

On Models for Object Lifetime Distributions

Darko Stefanović
Department of
Electrical Engineering
Princeton University
Princeton, NJ 08544
darko@acm.org

Kathryn S. McKinley
Department of
Computer Science
University of Massachusetts
Amherst, MA 01003
mckinley@cs.umass.edu

J. Eliot B. Moss
Department of
Computer Science
University of Massachusetts
Amherst, MA 01003
moss@cs.umass.edu

ABSTRACT

Analytical models of memory object lifetimes are appealing because having them would enable mathematical analysis or fast simulation of the memory management behavior of programs. In this paper, we investigate models for object lifetimes drawn from programs in object-oriented languages such as Java and Smalltalk. We present certain postulated analytical models and compare them with observed lifetimes for 58 programs. We find that observed lifetime distributions do not match previously proposed object lifetime models, but do agree in salient shape characteristics with the gamma distribution family used in statistical survival analysis for general populations.

Categories and Subject Descriptors

D.3.4 [Programming Languages]: Processors—*memory management (garbage collection)*; G.3 [Mathematics of Computing]: Probability and Statistics—*survival analysis*

General Terms

Measurement

Keywords

Object lifetimes, lifetime distributions, garbage collection modelling

1. INTRODUCTION

If we can develop accurate analytical models for object lifetimes in object-oriented programs, they will enable faster and more thorough exploration of memory management techniques. For instance, given a model of object lifetimes, we could compute an estimate of copying costs of a generational or some other garbage collector. If distribution models and garbage collector models are simple enough, we may even arrive at closed-form analytical descriptions; but even if both are quite complicated, we can use the lifetime distributions to drive simulations of a proposed garbage collector scheme.

Lifetime models are not sufficient for exploring garbage collection, because they do not account for heap pointer structure effects: the direct cost of pointer maintenance (including write barriers), and the copying cost increase owing to the excess retention of objects, both

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ISMM '00 Minneapolis MN USA
Copyright ACM 2000 1-58113-263-8/00/10...\$5.00

costs being present in generational and other heap-partitioning schemes. Nevertheless, the models can be useful as a tool for preliminary evaluation (and understanding) of collector performance.

Observed object lifetime behaviors are inherently discrete; we measure the lifetime of each object and arrive at a discrete distribution. Most performance-related propositions have dealt with mortality, which is a derivative form (in the sense of the calculus of functions $\mathbb{R} \rightarrow \mathbb{R}$); since obtaining a derivative of a discrete observed function involves inherently arbitrary smoothing decisions, it has been difficult to characterize the mortality of observed distributions, let alone to match it to an analytical model. The extremely fast decay of objects exacerbates the situation: most models developed in other domains are for much more slowly decaying populations.

If we knew which distribution family describes typical object behaviors, we could fit observed lifetimes to the model of that family, and find the best matching instance (i.e., its parameters). In fact, a running program could recognize the lifetime distribution of objects allocated (overall, or at a particular allocation site) and adjust collection policies accordingly. But it is not yet known which family this may be.

In the following, we first briefly introduce terms and notation related to lifetime distributions (with more details in the Appendix), then review what assumptions have been made implicitly (or stated explicitly) in past research in garbage collection. We develop models based on a plausible qualitative characterization of lifetimes; namely, that past lifetime is a strong predictor of future lifetime. Lastly, we put the models to the test of empirical evidence against actual lifetime distributions from object-oriented programs, using a graphical device recommended in statistical survival analysis.

We shall find that these particular analytical models are not a good match for actual distributions. In search of a good match, we tentatively consider several well-known distributions families, but must conclude that none is completely satisfactory: objects in programs are a much different population from those statisticians have examined. This indicates on the one hand the need for further modelling effort to achieve good matches that can be validated against experiment, and on the other, the need to derive such models from first principles, viz. from program semantics.

2. BACKGROUND MATERIAL

The lifetime of an object is defined as the amount of allocation that occurs between the allocation of the object and its demise.¹ We view

¹The actual point of demise depends on the accuracy of the memory management scheme. In the empirical data reported here, that scheme is an accurate-roots garbage collector performing full-heap collection

object lifetime as a random variable. Future studies may look at object lifetimes as stochastic processes, and in this context, distinguish each allocation site as generating a different process. Here, we do not attempt such fine distinctions.

Actual object lifetimes are natural numbers, thus discrete probability distributions are the obvious representation. However, continuous models are used for mathematical ease and convenience. Below we review some definitions and symbols from probability theory as they apply to survival analysis; further details appear in the Appendix along with a summary of properties of commonly used analytical distribution families. Although we shall write more formulae than are typically seen in garbage collection literature, the mathematical apparatus is elementary.

The *survival* function of a random variable \mathbf{L} is $s_{\mathbf{L}}(t) = \wp\{\mathbf{L} \geq t\}$. For object lifetimes, it expresses what fraction of original allocation volume is still live at age t . We usually drop the subscript \mathbf{L} . The survivor function is a monotone non-increasing function. The *probability density* function for object lifetime is $f(t) = -s'(t)$. Occasionally we also use the *cumulative distribution* function $F(t) = 1 - s(t)$. The *mortality* function is $m(t) = \frac{f(t)}{s(t)} = -\frac{d}{dt} \log s(t)$, and it expresses the age-specific death rate. Mortality is also known as the *hazard* function (and written $h(t)$) in the literature on lifetime analysis [4].

3. STATEMENTS ABOUT DISTRIBUTIONS

Object lifetime distributions have been of interest to researchers of garbage collection, especially generational collection: the success of a particular garbage collector organization or promotion policy depends on how well it is matched to the behavior of typical user programs. In fact, claims have been made about lifetimes.

Hayes introduced a distinction between a “weak” and a “strong” generational hypothesis [8]. Our understanding of his statement of the weak generational hypothesis is this: *newly created objects have a much higher mortality than objects that are older*. His statement of the strong generational hypothesis (which he in fact introduces) is that *even if the objects in question are not newly created, the relatively younger objects have a higher mortality than the relatively older objects*, or simply, that $m(t)$ is an everywhere decreasing function.

Baker clearly pointed out that an exponential distribution of lifetimes, with $m(t)$ constant, cannot be favorable to generational collection (as opposed to whole-heap collection), and that instead $m(t)$ *should* be decreasing [1]. Nevertheless, the exponential distribution has a unique cachet among survival distributions: its mathematical simplicity and the property of “lack of memory”. In a garbage collector this property assures that an object just discovered live by the collector has the same residual lifetime as the lifetime of a newly allocated object, and this greatly simplifies the analysis. Thus, the exponential distribution was used by Clinger and Hansen in the analysis (and to inspire the design) of a *non-predictive* collector, outside the generational realm [3]. In our examination of a generalized form of that collector [13], we decided to use not only the exponential distribution $s(t) = e^{-\rho t}$, but also a variation with decreasing mortality $s(t) = e^{-\sqrt{\rho t}}$ as being in agreement with the strong generational hypothesis, as well as a variation with increasing mortality $s(t) = e^{-(\rho t)^2}$ for control. In fact, these three are instances of the Weibull distribution family $s(t) = e^{-(\rho t)^c}$ [14, 10].

at each object allocation. Thus, demise is detected precisely at the point when the object becomes unreachable from the global roots.

4. “PAST-IS-FUTURE” MODELS

A multitude of models can be developed that have decreasing mortality. But developing them *ex vacuo*, just for the simplicity of their mathematical formulation (or their use in other domains) is not satisfactory. We can base models on a broad experimental study, and in Section 5 we make a first attempt at that. But, our understanding would be aided more if distribution models could be derived from certain principles that we expect to be naturally associated with program behavior. In this spirit, Appel suggested (in a personal communication to us) that plausible object lifetime distributions should satisfy the following property:

- (1) An object’s future expected lifetime is proportional to its current age.

Thus, past lifetime is a strong predictor of the future (residual) lifetime. This stands in stark contrast to the exponential distribution.

An object’s *future expected lifetime* $C(x)$ is the difference between its expected lifetime and its current age. The expected lifetime (once we know the current age) is the conditional expected value [12] of the lifetime random variable \mathbf{L} , $E[\mathbf{L} \mid \mathbf{L} \geq x]$, where x is the current age. It is calculated as:

$$\begin{aligned} E[\mathbf{L} \mid \mathbf{L} \geq x] &= \int_0^{\infty} t f(t \mid t \geq x) dt \\ &= \frac{\int_x^{\infty} t f(t) dt}{\int_x^{\infty} f(t) dt} \\ &= \frac{\int_x^{\infty} t f(t) dt}{s(x)}. \end{aligned}$$

Thus $C(x) = E[\mathbf{L} \mid \mathbf{L} \geq x] - x$, and statement (1) is:

$$(\exists \psi > 0)(\forall x \geq 0) C(x) = \psi x,$$

with ψ a proportionality constant between the current age x and the future expected lifetime $C(x)$. Unfortunately, with $x = 0$, we have $E[\mathbf{L}] = E[\mathbf{L} \mid \mathbf{L} \geq 0] = C(0) = \psi \cdot 0 = 0$ for any ψ . We look at two ways out of the quandary: first, by letting proportionality (1) hold only in the limit as $x \rightarrow \infty$; and second, by restricting the domain of definition of the distribution to an interval $[x_0, \infty)$.

We shall find use for an alternative form of (1): Let $G(x) = \int_x^{\infty} t f(t) dt$; then $E[\mathbf{L} \mid \mathbf{L} \geq x] = \frac{G(x)}{s(x)}$. Let $\Lambda(x) = \frac{G(x)}{xs(x)} = \frac{E[\mathbf{L} \mid \mathbf{L} \geq x]}{x} = \frac{C(x)+x}{x} = \frac{C(x)}{x} + 1$.

Then statement (1) is

$$(\exists \psi' > 1)(\forall x \geq 0) \Lambda(x) = \psi',$$

with $\psi' = \psi + 1$.

4.1 Past-is-future in the limit

Let us weaken the statement (1) so that the proportionality holds asymptotically, for large values of x . (A similar analysis was outlined previously by Pearlmuter in *comp. Lang. s. scheme*, October 1995.) We look for a distribution such that $\lim_{x \rightarrow \infty} \Lambda(x)$ exists and is strictly greater than 1. Here is one such distribution.

Let

$$f(t) = \frac{\beta}{(t + \tau)^\lambda},$$

so that

$$s(t) = \frac{\beta}{\lambda-1} (t+\tau)^{1-\lambda},$$

where $\lambda > 2$. Normalization:

$$1 = s(0) = \frac{\beta}{\lambda-1} \tau^{1-\lambda}$$

gives

$$\tau = \left(\frac{\lambda-1}{\beta} \right)^{1/(1-\lambda)}.$$

We find

$$G(x) = \beta \frac{1}{\lambda-2} (x+\tau)^{2-\lambda} - \tau \beta \frac{1}{\lambda-1} (x+\tau)^{1-\lambda}.$$

Expected value

$$E[L] = G(0) = \beta \frac{1}{(\lambda-2)(\lambda-1)} \left(\frac{\lambda-1}{\beta} \right)^{\frac{2-\lambda}{1-\lambda}}.$$

If we take this value as a free parameter $E[L] = V$ (as it is necessary to do in order to generate a trace for simulation, where we want to control the heap size in equilibrium), we have

$$\beta = (\lambda-1)(V(\lambda-2))^{\lambda-1}.$$

We can then simplify:

$$\tau = V(\lambda-2).$$

The ratio

$$\Lambda(x) = \frac{G(x)}{xs(x)} = \frac{\lambda-1}{\lambda-2} + \frac{1}{\lambda-2} \frac{\tau}{x} = \frac{\lambda-1}{\lambda-2} + \frac{V}{x},$$

hence

$$\lim_{x \rightarrow \infty} \Lambda(x) = \frac{\lambda-1}{\lambda-2} > 1,$$

as required. Varying λ changes the value $\Lambda(x)$ uniformly for all x ; lower values of λ produce distributions with heavier tails.

4.2 Past-is-future restricted to $x > x_0$

We make condition (1) strictly hold for $x > x_0$, where $x_0 > 0$, and we define the distribution (functions F , f , s , and m) in that interval, but we set $F = 0$, $f = 0$, $s = 1$, and $m = 0$ in the interval $[0, x_0]$. This formulation is intuitively appealing, since lifetimes in practice take on discrete values in \mathbb{N} , and hence setting $x_0 = 1$ is quite natural.

Condition (1):

$$(\exists \psi' > 1)(\forall x \geq x_0) \Lambda(x) = \psi',$$

that is,

$$G(x) = \psi' xs(x),$$

gives the integral equation,

$$\int_x^\infty t f(t) dt = \psi' x \int_x^\infty f(t) dt.$$

It is easy to obtain the corresponding differential equation:

$$F'(x) = \frac{\psi'(1-F(x))}{(\psi'-1)x}.$$

The solution, with the boundary value $F(x_0) = 0$, is

$$F(x) = 1 - \left(\frac{x_0}{x} \right)^{\frac{\psi'}{\psi'-1}},$$

$$f(x) = \frac{1}{x_0} \frac{\psi'}{\psi'-1} \left(\frac{x_0}{x} \right)^{\frac{2\psi'-1}{\psi'-1}},$$

$$s(x) = \left(\frac{x_0}{x} \right)^{\frac{\psi'}{\psi'-1}},$$

$$m(x) = \frac{\psi'}{\psi'-1} \frac{1}{x}.$$

The mortality is indeed an everywhere decreasing function.

$$G(x) = \psi' \left(\frac{x_0}{x} \right)^{\frac{1}{\psi'-1}},$$

$$G(x_0) = \psi' x_0.$$

The expected live amount in the heap is:

$$v(x) = \psi' x_0 - (\psi' - 1) \left(\frac{x_0}{x} \right)^{\frac{1}{\psi'-1}}.$$

Let us examine this family of distributions qualitatively. The steady-state heap volume $V = \lim_{x \rightarrow \infty} v(x)$ equals the expected value of L , which is $\psi' x_0$. What are reasonable parameter values? Suppose that we wish to set $V = 50000$.² Suppose also that we want x_0 to be 1. Then $\psi' = V/x_0 = 50000$. However, the live volume in the heap, $v(x)$, approaches its limit value V at the rate of decay of the second term, that is, as the 49999-th root of x . With such a slow approach, that is, with such a heavy tail in the distributions, one must allow a time $3.67 \cdot 10^{99997}$ to pass before the heap is within 1% of equilibrium; until the heap is in equilibrium, the distribution of objects in it does not reflect the heavy tail of the source distribution. Simulating that many objects is somewhat impractical. Moreover, can actual programs exhibit such extremely slow heap growth as with $\psi' = 50000$? Suppose that a program does. We can only observe executions of much shorter duration than 10^{99997} , say, up to 10^{10} ; but then we cannot *empirically* distinguish the postulated $\psi' = 50000$ past-is-future distribution from another distribution that agrees with it up to $t = 10^{10}$, but lacks the heavy tail beyond that age. Alternatively, to allow 99% of the steady-state volume to be reached with 10^7 objects simulated, one must have $\psi' < 2.5$ (approximately), but then $x_0 > 20000$. Thus, beyond the construction of an elegant analytical model of object lifetimes, we must keep in mind the need to be able to validate it against real data. This example shows that sometimes validation may be difficult to achieve.

5. MODELS VS. EMPIRICAL EVIDENCE

In validating lifetime models, we apply the tools of statistical analysis of survival data to the distributions of object lifetimes, to the extent that they are applicable to our problem. (The populations traditionally studied in statistical analyses are quite different from programs'

²This number is sufficiently large so that in simulation, even when a heap is divided into ~ 100 regions, each one contains at least ~ 100 objects, which allows us to vary the heap configuration widely without incurring significant fragmentation effects.

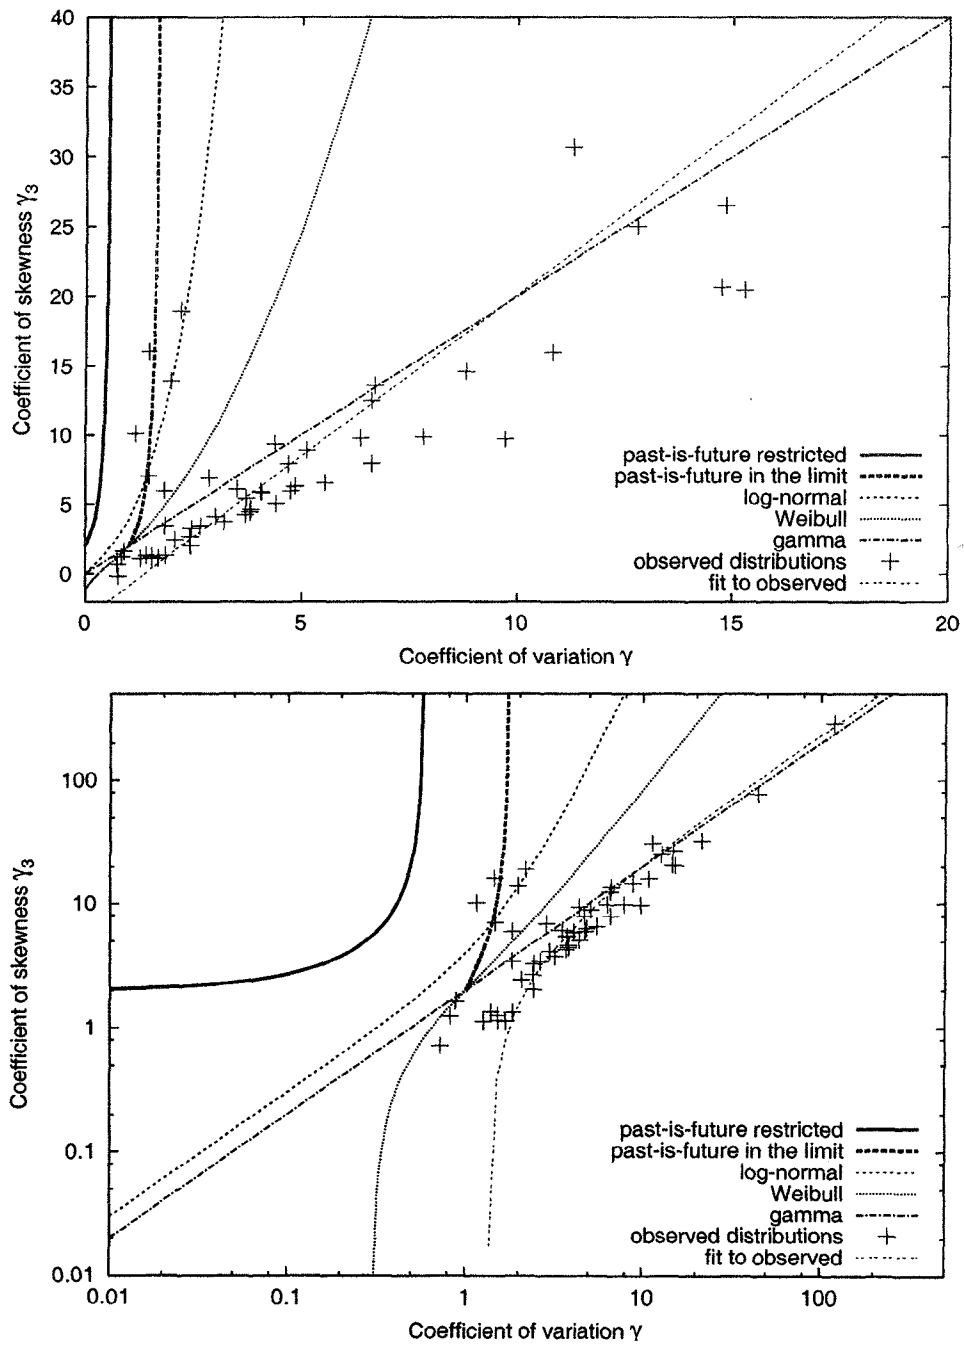


Figure 1: Coefficient of skewness, γ_3 , versus coefficient of variation for several analytical distribution families and for empirically observed object lifetime distributions. Top: linear scale; bottom: logarithmic scale. (See the appendix for definitions.)

objects. Moreover, the relevant literature concentrates mostly on statistical confirmation of explanatory variables (e.g., in clinical trials), which is different from our goal: finding distributions.)

One recommended test for comparing *families* of distributions is based on the graph of moment ratios: the coefficient of skewness γ_3 vs. the coefficient of variation γ . These moment ratios are dimensionless quantities independent of time-scale [4, p.24–28] that reflect only the *shape* of the distribution. The coefficient of variation is a measure of the spread of the distribution around its mean; the coefficient of skewness is a measure of the asymmetry of the distribution, and large positive values indicate heavy tails. Each distribution corresponds to a single point (γ, γ_3) ; a single-parameter family of distributions defines a parametric curve in the γ - γ_3 plane. (See Appendix for definition of γ and γ_3 .) Three families of distributions commonly used in survival analysis (log-normal (Section A.4.3), Weibull (Section A.4.2), and gamma (Section A.4.1)), are plotted in the γ - γ_3 plane in Figure 1. Note that the Weibull and gamma curves intersect at the point (1, 2); at this point each has degenerated into the exponential distribution. The two families introduced here are also shown (past-is-future in the limit (Section 4.1) and past-is-future restricted (Section 4.2)). Finally, the figure contains scatter points (γ, γ_3) of object lifetime distributions obtained empirically, and a line of least-squares fit for these points. These empirical distributions come from 58 Smalltalk and Java programs. (Complete object-level traces are available at <http://ali-www.cs.umass.edu/~stefanov/ISMM2000objecttraces-README.html>.)

The scatter points of empirical distributions show a trend of correlation between γ and γ_3 . This trend is somewhat surprising, since there is no *a priori* reason to expect it from a haphazard collection of benchmark programs. Perhaps the presence of this trend points to fundamental properties of program behavior, and it certainly ought to be studied further.

The scatter points lie for the most part well to the right and below the common analytical distribution families. The one exception is the gamma family: in fact, even though most scatter points are to the right and below the gamma curve, they are quite close to it.

To our chagrin, the two past-is-future families both have much lower coefficient of variation and much higher coefficient of skewness than the empirical distributions. Therefore, however intuitively plausible they are, they should not be employed to model object lifetimes. Indeed, we must concentrate the search for analytical distributions on those—not in standard literature—with much higher coefficient of variation; in the meantime, the gamma family is to be favored as a candidate.

We see that the γ - γ_3 diagram usefully summarizes the shape properties of distributions and allows us to *exclude* certain analytical distribution families as models for a set of observed distributions. Note, however, that the two moments displayed do not completely capture the shape of a distribution. For positive matching further statistical tests are necessary.

6. SUMMARY

Analytical modelling of object lifetimes is desirable for the design, analysis, and simulation of dynamic memory management systems, but it remains a difficult problem. We examined certain qualitative criteria that may be imposed on lifetime distributions, and demonstrated the use of a simple graphical technique for (in)validating postulated distribution models against empirical evidence.

7. ACKNOWLEDGMENTS

We thank Gilberto Contreras, Barak Pearlmitter, and the reviewers for their constructive advice. We are especially grateful to Andrew Appel for comments and discussions that initially spurred the developments presented.

8. REFERENCES

- [1] H. G. Baker. 'Infant Mortality' and generational garbage collection. *SIGPLAN Notices*, 28(4):55–57, 1993.
- [2] R. S. Burrington and D. C. May, Jr. *Handbook of Probability and Statistics with Tables*. McGraw-Hill, New York, 2nd edition, 1970.
- [3] W. D. Clinger and L. T. Hansen. Generational garbage collection and the radioactive decay model. *SIGPLAN Notices*, 32(5):97–108, May 1997. *Proceedings of the ACM SIGPLAN '97 Conference on Programming Language Design and Implementation*.
- [4] D. R. Cox and D. Oakes. *Analysis of Survival Data*. Chapman and Hall, London, 1984.
- [5] L. Devroye. *Non-Uniform Random Variate Generation*. Springer-Verlag, New York, 1986.
- [6] R. C. Elandt-Johnson and N. L. Johnson. *Survival Models and Data Analysis*. Wiley, New York, 1980.
- [7] M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. John Wiley, New York, 2nd edition, 1993.
- [8] B. Hayes. Using key object opportunism to collect old objects. In *Proceedings of the Conference on Object-Oriented Programming Systems, Languages, and Applications*, pages 33–46, Phoenix, Arizona, Oct. 1991. *SIGPLAN Notices* 26, 11 (Nov. 1991).
- [9] N. L. Johnson and S. Kotz. *Distributions in Statistics: Continuous Univariate Distributions (vols. 1 and 2)*. John Wiley, New York, 1970.
- [10] A. Lemmi. Empirical testing of the weibull distribution in the analysis of younger ages mortality. Technical report, Facoltà di Scienze Economiche e Bancarie, Università di Siena, Dec. 1982. Paper presented to a seminar on Demographic Models at the Department of Demography, University of Kinshasa, December 15–22, 1982.
- [11] I. Miller and J. E. Freund. *Probability and Statistics for Engineers*. Prentice-Hall, Englewood Cliffs, New Jersey, 1965.
- [12] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York, 3rd edition, 1991.
- [13] D. Stefanović, J. E. B. Moss, and K. S. McKinley. Oldest-first garbage collection. Computer Science Technical Report 98-81, University of Massachusetts, Amherst, Apr. 1998.
- [14] W. Weibull. *Fatigue testing and analysis of results*. Pergamon Press, for the Advisory Group for Aeronautical Research and Development, North Atlantic Treaty Organization, New York, 1961.

APPENDIX

A. ABOUT LIFETIME DISTRIBUTIONS

A.1 Basic definitions

The *survival* function of a random variable \mathbf{L} is $s_{\mathbf{L}}(t) = \mathcal{P}\{\mathbf{L} \geq t\}$. For object lifetimes, it expresses what fraction of original allocation volume is still live at age t . We usually drop the subscript \mathbf{L} . The survivor function is a monotone non-increasing function.

The *probability density* function is $f(t) = -s'(t)$. Occasionally we also use the *cumulative distribution* function $F(t) = 1 - s(t)$. The *mortality* function is $m(t) = \frac{f(t)}{s(t)} = -\frac{d}{dt} \log s(t)$, and it expresses the age-specific death rate. Mortality is also known as the *hazard* function (and written $h(t)$) in the literature on lifetime analysis. Occasionally we also use the integrated mortality: $M(t) = \int_0^t m(u)du$. Certain properties always hold: $s(0) = 1$; $\lim_{t \rightarrow \infty} s(t) = 0$; $\int_0^{\infty} f(t)dt = 1$ (normalization of density); $s(t) = e^{-M(t)}$ [4, p.14].

A.2 Moments

The *moments* of a distribution of random variable \mathbf{L} are defined as $m_k = E[\mathbf{L}^k] = \int_0^{\infty} t^k f(t)dt$. The *central moments* are defined as $\mu_k = E[(\mathbf{L} - m_1)^k] = \int_0^{\infty} (t - m_1)^k f(t)dt$. Here m_1 is the mean, or expected value, and m_2 is the variance, or dispersion; a common notation is $\sigma^2 = \mu_2$, where σ is called *standard deviation*. In calculation, we usually first find moments m_1 , m_2 , and m_3 , by integration in the case of analytical definitions, or by summation over observed discrete points in empirical distributions, and then compute central moments using the formulae $\mu_2 = m_2 - m_1^2$ and $\mu_3 = m_3 - 3m_1m_2 + 2m_1^3$.

The *coefficient of variation* is $\gamma = \frac{\sigma}{m_1}$. The *standardized third moment* or *coefficient of skewness* is $\gamma_3 = \frac{\mu_3}{\sigma^3}$; it is also written η_3 or $\sqrt{\beta_1}$.

A.3 Finiteness of expected value

It is a simple exercise to show that the expected value of the live amount in the heap at time x (that is, after an amount x has been allocated) is $v(x) = \int_0^x s(t)dt$, and that

$$V = \lim_{x \rightarrow \infty} v(x) = \int_0^{\infty} s(t)dt = \int_0^{\infty} t f(t)dt = E[\mathbf{L}],$$

when they exist.

We may impose on the object lifetime distribution an additional property of finiteness (existence) of expected value, to ensure that a heap equilibrium is reached in the limit. (Heap equilibrium has been the underlying assumption in some comparative analyses of garbage collection costs [3, 13]. A relative heap size parameter is used as the basis for comparison of two collection algorithms: heap size is a fixed multiple of a *steady-state* live data amount.) How essential is this requirement, and could we also consider distributions with unbounded expected value? On the one hand, the running time of real programs is finite, and thus f is finally-zero, hence $E[\mathbf{L}]$ is finite.

On the other hand, it is theoretically plausible that we are observing *initial segments* of potentially infinite computations, and so it is useful to investigate heaps that grow without bound. From a purist standpoint, many realistic programs that run indefinitely do use increasing amounts of space; for instance, counting requires logarithmically increasing space.

If the live data amount does not stabilize, but rather grows indefinitely, then the available heap size ought to grow in equal proportion—if one desires measurements in terms of the relative heap size parameter. This property must be ensured with due care in analysis and simulation.

A.4 Distribution families of Figure 1

Basic definitions of distribution families, compiled from textbooks [2, 5, 6, 7, 9, 11, 12].

A.4.1 Gamma

$$f(t) = \frac{c^{b+1}}{\Gamma(b+1)} t^b e^{-ct}$$

$$m_k = (b+1)(b+2) \cdots (b+k)c^{-k}$$

$$\mu_2 = \frac{b+1}{c^2}$$

$$\mu_3 = 2 \frac{b+1}{c^3}$$

$$\gamma = \frac{1}{\sqrt{b+1}}$$

$$\gamma_3 = \frac{2}{\sqrt{b+1}}$$

Therefore $\gamma_3 = 2\gamma$ is a straight line in Figure 1.

A.4.2 Weibull

$$f(t) = ct^{c-1} e^{-t^c}$$

$$m_1 = \frac{1}{c} \Gamma\left(\frac{1}{c}\right)$$

$$m_2 = \Gamma\left(1 + \frac{2}{c}\right)$$

$$m_3 = \Gamma\left(1 + \frac{3}{c}\right)$$

The curve γ_3 is plotted parametrically with respect to c in Figure 1.

A.4.3 Log-normal

$$f(t) = \frac{1}{t\delta\sqrt{2\pi}} e^{-\frac{(\log t - \xi)^2}{2\delta^2}}$$

It can be shown that, with abbreviation $\omega = e^{\delta^2}$, $\gamma = \sqrt{\omega - 1}$ and $\gamma_3 = (\omega + 2)\sqrt{\omega - 1}$, therefore $\gamma_3 = 3\gamma + \gamma^3$ in Figure 1.