Optimal File Splitting for Wireless Networks with Concurrent Access

Gerard Hoekstra[1], Rob van der Mei[2],
Yoni Nazarathy[3], Bert Zwart[4]

NET-COOP 2009,
EURANDOM, Eindhoven

[1] CWI Amsterdam, Thales Nederland B.V.

[2] CWI Amsterdam, VU University Amsterdam

[3] CWI Amsterdam, EURANDOM, Mechanical Engineering TU/e

[4] CWI Amsterdam, VU University Amsterdam, EURANDOM, Georgia Institute of Technology

Outline

Outline

Idea: Use Several Wireless Links for File Transfer

Example Story

- Police vehicle equipped with several cameras
- Photos are taken and files are to be uploaded to headquarters
- Vehicle has **several** slow wireless links
  Example: GSM, SATCOM, MANET
- Upon file transfer, FTP opens several concurrent connections one on each link
- Goal: minimize file transfer delay by utilizing all links

Similar Stories (perhaps)

- Jobs arriving to web-server farms are split
- Downloading of file fragments in peer to peer networks

Model a Link as a PS Queue

Processor Sharing Abstraction of A Link

- On the processor sharing of file transfers in wireless LANs (Hoekstra, van der Mei, 2009)

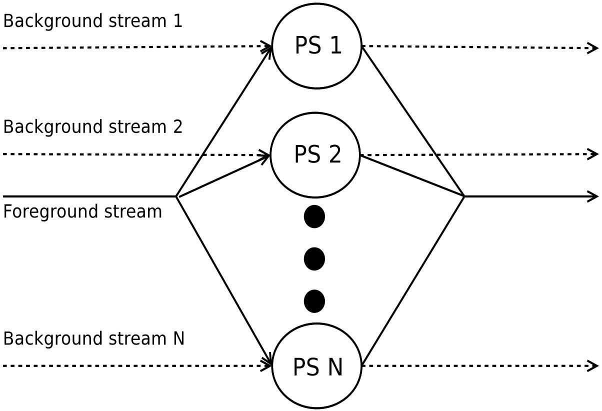

How to Split the File?

Some Options

1. Don't split. Instead route (e.g. JSQ)
2. Split (to 2 pieces):
   – Transmit piece 1 from first byte "forward"
   – Transmit piece 2 from last byte "backwards"
3. Use fixed static splitting:
   decide on $\alpha = (\alpha_1, \ldots, \alpha_N)$ and always use it
4. $\ldots$

In this work, we follow option 3

Queueing Model

- N processor sharing queues, rates $c_1, \ldots, c_N$
- N background Poisson file streams, $\lambda_1, \ldots, \lambda_N$
- 1 foreground Poisson stream, λ_0
- I.I.D. file sizes per stream, $B_0, B_1, \ldots, B_N$, means $\beta_0, \beta_1, \ldots, \beta_N$
- Foreground files split into fragments:
  Splitting rule: $\alpha = (\alpha_1, \ldots, \alpha_N)$, fragment i is of size $\alpha_i B_0$
- $\rho_i := \lambda_i \beta_i$. Assume $\alpha_i \rho_0 + \rho_i < c_i$
- Assume steady state
- Of interest:
  - Sojourn time of foreground files (maximum of fragments)
  - Choosing a "good" splitting rule α

Related Work: "JSQ to PS Farm", Gupta, Harchol Balter, Sigman, Whitt (2007), "Joining Games", Altman, Ayesta, Prabhu (2008), Chen, Marden, Wierman (2009), "Fork-Join FCFS Tail Asymptotics", Lelarge (2008, 2009)

M/G/1-PS Refresher

Denote:

- $V \equiv$ sojourn time of a job arriving to steady state
- $\tilde{B} \equiv$ service time random variable
- $R(x) := \int_0^x \frac{1}{1+Q(t)} dt$

Property 1

$E[V|\tilde{B}] = \frac{\tilde{B}}{c-\rho}$

Property 2

$P(V > x) = P(\tilde{B} > R(x))$

Property 3, Reduced Load Approximation (RLA) (Zwart, Boxma 2000)

Assume $P(\tilde{B} > x) = L(x)x^{-\nu}$, $\nu > 1$, $L(\cdot)$ slowly varying. Then,

$$P(V > x) \sim P(\tilde{B} > (c - \rho)x)$$

Outline

Main Result

Reminder of our Story

1. Foreground file B_0 arrives to steady state
2. File splits to N fragments $\alpha_i B_0$, with sojourn V_i
3. Sojourn time of whole file: $M_\alpha := \max\{V_1, \dots, V_N\}$

Theorem (Reduced Load Equivalence)

Assume,

$$\text{(i) } E[B_i^{1+\epsilon}] < \infty,\ i = 1, \dots, N$$
$$\text{(ii) } P(B_0 > x) = L(x)x^{-\nu},\ \nu > 1,\ L(\cdot) \text{ slowly varying}$$

Then,

$$P(M_\alpha > x) \sim P(B_0 > \gamma_\alpha x),$$

with,

$$\gamma_\alpha = \min_{i=1,\dots,N} \gamma_i, \qquad \text{with} \qquad \gamma_i = \frac{c_i - \rho_i - \alpha_i \rho_0}{\alpha_i}$$

Proof, cont.

Proof: $\Big(\; P(\max(V_1, V_2) > x) \sim P(B_0 > \min(\gamma_1, \gamma_2)x) \; \Big)$

- Assume $\min(\gamma_1, \gamma_2) = \gamma_1$.
- $\frac{P(M_\alpha > x)}{P(B_0 > \gamma_\alpha x)} = \frac{P(B_0 > \frac{R_1(x)}{\alpha_1})}{P(B_0 > \gamma_1 x)} + \frac{P(B_0 > \gamma_2 x)}{P(B_0 > \gamma_1 x)} \left(\frac{P(B_0 > \frac{R_2(x)}{\alpha_2})}{P(B_0 > \gamma_2 x)} - \frac{P(B_0 > \max(\frac{R_1(x)}{\alpha_1}, \frac{R_2(x)}{\alpha_2}))}{P(B_0 > \max(\gamma_1, \gamma_2)x)} \right).$
- $\frac{P(B_0 > \gamma_2 x)}{P(B_0 > \gamma_1 x)} = \frac{L(\gamma_2 x)}{L(\gamma_1 x)} \left(\frac{\gamma_2}{\gamma_1}\right)^{-\nu} \to \quad \left(\frac{\gamma_2}{\gamma_1}\right)^{-\nu},$
- Each queue is a 2-class queue: $P(B_0 > \frac{R_i(x)}{\alpha_i}) \sim P(B_0 > \gamma_i x)$ (Zwart 1999)
- Remaining to show: $P(B_0 > \max(\frac{R_1(x)}{\alpha_1}, \frac{R_2(x)}{\alpha_2})) \sim P(B_0 > \max(\gamma_1, \gamma_2)x)$
- Use Guillemin, Robert, Zwart, 2004. Thm1: Let $\lim_{x\to\infty} S(x)/x = \gamma$ a.s. and take $\tilde{B}$ independent regularly varying. Then $P(\tilde{B} > S(x)) \sim P(\tilde{B} > \gamma x)$, if we can find a c, such that $P\big(S(x) \le cx\big) = o\big(P(\tilde{B} > x)\big)$
- In our case $\tilde{B} = B_0$ and $S(x) = \max\left(\frac{R_1(x)}{\alpha_1}, \frac{R_2(x)}{\alpha_2}\right)$
- $P(\max\left(\frac{R_1(x)}{\alpha_1}, \frac{R_2(x)}{\alpha_2}\right) \le cx) \le P(\frac{R_1(x)}{\alpha_1} \le cx) = o\big(P(B_0 > x)\big)$
- Last step is as in analysis of Guillemin, Robert, Zwart for M/G/1

Outline

The Rule To Use

$$\alpha_i^* := \frac{c_i - \rho_i}{\sum_{j=1}^{N}(c_j - \rho_j)}$$

Interpretations

- Minimizes asymptotic tail (as we show in next slide)
- "Split proportional to free capacity"
- Equating conditional (on B_0) mean sojourn times in PS queues. Set,

$$\frac{\alpha_1^* B_0}{c_1 - \rho_1 - \alpha_1^* \rho_0} = \frac{(1 - \alpha_1^*)B_0}{c_2 - \rho_2 - (1 - \alpha_1^*)\rho_0}$$

  and solve for α_1^*

Minimizing the Tail

- Since we have $P(M_\alpha > x) \sim P(B_0 > \gamma_\alpha x)$, let's optimize γ_α:

$$\begin{aligned} \max_{\alpha} \quad & \min_{i=1,\dots,N} \left(\frac{c_i - \rho_i}{\alpha_i} \right) - \rho_0 \\ \text{s.t.} \quad & \sum_{i=0}^{N} \alpha_i = 1, \quad \alpha \geq 0. \end{aligned}$$

- The unique solution is given by $\alpha^* = \left(\frac{c_i - \rho_i}{\sum_{j=1}^{N} (c_j - \rho_j)} \right)$
- Example for $N = 2$:

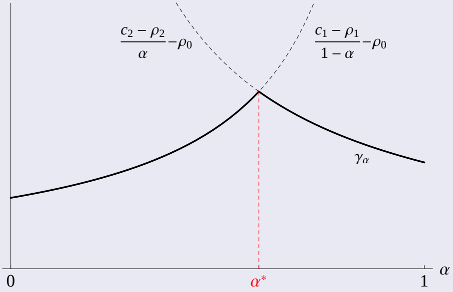

Simulated Examples

Some Examples

- $N = 2$, $c_1 = c_2 = 1$, $\beta_0 = \beta_1 = \beta_2 = 1$
- Parameterize by ρ, κ, η:

$$\rho = \frac{\lambda_0 + \lambda_1 + \lambda_2}{2}, \quad \kappa = \frac{1 - \lambda_1}{1 - \lambda_2}, \quad \eta = \frac{\lambda_0}{\lambda_1 + \lambda_2}.$$

- Examples:

System	ρ	κ	η	Distribution 0	Distribution 1	Distribution 2	$(\lambda_0, \lambda_1, \lambda_2)$	α^*
1	0.5	1.5	0.5	Pareto 3	Pareto 3	Pareto 3	$(\frac{1}{3}, \frac{1}{5}, \frac{7}{15})$	0.6
2	0.5	1.5	0.5	Pareto 3	Deterministic	Deterministic	as System 1	-
3	0.5	1.5	0.5	Pareto 3	Exponential	Exponential	as System 1	-
4	0.5	1.5	0.5	Pareto 3	Exponential	Deterministic	as System 1	-
5	0.5	1.5	0.5	Deterministic	Deterministic	Deterministic	as System 1	-
6	0.5	1.5	0.5	Erlang 2	Erlang 2	Erlang 2	as System 1	-
7	0.5	1.5	0.5	Exponential	Pareto 3	Erlang 2	as System 1	-
8	0.5	2.0	0.5	Pareto 3	Pareto 3	Pareto 3	$(\frac{1}{3}, \frac{1}{9}, \frac{5}{9})$	$\frac{2}{3}$
9	0.5	1.0	0.5	Exponential	Exponential	Exponential	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$	0.5

Finite Tail Optimality vs. Asymptotic Optimality

Simulation runs of System 4.

- Search for $\alpha^*(x) = \text{argmin}_\alpha \; P(M_\alpha > x)$,
  $x = 1, 2, 3, 5, 8, 11, 17, 25, 35, 48, 64, 85, 115, 160, 210, 270, 350, 500$
- Compare to $\alpha^* = \lim_{x \to \infty} \alpha^*(x) = 0.6$
- Plot also $E[M_\alpha]$ (to see $\text{argmin}_\alpha E[M_\alpha]$)

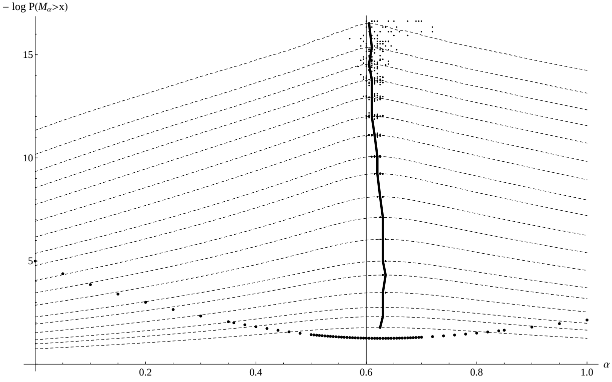

More Graphs...

$$\text{Error} \equiv \frac{P(M_{\alpha^*} > x) - P(M_{\alpha^*(x)} > x)}{P(M_{\alpha^*(x)} > x)}$$

Heavy Tailed Foreground

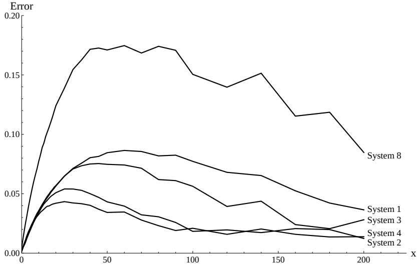

Light Tailed Foreground

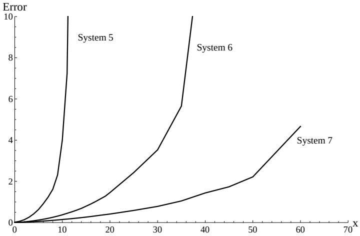

Comments

- For heavy tailed foreground files: Approximation is good for moderate x
- For light tailed foreground files: Our α^* is not tail optimal

Outline

Optimization of $E[M_\alpha]$

Plots of $E[M_\alpha]$, for Systems 1-9.

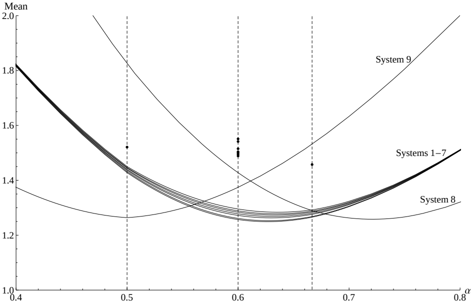

Note: The dots are plots of $E[M_{\text{JSQ}}]$.

Some Observations:

- $\min_\alpha E[M_\alpha] \approx E[M_{\alpha^*}]$
- Near insensitivity, similar to the JSQ results (Gupta et. al.)
- $E[M_{\alpha^*}] < E[M_{\text{JSQ}}]$ (not surprising)

Discussion on $\min_\alpha E[M_\alpha] \approx E[M_{\alpha^*}]$, 1'st Justification

Reason 1

- Denote $R(x) := \min_{i=1,\dots,N} \frac{R_i(x)}{\alpha_i}$
- Observe $\frac{R(x)}{x} \to \gamma_\alpha$ and $\frac{R^{-1}(x)}{x} \to \frac{1}{\gamma_\alpha}$, a.s.
- We have, $P(M_\alpha > x) = P(B > R(x))$
- Denote by $M(b)$ the sojourn time of a foreground file of size b
- We have that $M(b) = R^{-1}(b)$.
- Define $\mu(b) := E[M(b)]$
- We have $\frac{\mu(b)}{b} \to \frac{1}{\gamma_\alpha}$ as $b \to \infty$
- Thus, for large b: $\mu(b) \approx \frac{b}{\gamma_\alpha}$
- Thus selecting α such that γ_α is maximal minimizes $\mu(b)$ when b is large. and approximately minimizes the unconditional sojourn time $E[M] = E_B[\mu(B)]$

Discussion on $\min_\alpha E[M_\alpha] \approx E[M_{\alpha^*}]$, 2'nd Justification

General Setting

- Some "stochastic model" parameterized by scalar α
- α yields $1 - \overline{F}_\alpha(x)$. $\mu_\alpha = \int_0^\infty \overline{F}_\alpha(u)du$

Lemma

Assume that $\overline{F}_\alpha(x)$ is unimodal in α and that $\overline{F}_\alpha(x)$ and μ_α are differentiable in α, then there exists an $x > 0$ such that

$$argmin_\alpha \mu_\alpha = argmin_\alpha \overline{F}_\alpha(x).$$

In Our Model

- We saw empirically that $\alpha^*(x)$ does not vary much with x
- Combining with the above lemma, we get $\min_\alpha E[M_\alpha] \approx E[M_{\alpha^*}]$

Outline

Some Unanswered Questions

- Tail asymptotics for light tails of foreground (harder)
- Mathematical basis for $\min_\alpha E[M_\alpha] \approx E[M_{\alpha^*}]$
- Mathematical basis of "Near-Insensitivity"
- Tail Asymptotics for JSQ

Enjoy The Visit in Eindhoven...