Worst-case Bounds and Optimized Cache on M^{th} Request Cache Insertion Policies under Elastic Conditions

Niklas Carlsson, *Linköping University*
Derek Eager, *University of Saskatchewan*

LINKÖPING UNIVERSITY

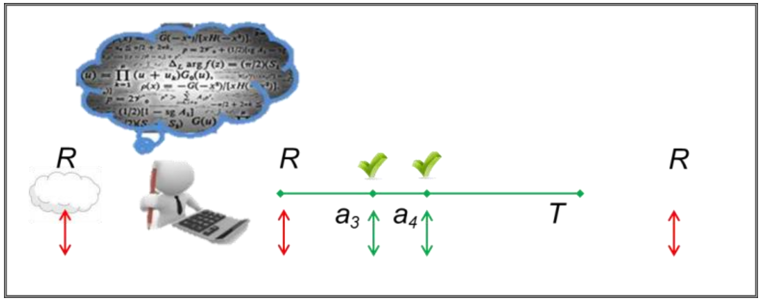

Proc. IFIP Performance, Toulouse, France, Dec. 2018.

Motivation and problem

- Cloud services and other shared infrastructures increasingly common
  - Typically third-party operated
  - Allow service providers to easily scale services based on current resource demands
- Content delivery context: Many content providers are already using third-party operated Content Distribution Networks (CDNs) and cloud-based content delivery platforms
- This trend towards using third-party providers on an on-demand basis is expected to increase as new content providers enter the market

Motivation and problem

- Cloud services and other shared infrastructures increasingly common
  - Typically third-party operated
  - Allow service providers to easily scale services based on current resource demands
- Content delivery context: Many content providers are already using third-party operated Content Distribution Networks (CDNs) and cloud-based content delivery platforms
- This trend towards using third-party providers on an on-demand basis is expected to increase as new content providers enter the market

Problem: Individual content provider that wants to minimize its delivery costs under the assumptions that

- the storage and bandwidth resources it requires are *elastic*,
- the content provider *only pays for the resources that it consumes*, and
- *costs are proportional* to the resource usage.

High-level picture

- Analyze the optimized delivery costs of different *cache on M^{th} request* cache insertion policies when using a Time-to-Live (TTL) based eviction policy
  - File object remains in the cache until a time *T* has elapsed
- Assuming elastic resources, cache eviction is not needed to make space for a new insertion
  - Rather to reduce cost by removing objects that are not expected to be requested again soon
  - A TTL-based eviction policy is a good heuristic for such purposes
  - Bonus: TTL provides approximation for fixed-size LRU caching
- Cloud service providers already provide elastic provisioning at varying granularities for computation and storage
  - Support for fine-grained elasticity likely to increase in the future

Contributions

Within this context, we

- derive worst-case bounds for the optimal cost and competitive cost ratios of different classes of *cache on M^{th} request* cache insertion policies,
- derive explicit average cost expressions and bounds under arbitrary inter-request distributions,
- derive explicit average cost expressions and bounds for short-tailed (deterministic, Erlang, and exponential) and heavy-tailed (Pareto) inter-request distributions, and
- present numeric and trace-based evaluations that reveal insights into the relative cost performance of the policies.

Contributions

Within this context, we

- derive worst-case bounds for the optimal cost and competitive cost ratios of different classes of *cache on M^{th} request* cache insertion policies,
- derive explicit average cost expressions and bounds under arbitrary inter-request distributions,
- derive explicit average cost expressions and bounds for short-tailed (deterministic, Erlang, and exponential) and heavy-tailed (Pareto) inter-request distributions, and
- present numeric and trace-based evaluations that reveal insights into the relative cost performance of the policies.

Our results show that a *window-based cache on 2^{nd} request* policy (using a single threshold parameter optimized to minimize the best worst-case costs) provides good average performance across the different distributions and the full parameter ranges of each considered distribution

System model

System model

- Assumptions:
  - storage and bandwidth resources it requires are *elastic*
  - content provider *only pays for the resources that it consumes*
  - *costs are proportional* to the resource usage

System model and problem

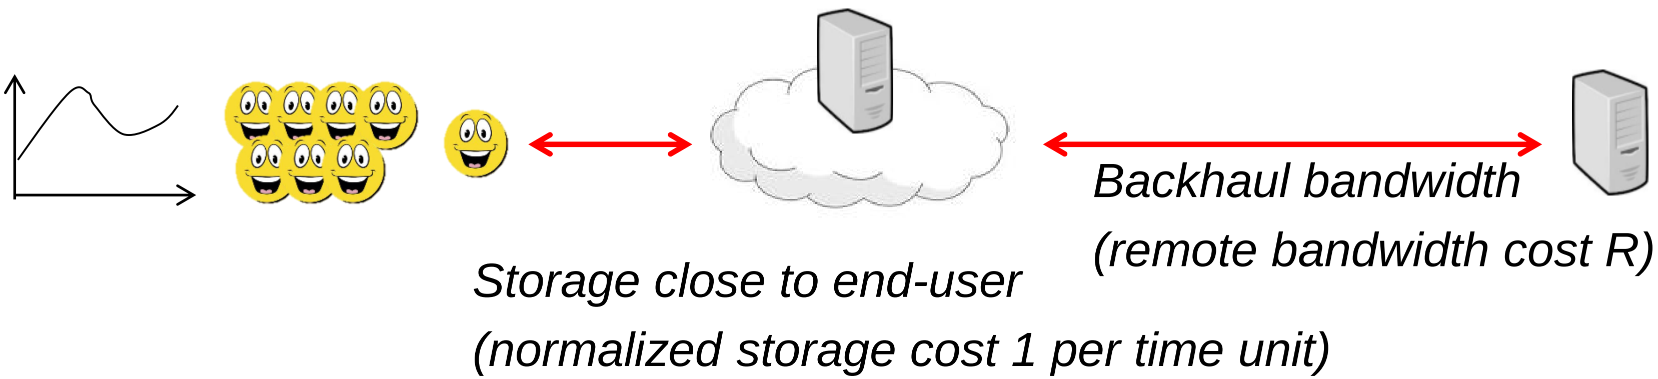

- Assumptions:
  - storage and bandwidth resources it requires are *elastic*
  - content provider *only pays for the resources that it consumes*
  - *costs are proportional* to the resource usage

- Policy decision: At the time a request is made for a file object not currently in the cache, the system must, in an online fashion, decide whether the object should be cached or not

Insertion policies

Insertion policies

Insertion policies

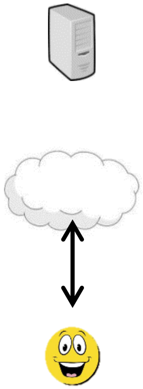

Insertion policies

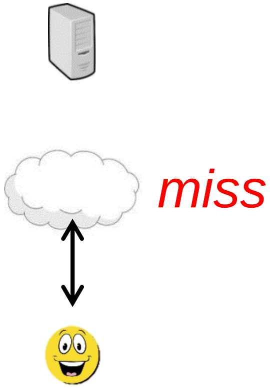

Insertion policies

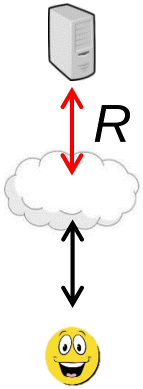

Insertion policies

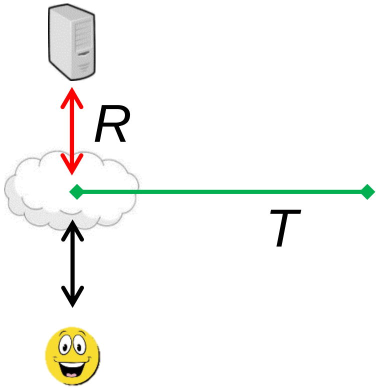

Insertion policies

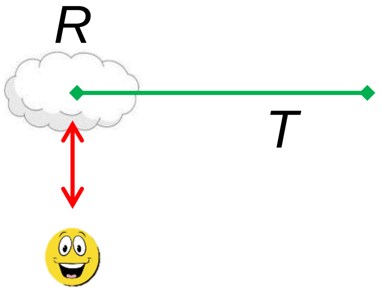

Insertion policies

Always on 1st (T)

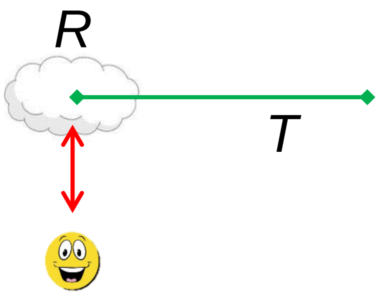

Insertion policies

Always on 1st (*T*)

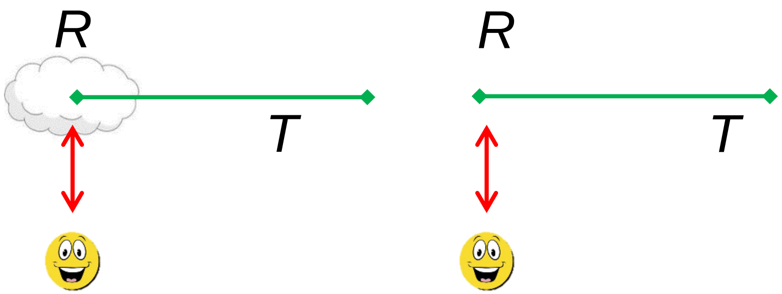

Insertion policies

Always on 1st (T)

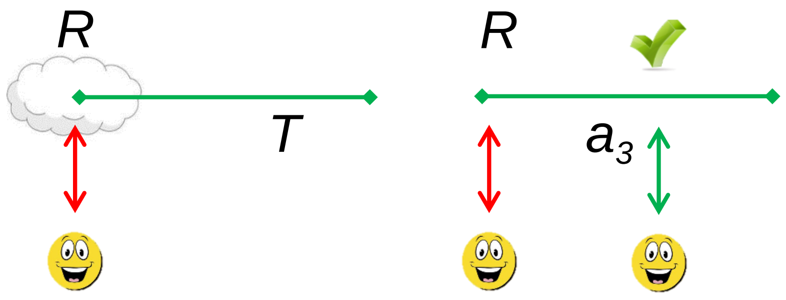

Insertion policies

Always on 1st (T)

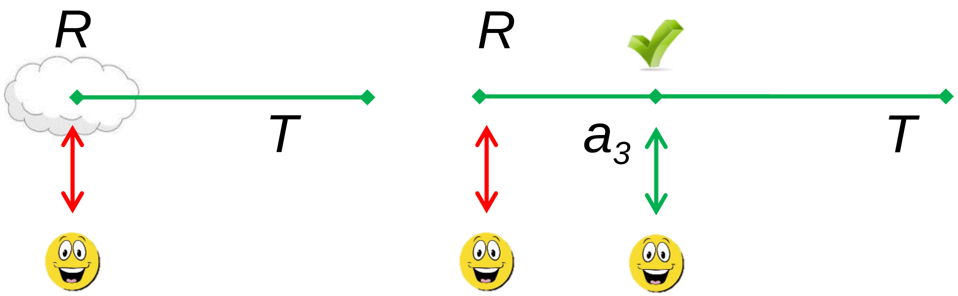

Insertion policies

Always on 1st (T)

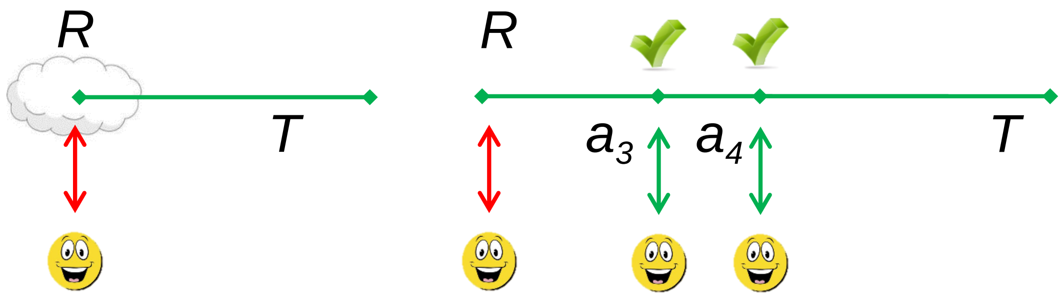

Insertion policies

Always on 1st (T)

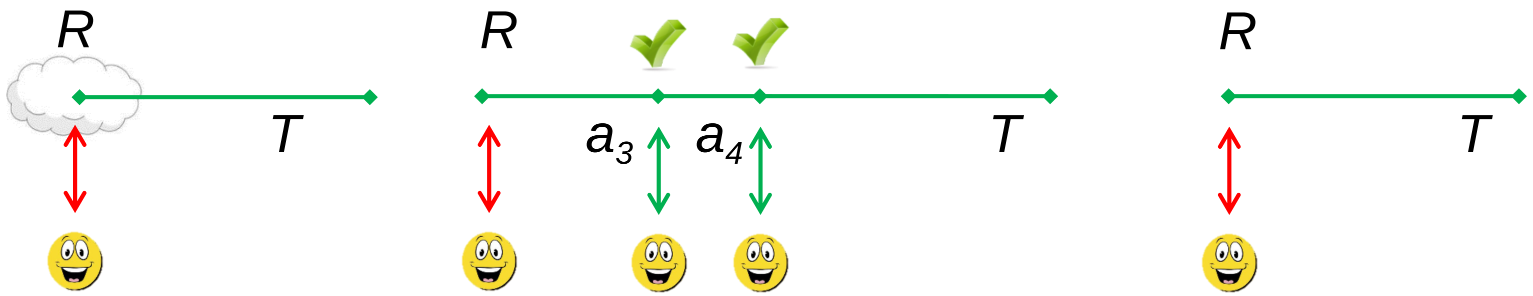

Insertion policies

Always on 1st (*T*)

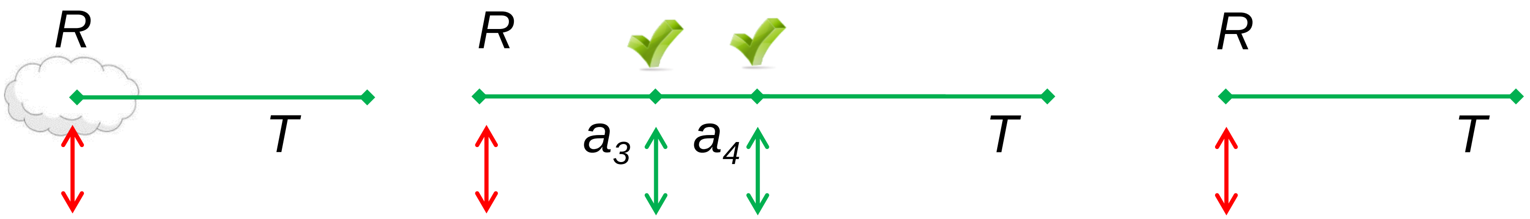

Insertion policies

Always on 1st (T)

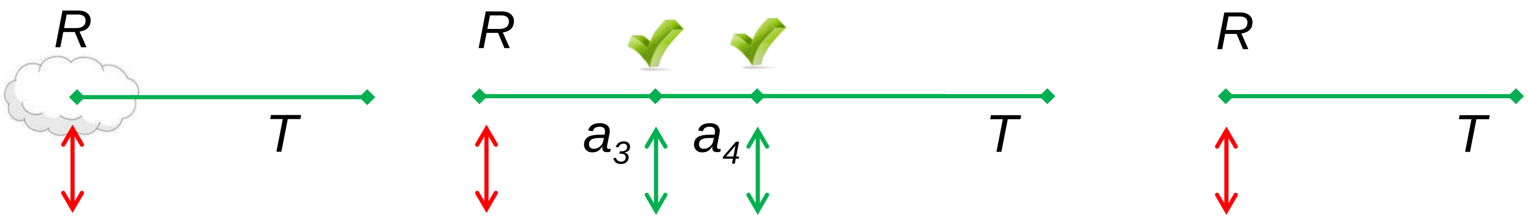

Always on 2nd (T)

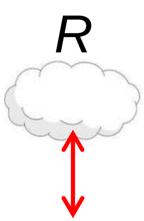

Insertion policies

Always on 1st (*T*)

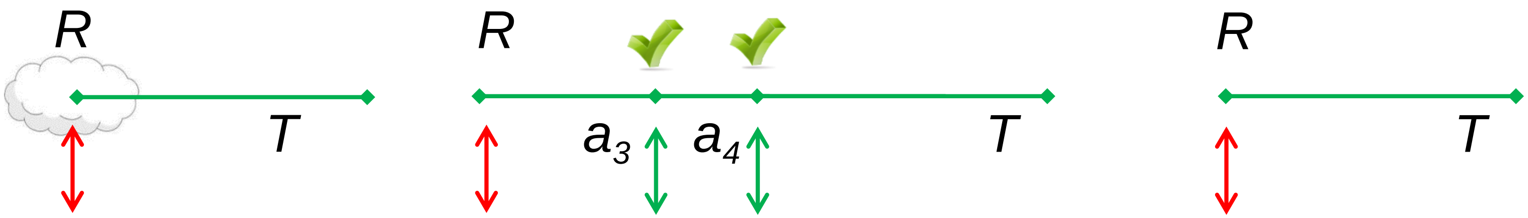

Always on 2nd (*T*)

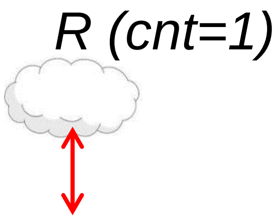

Insertion policies

Always on 1st (T)

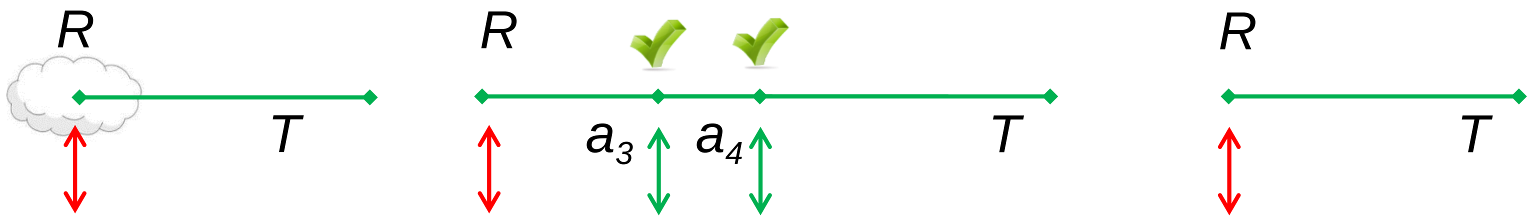

Always on 2nd (T)

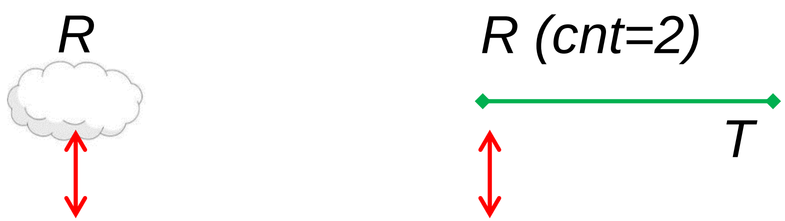

Insertion policies

Always on 1st (T)

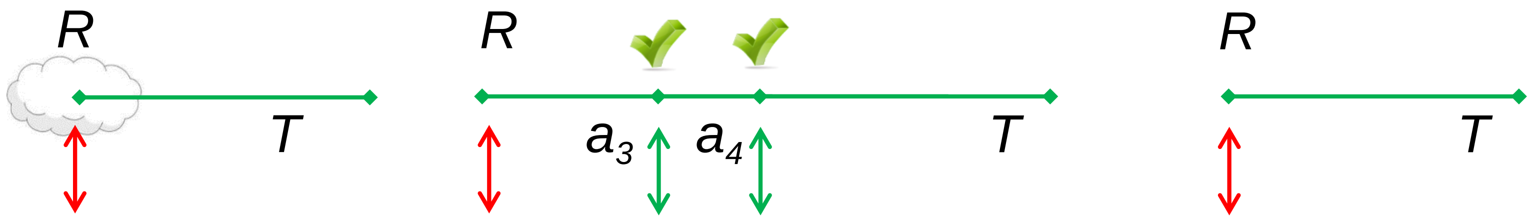

Always on 2nd (T)

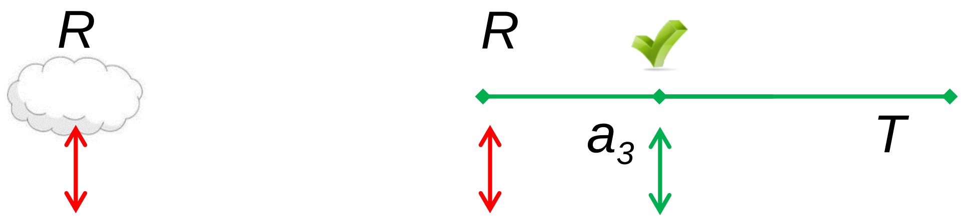

Insertion policies

Always on 1st (T)

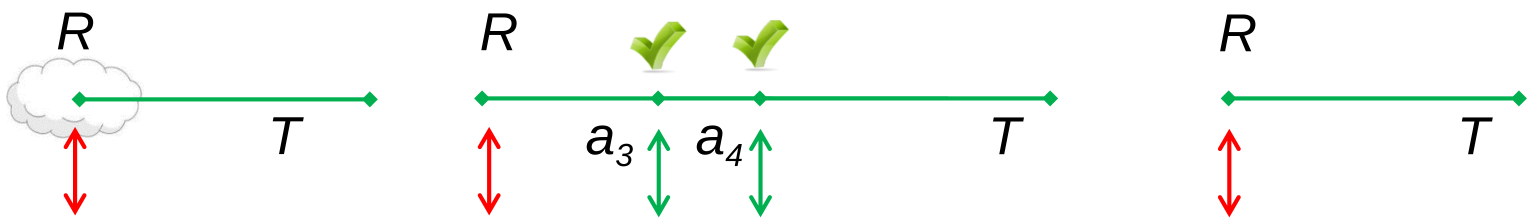

Always on 2nd (T)

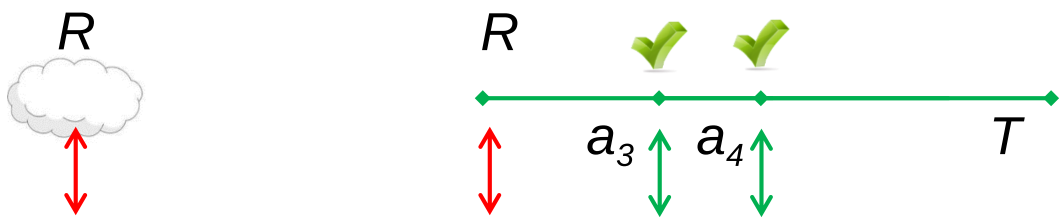

Insertion policies

Always on 1st (T)

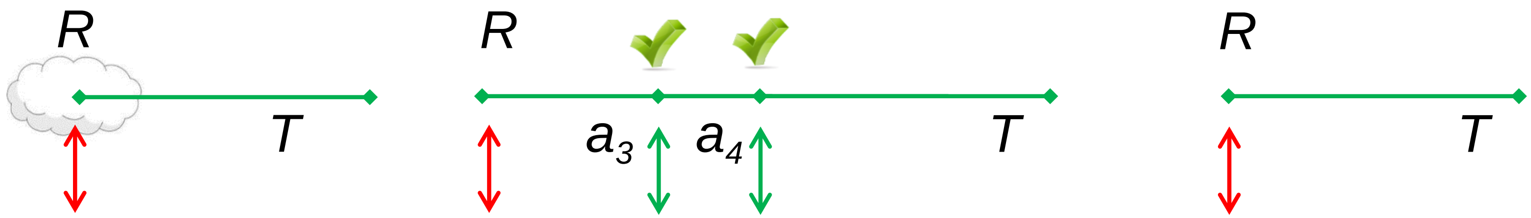

Always on 2nd (T)

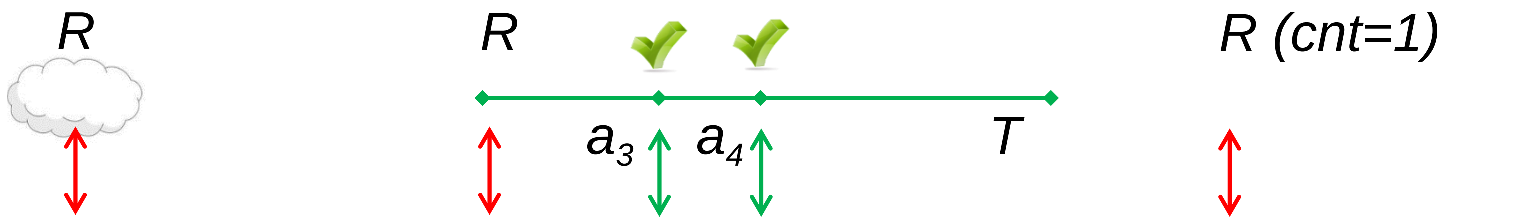

Insertion policies

Always on 1st (T)

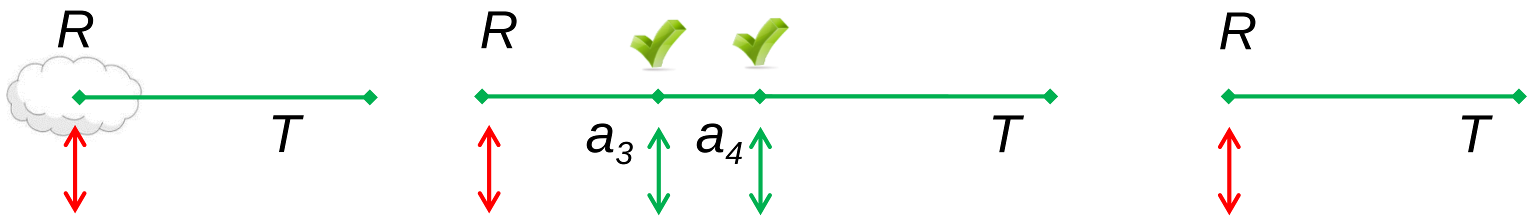

Always on 2nd (T)

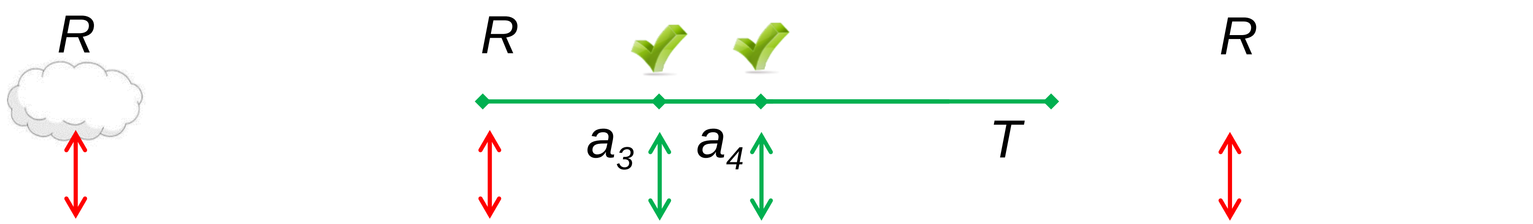

Single-window on 2nd (T)

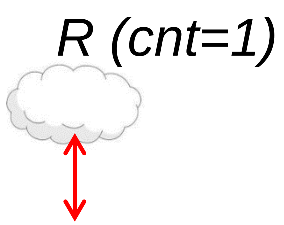

Insertion policies

Always on 1st (T)

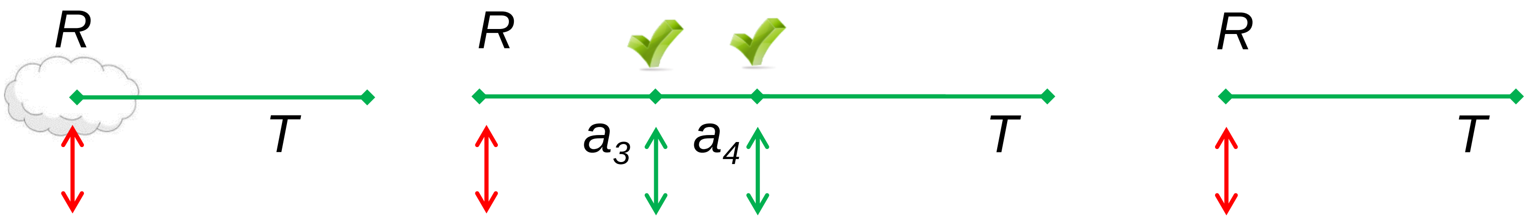

Always on 2nd (T)

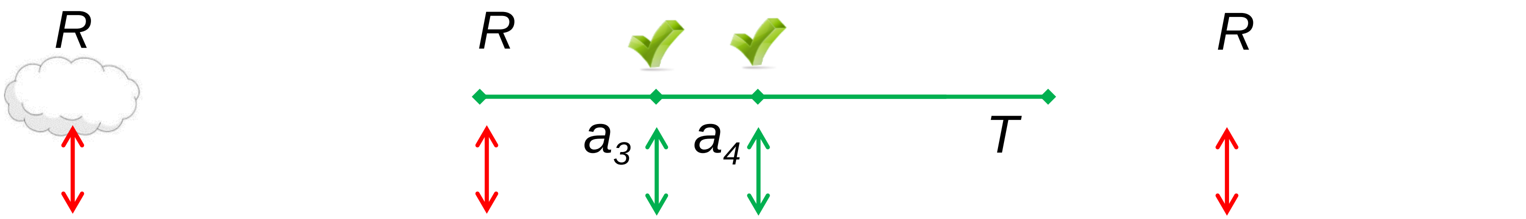

Single-window on 2nd (T)

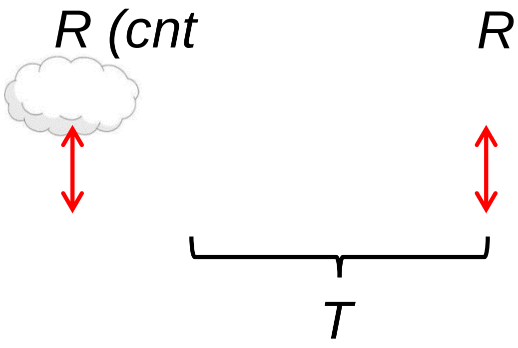

Insertion policies

Always on 1st (T)

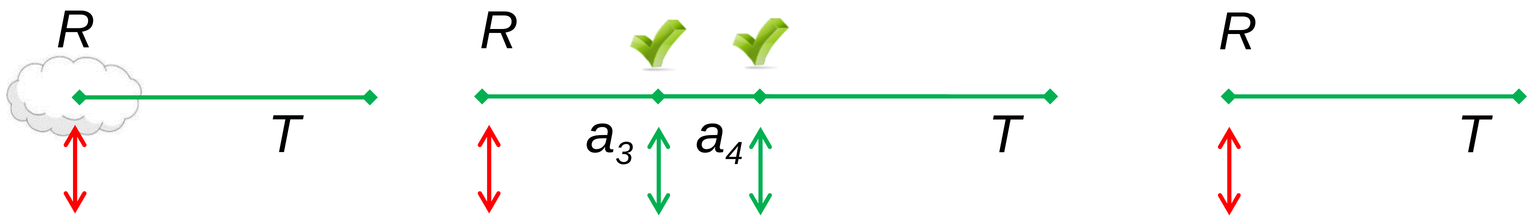

Always on 2nd (T)

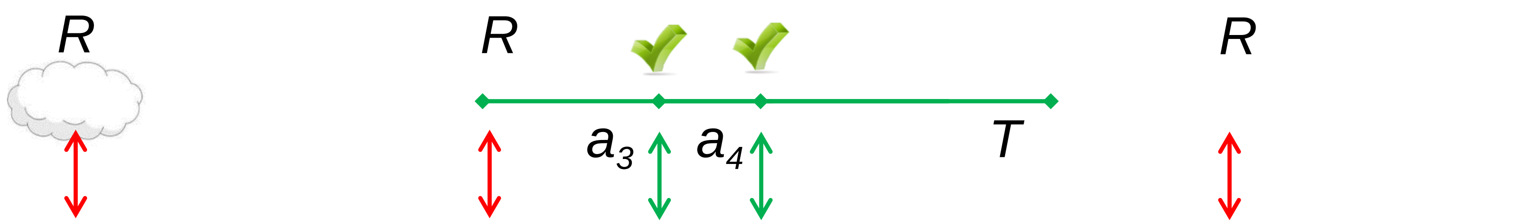

Single-window on 2nd (T)

Insertion policies

Always on 1st (T)

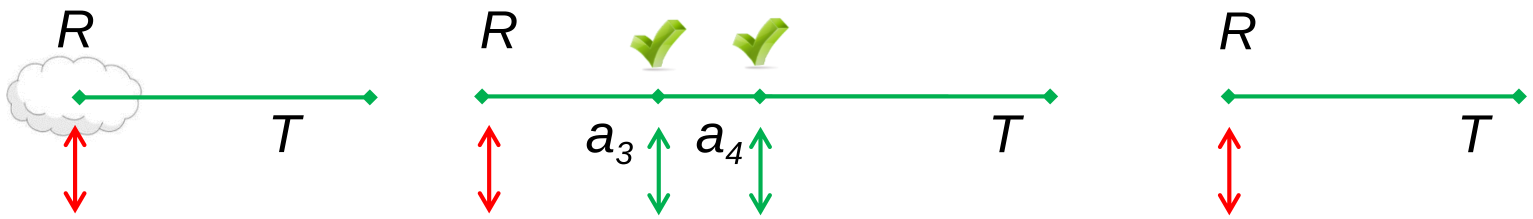

Always on 2nd (T)

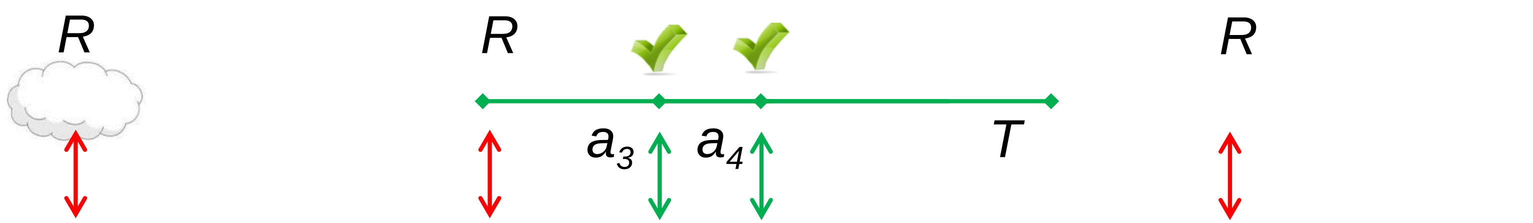

Single-window on 2nd (T)

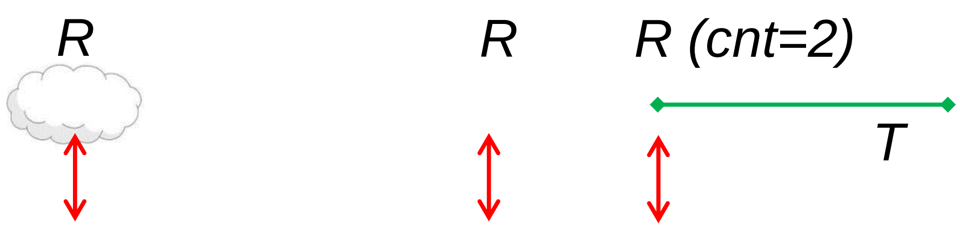

Insertion policies

Always on 1st (T)

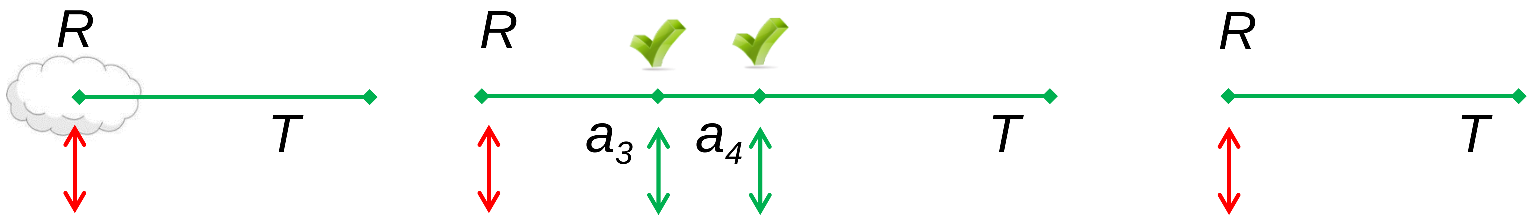

Always on 2nd (T)

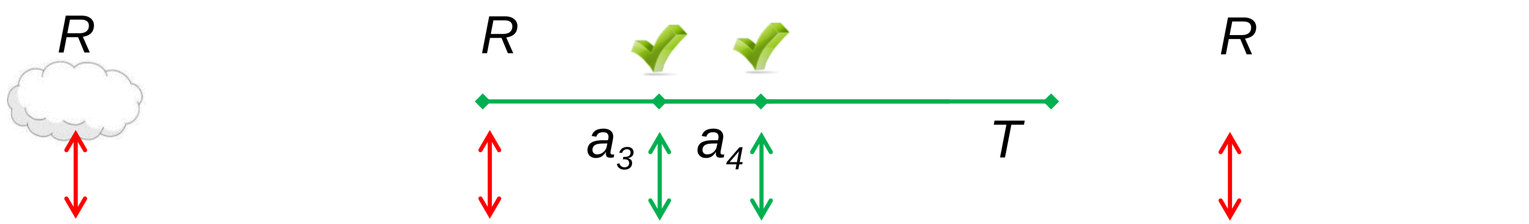

Single-window on 2nd (T)

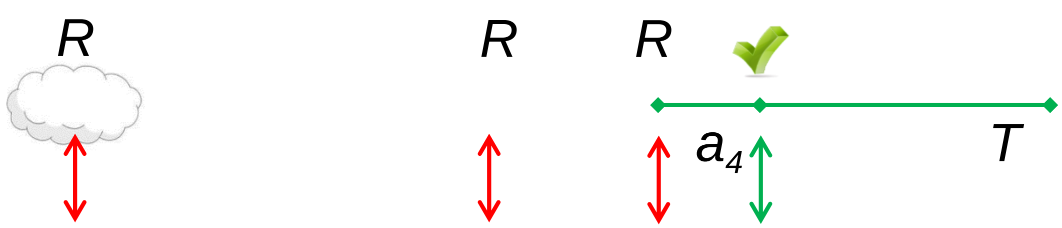

Insertion policies

Always on 1st (*T*)

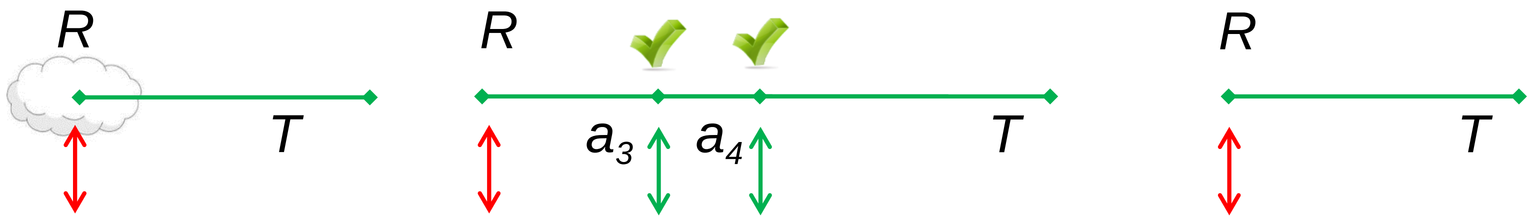

Always on 2nd (*T*)

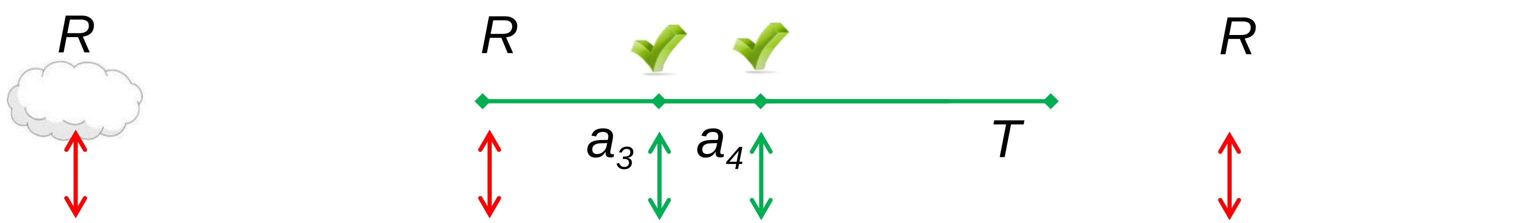

Single-window on 2nd (*T*)

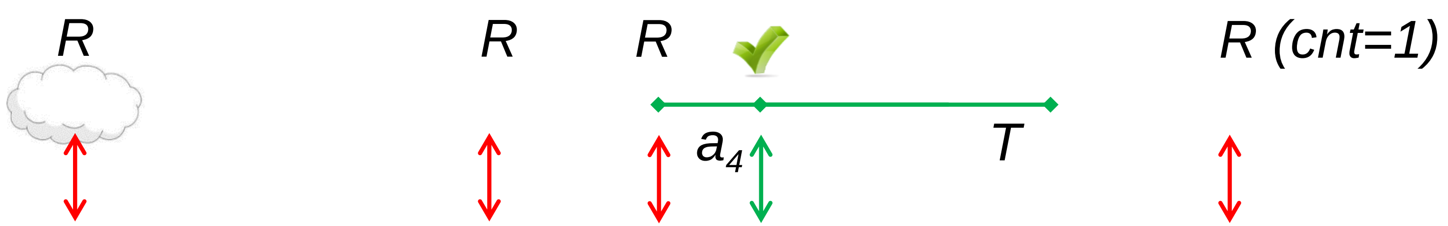

Insertion policies

Always on 1st (T)

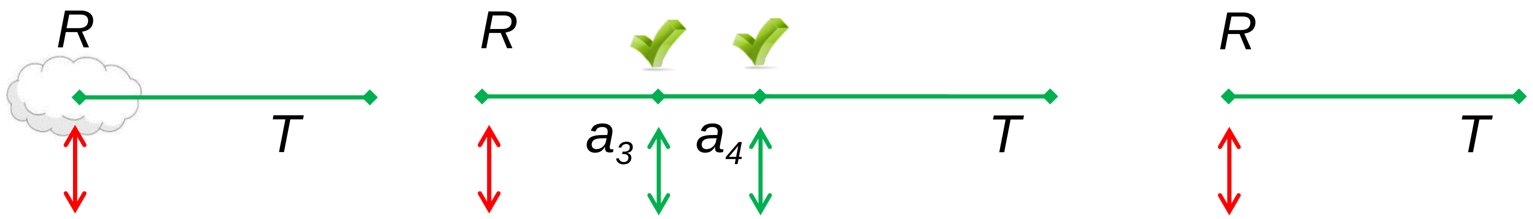

Always on 2nd (T)

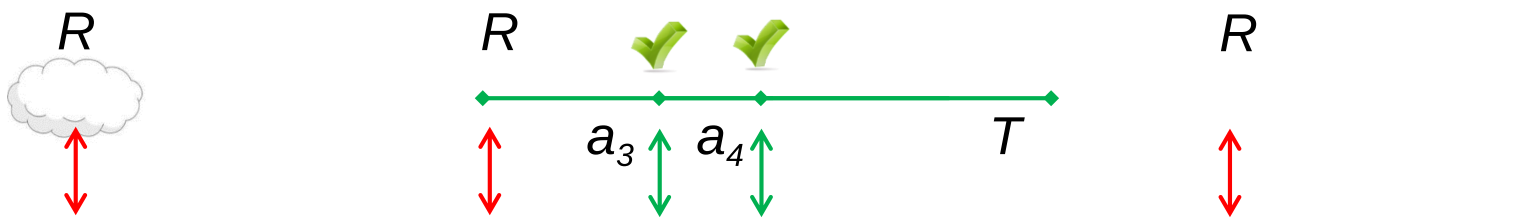

Single-window on 2nd (T)

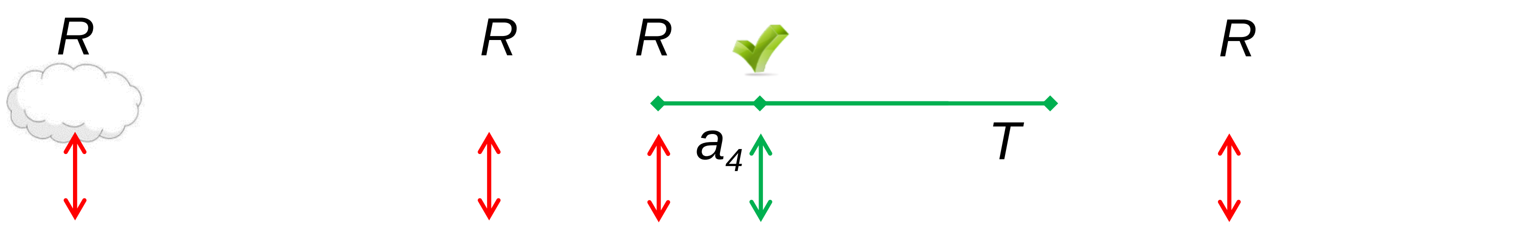

Insertion policies

Always on 1st (T)

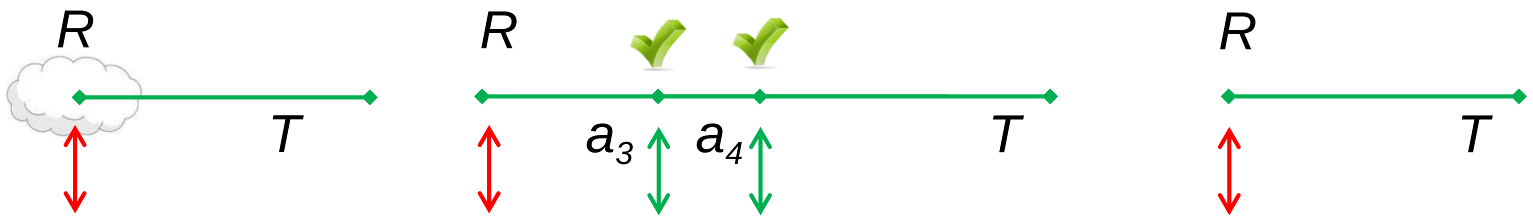

Always on 2nd (T)

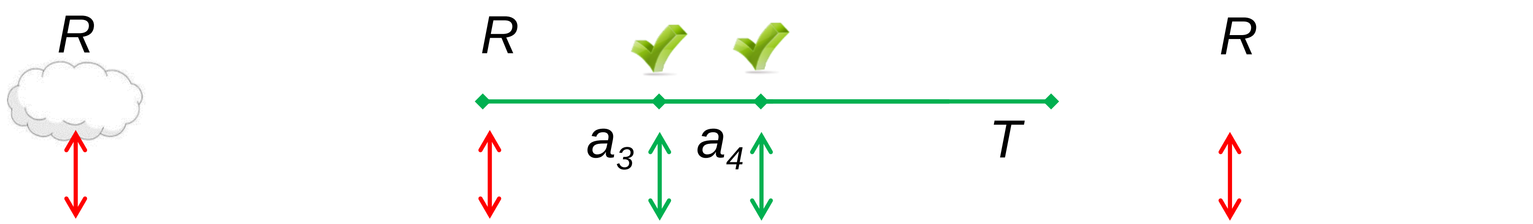

Single-window on 2nd (T)

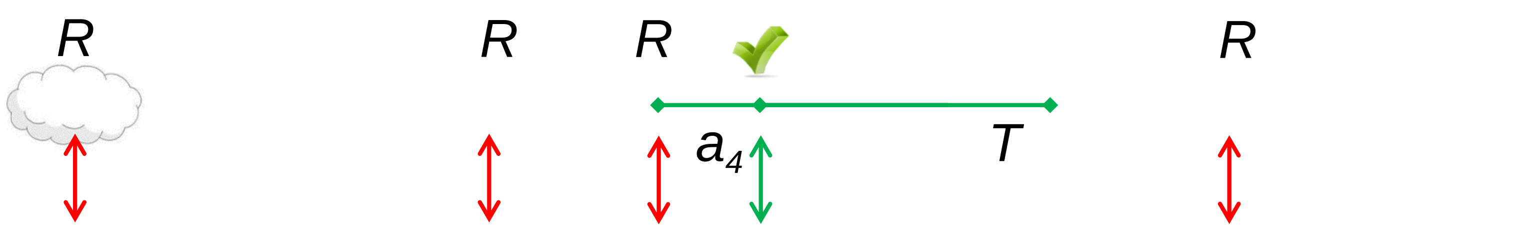

Dual-window on 2nd (W ≤ T), here W = T/2

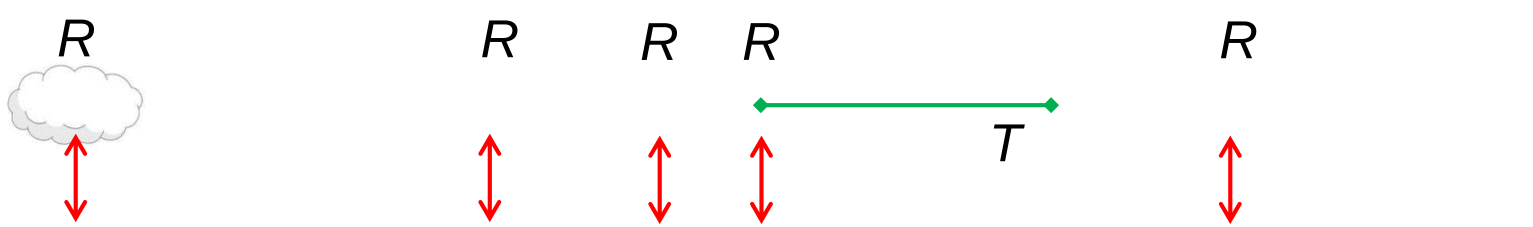

Insertion policies

Always on 1st (T)

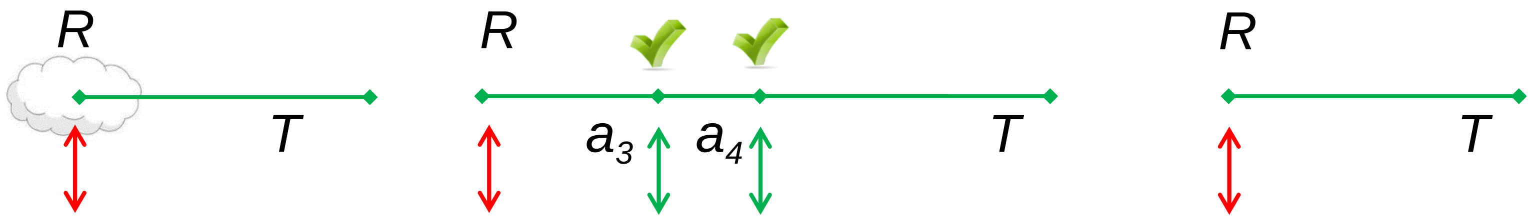

Always on 2nd (T)

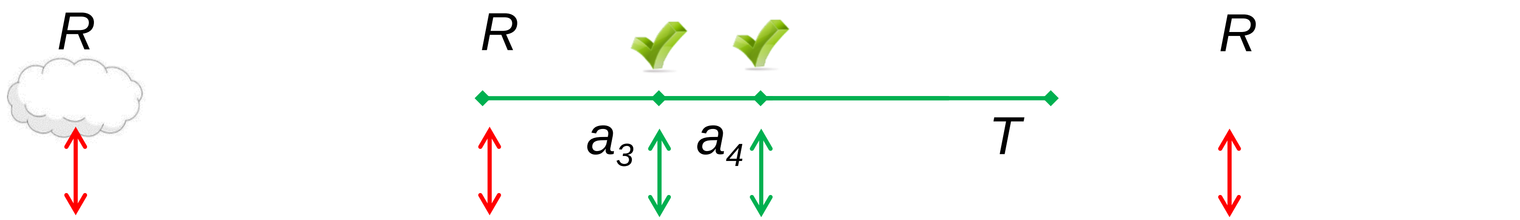

Single-window on 2nd (T)

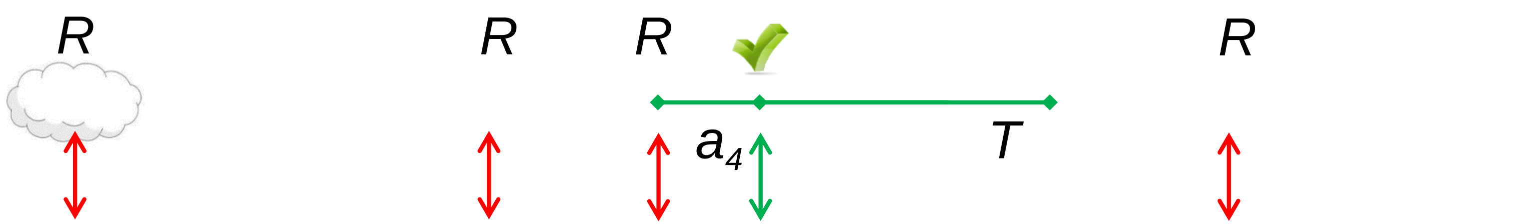

Single-window on 3rd (T)

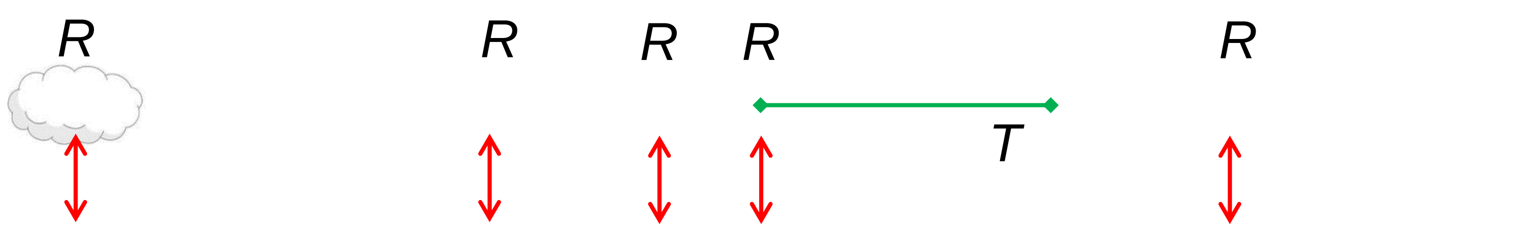

Worst-case bounds

Offline-optimal lower bound

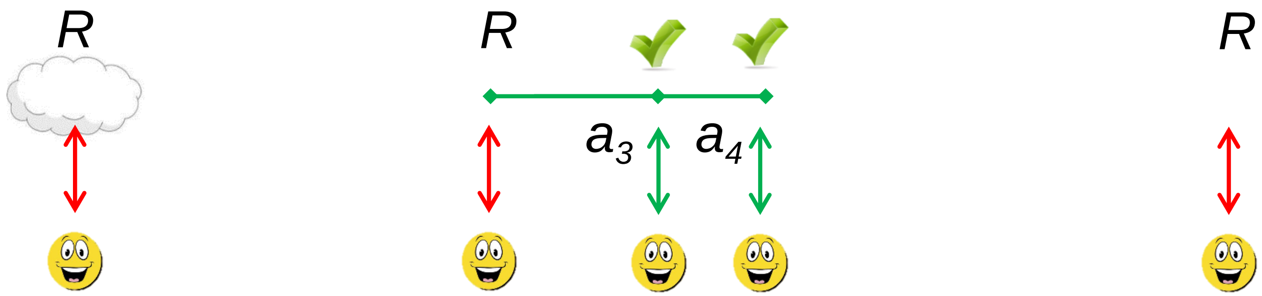

"Oracle" policy: *Keep in cache until (at least) the next inter-request arrival i whenever $a_i < R$; otherwise, do not cache.*

Offline-optimal lower bound

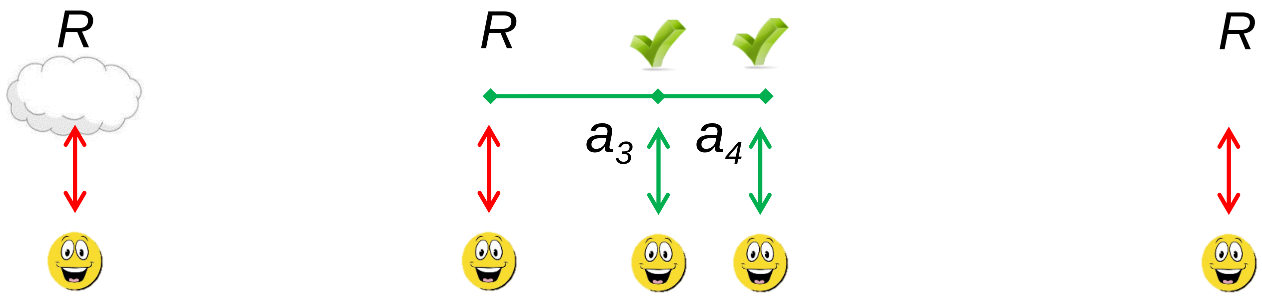

"Oracle" policy: *Keep in cache until (at least) the next inter-request arrival i whenever a_i < R; otherwise, do not cache.*

LEMMA 4.1. *Given an arbitrary request sequence $\mathcal{A}$, the minimum total delivery cost of the optimal offline policy is:*

$$C_{opt}^{offline} = R + \sum_{i=2}^{N} \min[a_i, R]. \qquad (1)$$

Example: Always on 1st

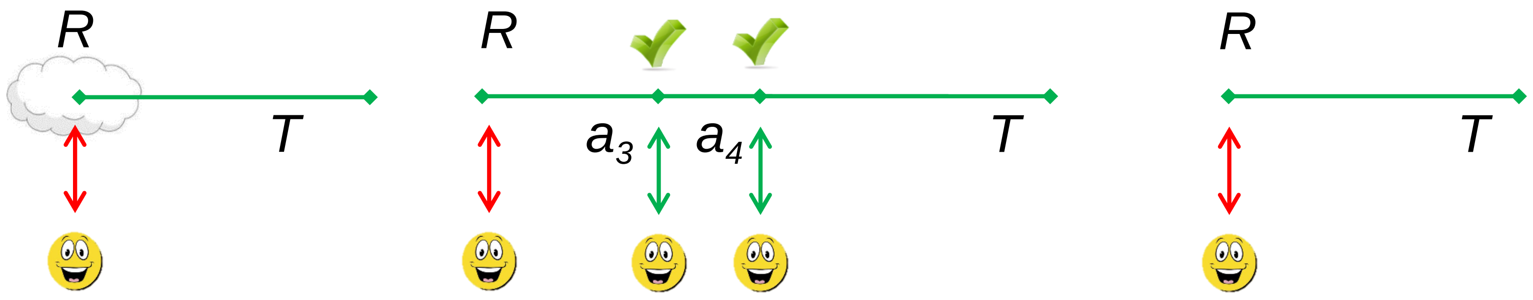

$$C_{M=1,T}^{always} = R + T + \sum_{i=2}^{N} x_i, \qquad (2)$$

$$x_i = \begin{cases} T + R, & \text{if } a_i > T \\ a_i, & \text{otherwise.} \end{cases} \qquad (3)$$

Worst-case ratio: Always on 1^{st}

THEOREM 4.2. *The best (optimal) competitive ratio using* always on 1^{st} *is achieved with* $T = R$ *and is equal to 2. More specifically,*

$$\max_{\mathcal{A}} \frac{C_{M=1,T=R}^{always}}{C_{opt}^{offline}} \le \max_{\mathcal{A}} \frac{C_{M=1,T}^{always}}{C_{opt}^{offline}} \tag{4}$$

for all T, *and* $\frac{C_{M=1,T=R}^{always}}{C_{opt}^{offline}} \le 2$ *for all possible sequences* $\mathcal{A} = \{a_i\}$.

Worst-case ratio: Always on 1st

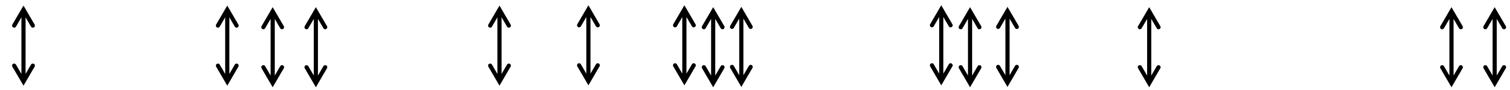

$$\frac{C_{M=1,T}^{always}}{C_{opt}^{offline}} =$$

?? Given arbitrary worst-case request sequence

$$C_{M=1,T}^{always} = R + T + \sum_{i=2}^{N} x_i, \qquad (2)$$

$$x_i = \begin{cases} T + R, & \text{if } a_i > T \\ a_i, & \text{otherwise.} \end{cases} \qquad (3)$$

$$C_{opt}^{offline} = R + \sum_{i=2}^{N} \min[a_i, R]. \qquad (1)$$

Worst-case ratio: Always on 1st

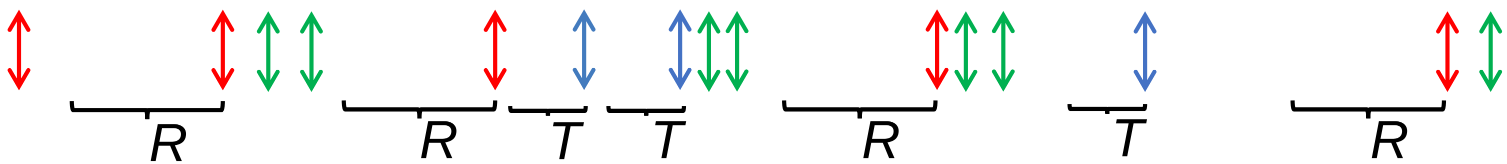

Case: $T \leq R$

$$S = \{i | a_i \leq T\},$$

$$S' = \{i | T < a_i \leq R\},$$

$$S'' = \{i | R < a_i\}.$$

$$\frac{C_{M=1,T}^{always}}{C_{opt}^{offline}} = \quad ??$$

$$C_{M=1,T}^{always} = R + T + \sum_{i=2}^{N} x_i, \qquad (2)$$

$$x_i = \begin{cases} T + R, & \text{if } a_i > T \\ a_i, & \text{otherwise.} \end{cases} \qquad (3)$$

$$C_{opt}^{offline} = R + \sum_{i=2}^{N} \min[a_i, R]. \qquad (1)$$

Worst-case ratio: Always on 1st

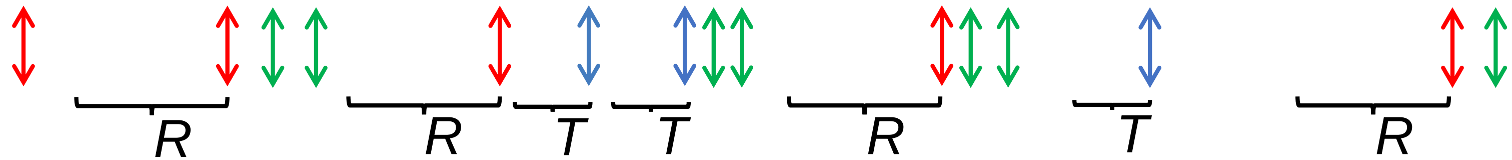

Case: $T \leq R$

$$S = \{i | a_i \leq T\},$$

$$S' = \{i | T < a_i \leq R\},$$

$$S'' = \{i | R < a_i\}.$$

$$\frac{C_{M=1,T}^{always}}{C_{opt}^{offline}} = \frac{R + \sum_{i \in S} a_i + (|S'| + |S''|)(T + R) + T}{R + \sum_{i \in S} a_i + \sum_{i \in S'} a_i + |S''| R}$$

$$\leq \frac{(R+T)(1+|S''|) + (R+T)|S'|}{R(1+|S''|) + \sum_{i \in S'} a_i}$$

... [some steps] ...

$$\leq \frac{(R+T)(1+|S''|) + (R+T)|S'|}{R(1+|S''|) + |S'| T} \leq \frac{R+T}{T}.$$

$$C_{M=1,T}^{always} = R + T + \sum_{i=2}^{N} x_i, \qquad (2)$$

$$x_i = \begin{cases} T + R, & \text{if } a_i > T \\ a_i, & \text{otherwise.} \end{cases} \qquad (3)$$

$$C_{opt}^{offline} = R + \sum_{i=2}^{N} \min[a_i, R]. \qquad (1)$$

Worst-case ratio: Always on 1st

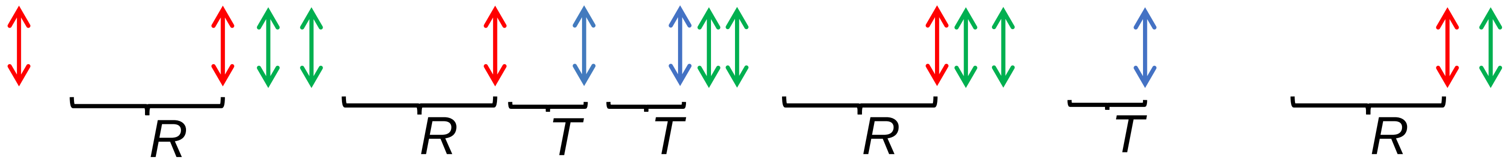

Case: $T \leq R$

$S = \{i | a_i \leq T\},$

$S' = \{i | T < a_i \leq R\},$

$S'' = \{i | R < a_i\}.$

$$\frac{C^{always}_{M=1,T}}{C^{offline}_{opt}} = \frac{R + \sum_{i \in S} a_i + (|S'| + |S''|)(T + R) + T}{R + \sum_{i \in S} a_i + \sum_{i \in S'} a_i + |S''|R}$$

$$\leq \frac{(R+T)(1+|S''|) + (R+T)|S'|}{R(1+|S''|) + \sum_{i \in S'} a_i}$$

... [some steps] ...

$$\leq \frac{(R+T)(1+|S''|) + (R+T)|S'|}{R(1+|S''|) + |S'|T} \leq \frac{R+T}{T}.$$

Worst-case ratio: Always on 1st

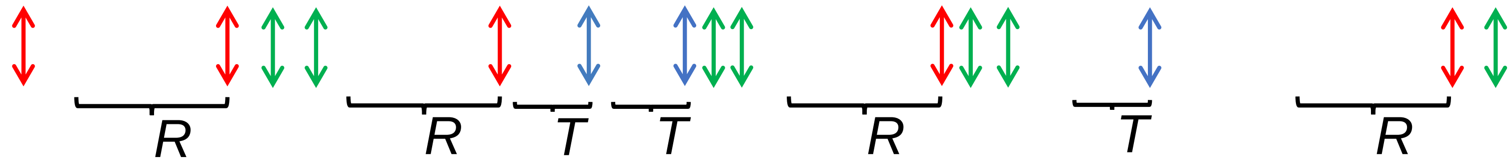

Case: $T \leq R$

$$S = \{i | a_i \leq T\},$$

$$S' = \{i | T < a_i \leq R\},$$

$$S'' = \{i | R < a_i\}.$$

$$\frac{C^{always}_{M=1,T}}{C^{offline}_{opt}} = \frac{R + \sum_{i \in S} a_i + (|S'| + |S''|)(T + R) + T}{R + \sum_{i \in S} a_i + \sum_{i \in S'} a_i + |S''| R}$$

$$\leq \frac{(R+T)(1+|S''|) + (R+T)|S'|}{R(1+|S''|) + \sum_{i \in S'} a_i}$$

… [some steps] …

$$\leq \frac{(R+T)(1+|S''|) + (R+T)|S'|}{R(1+|S''|) + |S'|T} \leq \frac{R+T}{T}.$$

Bound monotonically decreasing in range $0 \leq T \leq R$.

Bound tight when $T \rightarrow R$ (and equal to 2); achieved with $T+\varepsilon$ spacing

Similar approach for case when $R \leq T$

Worst-case bounds

Policy	Parameters	Optimal choice	Tight bound
Always 1st	T	$T = R$	2
Always Mth			
Single-window Mth			
Dual-window 2nd			

Worst-case bounds

Policy	Parameters	Optimal choice	Tight bound
Always 1st	T	$T = R$	2
Always Mth	T	$T = R$	M+1
Single-window Mth	T	$T = R$	M+1
Dual-window 2nd	W, T	$W = T = R$	3

Worst-case bounds

Policy	Parameters	Optimal choice	Tight bound
Always 1st	T	$T = R$	2
Always Mth	T	$T = R$	$M+1$
Single-window Mth	T	$T = R$	$M+1$
Dual-window 2nd	W, T	$W = T = R$	3

- Although $M+1$ worst-case bounds may seem discouraging, we will see that window-based policies are good on average (across different distributions and distribution parameters)

Steady-state analysis

Offline-optimal lower bound

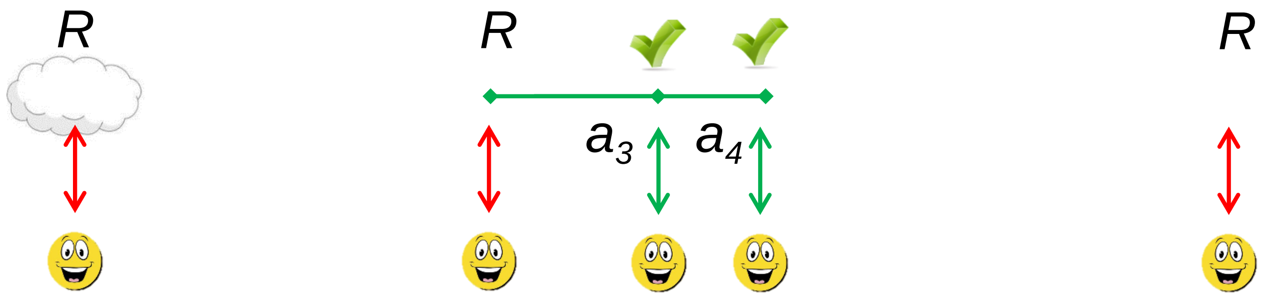

"Oracle" policy: *Keep in cache until (at least) the next inter-request arrival i whenever a_i < R; otherwise, do not cache.*

$$C_{opt}^{offline} = \frac{1}{E[a_i]}\left[\int_0^R t f(t)\mathrm{dt} + R\int_R^\infty f(t)\mathrm{dt}\right]$$

Offline-optimal lower bound

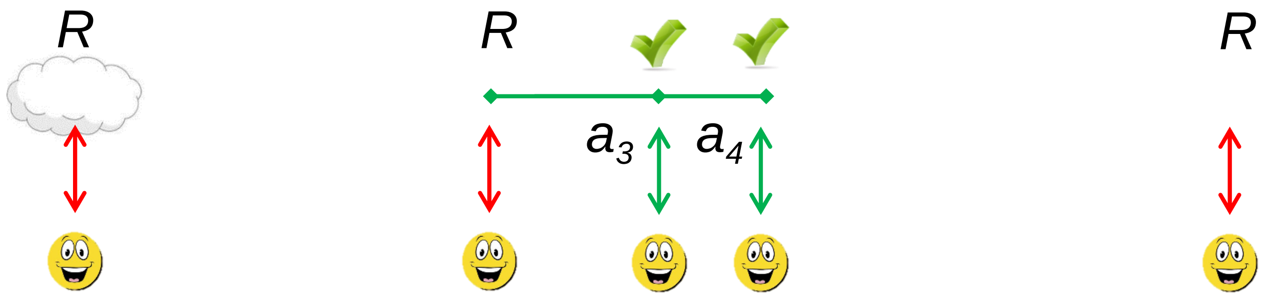

"Oracle" policy: *Keep in cache until (at least) the next inter-request arrival i whenever a_i < R; otherwise, do not cache.*

$$C_{opt}^{offline} = \underbrace{\frac{1}{E[a_i]}}_{\text{Rate of new requests}} \left[\underbrace{\int_0^R t f(t) \mathrm{dt}}_{\text{Cost } a_i} + \underbrace{R \int_R^\infty f(t) \mathrm{dt}}_{\text{Cost } R} \right]$$

Rate of new requests

Cost a_i (per request)

Cost R (per request)

Offline-optimal lower bound

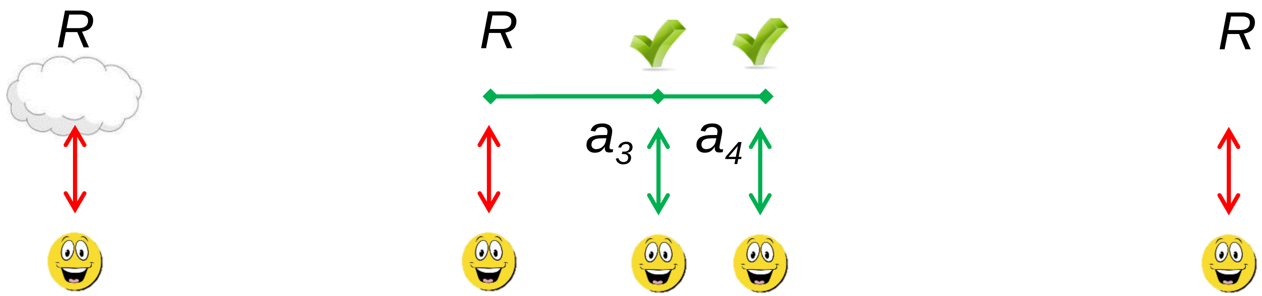

"Oracle" policy: *Keep in cache until (at least) the next inter-request arrival i whenever $a_i < R$; otherwise, do not cache.*

$$C_{opt}^{offline} = \frac{1}{E[a_i]}\left[\int_0^R t f(t)\mathrm{d}t + R\int_R^{\infty} f(t)\mathrm{d}t\right]$$

… [some steps] …

$$= \frac{1}{E[a_i]}\left[R - \int_0^R F(t)\mathrm{d}t\right]. \qquad (12)$$

Example distribution results

Table 1: Summary of costs for different distributions and insertion policies. To make room, for Erlang, we simplified expressions using $F(t) = 1 - \sum_{n=0}^{k-1} \frac{1}{n!} e^{-\lambda t} (\lambda t)^n$ and $\Phi(T) = \frac{e^{-\lambda T}}{\lambda} \sum_{m=1}^{k} \sum_{n=0}^{m-1} \frac{(\lambda T)^n}{n!}$.

Policy	Exponential	Erlang	Deterministic	Pareto
Offline	$1 - e^{-\lambda R}$	$1 - \frac{\lambda}{k} \Phi(R)$	$\min[\frac{R}{a}, 1]$	$\begin{cases} 1 - \frac{1}{\alpha} \left(\frac{t_m}{R}\right)^{\alpha - 1}, & \text{if } t_m \leq R \\ \frac{R(\alpha-1)}{\alpha t_m}, & \text{if } R < t_m \end{cases}$

Exponential *Erlang* *Deterministic* — *Short-tailed*

Pareto — *Heavy-tailed*

Example distribution results

Table 1: Summary of costs for different distributions and insertion policies. To make room, for Erlang, we simplified expressions using $F(t) = 1 - \sum_{n=0}^{k-1} \frac{1}{n!} e^{-\lambda t} (\lambda t)^n$ and $\Phi(T) = \frac{e^{-\lambda T}}{\lambda} \sum_{m=1}^{k} \sum_{n=0}^{m-1} \frac{(\lambda T)^n}{n!}$.

Policy	Exponential	Erlang	Deterministic	Pareto
Offline	$1 - e^{-\lambda R}$	$1 - \frac{\lambda}{k} \Phi(R)$	$\min[\frac{R}{a}, 1]$	$\begin{cases} 1 - \frac{1}{\alpha} \left(\frac{t_m}{R}\right)^{\alpha - 1}, & \text{if } t_m \le R \\ \frac{R(\alpha-1)}{\alpha t_m}, & \text{if } R < t_m \end{cases}$
Baseline	$\min[\lambda R, 1]$	$\min[\frac{\lambda}{k} R, 1]$	$\min[\frac{R}{a}, 1]$	$\min[\frac{\alpha - 1}{\alpha} \frac{R}{t_m}, 1]$

Exponential *Erlang* *Deterministic* *Pareto*

Short-tailed *Heavy-tailed*

“Static baseline” policy: *Either “always remote” or “always local”.*

Example distribution results

Table 1: Summary of costs for different distributions and insertion policies. To make room, for Erlang, we simplified expressions using $F(t) = 1 - \sum_{n=0}^{k-1} \frac{1}{n!} e^{-\lambda t} (\lambda t)^n$ and $\Phi(T) = \frac{e^{-\lambda T}}{\lambda} \sum_{m=1}^{k} \sum_{n=0}^{m-1} \frac{(\lambda T)^n}{n!}$.

Policy	Exponential	Erlang	Deterministic	Pareto
Offline	$1 - e^{-\lambda R}$	$1 - \frac{\lambda}{k}\Phi(R)$	$\min[\frac{R}{a}, 1]$	$1 - \frac{1}{\alpha}\left(\frac{t_m}{R}\right)^{\alpha-1}$, if $t_m \leq R$ $\frac{R(\alpha-1)}{\alpha t_m}$, if $R < t_m$
Baseline	$\min[\lambda R, 1]$	$\min[\frac{\lambda}{k} R, 1]$	$\min[\frac{R}{a}, 1]$	$\min[\frac{\alpha-1}{\alpha} \frac{R}{t_m}, 1]$

Exponential *Erlang* *Deterministic*

Short-tailed

"Static baseline" policy: *Either "always remote" or "always local".*

Example distribution results

Table 1: Summary of costs for different distributions and insertion policies. To make room, for Erlang, we simplified expressions using $F(t) = 1 - \sum_{n=0}^{k-1} \frac{1}{n!} e^{-\lambda t} (\lambda t)^n$ and $\Phi(T) = \frac{e^{-\lambda T}}{\lambda} \sum_{m=1}^{k} \sum_{n=0}^{m-1} \frac{(\lambda T)^n}{n!}$.

Policy	Exponential	Erlang	Deterministic	Pareto
Offline	$1 - e^{-\lambda R}$	$1 - \frac{\lambda}{k}\Phi(R)$	$\min[\frac{R}{a}, 1]$	$1 - \frac{1}{\alpha}\left(\frac{t_m}{R}\right)^{\alpha-1}$, if $t_m \leq R$; $\frac{R(\alpha-1)}{\alpha t_m}$, if $R < t_m$
Baseline	$\min[\lambda R, 1]$	$\min[\frac{\lambda}{k} R, 1]$	$\min[\frac{R}{a}, 1]$	$\min[\frac{\alpha-1}{\alpha} \frac{R}{t_m}, 1]$

Exponential *Erlang* *Deterministic*

Short-tailed

"Static baseline" policy: *Either "always remote" or "always local".*

THEOREM 6.1. Static baseline *achieves the minimum cost of any* online policy *when the inter-request distribution has an increasing or constant hazard rate.*

Example distribution results

Table 1: Summary of costs for different distributions and insertion policies. To make room, for Erlang, we simplified expressions using $F(t) = 1 - \sum_{n=0}^{k-1} \frac{1}{n!} e^{-\lambda t} (\lambda t)^n$ and $\Phi(T) = \frac{e^{-\lambda T}}{\lambda} \sum_{m=1}^{k} \sum_{n=0}^{m-1} \frac{(\lambda T)^n}{n!}$.

Policy	Exponential	Erlang	Deterministic	Pareto
Offline	$1 - e^{-\lambda R}$	$1 - \frac{\lambda}{k}\Phi(R)$	$\min[\frac{R}{a}, 1]$	$1 - \frac{1}{\alpha}\left(\frac{t_m}{R}\right)^{\alpha-1}$, if $t_m \le R$; $\frac{R(\alpha-1)}{\alpha t_m}$, if $R < t_m$
Baseline	$\min[\lambda R, 1]$	$\min[\frac{\lambda}{k} R, 1]$	$\min[\frac{R}{a}, 1]$	$\min[\frac{\alpha-1}{\alpha} \frac{R}{t_m}, 1]$

Exponential　*Erlang*　*Deterministic*

Short-tailed

"Static baseline" policy: *Either "always remote" or "always local".*

THEOREM 6.1. Static baseline *achieves the minimum cost of any* online policy *when the inter-request distribution has an increasing or constant hazard rate.*

… is ***online optimal*** *for these cases!!*

Gap between online and offline optimal

THEOREM 6.3. *Under Poisson requests we have*

$$\frac{C_{opt}^{online}}{C_{opt}^{offline}} = \frac{C_{opt}^{static}}{C_{opt}^{offline}} \leq \frac{1}{1 - 1/e}. \qquad (31)$$

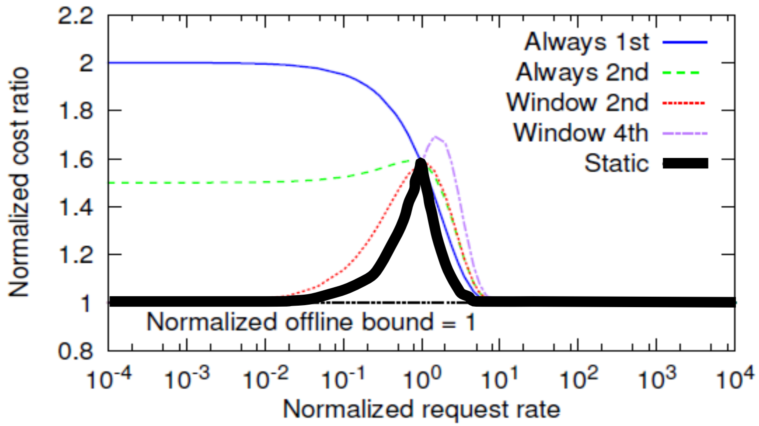

THEOREM 6.1. Static baseline *achieves the minimum cost of any* online policy *when the inter-request distribution has an increasing or constant hazard rate.*

Gap between online and offline optimal

THEOREM 6.3. *Under Poisson requests we have*

$$\frac{C_{opt}^{online}}{C_{opt}^{offline}} = \frac{C_{opt}^{static}}{C_{opt}^{offline}} \leq \frac{1}{1 - 1/e}. \tag{31}$$

THEOREM 6.4. *Under Erlang inter-request times, we have*

$$\frac{C_{opt}^{online}}{C_{opt}^{offline}} = \frac{C_{opt}^{static}}{C_{opt}^{offline}} \leq \frac{1}{1 - e^{-k} \frac{k^k}{k!}}. \tag{32}$$

THEOREM 6.5. *Under deterministic inter-request times, we have*

$$\frac{C_{opt}^{online}}{C_{opt}^{offline}} = \frac{C_{opt}^{static}}{C_{opt}^{offline}} = 1. \tag{36}$$

THEOREM 6.1. Static baseline *achieves the minimum cost of any* online policy *when the inter-request distribution has an increasing or constant hazard rate.*

However, not true for heavy-tailed …

… in fact, for Pareto the optimal static baseline can be far from optimal

THEOREM 6.6. *With Pareto inter-request times, the worst-case cost ratio for the optimal* static baseline *is unbounded. In particular,*

$$\frac{C_{opt}^{static}}{C_{opt}^{offline}} \to \infty \tag{37}$$

when $\alpha = \frac{1}{1-\frac{t_m}{R}}$ *and* $\frac{t_m}{R} \to 0+$.

Policy analysis: Always on 1st

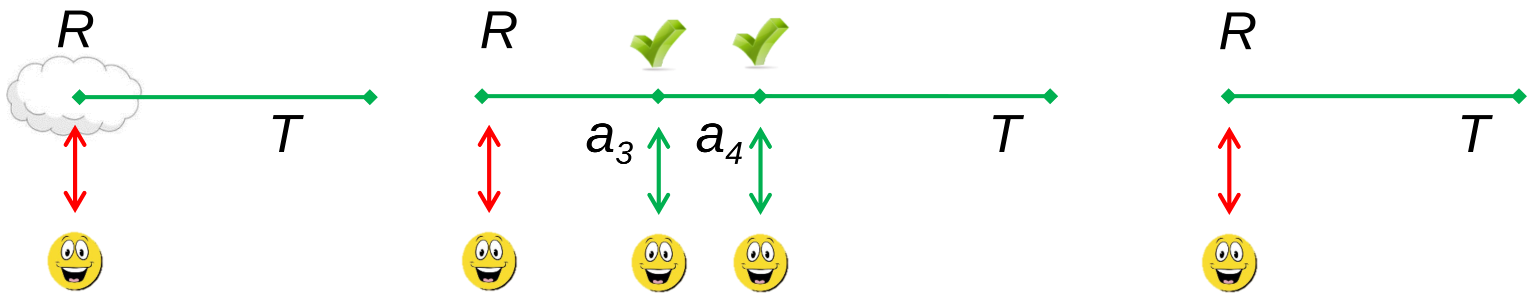

$$C_{M=1,T}^{always} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \quad (42)$$

Policy analysis: Always on 1st

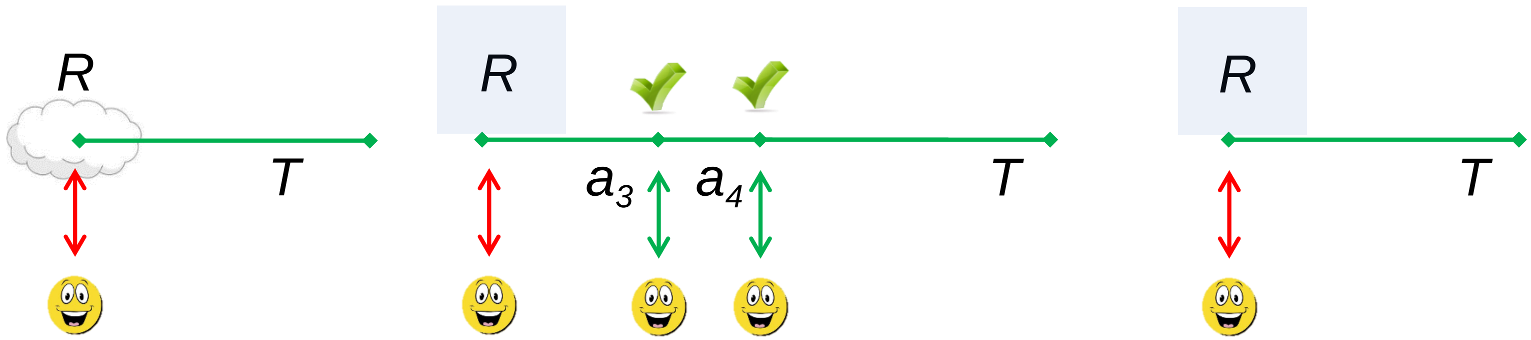

$$C^{always}_{M=1,T} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \quad (42)$$

Policy analysis: Always on 1st

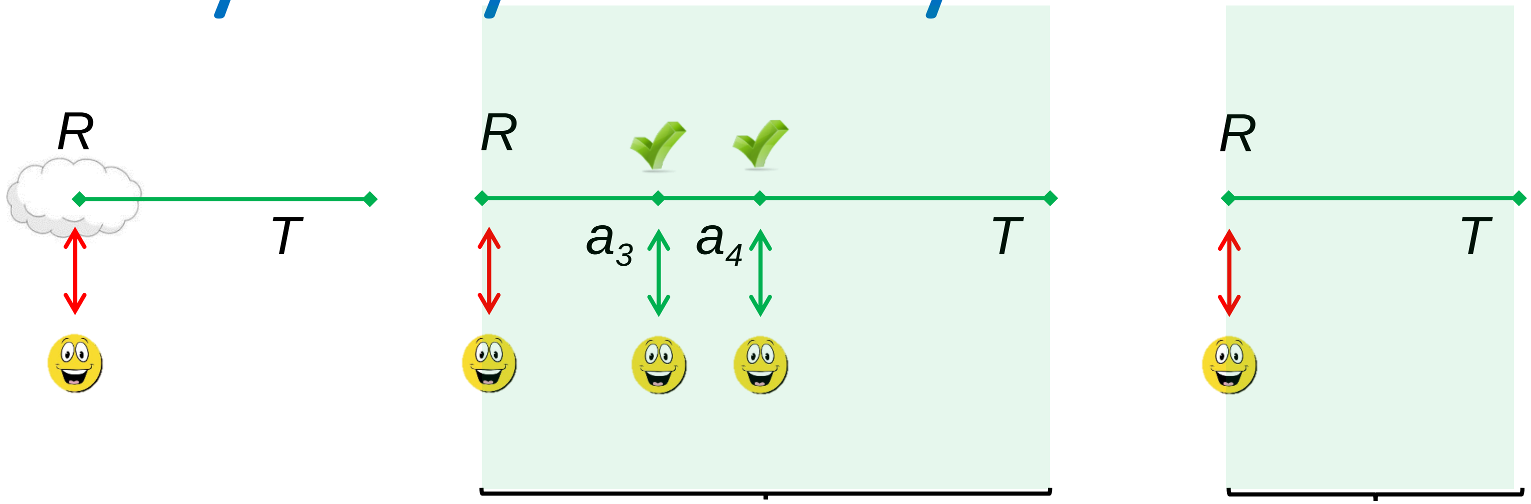

$$C^{always}_{M=1,T} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \tag{42}$$

Policy analysis: Always on 1st

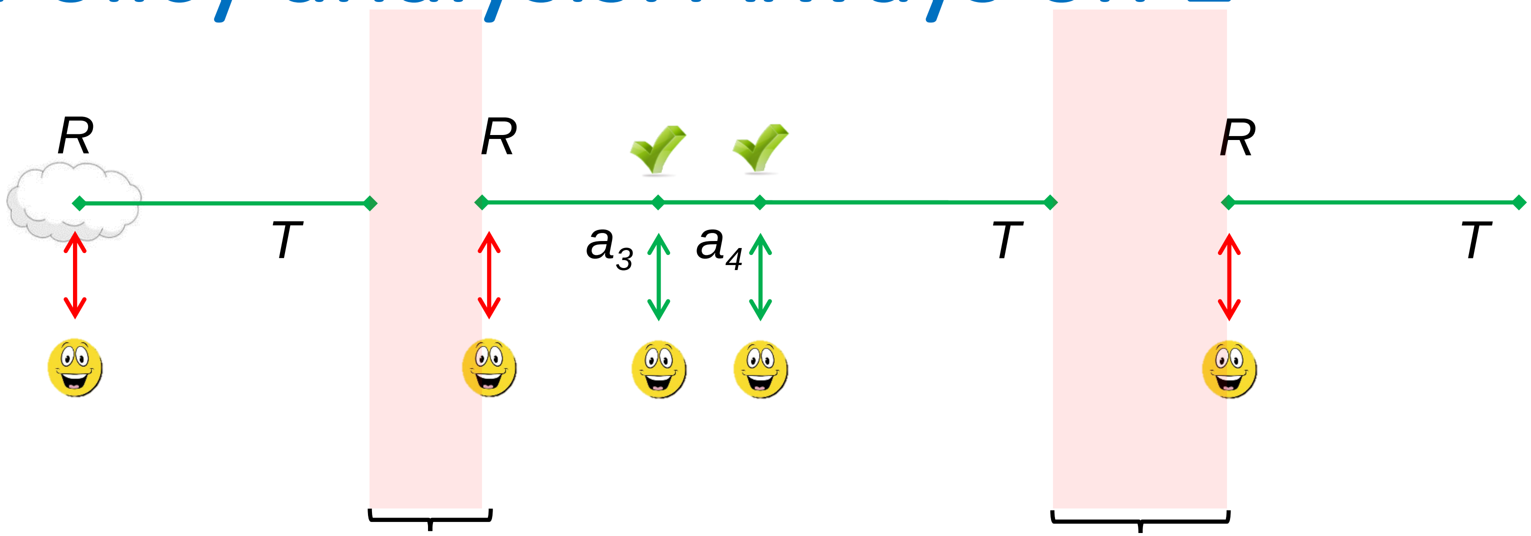

$$C^{always}_{M=1,T} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \qquad (42)$$

Policy analysis: Always on 1st

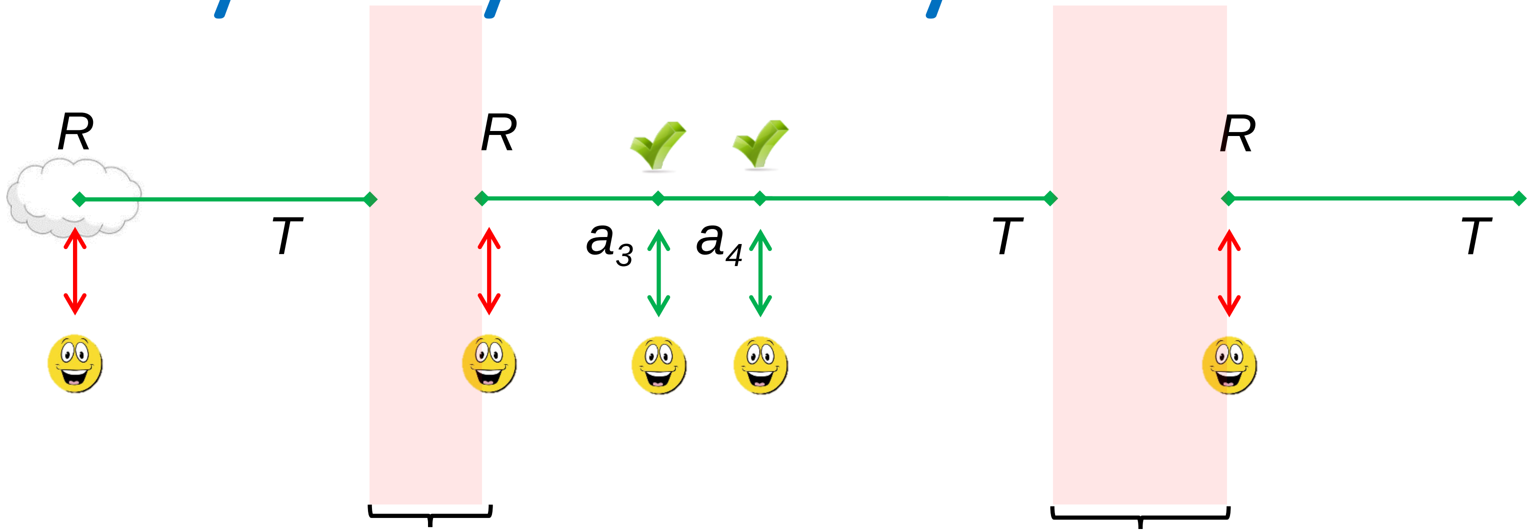

$$C^{always}_{M=1,T} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \quad (42)$$

$$E[\Delta_1] = E[a_i | a_i > T] - T = \frac{1}{1 - F(T)} \left(E[a_i] + \int_0^T F(t) \mathrm{d}t - T \right), \quad (43)$$

Policy analysis: Always on 1st

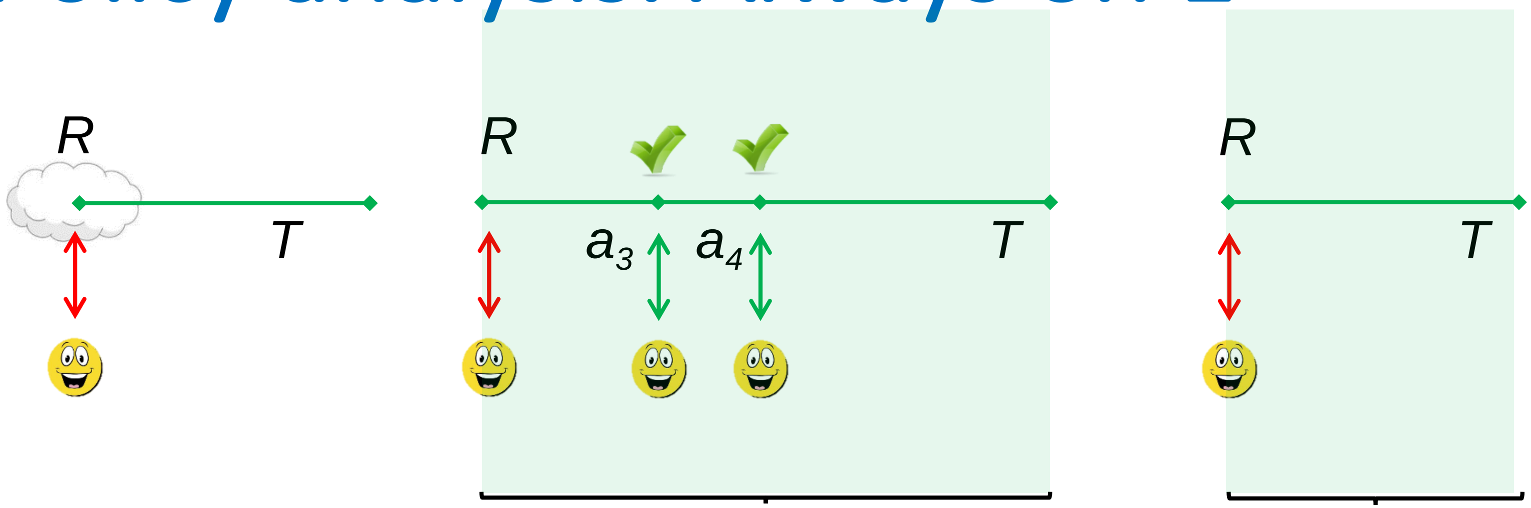

$$C^{always}_{M=1,T} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \qquad (42)$$

$$E[\Delta_1] = E[a_i | a_i > T] - T = \frac{1}{1 - F(T)} \left(E[a_i] + \int_0^T F(t)\mathrm{d}t - T \right), \qquad (43)$$

$$E[\Theta] = (1 - F(T))T + F(T)(E[a_i | a_i < T] + E[\Theta]), \qquad (44)$$

Policy analysis: Always on 1st

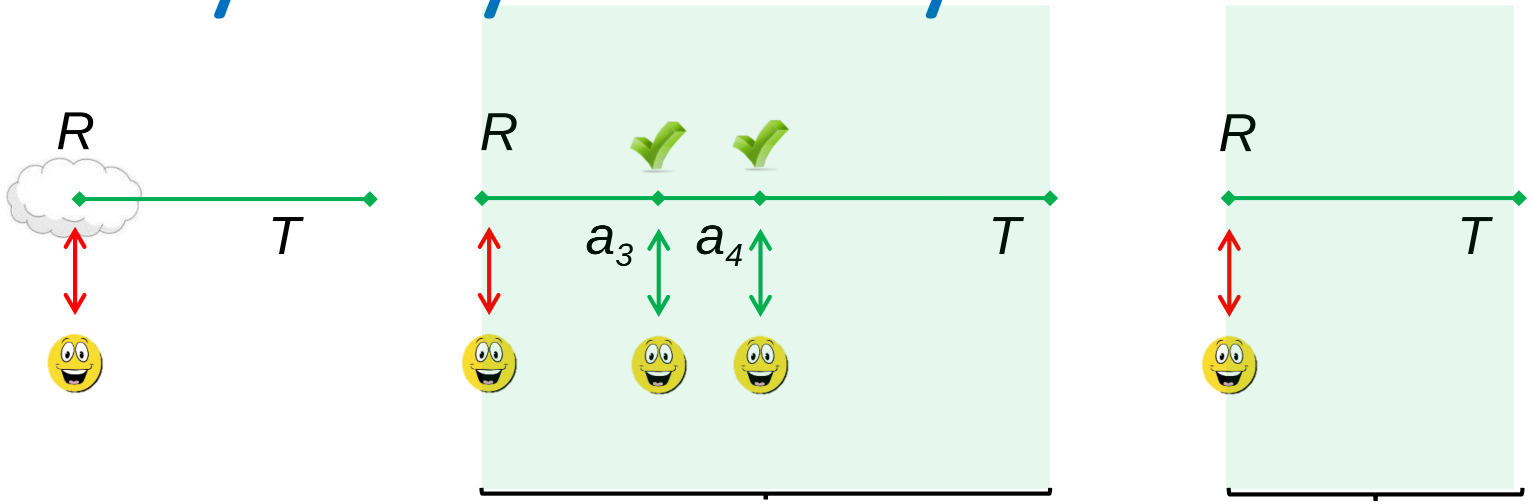

$$C^{always}_{M=1,T} = \frac{R + E[\Theta]}{E[\Delta_1] + E[\Theta]}, \quad (42)$$

$$E[\Delta_1] = E[a_i | a_i > T] - T = \frac{1}{1 - F(T)} \left(E[a_i] + \int_0^T F(t)\mathrm{d}t - T \right), \quad (43)$$

$$E[\Theta] = \underbrace{(1 - F(T))T}_{\text{No extension}} + \underbrace{F(T)(E[a_i | a_i < T] + E[\Theta])}_{\text{Extension case}}, \quad (44)$$

No extension *Extension case*

Results for example distributions

Example distributions: Summary of costs

Table 1: Summary of costs for different distributions and insertion policies. To make room, for Erlang, we simplified expressions using $F(t) = 1 - \sum_{n=0}^{k-1} \frac{1}{n!} e^{-\lambda t} (\lambda t)^n$ and $\Phi(T) = \frac{e^{-\lambda T}}{\lambda} \sum_{m=1}^{k} \sum_{n=0}^{m-1} \frac{(\lambda T)^n}{n!}$.

Policy	Exponential	Erlang	Deterministic	Pareto
Offline	$1 - e^{-\lambda R}$	$1 - \frac{\lambda}{k}\Phi(R)$	$\min[\frac{R}{a}, 1]$	$\begin{cases} 1 - \frac{1}{\alpha}\left(\frac{t_m}{R}\right)^{\alpha-1}, & \text{if } t_m \le R \\ \frac{R(\alpha-1)}{\alpha t_m}, & \text{if } R < t_m \end{cases}$
Baseline	$\min[\lambda R, 1]$	$\min[\frac{\lambda}{k} R, 1]$	$\min[\frac{R}{a}, 1]$	$\min[\frac{\alpha-1}{\alpha} \frac{R}{t_m}, 1]$
Always 1^{st}	$1 - e^{-\lambda T} + \lambda R e^{-\lambda T}$	$(1 - F(T))\frac{\lambda}{k} R + (1 - \frac{\lambda}{k}\Phi(T))$	$\begin{cases} 1, & \text{if } a \le T \\ \frac{R+T}{a}, & \text{if } T < a \end{cases}$	$\begin{cases} \frac{\alpha-1}{\alpha}\left(\frac{t_m}{T}\right)^{\alpha} \frac{R}{t_m} + \left(1 - \frac{1}{\alpha}\left(\frac{t_m}{T}\right)^{\alpha-1}\right), & \text{if } t_m \le T \\ \frac{(R+T)(\alpha-1)}{\alpha t_m}, & \text{if } T < t_m \end{cases}$
Always 2^{nd}	$\frac{1-e^{-\lambda T}+2\lambda R e^{-\lambda T}}{1+e^{-\lambda T}}$	$\frac{(1-F(T))\frac{\lambda}{k}2R+(1-\frac{\lambda}{k}\Phi(T))}{2-F(T)}$	$\begin{cases} 1, & \text{if } a \le T \\ \frac{2R+T}{2a}, & \text{if } T < a \end{cases}$	$\begin{cases} \frac{\frac{\alpha-1}{\alpha}\left(\frac{t_m}{T}\right)^{\alpha} \frac{2R}{t_m} + \left(1 - \frac{1}{\alpha}\left(\frac{t_m}{T}\right)^{\alpha-1}\right)}{1+\left(\frac{t_m}{T}\right)^{\alpha}}, & \text{if } t_m \le T \\ \frac{2R+T}{2} \frac{\alpha-1}{\alpha t_m}, & \text{if } T < t_m \end{cases}$
Single M^{th}	$\lambda e^{-\lambda T} \sum_{i=0}^{M-1} (1 - e^{-\lambda T})^i R + \left(1 - e^{-\lambda T}\right)^M$	$(1 - F(T))\frac{\lambda}{k} \sum_{i=0}^{M-1} F(T)^i R + \left(1 - \frac{\lambda}{k}\Phi(T)\right) F(T)^{M-1}$	$\begin{cases} 1, & \text{if } a \le T \\ \frac{R}{a}, & \text{if } T < a \end{cases}$	$\begin{cases} \frac{\alpha-1}{\alpha}\left(\frac{t_m}{T}\right)^{\alpha} \sum_{i=0}^{M-1} (1 - \left(\frac{t_m}{T}\right)^{\alpha})^i \frac{R}{T} + \left(1 - \frac{1}{\alpha}\left(\frac{t_m}{T}\right)^{\alpha-1}\right)(1 - \left(\frac{t_m}{T}\right)^{\alpha})^{M-1}, & \text{if } t_m \le T \\ \frac{R(\alpha-1)}{\alpha t_m}, & \text{if } T < t_m \end{cases}$
Dual 2^{nd}	$\frac{\lambda R e^{-\lambda T}\left(2-e^{-\lambda W}\right)+\left(1-e^{-\lambda T}\right)\left(1-e^{-\lambda W}\right)}{1-e^{-\lambda W}+e^{-\lambda T}}$	$\frac{(1-F(T))\left(2+\frac{1-F(W)}{F(W)}\right)R+\left(\frac{k}{\lambda}-\Phi(T)\right)}{\frac{k}{\lambda}(1+\frac{1-F(T)}{F(W)})}$	$\begin{cases} 1, & \text{if } a < W \le T \\ \frac{R}{a}, & \text{if } W < a \end{cases}$	$\begin{cases} \frac{(\alpha-1)\left(\frac{t_m}{T}\right)^{\alpha}(2-\left(\frac{t_m}{W}\right)^{\alpha})R+(1-\left(\frac{t_m}{W}\right)^{\alpha})(t_m\alpha-T\left(\frac{t_m}{T}\right)^{\alpha})}{\alpha t_m(1-\left(\frac{t_m}{W}\right)^{\alpha}+\left(\frac{t_m}{T}\right)^{\alpha})}, & \text{if } t_m \le W \\ \frac{R(\alpha-1)}{\alpha t_m}, & \text{if } W < t_m \end{cases}$

Example distribution: Exponential

Policy	Exponential
Offline	$1-e^{-\lambda R}$
Baseline	$\min[\lambda R, 1]$
Always 1^{st}	$1-e^{-\lambda T}+\lambda R e^{-\lambda T}$
Always 2^{nd}	$\frac{1-e^{-\lambda T}+2\lambda R e^{-\lambda T}}{1+e^{-\lambda T}}$
Single M^{th}	$\lambda e^{-\lambda T}\sum_{i=0}^{M-1}(1-e^{-\lambda T})^i R+\left(1-e^{-\lambda T}\right)^M$
Dual 2^{nd}	$\frac{\lambda R e^{-\lambda T}\left(2-e^{-\lambda W}\right)+\left(1-e^{-\lambda T}\right)\left(1-e^{-\lambda W}\right)}{1-e^{-\lambda W}+e^{-\lambda T}}$

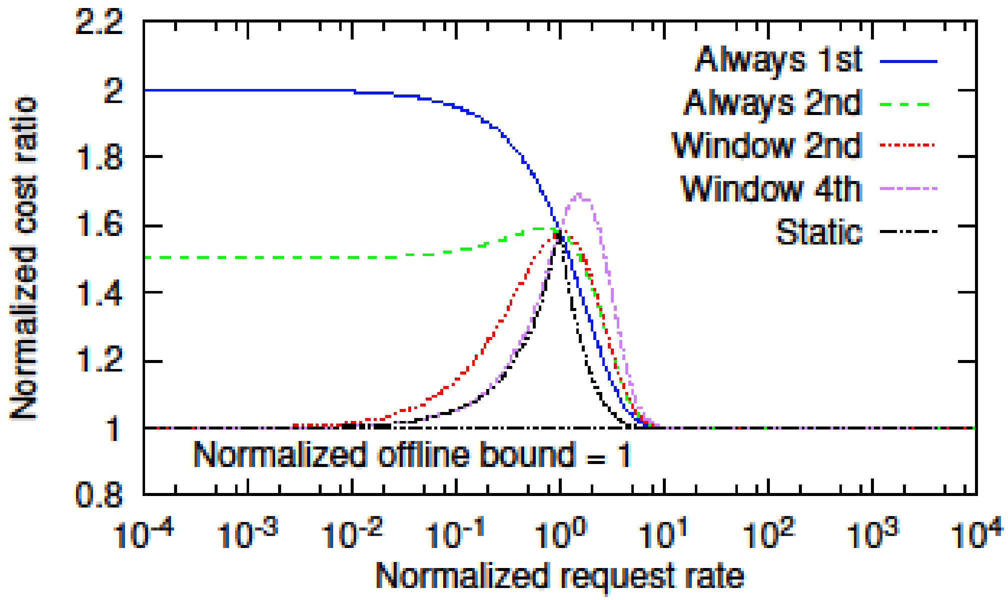

- Results with $W = T = R$

Example distribution: Exponential

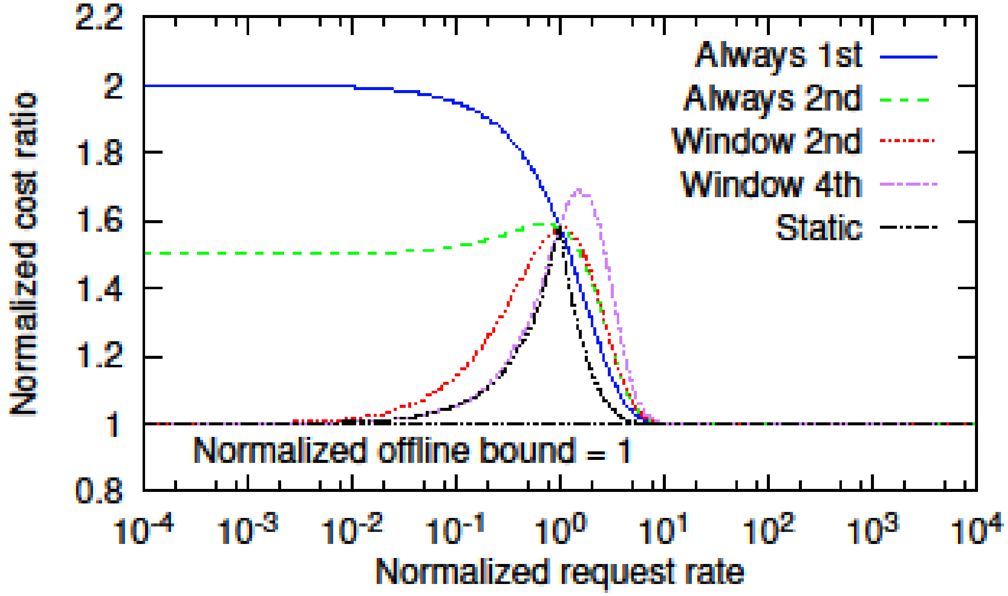

- Results with $W = T = R$

Example distribution: Exponential

Always on M^{th} asymptotes at $M/(M+1)$

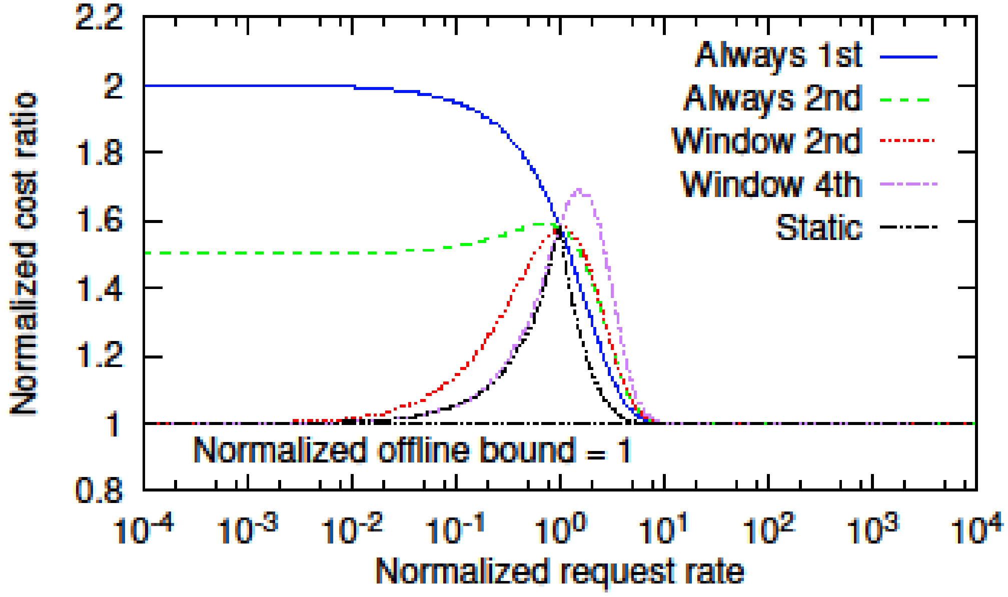

- Results with $W = T = R$

Example distribution: Exponential

Always on M^{th} asymptotes at $M/(M+1)$

Window on 2^{nd} peaks at (1.052,1.588)

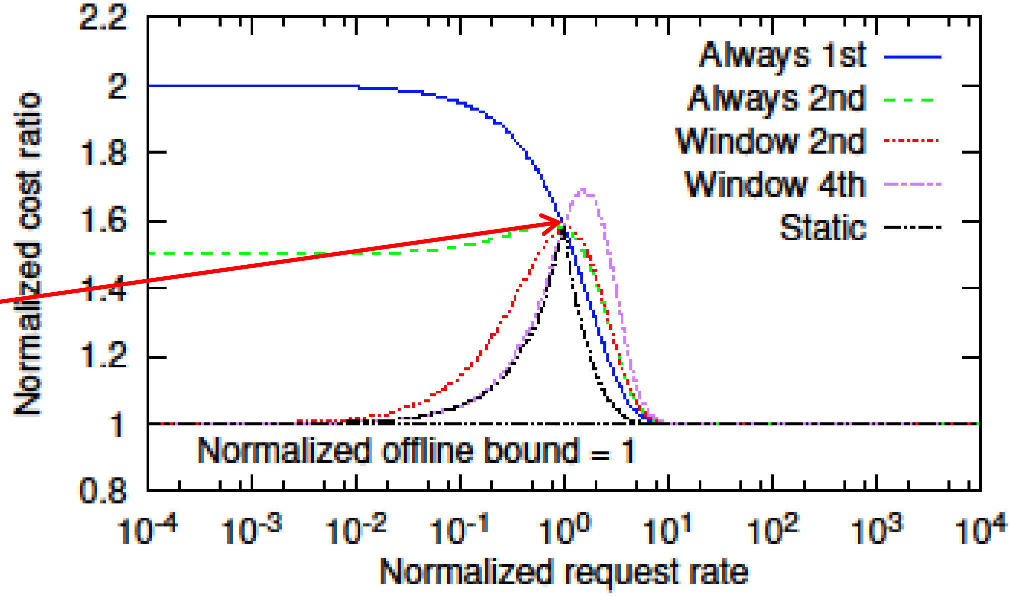

- Results with $W = T = R$

Example distribution: Exponential

Always on M^{th} asymptotes at $M/(M+1)$

Window on 2^{nd} peaks at (1.052,1.588)

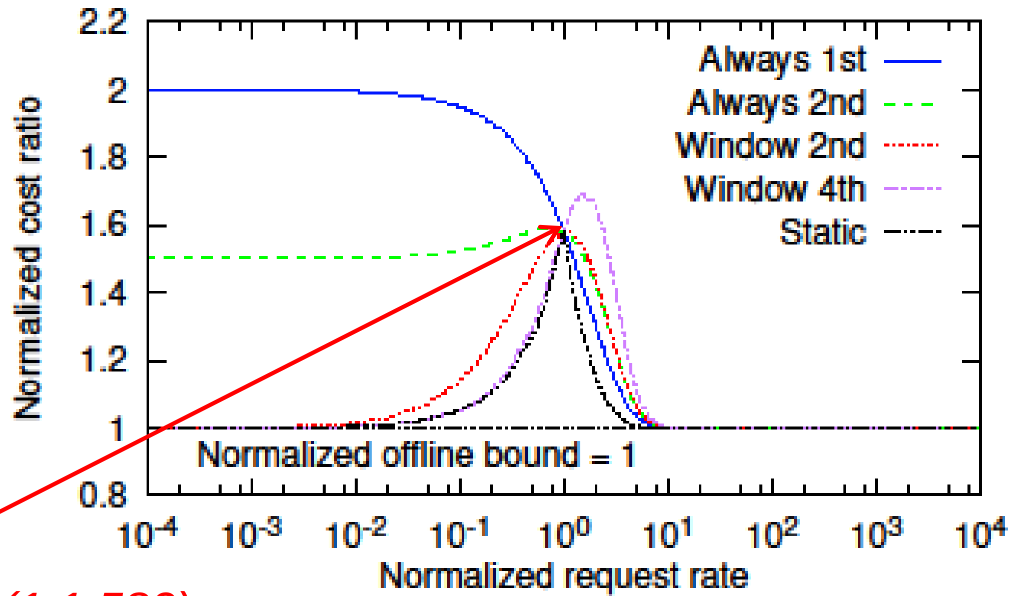

Static peaks at (1,1.582)

- Results with $W = T = R$

Example distribution: Exponential

Always on M^{th} asymptotes at $M/(M+1)$

Window on 2^{nd} peaks at (1.052,1.588)

Static peaks at (1,1.582)

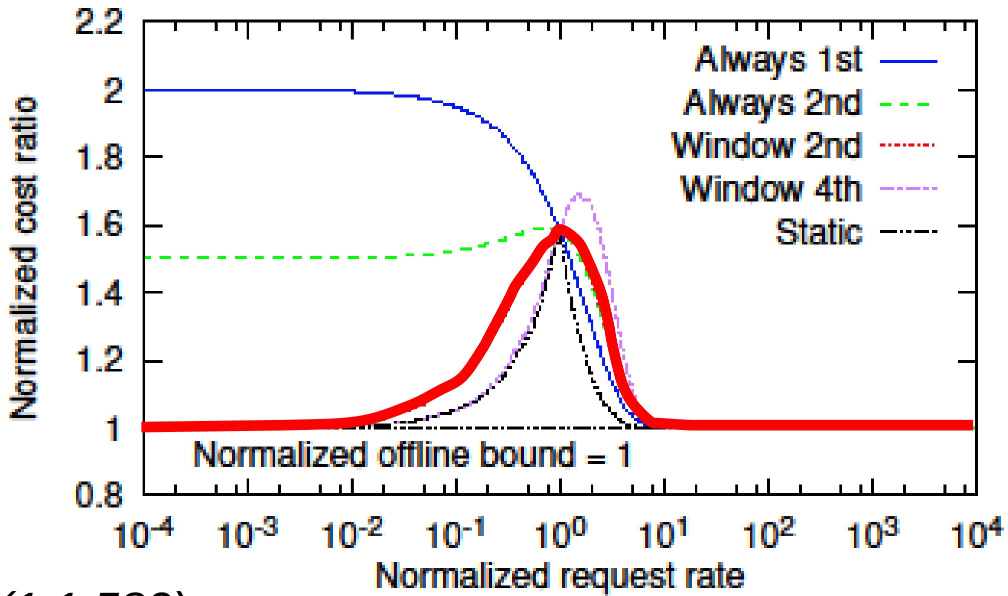

- Results with $W = T = R$
- *Window on 2^{nd} performs good throughout*
- *Window on 4^{th} performs somewhat better for lower request rates, but at an increased peak cost (at somewhat higher rates)*

Example distribution: Exponential

*Always on M^th^*

asymptotes at M/(M+1)

*Window on 2^nd^*

peaks at (1.052,1.588)

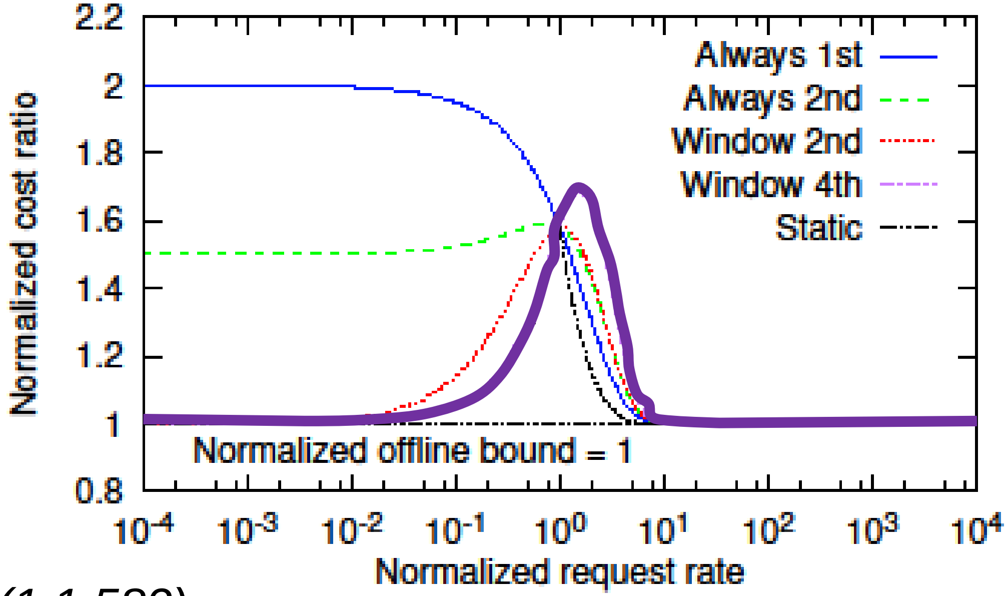

Static

peaks at (1,1.582)

- Results with $W = T = R$
- *Window on 2^{nd} performs good throughout*
- *Window on 4^{th} performs somewhat better for lower request rates, but at an increased peak cost (at somewhat higher rates)*

Example distribution: Exponential

Always on M^th asymptotes at M/(M+1)

Window on 2^nd peaks at (1.052,1.588)

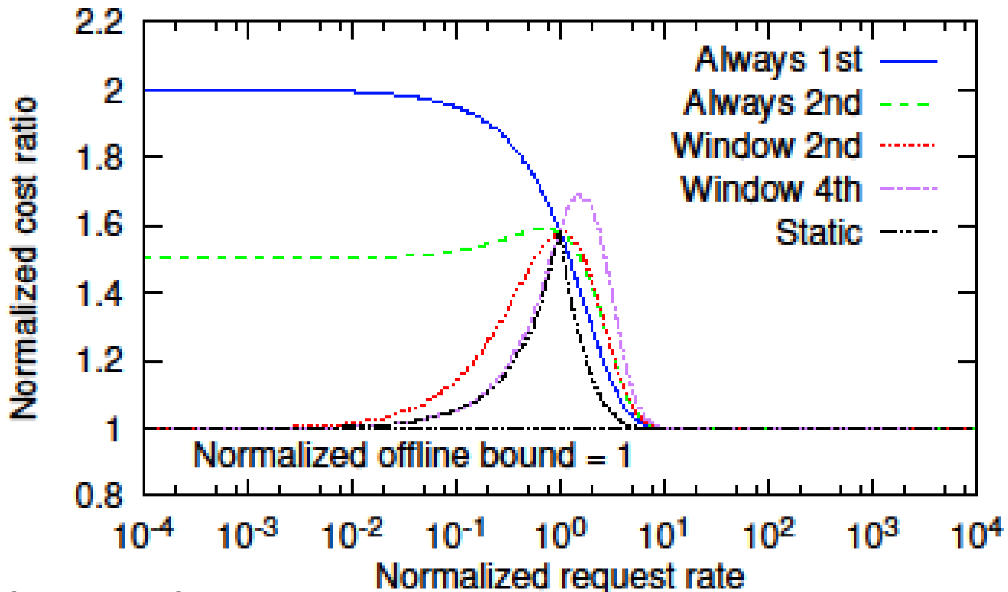

Static peaks at (1,1.582)

- Results with $W = T = R$
- *Window on 2^nd performs good throughout*
- *Window on 4^th performs somewhat better for lower request rates, but at an increased peak cost (at somewhat higher rates)*

Example distributions: Low variability distributions

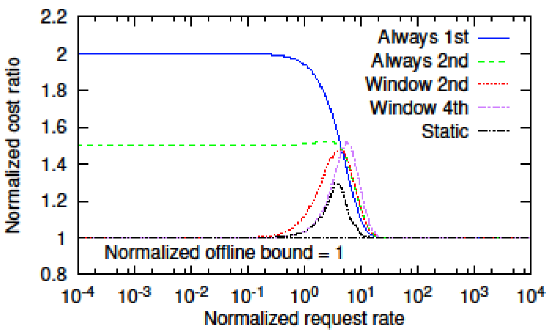

Erlang k=2

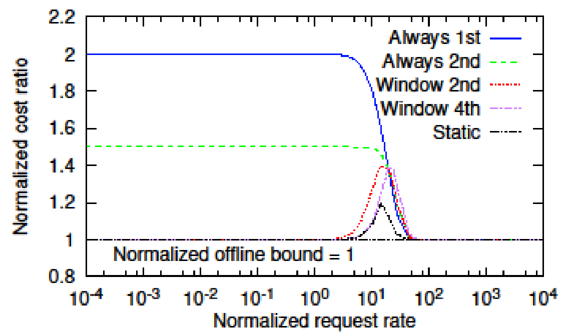

Erlang k=4

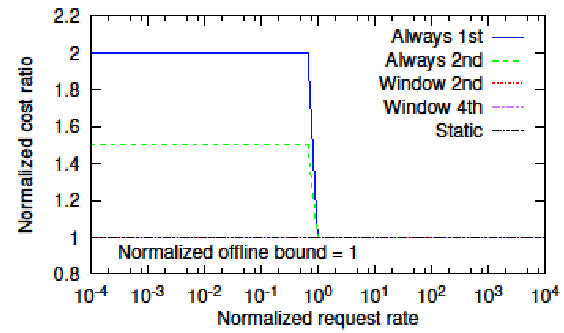

Deterministic

Increasingly deterministic inter-request times

- Peak cost ratio for single-window on *2nd* reduces as *k* increases and inter-request times become increasingly deterministic (rightmost fig)

Example distribution: Pareto

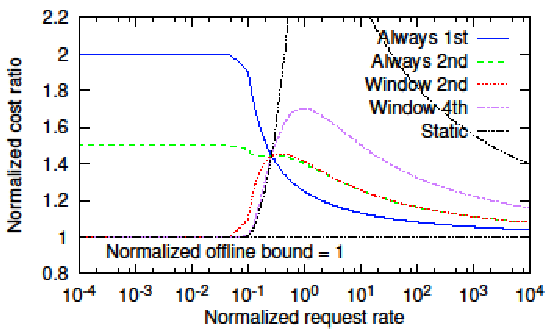

α =*1.1*

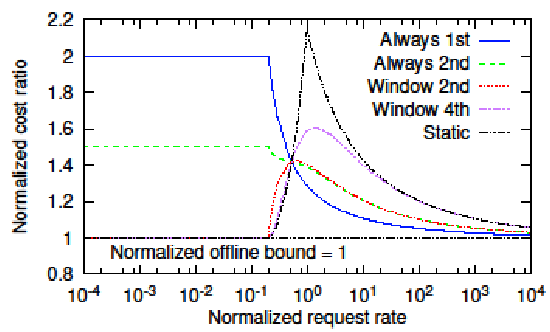

α =*1.25*

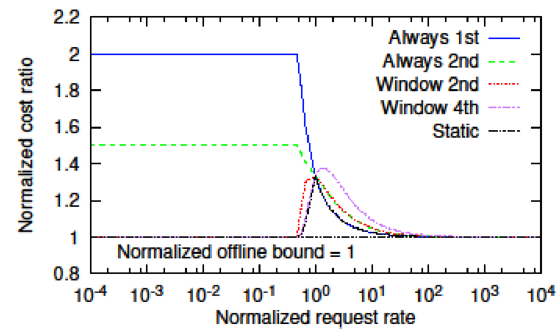

α = 2

- As per Theorem 6.6, *static baseline* performs very poorly when $\alpha \rightarrow 1$ (and t_m is small). E.g., large peak cost ratio in left-most fig
  - For larger α (e.g., α = 2), this peak reduces substantially.
- Otherwise, the results are similar as for the other inter-request distributions, suggesting that *single-window on 2nd* with $T = R$ is a good choice

Multi-file evaluation

Example distributions

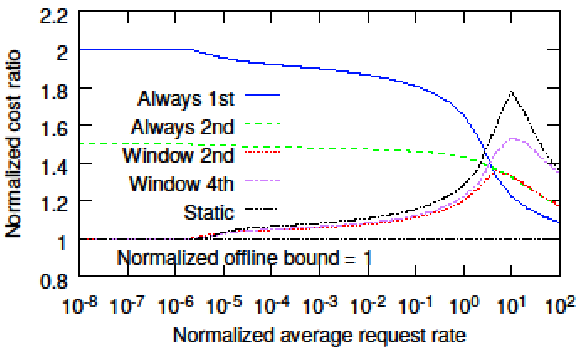

Pareto, α =1.25

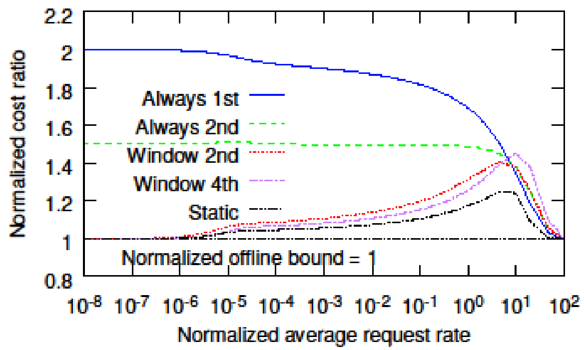

Exponential

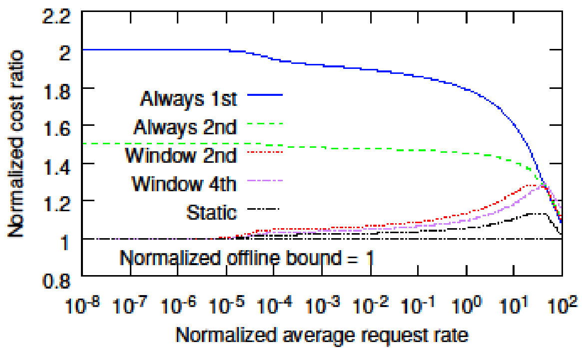

Erlang, k=4

- Setup: 1,000,000 objects with Zipf popularity
  - Here, γ=1 (but results with γ=0.75 and γ=1.25 similar)
  - *W = T = R*
- Significant benefits to being selective
  - *Window on 2nd* significantly outperforms *always on Mth*
- *Window on 2nd* good throughout
  - Close to static optimal when Exponential and Erlang
  - Outperform static when Pareto
  - Has a peak cost-ratio of 1.4

Trace-based simulations

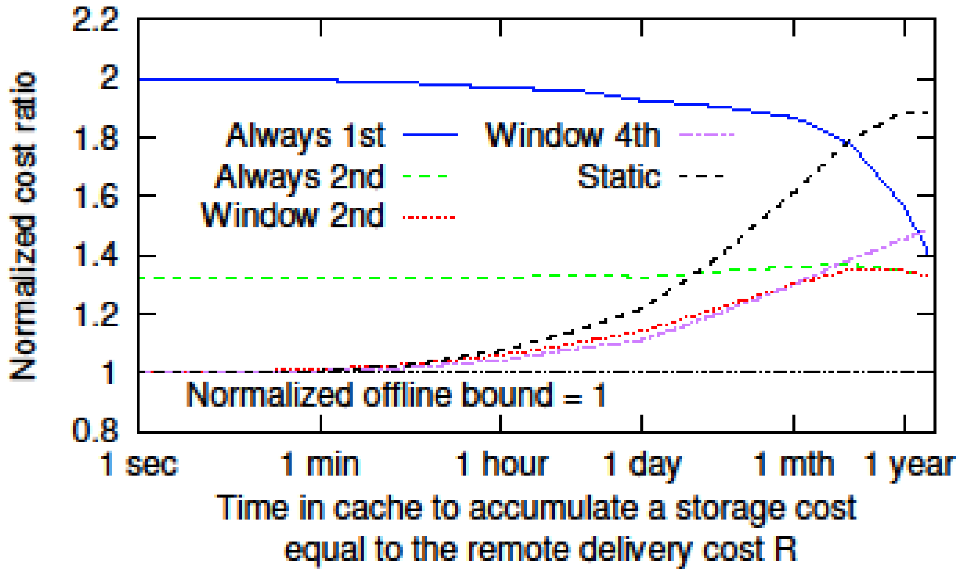

- Setup: 20-month long university trace with YouTube viewings
  - 5.5 M video request to 2.4 M unique videos
  - Long tail of less popular videos
  - $W = T = R$
- "Static" (highly optimistically) assumes "oracle" knowledge of which choice is better (*always local* or *always remote*) for each individual video ...
- Yet, *window on 2nd* outperform static
  - Highlights value of policy when request rates are unknown and variable

Break-down of cost contributions

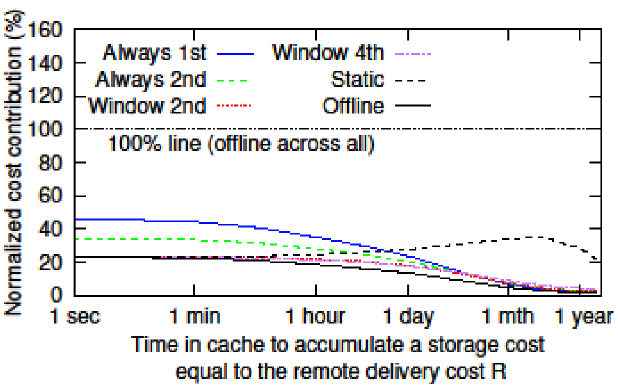

Top (more than 20)

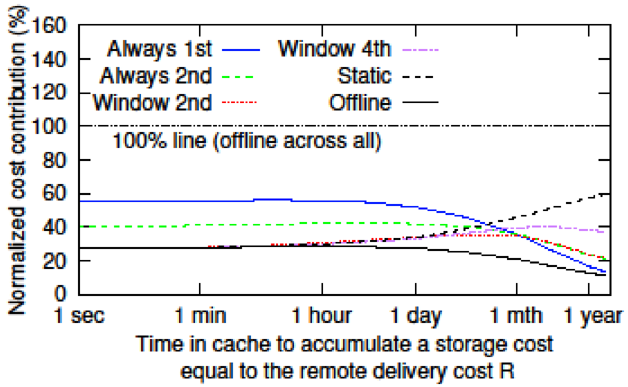

Middle (4-20 views)

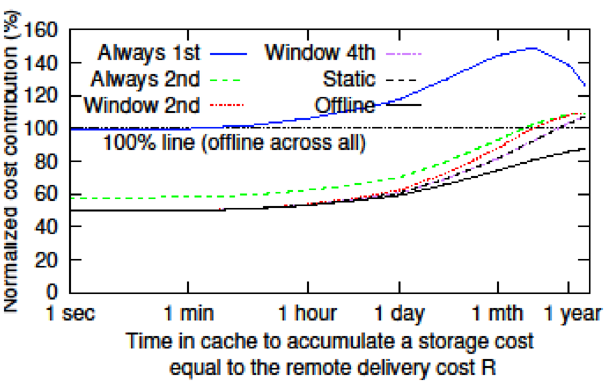

Tail (1-3 views)

- Tail contribute to most of the costs …

  … highlighting importance of selective insertions.

Conclusions

Conclusions

Worst-case bounds for the optimal cost and competitive cost ratios

- E.g., Best worst-case bounds of M+1 are achieved by selecting $W = T = R$

Average cost expressions and bounds

- Arbitrary inter-request distributions
- Example inter-request distributions (both short-tailed and heavy-tailed)
- *Static* is *online optimal* for constant and decreasing hazard rates, but can be arbitrarily bad when heavy tailed or request rates are not known

Numeric and trace-based evaluations reveal insights into the relative cost performance of the policies

- Substantial cost benefits of using window-based with intermediate M (e.g., 2-4) and the optimal worst-case parameter setting (i.e., $W = T = R$)

Window-based cache on 2nd request policy using a single threshold optimized to minimize worst-case costs provides good average performance

- Attractive choice for a wide range of practical conditions where request rates of individual file objects typically are not known and can change quickly ...

Thanks for listening!

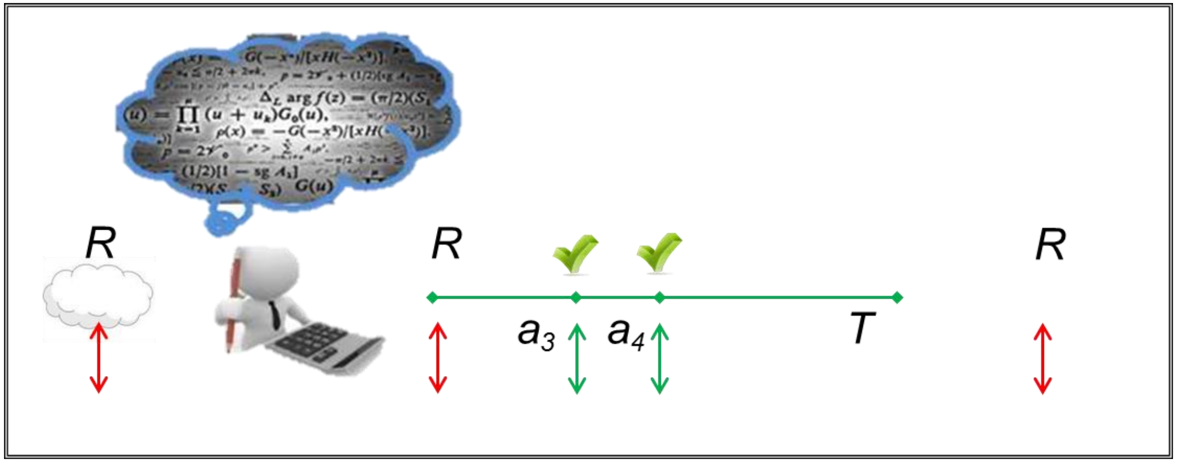

Worst-case Bounds and Optimized Cache on *M*th Request Cache Insertion Policies under Elastic Conditions

Niklas Carlsson and Derek Eager

LINKÖPING UNIVERSITY

Niklas Carlsson (niklas.carlsson@liu.se)