

# Improving Confidence in Network Simulations

Michele C. Weigle

Department of Computer Science  
Old Dominion University




Winter Simulation Conference  
December 5, 2006

## Introduction

- “Credibility crisis” in networking simulation analysis (Pawlikowski *et al.*, 2002)
  - effect of randomness
  - affects quality and repeatability of results
- Problems dealing with *heavy-tailed distributions*
  - describes Internet file sizes
- Problems dealing with *random number generators*

# Heavy-Tailed Distributions

- Most common
  - Pareto
  - Lomax
    - Pareto of the 2nd kind
- Parameters
  - scale,  $k$
  - shape,  $\alpha$ 
    - $\alpha \leq 2$ , infinite variance
    - $\alpha \leq 1$ , infinite mean



# Heavy-Tailed Distributions

## Usage in Network Simulation

- Network traffic is driven by transfers of files with a heavy-tailed size distribution
- Big question in network simulations is appropriate running time of simulation
  - typically, when sample mean approaches theoretical mean
- Workloads based on heavy-tailed distributions take a long time to converge

# Heavy-Tailed Distributions

## Problems in Network Simulation

- For heavy-tailed distributions with  $\alpha = 1.2$ , takes  $10^{10}$  samples for two-digit accuracy (Crovella and Lipsky, 1997)
- Swamping observation
  - one sample doubles the mean
  - with  $\alpha = 1.2$ , probability greater than 1 in 100
- When  $\alpha < 1.5$ , “simulation becomes infeasible” (Crovella and Lipsky, 1997)
- Practically, won't reach steady state, will always be transient


# Problems in Network Simulation

## Example

- PackMime-HTTP in ns-2
  - request sizes
    - Pareto,  $\mu = 50$  bytes,  $\alpha = 1.2$ ,  $k = 8.3$
  - response sizes
    - Pareto,  $\mu = 7000$  bytes,  $\alpha = 1.2$ ,  $k = 1400$
- 100,000 response size samples
- Took several hours to run
- Measure
  - offered load
  - cumulative mean of response sizes


# Heavy-Tailed Distributions

## Offered Load



# Heavy-Tailed Distributions

## Mean Response Size



# Heavy-Tailed Distributions

## Problems in Network Simulation

- Running until steady-state is not feasible
- Approaches
  - bound the distribution
  - approximate Pareto
  - treat as transient

## Approaches

### Bound the Distribution

- Crovella and Lipsky, 1997
- Let  $Y$  be the maximum observation
- $E[Y] = knB(n, 1-1/\alpha)$ 
  - $B$  is beta function
  - $n$  is number of samples
- Maximum possible observation,  $p$ 
  - e.g., maximum file size possible
- Run simulation until  $E[Y]$  approaches  $p$

## Bound the Distribution

### Example

- $k = 1400$ ,  $\alpha = 1.2$
- Number of samples needed for  $E[Y]$  to approach maximum  $p$

| Maximum File Size ( $p$ ) | Samples ( $n$ ) |
|---------------------------|-----------------|
| MP3 - 5 MB                | 2500            |
| TV episode - 350 MB       | 400,000         |
| CD image - 650 MB         | 850,000         |
| DVD image - 4.7 GB        | 9,000,000       |

## Bound the Distribution

### Bounded Pareto

- Harchol-Balter *et al.*, 1999
- Use Bounded Pareto distribution to compute mean and variance of similar Pareto
- Run simulation until sample mean approaches mean of Bounded Pareto
- Technique has been used in previous studies


# Approaches

## Approximate Pareto

- Fishman and Adan, 2006
- Approximate Lomax (Pareto of 2nd kind) with Lognormal
  - variance is always defined
- Given  $k$  and  $\alpha$  of Lomax, provide parameters for Lognormal, based on desired fit of the tail
  - $\zeta$  sets the point at which the tails diverge
  - smaller  $\zeta$ , the better the fit
- Provide sample path length given desired accuracy of the mean
  - high  $10^{-4}$ , moderate  $10^{-3}$ , gross  $10^{-2}$

## Approximate Pareto


### CCDF Generated File Sizes



# Approaches

## Approximate Pareto

- Lognormal  $\zeta = 10^{-3}$ 
  - 330,000 samples
- Lognormal  $\zeta = 10^{-5}$ 
  - 100,000,000 samples
  - 50 minutes on 2.8 GHz Xeon
- Lomax
  - 643,783,352 samples
  - $10^{12}$  required for accuracy



# Approaches

## Treat as Transient

- Weigle *et al.*, 2005
- Run simulation long enough to capture a certain number of completed HTTP request-response pairs
  - 250,000 pairs
- Require that at least one response from tail is drawn
  - 10 MB response size
- Use same RNG seeds to generate traffic in each experiment
  - testing behavior of different protocols



# Approaches

## Treat as Transient


- Running time, or number of samples generated, may not be the most important factor in how well the tail of the distribution fits
- Large samples can occur at any time
  - be aware of the number of large samples present in simulations
  - run simulations long enough to observe a reasonable number of these large samples

# Random Number Generation

- Added the MRG32k3a combined multiple recursive generator (L'Ecuyer) to ns-2
  - default since ns-2.1b9, April 2002
- Divides the period into  $2^{64}$  streams
  - each stream is guaranteed to be independent of every other stream
- Each stream consists of  $2^{51}$  substreams
  - each substream is guaranteed to be independent of every other substream
- Each substream has a period of  $2^{76}$

# RNG


## Streams and Substreams



# RNG

## In ns-2

- Users set only a starting seed
- Independent seeds for additional RNGs automatically computed
  - [new RNG]
- Multiple replications
  - use a separate stream for each random variable
  - use a separate substream for each replication
    - next-substream



# RNG

## Effects of Improper RNG Usage

- Decreased reproducibility
  - if you use time of day as seed, it is very difficult for others to reproduce your simulation
- PackMime-HTTP as example
  - HTTP request size
  - HTTP response size
  - time between start of new connections
  - server delay time
  - round-trip time

# RNG

## Effects of Improper RNG Usage

- When new connection starts,
  - request size, response size, time for next connection drawn
- When request received by server,
  - server delay time drawn
- Effects of change in network settings
  - could affect when requests reach the server
  - if using same RNG for all distributions, could cause different requests and responses to be drawn with the same starting seed


# Effects of Improper RNG Usage

## Example

- *A change in network settings should not change which request size and response sizes are drawn*
- Four simulations
  - uncongested network, all random variables use different, independent RNG (proper use of RNG)
  - congested network, proper use of RNG
  - uncongested network, all random variables use same RNG (improper use of RNG)
  - congested network, improper use of RNG


## Example

### Proper Use of RNGs



# Example

## Improper Use of RNGs



## RNG

### Effects of Improper RNG Usage

- Congested case has no response larger than 10,000 KB
- Uncongested case has 3 such responses
- Results in a large impact on the offered load in the system
  - could invalidate results

# Summary

- Heavy-tailed workloads are often used in network simulations
  - must be aware of issues in determining appropriate simulation running time
- Important to deal with RNGs properly in simulations
  - improper use could affect the validity of results
- Simulation scripts are available at <http://www.cs.odu.edu/~mweigle/research/netsim/>