

On Zipf Models for Probabilistic Piece Selection in P2P Stored Media Streaming

Carey Williamson
Department of Computer Science
University of Calgary
Calgary, AB, Canada
Email: carey@cpsc.ucalgary.ca

Niklas Carlsson
Dept of Computer and Information Science (IDA)
Linköping University
Linköping, Sweden
Email: niklas.carlsson@liu.se

Abstract—The Zipf distribution is widely used to model Web site popularity, video popularity, and file referencing behavior. In recent published work, we proposed and evaluated a Zipf-based policy for probabilistic piece selection in Peer-to-Peer (P2P) media streaming. In this current paper, we revisit this Zipf model in more detail, and identify two fundamentally different modeling approaches, namely regenerative versus degenerative Zipf models. We illustrate the differences between the two models, provide refined analytical models for each, and validate the models with simulations in the context of P2P media streaming. The results show that the regenerative model is more appropriate for P2P streaming, because of its stronger sequential progress.

Keywords—P2P streaming; piece selection; Zipf; modeling

I. INTRODUCTION

On-demand streaming of stored media files has generated significant interest in the networking research literature [2], [12], [29], [32]. A wide range of delivery mechanisms have been proposed, including centralized, distributed, and hybrid architectures. Centralized approaches typically rely on the classic client-server architecture, augmented with scalable streaming protocols, IP multicast, or combinations thereof. Decentralized approaches include the peer-to-peer (P2P) paradigm [1], [7], [12], [15], [19], [29], [31], as well as peer-assisted streaming [8], [9], [16]. Hybrid approaches using Content Delivery Networks (CDN) or cloud-based streaming have also emerged recently.

In this paper, we focus particularly on P2P approaches, which offer a scalable approach for Internet-based media streaming. The P2P paradigm has been used successfully for live media streaming [19], [30], [31], as well as for on-demand streaming [1], [7], [10], [12], [21], [29]. In either of these scenarios, the sequential playback requirements of media streaming applications are paramount. Unlike file download applications, which can obtain pieces in any order, on-demand video streaming applications ultimately require pieces in sequential order in order to facilitate playback.

In prior work [23], [24], we argued that the analysis of P2P media streaming is decomposable into *download progress* and *sequential progress*, which can be analyzed separately. Furthermore, improving one component can usually be done without compromising the other. We also

developed analytical models for four different piece selection policies: In-Order, Random, Portion, and Zipf. These models provide important insights into the efficiency of on-demand media streaming in P2P networks [24].

In this current paper, we further explore the dynamics of probabilistic piece selection policies, particularly for the Zipf piece selection policy. Our work is motivated by the limited accuracy of the Zipf modeling results in [24], and a desire to improve the model. The main insight that emerges in our work is that there are two fundamentally different ways of modeling Zipf referencing behavior, which we call *regenerative* and *degenerative* Zipf models. Both approaches have been used in the networking literature (e.g., a version of the regenerative approach in [7], [24], and degenerative approaches in [5], [6], [27]), with few details provided about the modeling assumptions or their performance impacts. We explore the details of these two Zipf models in this paper, as well as their implications.

The remainder of the paper is organized as follows. Section II provides some background on the Zipf distribution, and reviews recent literature on P2P media streaming. Section III presents our system model. Section IV presents our modeling results, using both analysis and simulation. Section V provides some additional simulation validation and comparison of our models. Finally, Section VI concludes the paper.

II. BACKGROUND AND RELATED WORK

This section provides background material on the Zipf distribution, and summarizes prior work on piece selection policies for P2P media streaming.

A. Zipf Distribution

The Zipf distribution exhibits a power-law structure that is seen in many Internet traffic measurement studies [4]. The primary feature is a non-uniform pattern of referencing, wherein a subset of the items garners most of the interest. Furthermore, this non-uniform referencing pattern continues to hold recursively, even when the most frequent items are removed from consideration. Web site popularity, file system usage, Web page referencing, and video object popularity are

Table I
EXAMPLES OF HARMONIC NUMBERS $H_{M,\theta}$

M	1	2	3	4	5
$\theta = 1$	1.000	1.500	1.833	2.083	2.283
$\theta = 2$	1.000	1.250	1.361	1.424	1.464

examples of domains in which Zipf or Zipf-like referencing patterns are seen [3], [4], [6].

Mathematically, the Zipf distribution is expressed as $Prob(i) = c/i^\theta$, where $Prob(i)$ is the probability of referencing the item with rank i , and c is a normalizing constant. When $\theta = 0$, $c = \frac{1}{M}$, and all pieces are equally likely, just as in the Uniform distribution. When $\theta = 1$, c is the M -th Harmonic number, $H_M = \sum_{i=1}^M \frac{1}{i}$. When $\theta > 1$, c is the M -th Generalized Harmonic number, $H_{M,\theta} = \sum_{i=1}^M \frac{1}{i^\theta}$. Table I shows the first few Harmonic numbers for $\theta = 1$ and $\theta = 2$.

Table II shows examples of the Zipf probability distribution for $1 \leq M \leq 5$ items, when $\theta = 1$. Table III shows the corresponding values for $\theta = 2$. In general, as θ increases, the Zipf distribution becomes even more concentrated on the most popular items. Figure 1 shows the rank-frequency plot for $\theta = 1$ and $\theta = 2$ when $M = 10$. On a log-log scale, the power-law structure of the Zipf distribution manifests itself as a straight line, with (negative) slope θ .

Table II
ZIPF PROBABILITY DISTRIBUTION ($\theta = 1$)

i	$M = 1$	$M = 2$	$M = 3$	$M = 4$	$M = 5$
1	1.0	0.6667	0.5454	0.4800	0.4380
2	-	0.3333	0.2727	0.2400	0.2190
3	-	-	0.1818	0.1600	0.1460
4	-	-	-	0.1200	0.1095
5	-	-	-	-	0.0875

The Zipf distribution is widely used in synthetic workload models for Web proxy workloads, media object popularity, and file referencing behavior [3], [4], [5]. In this paper, we restrict our focus to understanding Zipf models in the context of P2P media streaming applications.

B. Media Streaming

Many P2P streaming approaches assume a BitTorrent-like system as the underlying content delivery mechanism. BitTorrent [11] is a popular peer-to-peer file sharing system

Table III
ZIPF PROBABILITY DISTRIBUTION ($\theta = 2$)

i	$M = 1$	$M = 2$	$M = 3$	$M = 4$	$M = 5$
1	1.0	0.8000	0.7347	0.7024	0.6833
2	-	0.2000	0.1837	0.1756	0.1708
3	-	-	0.0816	0.0781	0.0759
4	-	-	-	0.0439	0.0427
5	-	-	-	-	0.0273

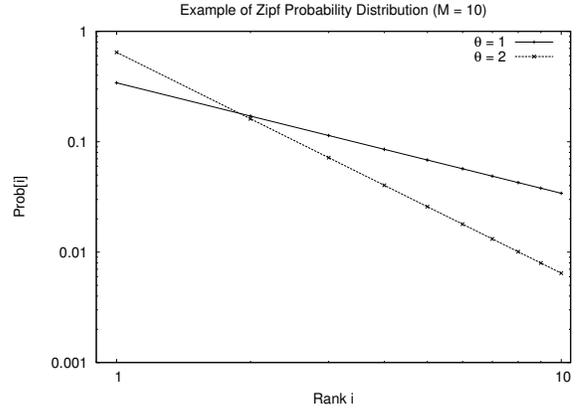


Figure 1. Example of Zipf probability distribution ($M = 10$)

designed for efficient file downloads. The set of peers participating in the download of a particular file is known as a swarm. BitTorrent splits files into pieces, which can be downloaded concurrently from multiple peers in parallel. Furthermore, the pieces can be downloaded in any order. There are several analytical models characterizing BitTorrent and the download efficiency of swarm-based systems [14], [22], [25], [26], [28].

Piece diversity is one of the central considerations in these systems. In BitTorrent, peers use a Rarest-First piece selection policy when deciding which piece to obtain next from a providing peer. This policy gives preference to under-represented pieces in the system, at least within a neighbourhood of cooperating peers. This approach ensures adequate piece diversity to maintain good download progress [17], [18], [25].

Compared to download progress, sequential progress is less well-studied [12], [23], [24]. Fan *et al.* [12] explicitly address the tradeoffs between optimality, robustness, and sequentiality in P2P streaming. One counter-intuitive result is that simple In-Order piece selection, which is clearly preferable for media streaming, has an adverse effect on BitTorrent, halving the system efficiency [24]. One reason for this is that In-Order selection breaks the “tit-for-tat” reciprocity of BitTorrent. Another reason is the limited piece diversity, which leads to non-uniform load in the system [13]. Carlsson *et al.* [7] proposed several probabilistic piece selection policies as a means to enhance piece diversity, without compromising sequential progress. Yang *et al.* [29] also study the effect of the request scheduling discipline on streaming latency.

In this paper, we consider the efficiency of probabilistic piece selection policies for on-demand stored media streaming in P2P systems.

III. SYSTEM MODEL AND PRIOR RESULTS

A. Problem Statement and Assumptions

The focus in this paper is on probabilistic piece selection policies in P2P media streaming systems. In particular, we are interested in the sequential progress of these policies.

The notion of sequential progress is defined in [24]. The idea is as follows: ‘‘After having retrieved k pieces, what is the probability that a peer possesses pieces 1 through j (inclusive) to enable streaming?’’ Figure 2 shows an example, in which a peer has retrieved $k = 5$ of M pieces, but has only obtained $j = 3$ of the initial pieces. Thus the sequential progress is 3. Table IV summarizes the notation used in our analysis.

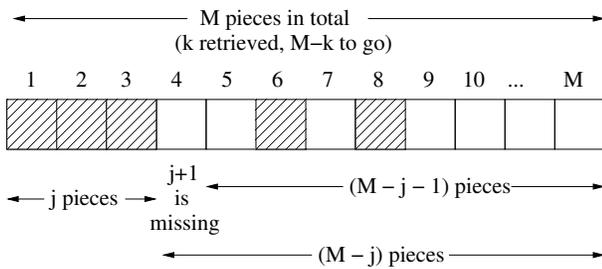


Figure 2. Example of sequential progress concept ($j = 3$, $k = 5$) [24]

Table IV
MODEL NOTATION

Symbol	Description
M	Number of pieces in media object
k	Number of pieces retrieved so far
j	Number of consecutive initial pieces of object retrieved
$E[j k]$	Expected sequential progress j after obtaining k pieces
p	Probability of In-Order selection in Portion policy
q	Probability of selecting piece $j + 1$ on next retrieval
θ	Exponent of Zipf distribution
H_M	M -th Harmonic number
$H_{M,\theta}$	M -th Generalized Harmonic number with exponent θ

B. Sequential Progress Analysis

Figure 3 illustrates, for $M = 10$, the expected sequential progress $E[j|k]$ of the four piece selection policies (In-Order, Random, Portion, and Zipf) that were analyzed in [24].

By definition, the In-Order policy is ideal in terms of sequential progress. Each peer simply retrieves the file pieces in sequential order from 1 to M . This pattern is illustrated by the straight diagonal line in the sequential progress plot.

The Random piece selection policy chooses pieces uniformly at random. As illustrated in Figure 3, this policy provides poor sequential progress. Specifically,

$$E[j|k] = \frac{k}{M - k + 1}. \quad (1)$$

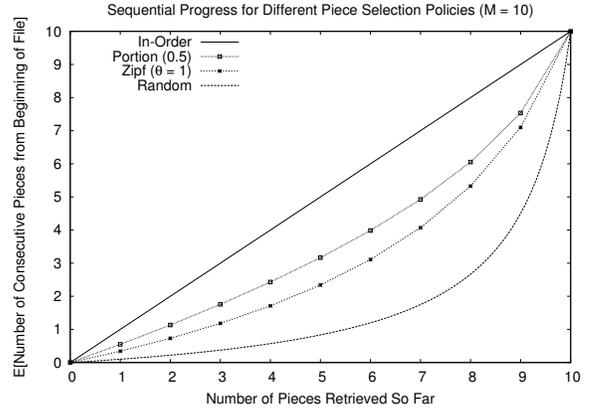


Figure 3. Sequential progress for piece selection policies ($M = 10$)

The sequential progress is slow initially, but increases monotonically and accelerates throughout the download. Additional analysis shows that:

$$\text{Var}[j|k] = \frac{k(M+1)(M-k)}{(M-k+2)(M-k+1)^2}. \quad (2)$$

The variance is a monotonically increasing function of k , though it degenerates to 0 when all M pieces are retrieved.

C. Portion Piece Selection Policy

The Portion policy [7] is a hybrid between In-Order and Random. It has a single configuration parameter p . At each step, it chooses pieces according to the In-Order policy with probability p , and according to the Rarest-First (i.e., Random) policy with probability $1 - p$. This section summarizes the analysis of the Portion policy from [24].

Let $P(j, k)$ be the probability that a peer has obtained exactly the first j in-order pieces after retrieving k of the M total pieces. By definition, the $P(j, k)$ values must satisfy $\sum_{j=0}^k P(j, k) = 1$, for a given k . Furthermore, the expected sequential progress $E[j|k]$ can be calculated as:

$$E[j|k] = \sum_{i=0}^k iP(i, k). \quad (3)$$

Let $q_{j,k}$ be the probability that the first missing piece ($j+1$) is selected for retrieval next. With the Random policy, $q_{j,k} = \frac{1}{M-k}$. With the Portion policy,

$$q_{j,k} = p + (1-p)\frac{1}{M-k}. \quad (4)$$

The first term represents In-Order selection based on the parameter p , while the second term represents the chance of ‘‘accidentally’’ selecting the next needed piece at random. This expression assumes that the missing pieces (excluding the first missing piece) are uniformly distributed, and are equally likely.

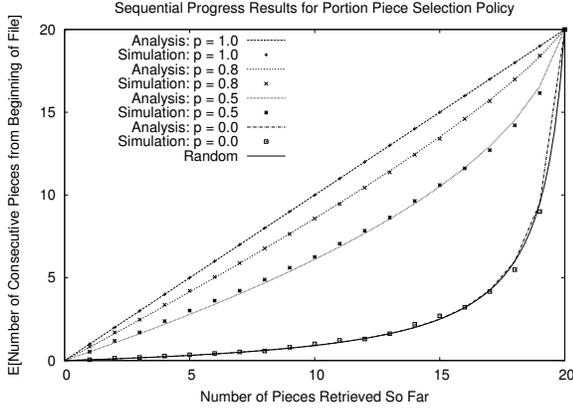


Figure 4. Validation results for Portion piece selection policy [24]

With this notation, probability $P(j, k)$ is expressed as:

$$P(j, k) = P(j, k-1)(1 - q_{j,k-1}) + \sum_{i=0}^{j-1} \left[P(i, k-1)q_{i,k-1}A_{i,k-1}^{(j-i-1)} \right]. \quad (5)$$

The first term in this expression represents the case when the needed piece $j+1$ is *not* selected; as a result, the sequential progress stays the same. The second term represents the case in which piece $j+1$ is successfully obtained. In this case, the sequential progress advances by at least one piece, and possibly more, as detailed next.

Let $A_{j,k}^n$ be the probability of having exactly n consecutive in-order pieces starting at position $j+2$ when the missing piece $j+1$ is selected. Assuming $k-j$ uniformly distributed pieces in the range $[j+2, M]$, we have:

$$A_{j,k}^n = \begin{cases} \left(1 - \frac{k-j-n}{M-j-n-1}\right) \prod_{i=0}^{n-1} \frac{k-j-i}{M-j-i-1}, & \text{if } n < k-j \\ \prod_{i=0}^{n-1} \frac{k-j-i}{M-j-i-1}, & \text{if } n = k-j \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

The factors in these expressions correspond to conditional probabilities, expressing the likelihood that each piece to the right of $j+1$ is (or is not) selected.

The Portion model is highly accurate, and has been validated by simulation, as illustrated in Figure 4. Two special cases of this model deserve mention. When $p \rightarrow 0$, the results match Random piece selection. When $p \rightarrow 1$, the results match In-Order. These observations are evident from Equation 4.

D. Zipf Piece Selection Policy

In the Zipf probabilistic piece selection policy, a Zipf distribution is used to skew the probability bias towards selecting lower-numbered pieces [7]. This section summarizes the analysis of the Zipf policy from [24].

The Zipf analysis mirrors that for the Portion policy. For simplicity, we again assume that the missing pieces (except

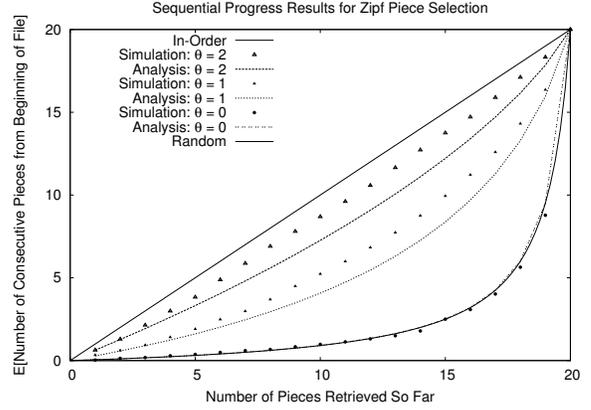


Figure 5. Validation results for Zipf piece selection policy [24]

the first missing piece) are uniformly distributed. Under this assumption, the corresponding expression for $q_{j,k}$ is:

$$q_{j,k} \approx \frac{1}{1 + \frac{M-k-1}{M-j-1} \sum_{i=2}^{M-j} \frac{1}{i^\theta}}. \quad (7)$$

The interpretation of Equation 7 is as follows. The summation term in the denominator represents the cumulative weight for selecting pieces to the right of piece $j+1$ in Figure 2, while the preceding coefficient prorates the sum based on the proportion¹ of these pieces that have already been retrieved. The numerator and the additional term in the denominator represent the unit weight assigned to piece $j+1$ in the Zipf model. The rest of the analysis is the same as for the Portion policy.

While the Zipf model is reasonably accurate, as shown in Figure 5, the simulation results show that the model underestimates sequential progress. As with the Portion policy, there are two special cases. When $\theta \rightarrow 0$, the model matches Random piece selection. When $\theta \rightarrow \infty$, it matches the In-Order policy. These observations follow directly from Equation 7.

Nonetheless, there are several deficiencies with this model. First, the model is complicated, and does not have a closed form. Rather, the model must be evaluated numerically using the foregoing equations. Second, the model is not very accurate. There are perceptible gaps between the simulation and analytical results in Figure 5. These modeling issues are addressed in the remainder of the paper.

IV. NEW MODEL AND RESULTS

This section presents new analytical modeling results for the Zipf probabilistic piece selection policy. We start with a statement of our new closed-form model, and then provide the details and insights behind the model development.

¹Note that equal weights are assumed for these pieces, despite the monotonically decreasing probabilities in a Zipf distribution.

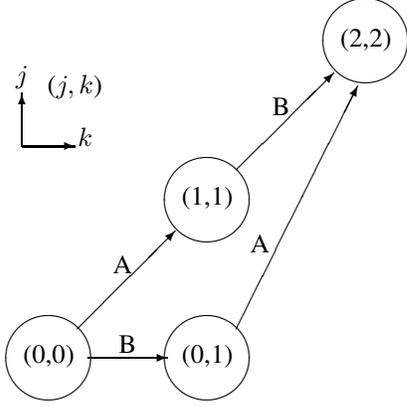


Figure 6. State Space Example for Probabilistic Piece Selection ($M = 2$)

A. Model Summary

Our primary result is for the Degenerative Zipf piece selection policy: when $\theta = 1$, the expected sequential progress after retrieving k of M pieces is:

$$E[j|k] \approx M \frac{(H_M - H_{M-k})}{H_M}. \quad (8)$$

B. Modeling Details

The initial insights for our model came from studying very small cases, such as the $M = 2$ case illustrated in Figure 6. In this scenario, we refer to the pieces of the media object as A and B. Two retrieval sequences are possible for this object, namely AB and BA.

Both sequences are illustrated in Figure 6, using the (j, k) state space notation indicated in the upper left of the diagram. Piece selection, with increasing k values, is indicated horizontally in the diagram, while sequential progress, with increasing j values, is indicated vertically. Clearly, sequential progress j can never exceed k , but it can match k at intermediate points, and always reaches k when the entire object is retrieved. Each of the states and arrows in Figure 6 has an associated probability, though these values are excluded from the diagram for simplicity. These probabilities depend on the piece selection policy in use (e.g., Random, Portion, Zipf), as will be explained shortly.

The next useful case to consider is $M = 3$, as illustrated in Figure 7. In this example, there are 6 possible retrieval orders: ABC, ACB, BAC, BCA, CAB, and CBA. Each of these sequences is represented in the diagram, with the latter two using italic fonts to reduce ambiguities.

A recursive relationship is apparent when comparing Figure 7 and Figure 6. For example, if the first piece retrieved is A, then the remaining states in the upper right of Figure 7 (for the remaining pieces B and C) match Figure 6 (for pieces A and B). Similarly, if B is the first piece chosen, then the remaining states represent the retrieval of A and

C (in either order), analogous to Figure 6. Finally, if C (in italics) is chosen first, then the remaining states (for pieces A (in italics) and B) correspond to those in Figure 6 (for A and B).

The recursive relationship between Figure 7 and Figure 6 suggests a recurrence relation as a natural way to formulate a Zipf analytical model. For $\theta = 1$, the proposed recurrence is:

$$E[j|k] = \frac{M q}{M - k + 1} + E[j|k - 1]. \quad (9)$$

Intuitively, the sequential progress after k piece selections is at least as large as it was after $k - 1$ pieces, and possibly larger if the needed piece $j + 1$ is selected (with probability q). Using Figure 7 as an example, a sequential progress of $j = 2$ can be achieved either by choosing piece B after the initial selection of A (increasing sequential progress by 1), or selecting piece A while already having piece B in hand (increasing sequential progress by 2). The likelihood of these transitions depends on the state probabilities.

This formulation is very similar to the Portion policy model in Equation 6, except that the recurrence relation expresses sequential progress directly, rather than via the state probabilities $P(j, k)$ and the (recursive) sequential progress $A_{j,k}^n$. Furthermore, our new expression is amenable to a closed form solution, as indicated earlier.

C. Zipf Model Variants

An important issue in the Zipf model is how to determine the probabilities for states and transitions. When exploring our model, we identified two fundamentally different approaches, which we refer to as Regenerative and Degenerative Zipf models. Figure 8 illustrates the key differences between these two models, using an example scenario with $M = 3$ and $\theta = 1$.

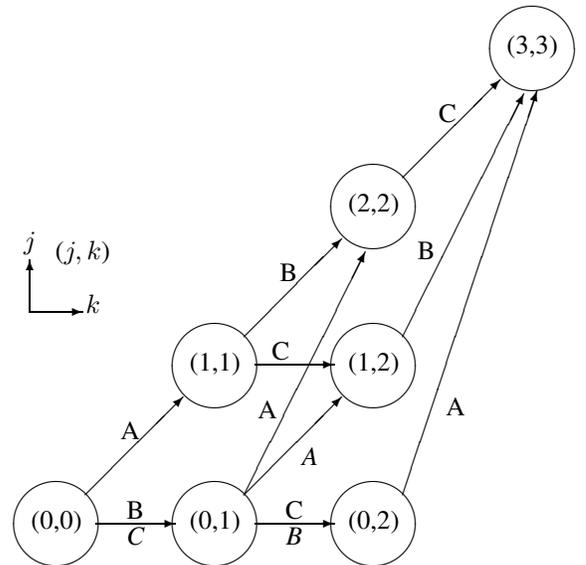


Figure 7. State Space Example for Probabilistic Piece Selection ($M = 3$)

In the Regenerative Zipf model, the Zipf distribution is re-instantiated after each piece selection, albeit with fewer pieces remaining to be retrieved. This model is illustrated in the upper row of Figure 8. The table on the left shows the $M = 3$ pieces initially available for selection, and the probability value associated with each piece. The three tables on the right show the resulting Zipf configurations for $M = 2$ pieces, based on the selection made from the leftmost table. As can be observed, the probability values are *identical* in each of the three tables. Specifically, the topmost item is twice as likely as the second item, as defined by the Zipf distribution for $M = 2$ and $\theta = 1$. In this model, only the *number of pieces* matters, and not what piece was chosen at each step.

In the Degenerative Zipf model, the structure of the Zipf distribution degrades after each piece selection, and becomes less and less Zipf-like. In this model, it is not only the number of pieces that matter, but also *which piece* was selected at each step. This model is illustrated in the lower row of Figure 8. Once again, the table on the left shows the $M = 3$ pieces initially available for selection, and the probability value associated with each piece. The three tables on the right show the resulting configurations for $M = 2$ pieces, based on the selection made from the leftmost table. If item A is selected (with probability $\frac{6}{11}$), then the remaining two items have their probabilities renormalized (relative to $\frac{5}{11}$), while *preserving the relative probabilities*. For example, item B remains 50% more likely than item C, just as it was in the leftmost table. Note that this distribution is no longer Zipf. Similar normalization principles apply for the other two tables shown, based on the initial item selected from the leftmost table. If item C is selected first (with probability $\frac{2}{11}$), then the remaining items A and B still have a Zipf distribution, as shown in the rightmost table. This case is identical to that for the Regenerative Zipf distribution above, but the other two cases are distinctly different.

The foregoing discussion is relevant because it affects the probabilities associated with the states and transitions in Figure 7. For example, if item A is selected first, then does item B have a 66.7% chance of being selected next (Regenerative model), or is it only a 60% chance (Degenerative model)?

Figure 9 shows a more elaborate example for $M = 4$, where the differences in piece selection probabilities between the two models are even more apparent. Note that the updated probability values in the Degenerative model could be lower, higher, or the same as those in the Regenerative model.

D. Model Implications

We conducted simulation experiments to better understand the differences between these two models, and their performance implications.

Figure 10 shows histogram plots illustrating the behavior of these models. The plot in Figure 10(a) is for $M = 3$ pieces. In this scenario, there are $3! = 6$ possible orderings in which the pieces can be retrieved: ABC, ACB, BAC, BCA, CAB, and CBA. We list these sequences in canonical (lexicographic) order from left to right in the plot. The upper part of the plot shows the relative frequencies of these retrieval sequences observed over 1000 iterations of the simulation, for the Degenerative Zipf model. The lower part of the plot shows the corresponding results for the Regenerative Zipf model. If the two models were identical, the plots would mirror each other symmetrically along the horizontal axis. However, several asymmetries can be observed. For example, the sequence BAC is more prevalent than ACB in the Regenerative Zipf model, while the opposite is true in the Degenerative Zipf model. As another example, the sequences BCA and CBA are almost equally prevalent in the Degenerative Zipf model, while BCA is far more prevalent in the Regenerative Zipf model.

The two other plots in Figure 10 show results for $M = 4$ and for $M = 5$ pieces, respectively. As the number of possible retrieval orderings grows, more differences between the models become apparent. In particular, the Regenerative Zipf model has a greater propensity to generate In-Order substrings. There are pronounced peaks in the histogram for the Regenerative model, induced by restoring the Zipf structure after each piece selection. These peaks are not nearly as prominent in the Degenerative Zipf model.

The recursive structure of the Zipf models is apparent when comparing the plots in Figure 10. For example, the 3-piece subsequence depicted in Figure 10(a) shows up four times from left to right in Figure 10(b), depending on the first piece retrieved. The relative magnitude and profile of these histograms matches with the models described. For example, the prominence of ABCDE in the Regenerative model in Figure 10(c) is consistent with the product of the first row of probability values (for $i = 1$) in Table II. One would expect the pattern ABCDE to occur 7.6% of the time in the Regenerative model, and BACDE would occur half as often. In the Degenerative model, ABCDE would occur about 4% of the time, with BACDE occurring only slightly less frequently than this.

The impact of these models on sequential progress are illustrated in Figure 11. Figure 11(a) shows the sequential progress for both models when $M = 4$. At this number of pieces, there is little difference between the two models. Figure 11(b) shows the sequential progress results for $M = 8$ pieces. For this setting, the two models are similar, although the Regenerative one tends to be slightly above the Degenerative one in terms of sequential progress. For $M = 16$ pieces in Figure 11(c), there is a noticeable difference between the two models (based on the 95% confidence intervals shown from 100 simulation replications). This observation is even more pronounced in Figure 11(d), for $M = 32$ pieces.

Regenerative Zipf Model

Item	Prob	Total
A	0.5454	0.5454
B	0.2727	0.8181
C	0.1818	1.0

if item A is chosen

Item	Prob	Total
B	0.6667	0.6667
C	0.3333	1.0

if item B is chosen

Item	Prob	Total
A	0.6667	0.6667
C	0.3333	1.0

if item C is chosen

Item	Prob	Total
A	0.6667	0.6667
B	0.3333	1.0

Degenerative Zipf Model

Item	Prob	Total
A	0.5454	0.5454
B	0.2727	0.8182
C	0.1818	1.0

if item A is chosen

Item	Prob	Total
B	0.6000	0.6000
C	0.4000	1.0

if item B is chosen

Item	Prob	Total
A	0.7500	0.7500
C	0.2500	1.0

if item C is chosen

Item	Prob	Total
A	0.6667	0.6667
B	0.3333	1.0

Figure 8. Example of Regenerative (upper row) and Degenerative (lower row) models for Zipf piece selection ($M = 3, \theta = 1$)

Regenerative Zipf Model

Item	Prob	Total
A	0.48	0.48
B	0.24	0.72
C	0.16	0.88
D	0.12	1.0

if item A is chosen

Item	Prob	Total
B	0.5454	0.5454
C	0.2727	0.8182
D	0.1818	1.0

if item B is chosen

Item	Prob	Total
A	0.5454	0.5454
C	0.2727	0.8182
D	0.1818	1.0

if item C is chosen

Item	Prob	Total
A	0.5454	0.5454
B	0.2727	0.8182
D	0.1818	1.0

if item D is chosen

Item	Prob	Total
A	0.5454	0.5454
B	0.2727	0.8182
C	0.1818	1.0

Degenerative Zipf Model

Item	Prob	Total
A	0.48	0.48
B	0.24	0.72
C	0.16	0.88
D	0.12	1.0

if item A is chosen

Item	Prob	Total
B	0.4615	0.4615
C	0.3077	0.7692
D	0.2308	1.0

if item B is chosen

Item	Prob	Total
A	0.6316	0.6316
C	0.2105	0.8421
D	0.1579	1.0

if item C is chosen

Item	Prob	Total
A	0.5714	0.5714
B	0.2857	0.8571
D	0.1429	1.0

if item D is chosen

Item	Prob	Total
A	0.5454	0.5454
B	0.2727	0.8182
C	0.1818	1.0

Figure 9. Example of Regenerative (upper row) and Degenerative (lower row) models for Zipf piece selection ($M = 4, \theta = 1$)

Figure 11 shows that the Zipf analytical model in [24] does not match *either* of the two Zipf approaches very well. The Regenerative Zipf model in the simulation provides much stronger sequential progress than predicted by the Zipf analytical model, while the Degenerative Zipf model provides much weaker sequential progress. In the “end game”, once all the higher probability pieces have been chosen, the Degenerative model resembles Random piece selection. Since the differences between piece probabilities are small in the tail of the Zipf distribution, the results are similar to uniform random piece selection.

E. Degenerative Zipf Model

Our initial efforts in refining the Zipf model focused on the Degenerative model, for which deeper understanding was sought. Formulating this model using a recurrence relation was a critical step, with the recurrence based on k , rather than M .

The intuition behind the particular recurrence relation given in Equation 9 is the following. M is a normalizing factor. q represents the probability of selecting a piece that advances the sequential progress. The factor $\frac{1}{M-k+1}$ reflects the relative contribution made to the (recursively remaining) sequential progress by each successful piece selection. For the $M = 3$ example in Figure 7, these contributions would

be $\frac{1}{3}$, $\frac{1}{2}$, and 1, respectively, for the transitions along the diagonal from bottom left to top right.

For $\theta = 1$, we obtained the following closed-form result:

$$E[j|k] \approx M \frac{(H_M - H_{M-k})}{H_M}. \quad (10)$$

This is an *approximate* model, obtained via repeated expansion of the recurrence relation in Equation 9, assuming a *fixed* probability $q = q_{0,0}$ throughout the entire download.

Equation 10 is compact, simple, and intuitively appealing. It is also reasonably accurate, as will be demonstrated in Section V. The boundary cases are easy to verify. Specifically, $E[j|k] = 0$ when $k = 0$, $E[j|k] = \frac{1}{H_M} = q_{0,0}$ when $k = 1$, and $E[j|k] = M$ when $k = M$.

Based on the structural insights from the $\theta = 1$ model, we developed a corresponding expression for $\theta = 0$:

$$E[j|k] \approx \frac{M}{M-k+1} \frac{(H_{M,\theta} - H_{M-k,\theta})}{H_{M,\theta}}. \quad (11)$$

Since $H_{i,\theta} = i$ when $\theta = 0$, this expression reduces to the Random model in Equation 1. This formula also provides another way to interpret Equation 10, namely that the probability q plays an important role in accelerating sequential progress compared to Random.

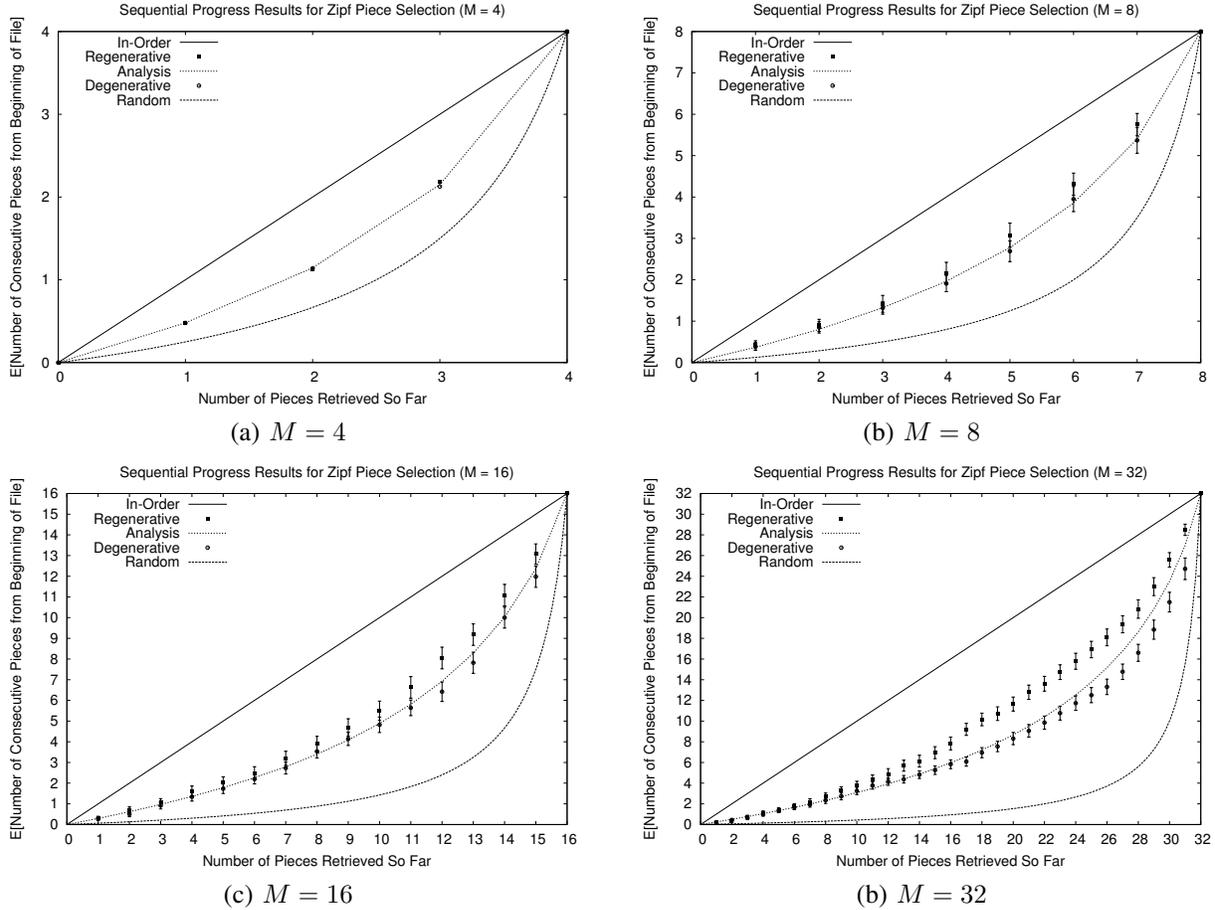


Figure 11. Sequential progress simulation results for Zipf piece selection policies

F. Regenerative Zipf Model

The Regenerative model is more complicated, since it requires different $q = q_{j,k}$ values in each step of the recurrence. We formulate this model generically as $E[j|k] = \sum_{i=1}^k A_i q_i$, where q_i is the probability of selecting the next needed piece, and A_i is how much the sequential progress advances when this piece is selected. In the Regenerative model, $q_i = \frac{1}{H_{M-i+1,\theta}}$, but A_i depends on θ .

For $\theta = 0$, we exploit the fact that $E[j|i] - E[j|i-1] = \frac{i}{M-i+1} - \frac{i-1}{M-i+2}$ to obtain $A_i = \frac{M+1}{M-i+2}$, which leads to an exact sequential progress formula as follows:

$$E[j|k] = \sum_{i=1}^k \frac{M+1}{M-i+2} \frac{1}{H_{M-i+1,\theta}}. \quad (12)$$

Two boundary cases for the latter expression are easy to verify. Specifically, $E[j|k] = 0$ when $k = 0$, and $E[j|k] = \frac{1}{H_{M,\theta}}$ when $k = 1$.

For larger values of θ , the shape of the A_i function varies. By definition, we know that $A_1 = 1$. We have also calculated an exact expression for $A_2 = 1 + \frac{1}{2^\theta H_{M,\theta}}$. The general form

appears to be:

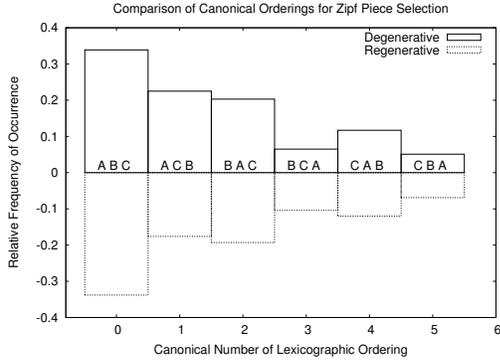
$$A_k \approx 1 + \sum_{i=1}^{k-1} \frac{1}{a_i^\theta H_{M,\theta}} + o\left(\frac{1}{H_{M,\theta}^2}\right), \quad (13)$$

where each a_i is a small positive integer. We do not yet have a compact closed form for this model, and thus must evaluate the model results numerically. When $\theta = 1$, A_i is very close to linear. For $\theta < 1$, the curve is convex (i.e., concave upwards), similar to the Random ($\theta = 0$) case. For $\theta > 1$, the curve is concave (i.e., concave downwards), and is (tightly) upper bounded by the Harmonic series $H_{k,\theta}$.

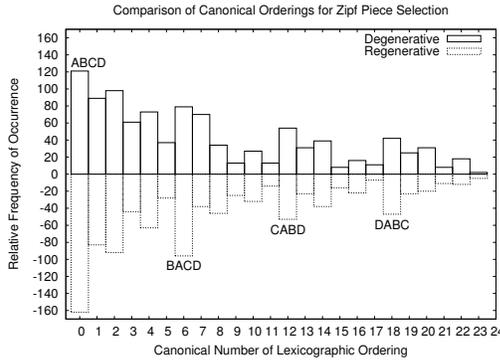
V. SIMULATION RESULTS

Figure 12 shows results to validate our revised models, and compare their properties.

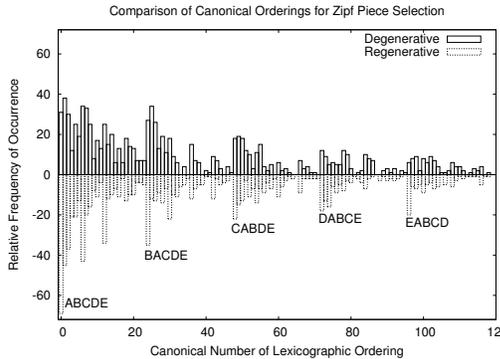
Figure 12(a) shows results for the Regenerative model, with $M = 32$. The value of θ is varied between 0 and 2. When $\theta = 0$, both the simulation results and the analytical results match those for the Random policy, as expected. As θ increases, the sequential progress curves move upward, with the Regenerative policy behaving more and more like In-Order. The 95% confidence intervals also become



(a) $M = 3$



(b) $M = 4$



(c) $M = 5$

Figure 10. Canonical sequence histograms for Zipf piece selection policies

tighter, since there is less variability in the piece selections generated.

Figure 12(b) shows the corresponding results for the Degenerative model. For $\theta = 0$, the simulation and analytical results match the Random policy, as expected. As θ increases, the sequential progress curves move upward, but more slowly than with the Regenerative model. The simulation and analytical results for a given θ are reasonably close together in the Degenerative model, though the analytical model is still somewhat conservative. For the largest θ value considered ($\theta = 2$), the Degenerative Zipf behaves quite differently from the In-Order policy. The slow “end game”

progress is a distinctive characteristic of the Degenerative Zipf policy.

One key assumption in the derivation of the compact formula for Degenerative sequential progress was a fixed probability q for successful retrieval of piece $j + 1$ at each step. This assumption is reasonable for the Degenerative model, but not suitable for the Regenerative model, which has monotonically increasing probability values for the most-needed piece.

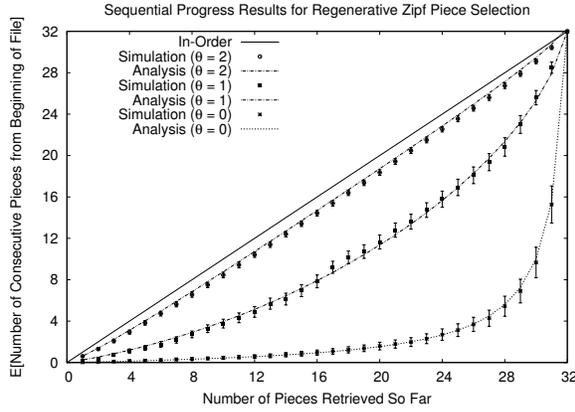
Figure 13(a) shows an example of how the Zipf selection probabilities for piece $j + 1$ evolve in the two models, when $M = 32$. For the Regenerative model (upper line), this probability increases monotonically, since it depends only on the number of pieces remaining. For the Degenerative model, Figure 13 shows the average probabilities (middle line) for missing piece $j + 1$ from 1000 independent simulation runs of the Degenerative model. This probability is non-monotonic. For most of the download, the average probability for piece $j + 1$ is about midway between those of the Regenerative and Random models, and its overall mean is very close to $q_{0,0} = 0.246$. This graph illustrates why a single fixed value of q is a reasonable approximation for the Degenerative model, at least for $M = 32$ and $\theta = 1$. The approximation is less accurate for other values of θ , since the mean probability for piece $j + 1$ differs substantially from $q_{0,0}$ (e.g., 0.1918 versus 0.1006 for $\theta = 0.5$, and 0.3313 versus 0.6195 for $\theta = 2$).

Figure 13(b) shows the average sequential progress advance A_i in each of the models, again for $M = 32$. For the Random policy, the simulation results match the analytical model. For the Regenerative model, the advances are lower, and are always upper bounded by the Harmonic number sequence. The Degenerative model has advances that are in between those of the other two models.

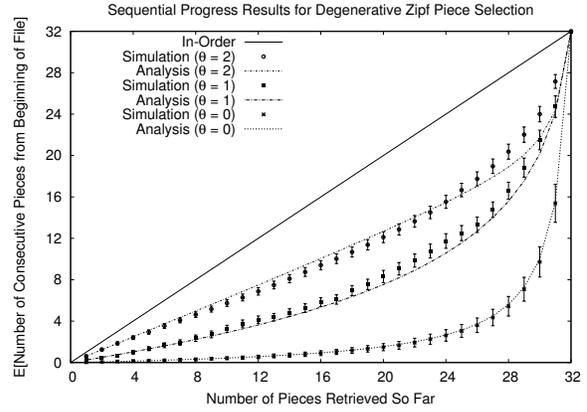
Figure 14 shows one final validation experiment for our models, for $M = 64$ and $\theta = 1$. For these parameter settings, the mean relative error of the Degenerative model is 6.0%, while that of the Regenerative model is 2.6%. These two models provide more accurate results than the original model from [24], which is shown as “Old Analysis” in Figure 14.

VI. CONCLUSIONS

In this paper, we evaluated two fundamentally different approaches to modeling Zipf referencing patterns, namely Regenerative and Degenerative Zipf models. Both models have been used in the previous literature for synthetic workload modeling, but without much discussion of their behavioral properties. In our work, we use analysis and simulation to highlight the differences between the two models. In particular, we provide revised analytical models for each, and validate the models with simulations of P2P stored media streaming. The results show that the Regenerative model is more suitable for P2P streaming, because of its stronger sequential progress.

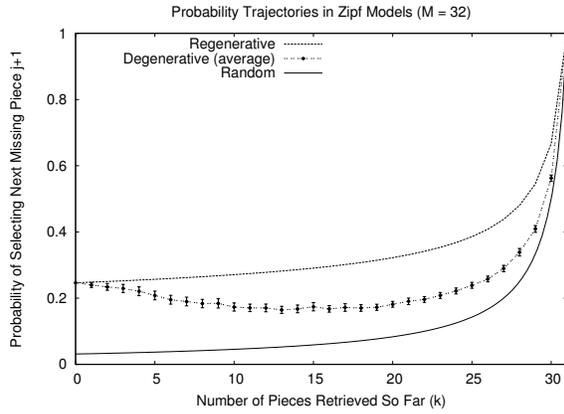


(a) Regenerative Zipf Model

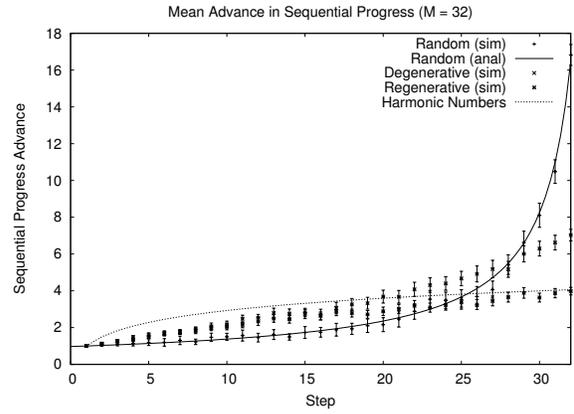


(b) Degenerative Zipf Model

Figure 12. Modeling results for different piece selection policies ($M = 32$)



(a) Probability trajectories



(b) Advances in sequential progress

Figure 13. Behavioral characteristics of Zipf piece selection policies ($M = 32, \theta = 1$)

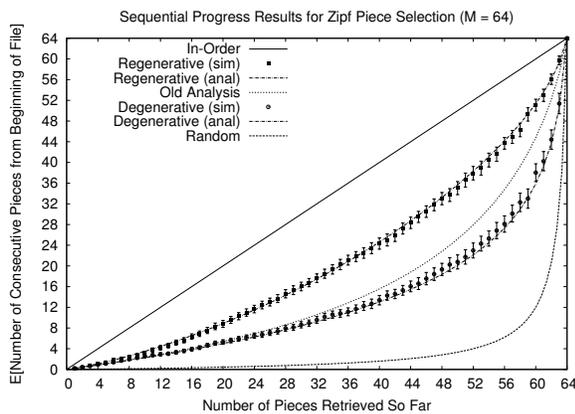


Figure 14. Model validation for Zipf piece selection policies ($M = 64$)

ACKNOWLEDGEMENTS

Financial support for this work was provided by Canada's Natural Sciences and Engineering Research Coun-

cil (NSERC). The authors are grateful to Pierre Beldor for suggesting the initial recurrence relation for the sequential progress of the Zipf piece selection policy. The authors thank the IEEE MASCOTS reviewers, and our assigned shepherd, for their help in improving the final version of our paper.

REFERENCES

- [1] S. Annapureddy, C. Gkantsidis, and P. Rodriguez, "Providing Video-on-Demand using Peer-to-Peer Networks", *Proceedings of Workshop on Internet Protocol TV (IPTV)*, Edinburgh, Scotland, May 2006.
- [2] A. Bestavros and S. Jin, "OSMOSIS: Scalable Delivery of Real-time Streaming Media in Adhoc Overlay Networks", *Proceedings of ICDCS Workshops*, pp. 214–219, Providence, RI, May 2003.
- [3] Y. Borghol, S. Mitra, S. Ardon, N. Carlsson, D. Eager, and A. Mahanti, "Characterizing and Modeling Popularity of User-generated Videos", *Performance Evaluation*, Vol. 68, No. 11, pp. 1037-1055, November 2011.

- [4] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web Caching and Zipf-like Distributions: Evidence and Implications", *Proceedings of IEEE INFOCOM*, pp. 126-134, New York, NY, March 1999.
- [5] M. Busari and C. Williamson, "On the Sensitivity of Web Proxy Cache Performance to Workload Characteristics", *Proceedings of IEEE INFOCOM*, pp. 1225-1234, Anchorage, AL, March 2001.
- [6] M. Busari and C. Williamson, "ProWGen: A Synthetic Workload Generation Tool for the Simulation Evaluation of Web Proxy Caches", *Computer Networks*, Vol. 38, No. 6, pp. 779-794, June 2002.
- [7] N. Carlsson and D. Eager, "Peer-assisted On-demand Streaming of Stored Media using BitTorrent-like Protocols", *Proc. IFIP/TC6 Networking*, pp. 570-581, Atlanta, GA, May 2007.
- [8] N. Carlsson and D. Eager, "Modeling Priority-based Incentive Policies for Peer-assisted Content Delivery Systems", *Proc. IFIP/TC6 Networking*, pp. 421-432, Singapore, May 2008.
- [9] N. Carlsson, D. Eager, and A. Mahanti, "Peer-assisted On-demand Video Streaming with Selfish Peers", *Proc. IFIP/TC6 Networking*, pp. 586-599, Aachen, Germany, May 2009.
- [10] Y. Choe, D. Schuff, J. Dyaberi, and V. Pai, "Improving VoD Server Efficiency with BitTorrent", *Proc. ACM Multimedia*, pp. 117-126, Augsburg, Germany, September 2007.
- [11] B. Cohen, "Incentives Build Robustness in BitTorrent", *Proceedings of Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, June 2003.
- [12] B. Fan, D. Andersen, M. Kaminsky, and K. Papagiannaki, "Balancing Throughput, Robustness, and In-Order Delivery in P2P VoD", *Proceedings of ACM CoNEXT*, pp. 570-581, Philadelphia, PA, December 2010.
- [13] A. Gai, F. Mathieu, F. de Montgolfier, and J. Reynier, "Stratification in P2P Networks: Application to BitTorrent", *Proceedings of ICDCS*, Toronto, Canada, June 2007.
- [14] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang, "Measurement, Analysis, and Modeling of BitTorrent-like Systems", *Proceedings of ACM Internet Measurement Conference*, pp. 35-48, Berkeley, CA, October 2005.
- [15] R. Kumar, Y. Liu, and K. Ross, "Stochastic Fluid Theory for P2P Streaming Systems", *Proceedings of IEEE INFOCOM*, pp. 919-927, Anchorage, AK, May 2007.
- [16] R. Kumar and K. Ross, "Peer Assisted File Distribution: The Minimum Distribution Time", *Proceedings of IEEE Workshop on Hot Topics in Web Systems and Technologies*, Boston, MA, November 2006.
- [17] A. Legout, N. Liogkas, E. Kohler, and L. Zhang, "Clustering and Sharing Incentives in BitTorrent Systems", *Proc. ACM SIGMETRICS*, pp. 301-312, San Diego, CA, June 2007.
- [18] A. Legout, G. Urvoy-Keller, and P. Michiardi, "Rarest First and Choke Algorithms Are Enough", *Proceedings of ACM Internet Measurement Conference*, pp. 203-216, Rio de Janeiro, Brazil, October 2006.
- [19] X. Liao, H. Jin, Y. Liu, L. Ni, and D. Deng, "AnySee: Peer-to-Peer Live Streaming", *Proceedings of IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [20] M. Lin, B. Fan, D. Chiu, and J. Lui, "Stochastic Analysis of File-Swarming Systems", *Performance Evaluation*, Vol. 64, No. 9-12, pp. 856-875, October 2007.
- [21] J. Luo, Y. Tang, and S. Yang, "Chasing: An Efficient Streaming Mechanism for Scalable and Resilient Video-on-Demand Service over Peer-to-Peer Networks", *Proceedings of IFIP Networking*, pp. 642-653, Coimbra, Portugal, May 2006.
- [22] L. Massoulié and M. Vojnovic, "Coupon Replication Systems", *Proceedings of ACM SIGMETRICS*, pp. 2-13, Banff, Canada, June 2005.
- [23] N. Parvez, C. Williamson, A. Mahanti, and N. Carlsson, "Analysis of BitTorrent-like Protocols for On-Demand Stored Media Streaming", *Proceedings of ACM SIGMETRICS*, pp. 301-312, Annapolis, MD, June 2008.
- [24] N. Parvez, C. Williamson, A. Mahanti, and N. Carlsson, "Insights on Media Streaming Progress using BitTorrent-like Protocols for On-Demand Streaming", *IEEE/ACM Transactions on Networking*, Vol. 20, No. 1, pp. 637-650, June 2012.
- [25] D. Qiu and R. Srikant, "Modeling and Performance Analysis of BitTorrent-Like Peer-to-Peer Networks", *Proceedings of ACM SIGCOMM*, pp. 367-378, Portland, OR, August 2004.
- [26] Y. Tia, D. Wu, and K. Ng, "Modeling, Analysis and Improvement for BitTorrent-Like File Sharing Networks", *Proceedings of IEEE INFOCOM*, Barcelona, Spain, April 2006.
- [27] C. Williamson and R. Bunt, "Characterizing Short Term File Referencing Behaviour", *Proceedings of IEEE IPCCC*, pp. 651-660, Phoenix, AZ, March 1986.
- [28] X. Yang and G. Veciana, "Service Capacity of Peer to Peer Networks", *Proceedings of IEEE INFOCOM*, pp. 2242-2252, Hong Kong, March 2004.
- [29] Y. Yang, A. Chow, L. Golubchik, and D. Bragg, "Improving QoS in BitTorrent-like VoD Systems", *Proceedings of IEEE INFOCOM*, San Diego, CA, March 2010.
- [30] M. Zhang, L. Zhao, Y. Tang, J. Luo, and S. Yang, "Large-scale Live Media Streaming over Peer-to-Peer Networks through Global Internet", *Proc. Workshop on Advances in Peer-to-Peer Multimedia Streaming*, pp. 21-28, Singapore, November 2005.
- [31] X. Zhang, J. Liu, B. Li, and T. Yum, "CoolStreaming/DONet: A Datadriven Overlay Network for Peer-to-Peer Live Media Streaming", *Proceedings of IEEE INFOCOM*, pp. 2102-2111, Miami, FL, March 2005.
- [32] Y. Zhou, D. Chiu, and J. Lui, "A Simple Model for Analyzing P2P Streaming Protocols", *Proceedings of ICNP*, pp. 226-235, Beijing, China, October 2007.