Talagrand, M., "Exact Bounds for the Stochastic Upward Matching Problem," *Transactions of the American Mathematical Society*, **307** (1988), 109–125.

[Sho] Shor, P. W., "The Average-Case Analysis of Some On-Line Algorithms for Bin Packing," *Combinatorica*, **6** (1986), 179–200.

[Tal] Talagrand, M., "Matching Theorems and Empirical Discrepancy Computations using Majorizing Measures," The Ohio State University, Columbus, OH (1991).