Cost-Aware Live Migration of Services in the Cloud

David Breitgand -- IBM Haifa Research Lab
Gilad Kutiel, Danny Raz -- Technion, Israel Institute of Technology

The research leading to these results has received funding from the European Union's Seventh Framework Programme under grant agreement n° 257448, 215605

Agenda

- Introduction
- The Cost of Live Migration
- Fixed Bandwidth Migration
- Variable Bandwidth Migration
  - Related Work
  - The CALM (Cost Aware Live Migration) Algorithm
  - Evaluation Study
- Conclusions

Introduction

- We consider pre-copy live migration model (but results hold for post-copy approach as well)
- We consider in-band migration
- We focus on network bandwidth as primary bottleneck (but the presented framework is general)
- We provide analytical study of our approach
- We validate our proposal using trace-driven simulations

The Cost of Live Migration 1/2

- Clearly no service is available during **downtime**
- If migration is done in-band then some of the bandwidth used to serve clients is used now for the migration
- We define the cost to be the probability to violate the SLA at a given time
- It is a function of the available bandwidth for the service and we denote it by $F(B_s)$

The Cost of Live Migration 2/2

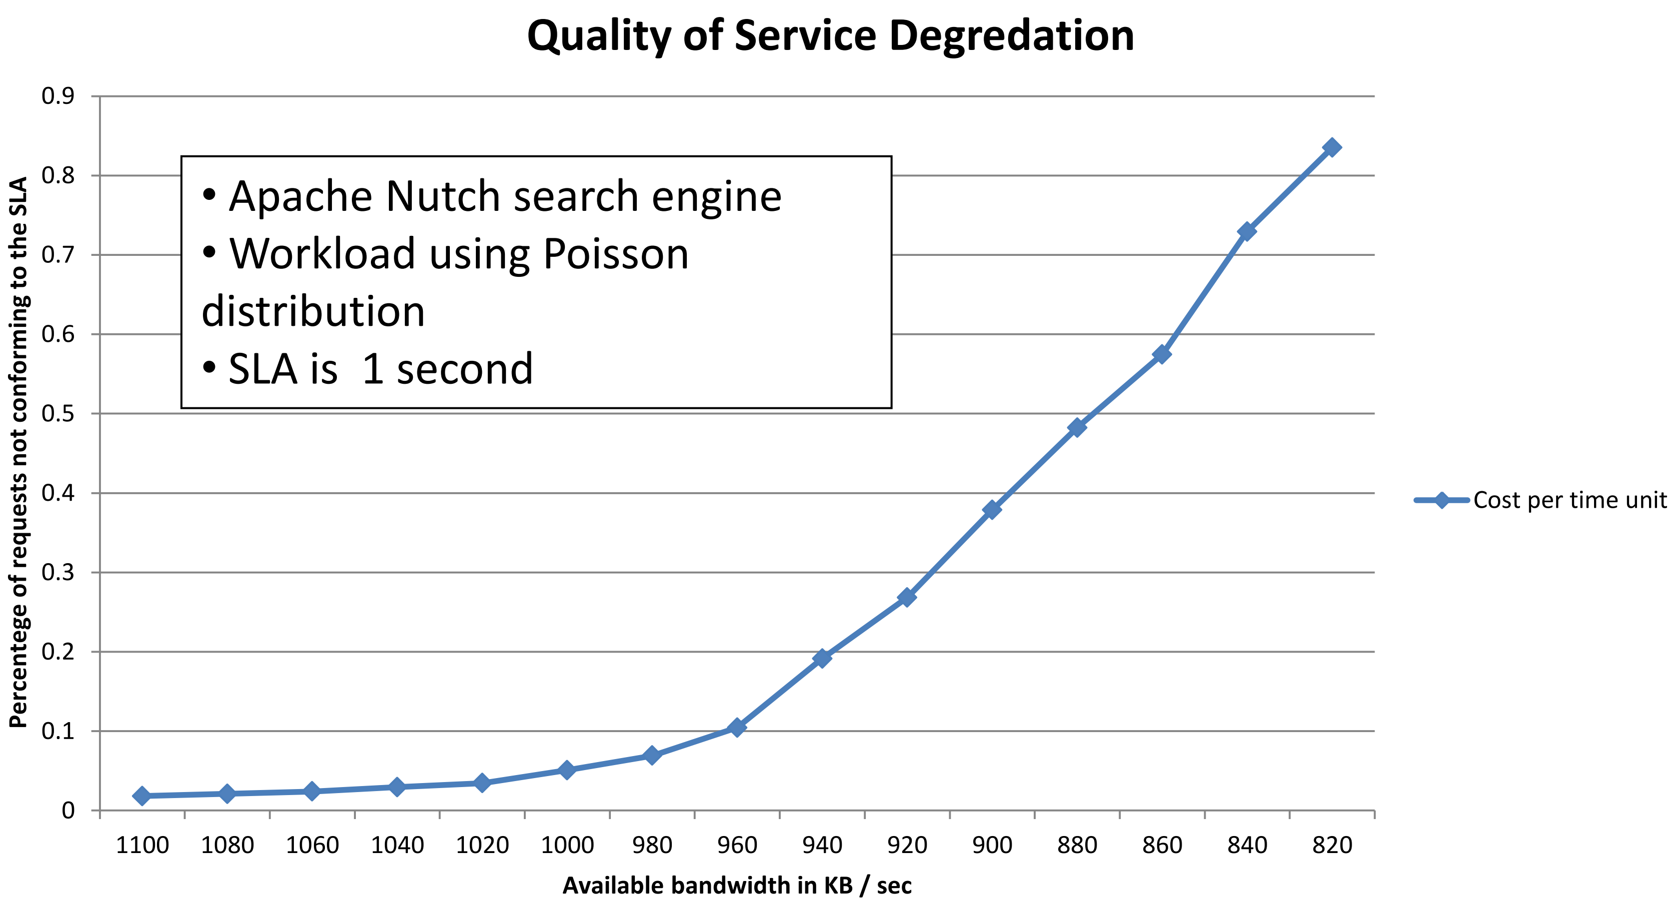

Fixed Bandwidth Migration 1/5

- We start with a simple case
- The bandwidth for the migrations is predefined
  - and fixed through the migration process
- Recall that memory is updated during the migration process, how much bandwidth should we use?
  - More, faster but more degradation
  - Less, better service while migrating but we might need to transfer pages again and again
- The optimal bandwidth depends on the cost function and other factors

Fixed Bandwidth Migration 2/5

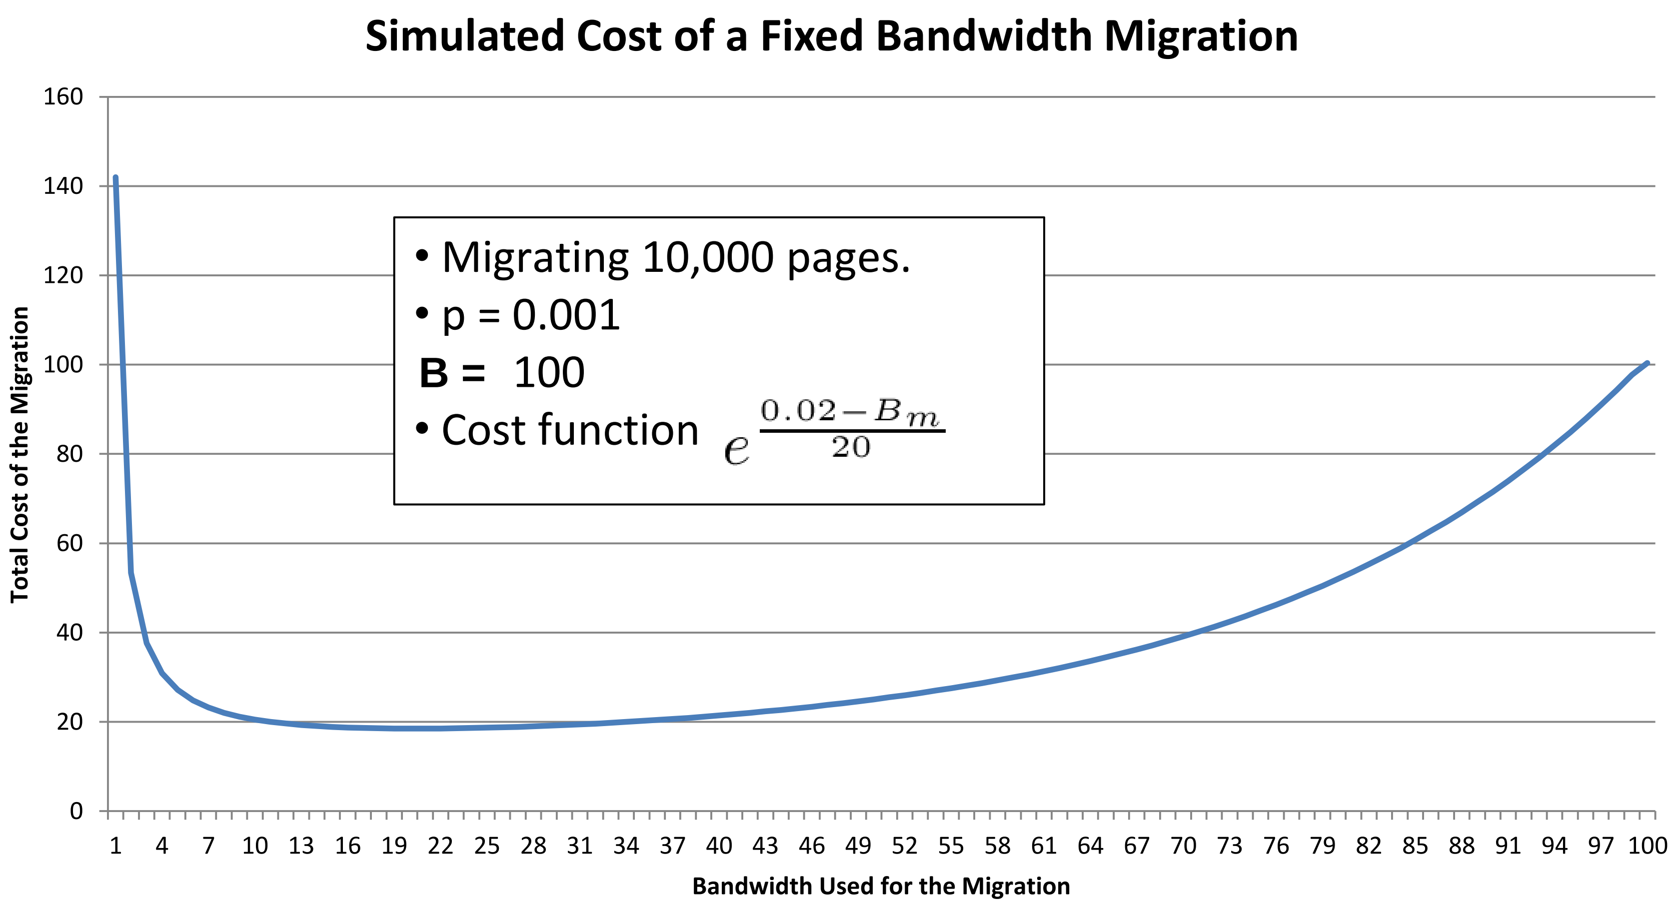

Fixed Bandwidth Migration 3/5

- Formulation:
  - Virtual machine with *M* pages
  - Total available bandwidth (service + migration) is *B*
  - B_m is the bandwidth used for the migration
  - B_S is the bandwidth available for the service
  - *p* is the probability for a page to be updated during a single time unit, we assume that it is uniform and independent ($q = 1 - p$)
  - A *clean* page is one that was copied and hasn't been updated since then
  - *N* pages are transferred during the **pre-copy** phase (and the rest during the copy phase)

Fixed Bandwidth Migration 4/5

- The expected cost of the migration process is given by:

$$E(cost) = E(cost_{pc}) + cost_{copy}$$

- Where:

$$E(cost_{pc}) = E(T_{pc}(B_S)) \cdot F(B_S)$$

- And:

$$E(T_{pc}(B_s))) = \frac{1}{ln(q)} \cdot ln(1 - N \cdot (1 - q^{\frac{1}{B - B_s}}))$$

- Optimal bandwidth can be found by minimizing the cost function (analytically / numerically).

Fixed Bandwidth Migration 5/5

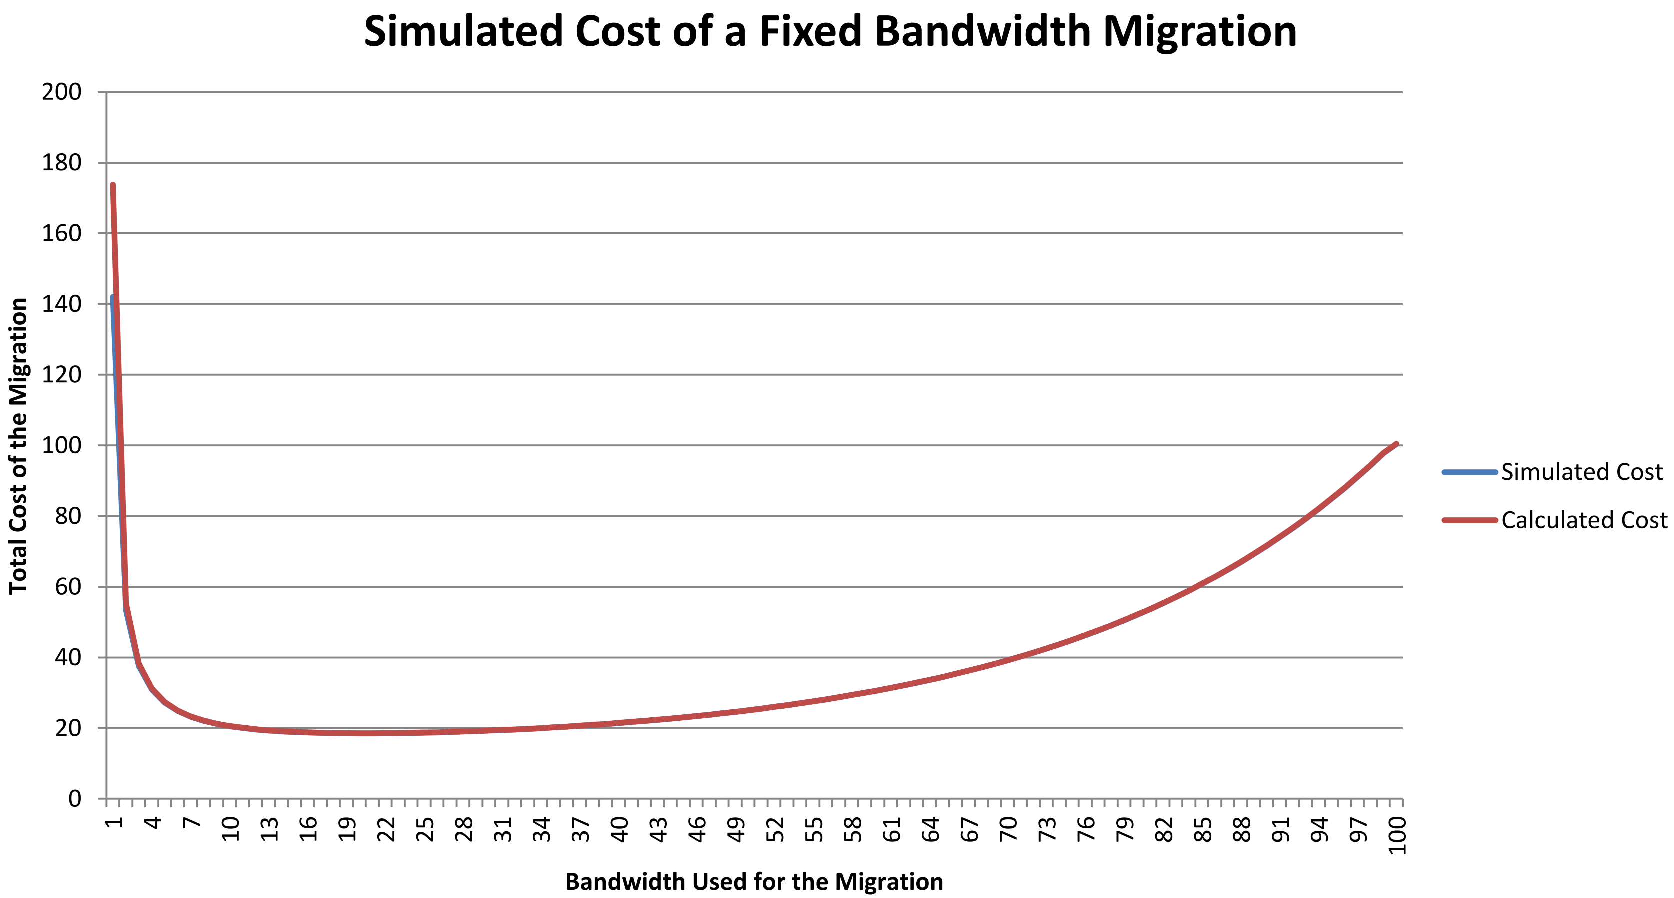

Trace Driven Simulation

- We found an optimal migration when the bandwidth is fixed and the dirtying probability is uniform
- What is the dirtying probability for a “real-world” application?
- We generated traces of dirtying patterns for several services and used those traces to simulate migration of a “real-world” services

Dirtying Probability of Real Application

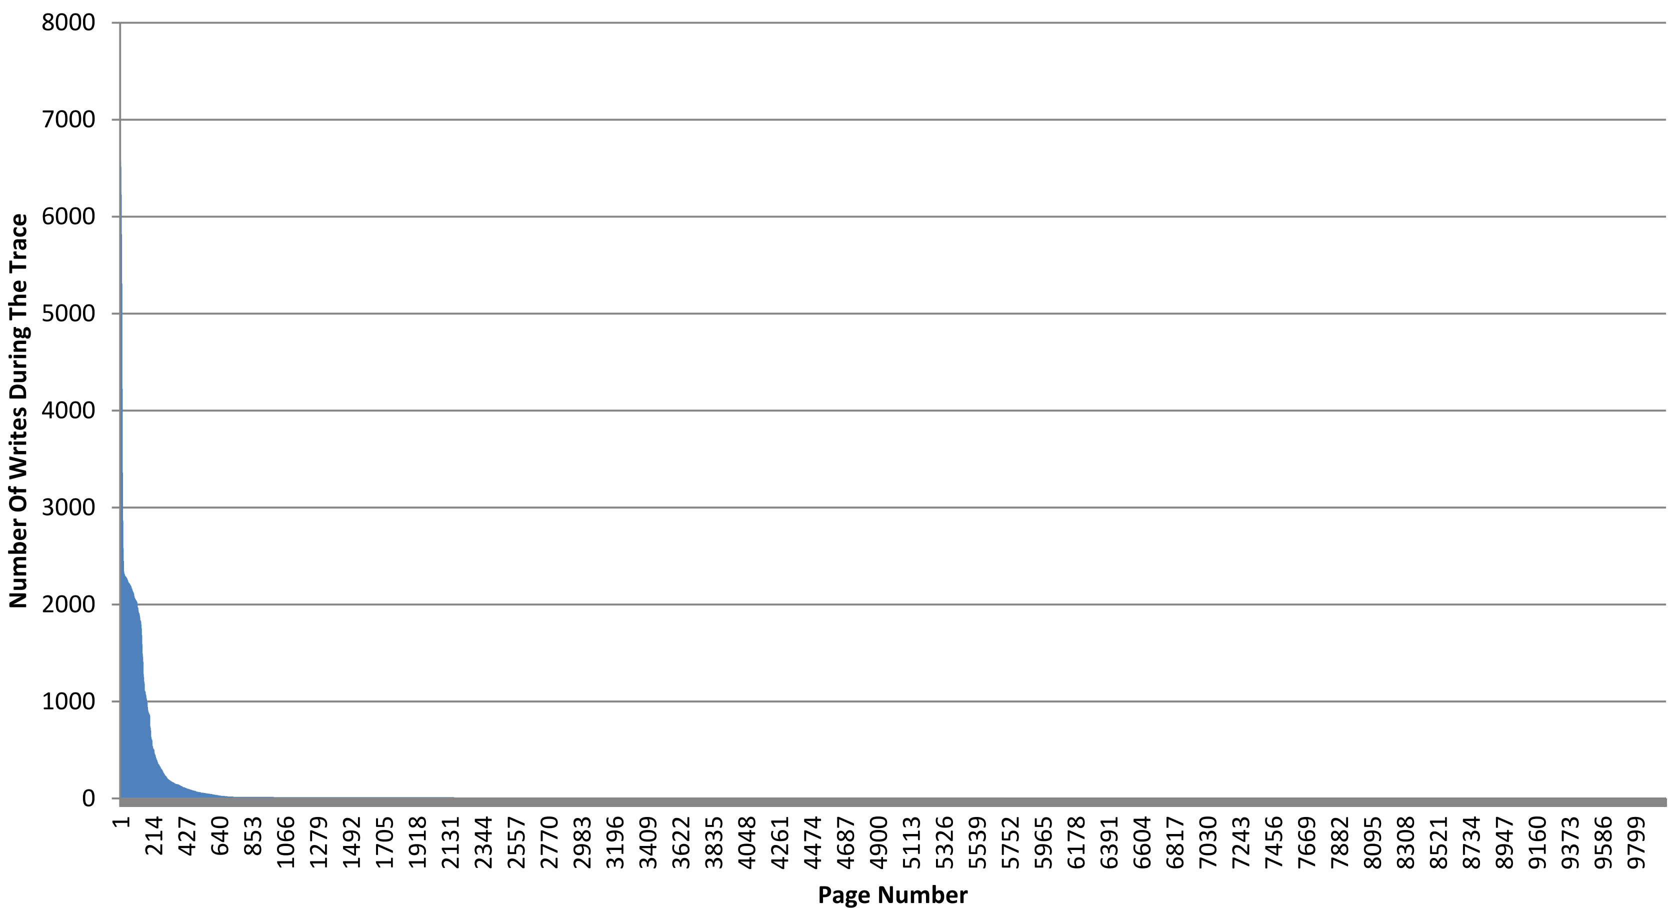

Dirtying Probability (top 100)

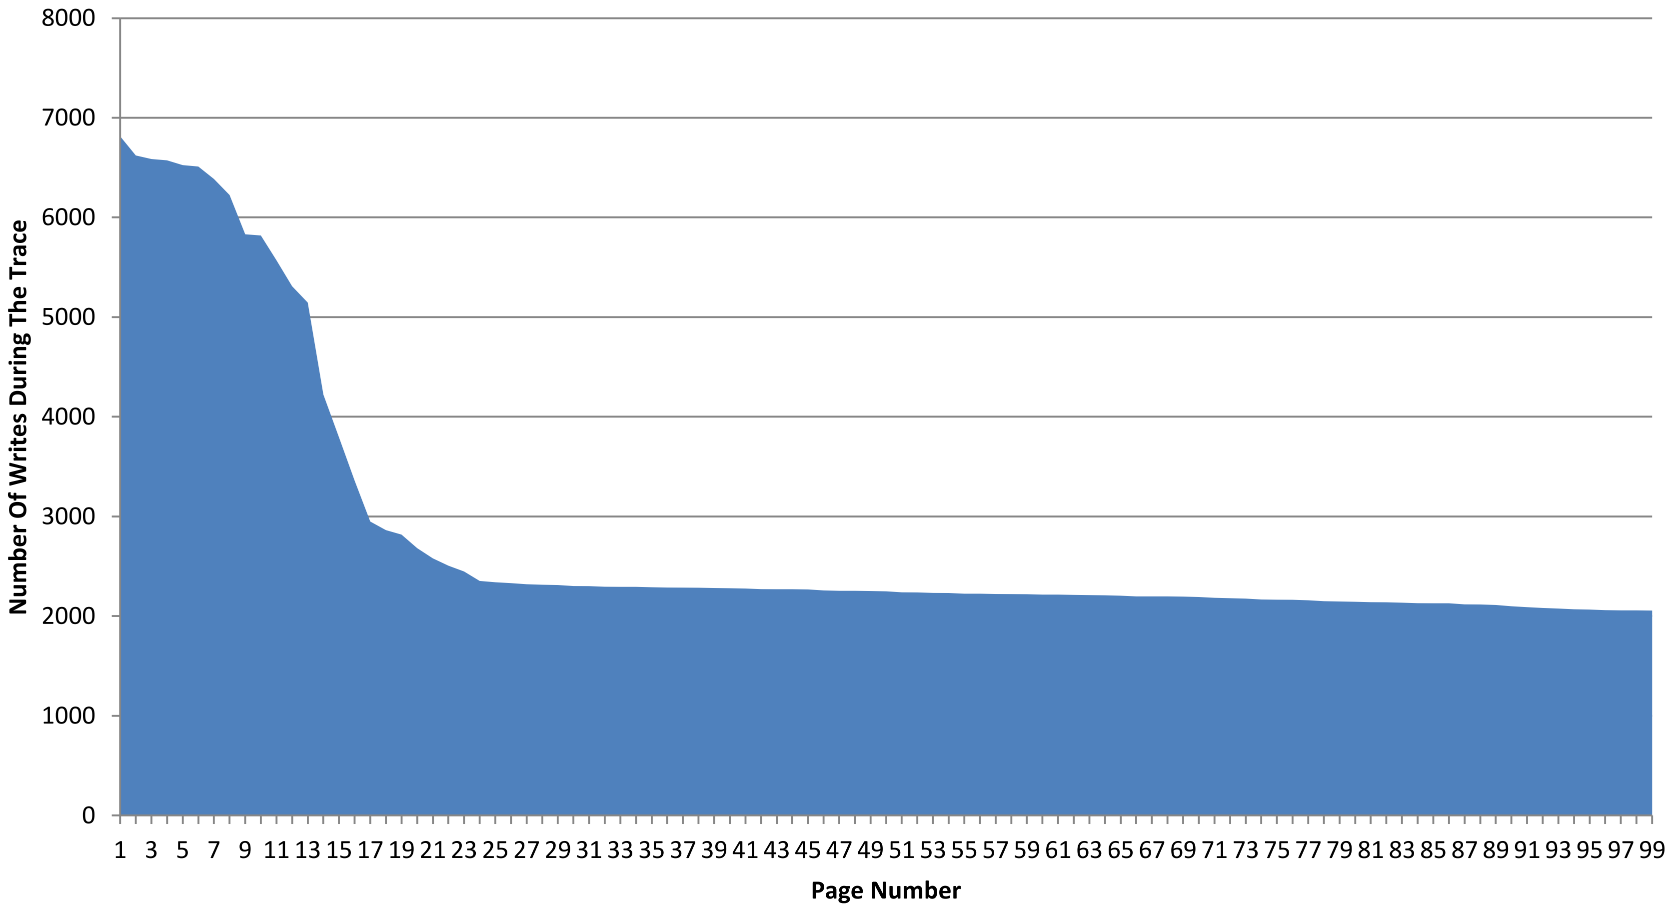

Dirtying Probability (top 500)

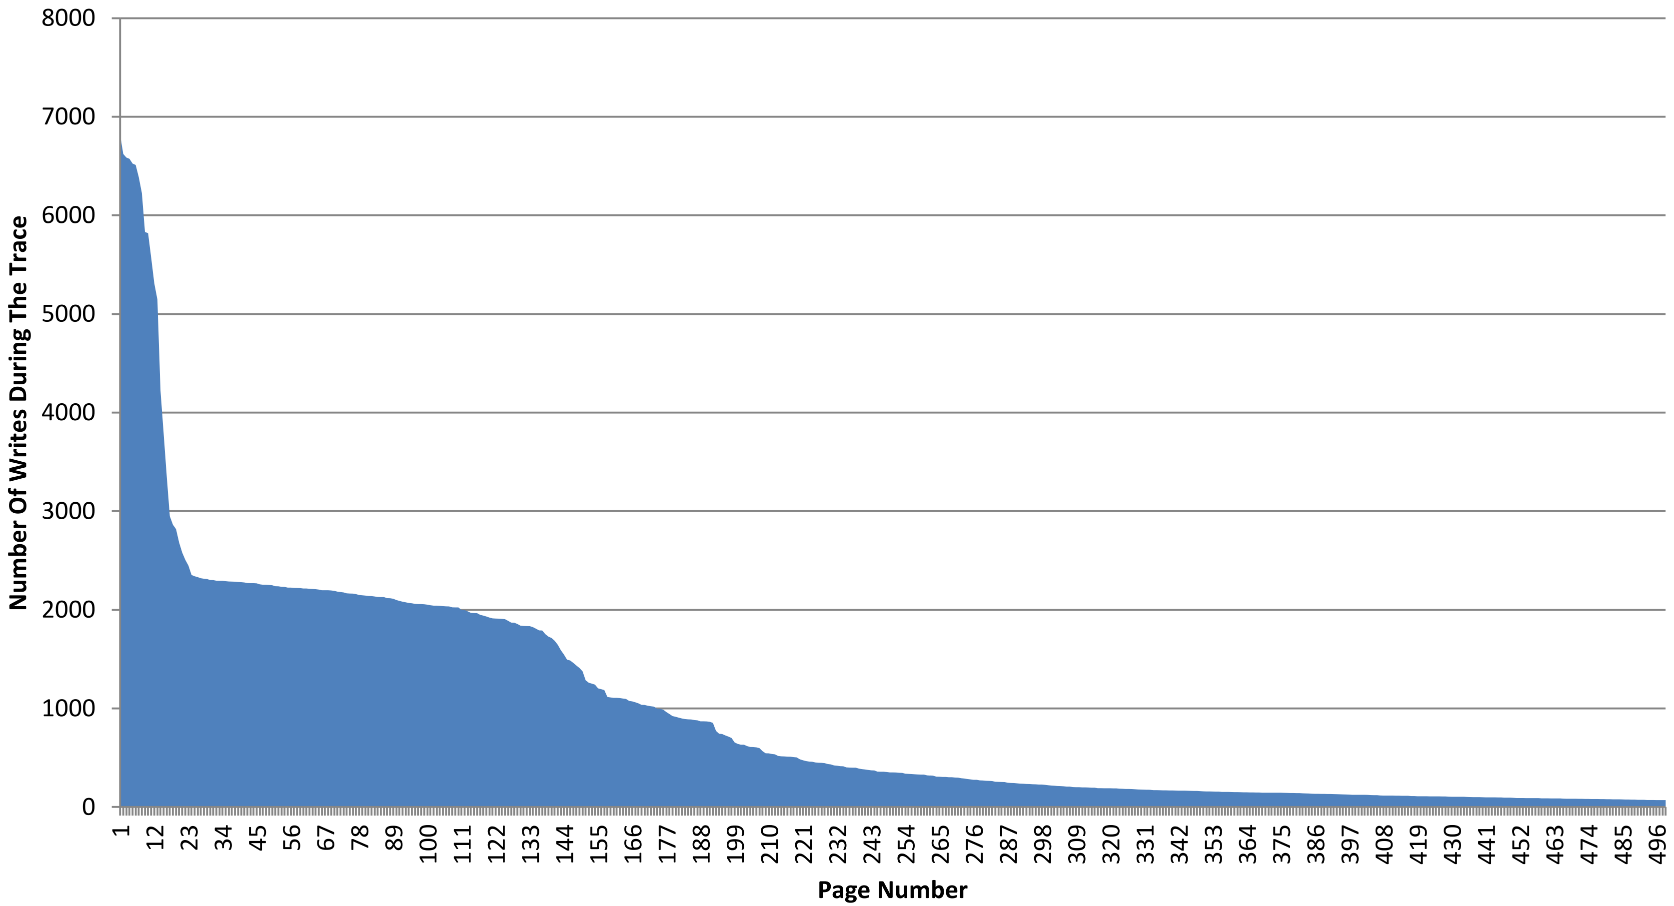

Dirtying Probability (Over Time)

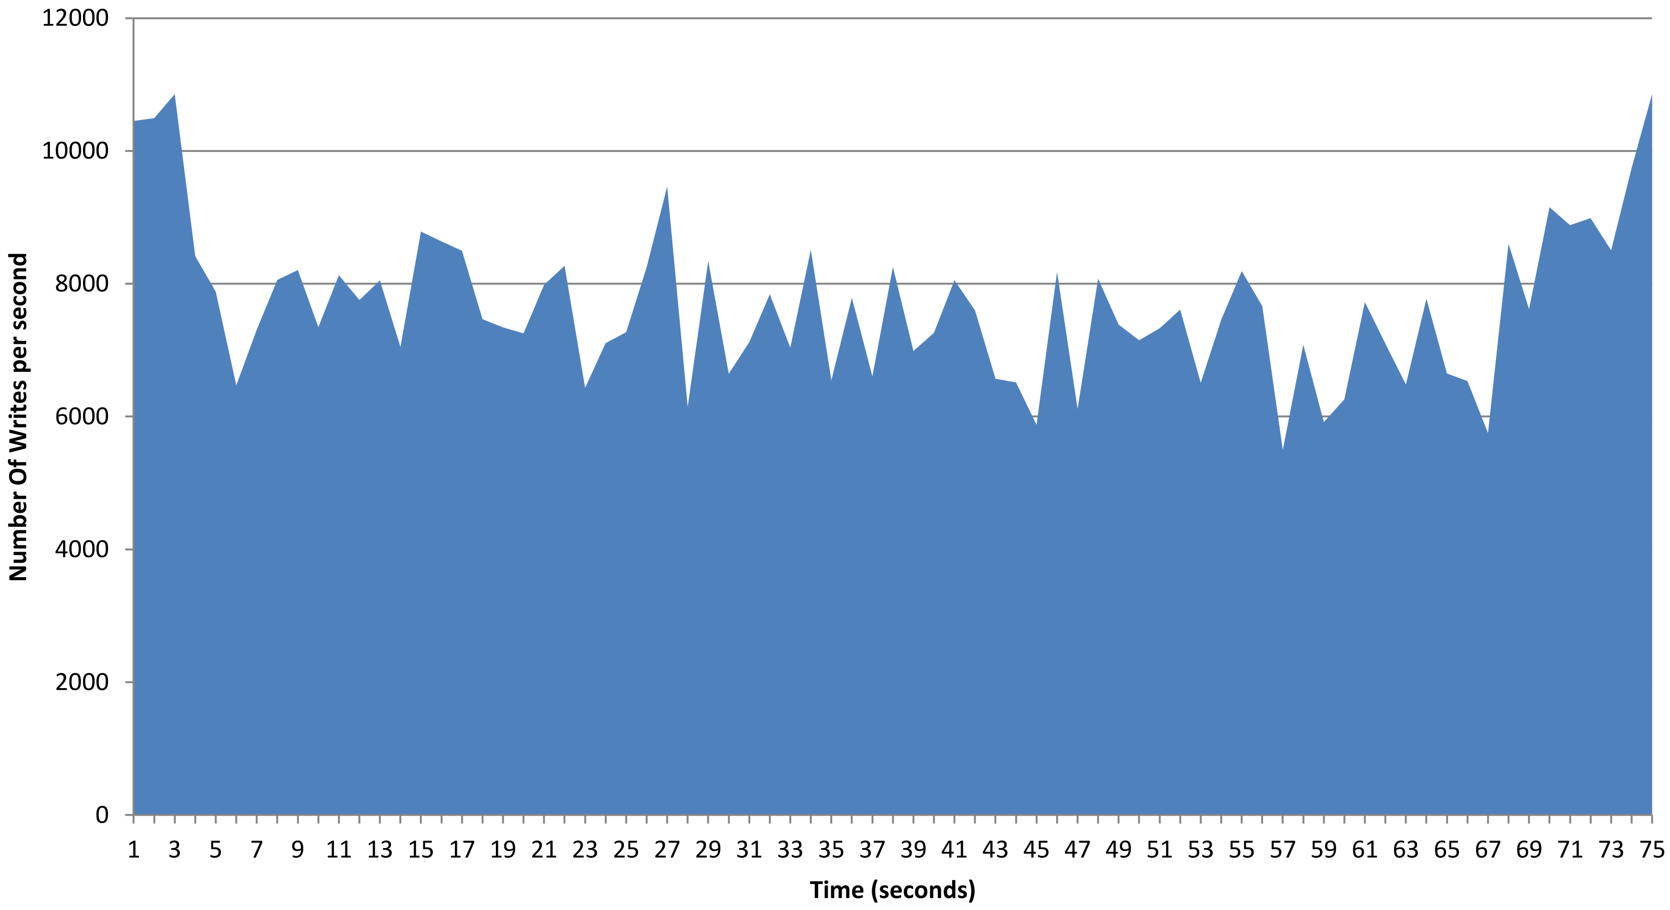

Migration's Cost With Real Dirtying Trace

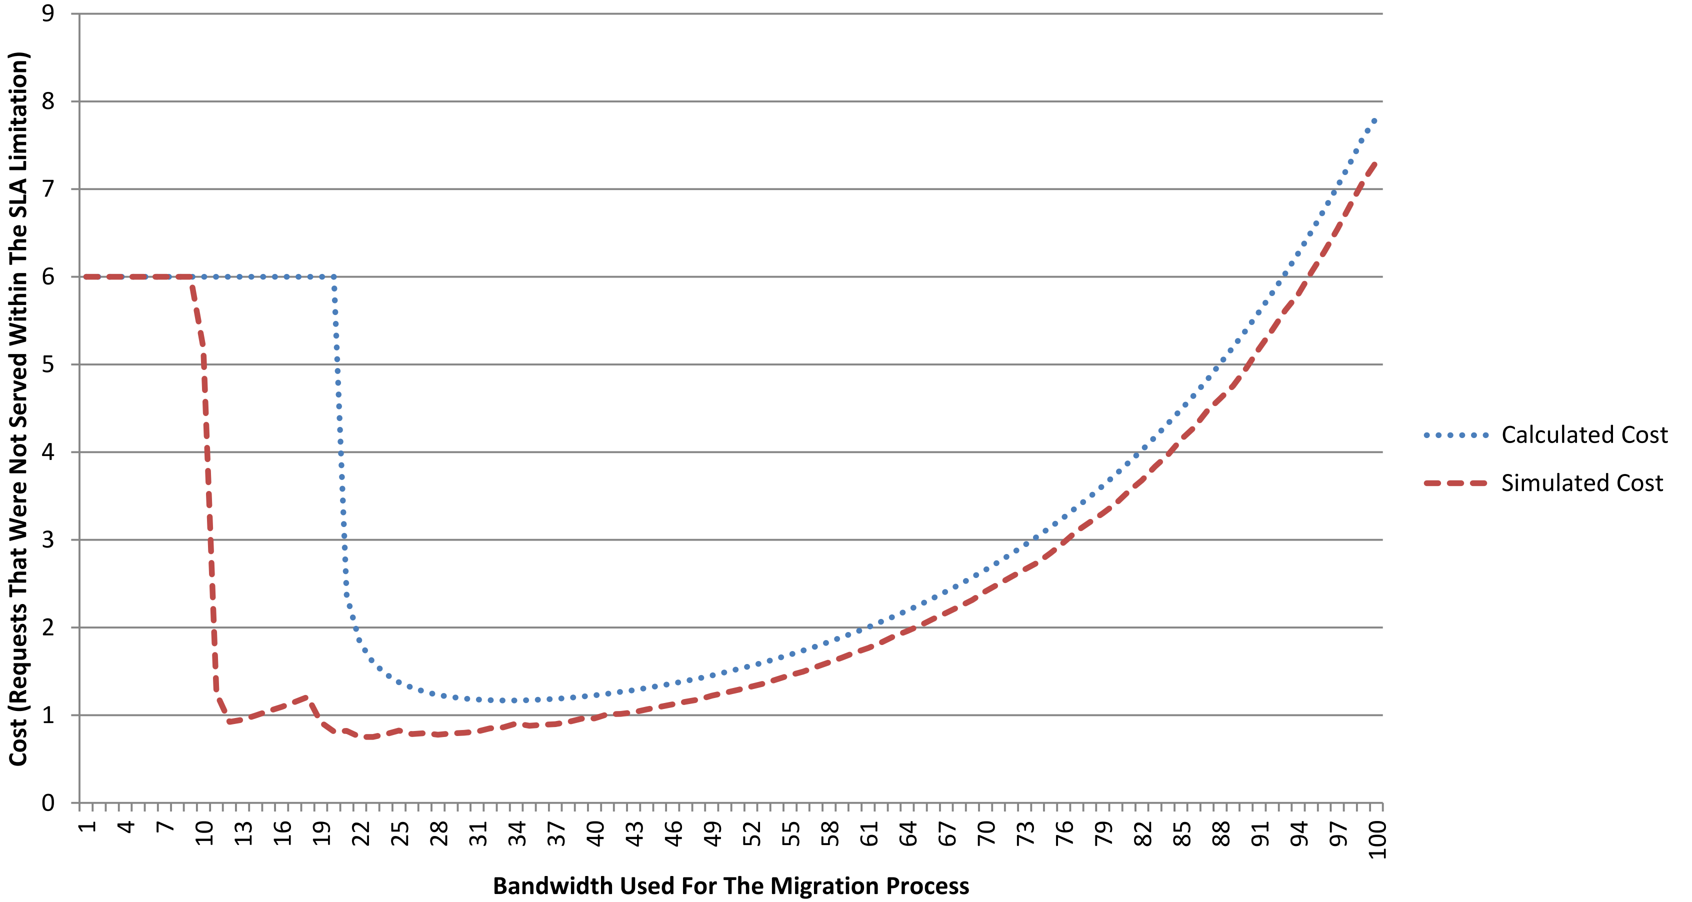

Variable Bandwidth Migration

- In reality we are not limited to a predefined, fixed, bandwidth.
- Bandwidth can be dynamically adjusted during the migration process.
- Intuitively, we should use a low bandwidth at the beginning of the process and increase it as we proceed, why ?

Related Work 1/2

- Clark et al. suggested the following algorithm:
  - The first pre-copy round copies all the pages from the source host to the destination host using initial bandwidth defined by the system administrator.
  - Each subsequent round copies only the dirty pages using a bandwidth equal to

$$\frac{\#\ dirty\ pages}{time\ of\ the\ last\ round} + \delta$$

    where δ is a fixed addition defined by the administrator.
  - Continue until bandwidth calculation exceed a maximum limit defined by the administrator or when there are less than 256KB to transfer.

Related Work 2/2

- Xen's migration algorithm uses similar principles, the stop conditions are defined as follows:
  - Less than 50 pages were dirtied during the last pre-copy iteration.
  - 29 pre-copy iterations have been carried out.
  - More than 3 times the total amount of RAM allocated to the VM has been copied to the destination host.
- What are the problems with the above algorithms ?

The CALM (Cost Aware Live Migration) Algorithm 1/2

- The bandwidth for the migrations can change over time
  - the algorithm determines
    - the bandwidth to be used at each phase of the migration process.
    - the end of the pre-copy phase.
- works in steps
  - each step moving from i clean pages to i+1 clean pages.
  - decides whether to continue or move to copy phase.

The CALM (Cost Aware Live Migration) Algorithm 2/2

- During step bandwidth is fixed so we can use previous results, the cost of the *i*th step is given by: $C_i = F(C - B_i) \cdot T_i$

  where $T_i = \frac{1}{ln(q)} \cdot ln(\frac{i \cdot (1-q^{\frac{1}{B_i}})-1}{i-1 \cdot (1-q^{\frac{1}{B_i}})-1})$

- Find best B_i (with minimal cost)
- Move to copy phase when: $C_i > \frac{1}{C}$

Evaluation Study 1/2

- We compare the CALM algorithm against the one suggested by Clark et al by simulating a live migration of a real-world services.
- Our simulations show that the CALM algorithm outperform Clark's algorithm even when used to migrate a real-world services.

Evaluation Study 2/2

RAM	Bandwidth	Algorithm	Cost	Total Time (sec)	Down Time (sec)
1GB	1GBits / sec	Clark	0.04	105.2	0.0424
		CALM	0.01	85.78	0.0078
		Fixed	0.01	19.93	0.0035
512MB	512MBits / sec	Clark	0.27	67.99	0.0039
		CALM	0.01	79.03	0.0078
		Fixed	0.02	142.55	0.0069
512MB	256MBits / sec	Clark	48.71	67.22	0.6412
		CALM	12.05	56.58	11.940
		Fixed	12.42	31.90	12.232
256MB	256MBits / sec	Clark	36.86	40.38	2.1160
		CALM	4.05	56.37	0.3940
		Fixed	4.52	35.60	4.2325

Conclusions

- We presented a novel model that accounts for the total cost of migration: pre-copy & copy phases
- Optimal migration strategy depends on various factors (available bandwidth, memory size, type of the service etc…).
- Cost-Aware migration algorithm is beneficial.
- CALM algorithm performs well also on real-world applications.
- The fixed algorithm performs well in certain cases.
- Future work is needed in order to better adjust the CALM algorithm to a real-world page dirtying pattern.

Thank You.

Bandwidth Usage on Different Scenarios

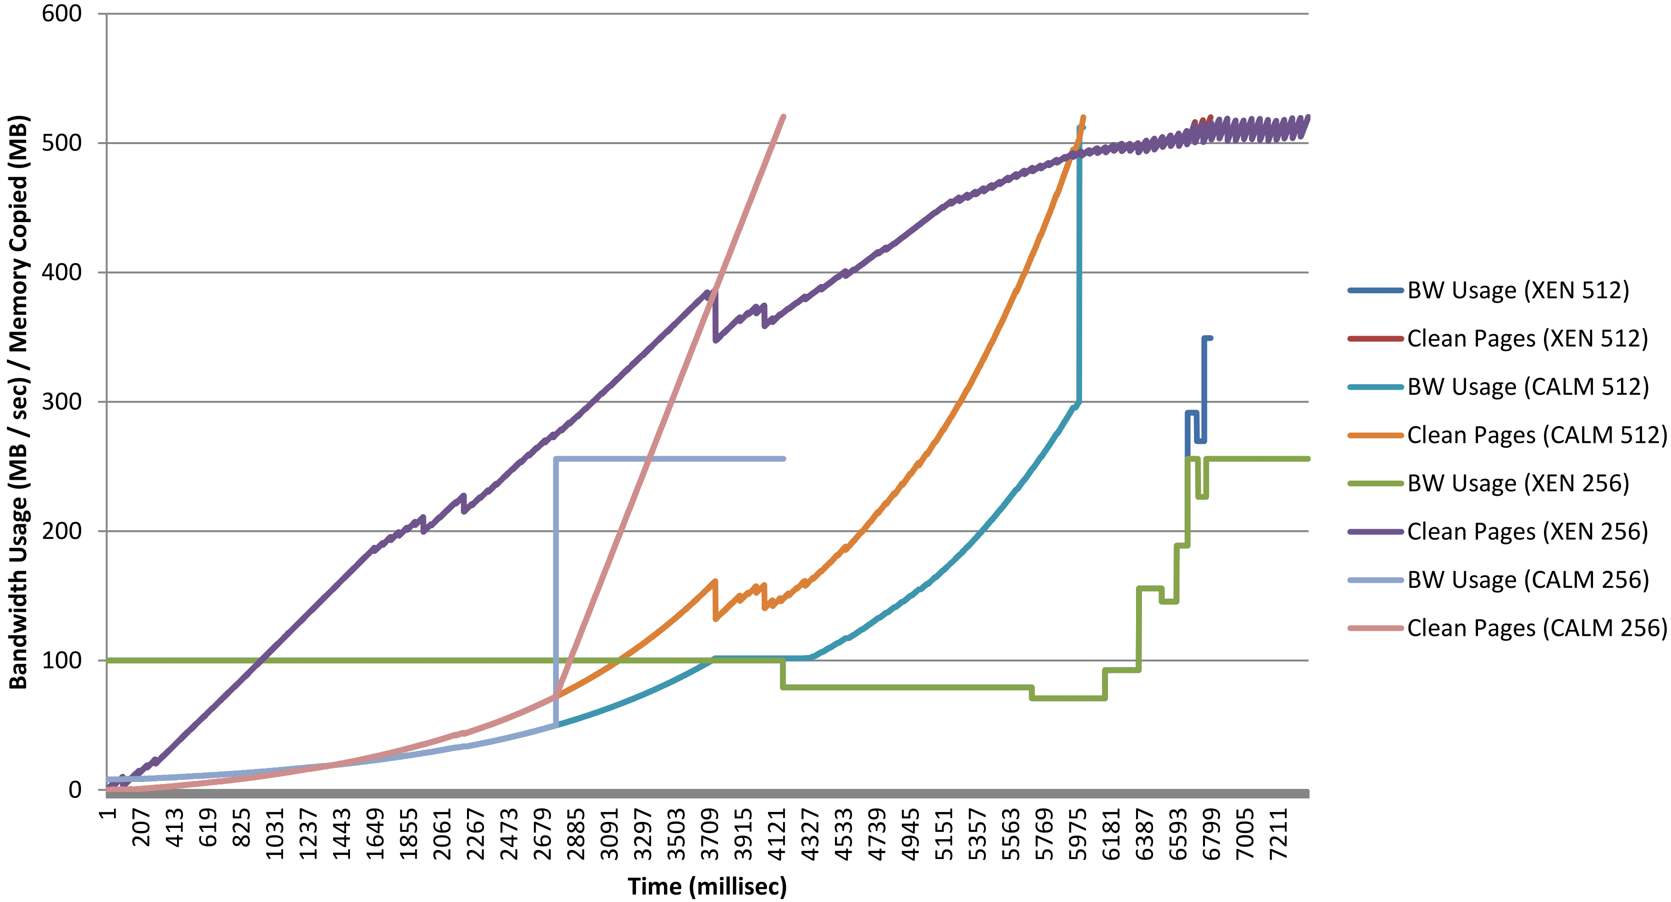

Fixed Bandwidth Migration 4/7

- We would like to calculate the expected number of *clean* pages (N_2) after t time units when at time = 0 the number of *clean* pages was N_1. This give us the following:

$$N_2 = N_1 \cdot q^t + \sum_{i=0}^{B_m \cdot t-1} q^{\frac{i}{B_m}} = N_1 \cdot q^t + \frac{1-q^t}{1-q^{\frac{1}{B_m}}}$$

$$T = \frac{1}{ln(q)} \cdot ln(\frac{N_2 \cdot (1-q^{\frac{1}{B_m}})-1}{N_1 \cdot (1-q^{\frac{1}{B_m}})-1})$$

Fixed Bandwidth Migration 5/7

- Using the formula above we can calculate the expected time (*T*) it takes until there are N_2 *clean* pages, when in time = 0 there were N_1 *clean* pages, we get: $E(f_g)$

Fixed Bandwidth Migration 6/7

- Finally, the cost of the **pre-copy** phase is given by:

$$cost_{pc} = \frac{1}{ln(q)} \cdot ln(1 - N \cdot (1 - q^{\frac{1}{B_m}})) \cdot F(C - B_m)$$

- And the total cost by:

$$cost_{pc} + \frac{M-N}{C} \cdot F(0) = cost_{pc} + \frac{M-N}{C}$$

- Optimal bandwidth can be found by minimizing the cost function (analytically / numerically).