Paging and k-Server revisited : New approaches in online algorithms

Sushmita Gupta

Kyoto University, Japan

28th April, 2014

OFFLINE vs ONLINE

OFFLINE

Input arrives in one piece

ONLINE

Input arrives piece by piece

OFFLINE vs ONLINE

OFFLINE	ONLINE
Input arrives in one piece	Input arrives piece by piece

Offline Algorithm : Receives its input in entirety, can make processing decision after reading the whole input or only parts of it.

Online Algorithm : Receives its input as a sequence of items, one at a time, and for every item, without knowing subsequent items must make a processing decision for the current item.

Paging and k-Server

$p_1, \ldots, p_N$

LARGE MEMORY

$p_{i_1}, \ldots, p_{i_k}$

SMALL MEMORY

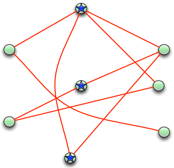

N vertices ★ k Servers

Paging and k-Server are examples of online problems

Paging and k-Server

$p_1, \ldots, p_N$

LARGE MEMORY

$p_{i_1}, \ldots, p_{i_k}$

SMALL MEMORY

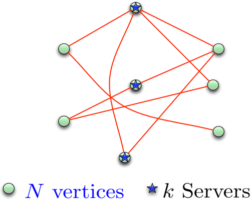

Paging and k-Server are examples of online problems

PAGING and k-SERVER

$p_1, \ldots, p_N$

LARGE MEMORY

$p_{i_1}, \ldots, p_{i_k}$

SMALL MEMORY

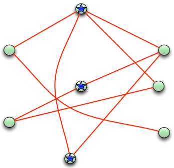

N vertices ★ k Servers

PAGING and k-SERVER are examples of online problems

BIN PACKING, MACHINE SCHEDULING and LOAD BALANCING

are interesting in both offline and online settings.

Computation Models

Computation Models

Distributional/Average-case Complexity Model : A distribution on events (or sequence of events) is hypothesized and then the expected total cost or expected cost per event is analyzed.

Distributional/Average-case Complexity Model : A distribution on events (or sequence of events) is hypothesized and then the expected total cost or expected cost per event is analyzed.

A good hypothesis is hard to formulate!

Computation Models

Distributional/Average-case Complexity Model : A distribution on events (or sequence of events) is hypothesized and then the expected total cost or expected cost per event is analyzed.

A good hypothesis is hard to formulate!

Worst-case Complexity Model : Nothing is assumed about the events (or sequence of events) that give rise to an input instance. One seeks the worst performance of an algorithm over all input instances of the same length.

Computation Models

Distributional/Average-case Complexity Model : A distribution on events (or sequence of events) is hypothesized and then the expected total cost or expected cost per event is analyzed.

A good hypothesis is hard to formulate!

Worst-case Complexity Model : Nothing is assumed about the events (or sequence of events) that give rise to an input instance. One seeks the worst performance of an algorithm over all input instances of the same length.

Competitive analysis is the canonical example.

Worst-Case Complexity Models

Competitive Analysis ($\mathcal{CA}$)

The comparison of the worst case performance of the online algorithm to that of the optimal offline algorithm, OPT on the same sequence.

Competitive Analysis ($\mathcal{CA}$)

The comparison of the worst case performance of the online algorithm to that of the optimal offline algorithm, OPT on the same sequence.

$\exists$ a constant α such that $\forall$ sequence I : $\mathcal{A}(I) \leqslant c.\text{OPT}(I) + \alpha$

Competitive Ratio *is the online analog of Approximation Ratio, used in the offline setting.*

Issues with Competitive Analysis

Strengths and Praises

- Conceptually simple and widely applicable to a wide array of problems.
- Has driven a lot of research in the area of online algorithms.

Limitations And Criticisms

Coarse Classification : Several algorithms are deemed equivalent by $\mathcal{CA}$ even though their performance differs significantly in practice.

Overly Pessimistic : The performance of an online algorithm as indicated by $\mathcal{CA}$ is significantly worse than its performance on most input instances encountered in practical applications.

Issues with Competitive Analysis

Limitations And Criticisms

Use of offline optimal, OPT

Coarse Classification : Several algorithms are deemed equivalent by $\mathcal{CA}$ even though their performance differs significantly in practice.

Worst-case complexity model

Overly Pessimistic : The performance of an online algorithm as indicated by $\mathcal{CA}$ is significantly worse than its performance on most input instances encountered in practical applications.

PAGING: **LRU,FIFO, FWF illustrate these properties**

“Fixing” Competitive Analysis

When comparing performances of algorithms, what if we

Avoid OPT as the medium of comparison

and compare algorithms directly to each other.

Compare performances on *related sequences*

and not the exact same sequence.

"Fixing" Competitive Analysis : Some Approaches

Avoiding OPT as the medium of comparison

Use of direct performance measures.

Comparing performances on *related sequences*

Model locality of reference (LOR) : both deterministic and probabilistic.

“Fixing” Competitive Analysis : Some Approaches

Avoiding OPT as the medium of comparison

Use of direct performance measures. Mixed success when applied to PAGING.

Comparing performances on *related sequences*

Model locality of reference (LOR) : both deterministic and probabilistic.

Promising success.

Many alternate performance measures have been proposed

EXTRA-RESOURCE, BIJECTIVE, RELATIVE WORST ORDER, RELATIVE INTERVAL ANALYSIS are some examples.

Going beyond Competitive Analysis

- Induce a partition on the input space.
  - Given by **locality of reference** modelled by access graphs, concave analysis, Markov chain etc.

BIRS95, IKP96, KPR00, AFG05, ADL07

- Imply, if not explicitly hypothesize a probability distribution on the input space.
  - Given by diffuse adversary model.

You98, You00, KP94, Bec04.

Locality of Reference

Locality can be modeled by an *Access Graph*, defined by Borodin, Irani, Raghavan and Schieber, STOC'91.

A graph $G = (V, E)$ whose vertex set corresponds to the set of pages that can be requested in a sequence. A sequence is said to *respect* an access graph, if the sequence of requests constitute a walk in that access graph.

Locality of Reference with Competitive Analysis

Locality can be modeled by an *Access Graph*, defined by Borodin, Irani, Raghavan and Schieber, STOC'91.

A graph $G = (V, E)$ whose vertex set corresponds to the set of pages that can be requested in a sequence. A sequence is said to *respect* an access graph, if the sequence of requests constitute a walk in that access graph.

- Borodin, Irani, Raghavan and Schieber, JCSS '95 showed that
  - LRU is offline optimal for a path access graph.
  - LRU is online optimal w.r.t $\mathcal{CA}$ for a tree access graph.
  - Introduced a new algorithm FAR, which is strongly competitive i.e, it is within a constant factor of the online optimal.
- Chrobak and Noga , SODA'98 showed that LRU is at least as good as FIFO on every access graph.

Relative Worst Order Analysis

Relative Worst Order Analysis with Locality Of Reference

- LRU and FIFO are equivalent.
  - LRU is at least as good as FIFO on all paths and cycles. Strictly better if length $\geqslant k+1$.
    Separation in cycles cannot be shown under $\mathcal{CA}$.
  - But incomparable on general access graphs.
- FWF and any conservative algorithm are comparable in latter's favor.
  - LRU, FIFO strictly better than FWF on any graph containing P_{k+1}.
- FAR and LRU are identical on all paths and cycles of length $\geqslant 2k$. For C_l, comparable in FAR's favor if $l = k+1$ and only weakly comparable if $k+2 \leqslant l \leqslant 2k-1$.

Results by Boyar, Favroldt, Larsen, JCSS'07 and Boyar, Gupta, Larsen, SWAT'12.

Relative Interval Analysis

Relative Interval Analysis

- $\mathcal{RIA}$ can separate LRU from FWF, while $\mathcal{CA}$ cannot.
- But $\mathcal{RIA}$ cannot separate LRU from FIFO.
  In fact, FIFO *performs better but does not dominate* LRU.

Results by Dorrigiv , Lopez-Ortiz, Munro ISAAC'09, Boyar, Gupta, Larsen WADS'13.

Relative Interval Analysis with Locality Of Reference

- $\mathcal{RIA}$ can separate LRU from FWF, while $\mathcal{CA}$ cannot.
- But $\mathcal{RIA}$ cannot separate LRU from FIFO.

  In fact, FIFO *performs better but does not dominate* LRU.

  - Excluding P_N, this is true.

    Paths : All sequences are served by $\{\text{FAR}, \text{LRU}\}$ at least as well as FIFO.

    Stars : There are sequences which are better served by FIFO than $\{\text{LRU}, \text{FAR}\}$ and vice versa. But skews towards FIFO.

    Cycles : There are sequences which are better served by FIFO than $\{\text{LRU}, \text{FAR}\}$ and vice versa. But skews towards FAR.

Results by Dorrigiv , Lopez-Ortiz, Munro ISAAC'09, Boyar, Gupta, Larsen WADS'13.

A question : Role of OPT

In the light of these improved results, we have to wonder....

Is there a meaningful role for OPT in analysis of online algorithms?

A question : Role of OPT

In the light of these improved results, we have to wonder....

Is there a meaningful role for OPT in analysis of online algorithms?

We take a step back and reevaluate the essential role of OPT.

Dual Role

- Adversary: which generates sequences that the ONL algorithm has to serve.
- Standard: against which the ONL solution is assessed.

A question : Role of OPT

In the light of these improved results, we have to wonder....

Is there a meaningful role for OPT in analysis of online algorithms?

We take a step back and reevaluate the essential role of OPT.

Dual Role

- Adversary which generates sequences that the ONL algorithm has to serve.
- Standard against which the ONL solution is assessed.

OPT gives a succinct description of the perfect solution. What if we use it as an advisor?

Advice Complexity Model

Dobrev, Královič and Pardubská SOFSEM'08 and Hromkovič, Královič and Královič

- There is an oracle, which has access to the whole input in advance.
- It computes the optimal solution ahead of ONL.
- It then provides some advice to the ONL when it serves the request sequence.
- The goal is to minimise the amount of help provided to improve the performance of the online algorithm.

The advice is encoded as a string of $\{0, 1\}$ bits and the amount of advice quantifies the advice complexity.

Advice Complexity Model

Dobrev, Královič and Pardubská SOFSEM'08 and Hromkovič, Královič and Královič

- There is an oracle, which has access to the whole input in advance.
- It computes the optimal solution ahead of ONL.
- It then provides some advice to the ONL when it serves the request sequence.
- The goal is to minimise the amount of help provided to improve the performance of the online algorithm.

The advice is encoded as a string of $\{0, 1\}$ bits and the amount of advice quantifies the advice complexity.

Advice complexity quantifies the *information gap* separating ONL and OPT.

Advice Complexity Model : Paging

Optimal : Advice complexity $s(n) = n + k$.

1-competitive : Advice complexity $s(n) = n$.

Lower bound for optimality : $s(n) \geqslant n\left(1 - O(\frac{\log(k-1)}{4(k-1)}) - f(k)\right)$, $O(1)$ is independent of k.

Advice complexity converges to linear (1 bit per request) as k grows.

Results are by DKP08, BKKKM09, HKK10

Advice Complexity Model : Paging

Optimal : Advice complexity $s(n) = n + k$.

1-competitive : Advice complexity $s(n) = n$.

Lower bound for optimality : $s(n) \geqslant n\left(1 - O(\frac{\log(k-1)}{4(k-1)}) - f(k)\right)$, $O(1)$ is independent of k.

Advice complexity converges to linear (1 bit per request) as k grows.

Tradeoff : There is an algorithm with advice complexity $s(n) = bn$ attains competitive ration $r \leqslant 3b + \frac{2(k+1)}{2^b}$.
For $N = k + 1$, any deterministic online algorithm with advice complexity $s(n) = bn$ has competitive ratio $r \geqslant k/2^b$.

Results are by DKP08, BKKKM09, HKK10

Advice Complexity : k-Server

Lots of work on different types of Metric Spaces :
2-dimensional Euclidean space, finite trees, bounded treewidth, μ collective (q, r)-spanners, etc along with general metric space.

Very small sample of results...

General Metric Space For every $b \geqslant 3$, there is ONL with advice complexity $s(n) = bn$ that attains competitive ratio of $\left\lceil \frac{\lceil \log k \rceil}{b-1} \right\rceil$.

Advice complexity $s(n) = \Theta(n)$ is necessary and sufficient for an optimal algorithm.

Some concluding thoughts

- A performance measure alone, cannot be expected to capture the entire spectrum of observed behavior of all online algorithms.
- It seems that the use of locality models such as LOR or probabilistic models which are amenable to being tailored and fine tuned to model scenarios present in specific real-life applications are likely to give more satisfactory results.
- A more nuanced role for OPT may be desirable and useful for further development of the subject, one such possible direction is now being explored under *advice complexity*.

Thank you for your attention!

References I

- A. Borodin, S. Irani, P. Raghavan and B. Schieber, *Competitive paging with locality of reference*, JCSS, 50(2), 1995.
- J. Boyar, L.M. Favrholdt and K.S. Larsen, *The relative worst order ratio applied to paging*, JCSS, 73: 818-843, 2007.
- R. Dorrigiv, A. López-Ortiz and J. Ian Munro, *On the Relative Dominance of Paging Algorithms*, TCS, 410(38-40): 3694–3701, 2009.
- M. Chrobak and J. Noga, *LRU is Better than FIFO*, Algorithmica, 23:180–184, 1999.
- J. Boyar and L.M. Favrholdt, *The relative worst order ratio for on-line algorithms*, ACM TALG, 3(2):1549-6325, 2007.
- L. Epstein, L.M. Favrholdt and J.S. Kohrt, *Comparing online algorithms for bin packing problems*, Journal of Scheduling, 15(1): 13-2, 2009.