# Fully-dynamic-to-incremental reductions with known deletion order (e.g. sliding window) 

Binghui Peng*<br>Columbia University<br>bp2601@columbia.edu

Aviad Rubinstein ${ }^{\dagger}$<br>Stanford University<br>aviad@cs.stanford.edu


#### Abstract

Dynamic algorithms come in three main flavors: incremental (insertions-only), decremental (deletions-only), or fully dynamic (both insertions and deletions). Fully dynamic is the holy grail of dynamic algorithm design; it is obviously more general than the other two, but is it strictly harder?

Several works managed to reduce fully dynamic to the incremental or decremental models by taking advantage of either specific structure of the incremental/decremental algorithms (e.g. HK99, HdLT01, BKS12, $\mathrm{ADK}^{+} 16$, BS80, OL81, OvL81), or specific order of insertions/deletions (e.g. AW14, HKNS15, KPP16). Our goal in this work is to get a black-box fully-to-incremental reduction that is as general as possible. We find that the following conditions are necessary:


- The incremental algorithm must have a worst-case (rather than amortized) running time guarantee.
- The reduction must work in what we call the deletions-look-ahead model, where the order of deletions among current elements is known in advance. A notable practical example is the "sliding window" (FIFO) order of updates.
Under those conditions, we design:
- A simple, practical, amortized-fully-dynamic to worst-case-incremental reduction with a $\log (T)$-factor overhead on the running time, where $T$ is the total number of updates.
- A theoretical worst-case-fully-dynamic to worst-case-incremental reduction with a polylog $(T)$ factor overhead on the running time.

[^0]
## 1 Introduction

A dynamic algorithm is a data structure that maintains certain properties (e.g. shortest paths) of a ground set that is subject to a sequence of updates (e.g. insertions/deletions of edges), and the goal is to minimize the (total) update time. A fully dynamic algorithm supports both insertions and deletions of the ground set elements, while an incremental algorithm restricts the updates to be insertion-only and a decremental algorithm handles deletion operations only.

A fully dynamic algorithm clearly benefits from handling more general updates, but the at the same time, it is expected to be much harder than incremental/decremental algorithms. Meanwhile, several existing works have exploited special structure (HK99, HdLT01, BKS12, ADK ${ }^{+}$16, BS80, OL81, OvL81]) or specific order of update sequence ([AW14, HKNS15, KPP16]) and reduce fully dynamic to incremental or decremental algorithm. This motivates one to ask

Can a generic reduction transform an incremental algorithm into one that handles both insertions and deletions?

Perhaps surprisingly, we find that once the order of deletions of current elements is known to the algorithm (deletions-look-ahead), one can translate an incremental algorithm with worst case guarantee to a fully dynamic algorithm with worst case guarantee, with only polylogarithmic overhead.

Theorem 1.1 (Reduction, worst case to worst case). Let $T \geq 1$ be the total number of updates. Suppose there exists a dynamic algorithm in the incremental setting with query time $\Gamma_{q}$ and worst case update time $\Gamma_{u}$, then there is a dynamic algorithm for deletions-look-ahead setting with query time $\Gamma_{q}$ and worst case update time 1

Our reduction requires the incremental algorithm to have worst case (rather than amortized) runtime guarantee, and most importantly, the relative order of deletions of current elements must be known 2 The latter assumption is satisfied by well-studied models including sliding window (FIFO) model ELVZ17, EMMZ22, DGIM02, $\mathrm{BDM}^{+} 20$, WZ22, where only most recent insertion updates are of interests, as well as offline/look-ahead model KMW98, SM10, vdBNS19, CGH ${ }^{+}$20, AW14, where the entire sequence of updates is known in advance. These models are of interests to both theoretical and empirical community [HHS22]. Furthermore, the sliding window model is often used as a benchmark for empirical investigation of fully dynamic algorithms (e.g. WLFT18, LMNF ${ }^{+}$20, HN21). To complement our results, we also prove that both conditions are indeed indispensable for a black box reduction (see Section (5).

If one only aims for algorithms with amortized runtime guarantee, we have a reduction with improved update time. We believe it could be of independent interests due to its simplicity and could be beneficial for empirical implementation.

Theorem 1.2 (Reduction, amortized to worst case). Let $T \geq 1$ be the total number of updates. Suppose there exists a dynamic algorithm for incremental setting with query time $\Gamma_{q}$ and worst case update time $\Gamma_{u}$, then there is a dynamic algorithm for deletions-look-ahead setting with query time $\Gamma_{q}$ and amortized update time $O\left(\Gamma_{u} \cdot \log (T)\right)$.

On the technical side, Theorem 1.2 exploits the ideas of rewinding the incremental algorithm (e.g. AW14, HKNS15, KPP16), and in particular achieves only logarithmic overhead by rewinding

[^1]$2^{i}$ insertions every $\Theta\left(2^{i}\right)$ updates (e.g. [HK99, HdLT01, BKS12, ADK $\left.\left.{ }^{+} 16\right]\right)$. To achieve the worst case guarantee, Theorem 1.1 additionally requires amortizing this rewinding of $2^{i}$ insertions over $\Theta\left(2^{i}\right)$ updates; this is a little trickier because we have to start re-inserting elements in advance before all the elements we would want to insert have arrived.

We demonstrate the power of our reduction in Section 4 by providing applications on dynamic submodular maximization, dynamic Depth first search (DFS) tree and dynamic All-Pair ShortestPaths (APSP).

Related work A systematic study of black box reduction for dynamic algorithms has been initiated by the seminal work of [BS80], which shows how to make a static data structure support insertions for "decomposable search problem". The ideas of rewinding incremental algorithm and the logarithmic scheduling have been studied in different problem specific context AW14, HKNS15, KPP16, HK99, HdLT01, BKS12, ADK ${ }^{+}$16, BS80, OL81, OvL81, DS91, Cha12. Our deletion-lookahead model has also been considered in the computational geometry literature, and it is one variant of the semi-online model (see DS91] for a detailed discussion). The work of Cha12 is closely related to us, it presents a worst-case to amortized-case reduction, when the exact deletion time is known. The result is (almost) equivalent to Theorem $1.2^{3}$ It differs from Theorem 1.1 as our reduction is worst-case to worst-case.

Notation Let $[n]=\{1,2, \ldots, n\}$ and $\left[n_{1}: n_{2}\right]=\left\{n_{1}, n_{1}+1, \ldots, n_{2}\right\}$. For any ground set elements $e_{1}$ and $e_{2}$, we say $e_{1}$ is younger than $e_{2}$ if $e_{1}$ would be deleted earlier than $e_{2}$. An element is of rank $r$ if it is the $r$-th youngest of the current elements. For an ordered set $A$, we use $A\left[n_{1}: n_{2}\right]$ to denote the youngest $n_{1}$-th to $n_{2}$-th elements of $A$. For any $t \in[T]$, let $k(t)$ be the largest integer such that $t$ is exactly a multiple of $2^{k(t)}$. In our pseudocode, InSERT refers to the insertion procedure of the incremental algorithm (hence taking $O\left(\Gamma_{u}\right)$ time), while adding/removing elements to a set $A$ operates only over the the set (hence taking $O(1)$ time).

## 2 Reduction - Amortized to worst case

We start by providing the amortized to worst case reduction and prove Theorem 1.2. It also serves as a warm up for the worst case to worst case reduction. Our key idea is to re-order current elements in reverse order of deletions, this ensures a deletion update takes only $O\left(\Gamma_{u}\right)$ time by rewinding the computation. Of course, maintaining the reverse order could be expensive as the new-coming element could be deleted last (e.g. the sliding window model). Instead, we only maintain a partial reverse order - we re-insert the last $O\left(2^{i}\right)$ elements in reverse order in every $2^{i}$ updates. This suffices to handle fully-dynamic updates and brings $O(\log (T))$ computation overhead on average.

The reduction is formally depicted in Algorithm 1 The current elements are distributed to $m+1=\left\lceil\log _{2} T\right\rceil+1$ buckets $B_{0}, \ldots, B_{m}$. At the $t$-th update $(t \in[T])$, the algorithm first handles the update (insertion or deletion) on $B_{0}$ and it is guaranteed by our algorithm that a deletion occurs on $B_{0}$. The algorithm then rewinds the computation to $B=B_{m} \ldots B_{k(t)+2}$ (Line 3) and then Insert elements of $B_{0} \cup \cdots \cup B_{k(t)+1}$ in the reverse order (Line 8). Algorithm [1] re-arranges buckets as follows: It keeps buckets $B_{m}, \ldots, B_{k(t)+2}$ unchanged, it puts the youngest $\left[2^{i}: 2^{i+1}-1\right]$ elements in $B$ to bucket $B_{i}(i \in[0: k(t)])$ and the remaining elements to $B_{k(t)+1}$ (see Line 6] - 7).

For any $t \in[T]$, let $E_{t}$ be all existing elements at the end of $t$-th update, and let $E_{t, r}$ be the youngest $r$ elements in $E_{t}$ (when there are less than $r$ elements in $E_{t}$, we take $E_{t, r}=E_{t}$ ). The following lemma formalizes the main invariant for Algorithm 1 ,

[^2]```
Algorithm 1 Reduction - Amortized to worst case
    Initialize \(B_{i} \leftarrow \emptyset(i \in[0: m]) \quad \triangleright m=\left\lceil\log _{2} T\right\rceil\)
    for \(t=1,2, \ldots, T\) do
        Rewind the computation to \(B_{m} \ldots B_{k(t)+2}\)
        Add/remove the element in \(B_{0} \quad \triangleright\) Deletion is guaranteed to occur in \(B_{0}\)
        \(B \leftarrow B_{0} \cup \cdots \cup B_{k(t)+1}\)
        \(B_{i} \leftarrow B\left[2^{i}: 2^{i+1}-1\right](i \in[0: k(t)]) \quad \triangleright B_{i}\) contains the youngest \(\left[2^{i}: 2^{i+1}-1\right]\) elements
        \(B_{k(t)+1} \leftarrow B \backslash\left(B_{0} \cup \cdots \cup B_{k(t)}\right)\)
        Insert \(B_{k(t)+1}, \ldots, B_{0} \quad \triangleright\) Elements are inserted in reverse order of deletion
    end for
```

Lemma 2.1. At the end of $t$-th update $(t \in[T]$ ), one has

- $\left|B_{0} \cup \cdots \cup B_{k(t)+1}\right| \leq 2^{k(t)+2}+2^{k(t)} ;$
- $E_{t, 2^{i+1}-1} \subseteq B_{0} \cup \cdots \cup B_{i}$ for any $i \in[0: k(t)]$.

Proof. For the first claim, at the end of $\left(t-2^{k(t)}\right)$-th update, Line 6 of Algorithm 1 guarantees that $\left|B_{i}\right| \leq 2^{i}$ holds for any $i \in[0: k(t)+1]$ as $\left(t-2^{k(t)}\right)$ is a multiple of $2^{k(t)+1}$. Since there are at most $2^{k(t)}$ insertions between the $\left(t-2^{k(t)}\right)$-th and $t$-th update, we have

$$
\left|B_{0} \cup \cdots \cup B_{k(t)+1}\right| \leq 2^{k(t)}+\sum_{i=0}^{k(t)+1} 2^{i} \leq 2^{k(t)+2}+2^{k(t)}
$$

We prove the second claim by induction. The case of $t=1$ holds trivially and suppose the induction holds up to $t-1$. Then we know that

$$
E_{t-2^{k(t)}, 2^{k(t)+2}-1} \subseteq B_{0} \cup \cdots \cup B_{k(t)+1}
$$

at the end of $\left(t-2^{k(t)}\right)$-th update as $\left(t-2^{k(t)}\right)$ is a multiple of $2^{k(t)+1}$. As a consequence, one has

$$
E_{t-2^{k(t)}} \backslash E_{t-2^{k(t), 2^{k(t)+2}-1}} \cap E_{t, 2^{k(t)+1}-1}=\emptyset
$$

as there are at most $2^{k(t)}$ deletions between the $\left(t-2^{k(t)}\right)$-th and $t$-th update, and $2^{k(t)+2}-1-2^{k(t)} \geq$ $2^{k(t)+1}-1$. At the same time, new elements arrive between the $\left(t-2^{k(t)}\right)$-th and $t$-th update and they are all contained in $B_{0} \cup \cdots \cup B_{k(t)+1}$, hence, we conclude that $E_{t, 2^{k(t)+1}-1} \subseteq B_{0} \cup \cdots \cup B_{k(t)+1}$. By Line 6 of Algorithm $\mathbb{\square}$, we have $E_{t, 2^{i+1}-1} \subseteq B_{0} \cup \cdots \cup B_{i}$ for any $i \in[0: k(t)]$. We conclude the proof here.

The correctness of Algorithm $\mathbb{1}$ follows immediately from the correctness of the incremental algorithm because the set $B_{0}$ always contains the youngest element. It remains to bound the amortized running time.

Lemma 2.2. The amortized update time of Algorithm $\mathbb{1}$ is at most $O\left(\Gamma_{u} \cdot \log (T)\right)$.
Proof. Line 4 takes constant time since $B_{0}$ has constant size. By Lemma 2.1, $\left|B_{0} \cup \cdots \cup B_{k(t)+1}\right| \leq$ $2^{k(t)+2}+2^{k(t)}$, the rewinding step (Line (3) takes at most $O\left(2^{k(t)} \Gamma_{u}\right)$ time and we make $O\left(2^{k(t)}\right)$ calls to Insert at Line 8. The allocation step (Line 6) takes no more than $O\left(2^{t(k)}\right)$ time, since the bucket
$B_{i}$ has already been sorted $(i \in[0: k(t)+1])$ and it remains to merge them. The total update time equals

$$
\sum_{t=1}^{T} O\left(2^{k(t)} \Gamma_{u}\right)=\sum_{k=0}^{m} O\left(2^{k} \Gamma_{u}\right) \cdot T / 2^{k}=O\left(\Gamma_{u} \cdot T \log (T)\right)
$$

We conclude the proof here.
Remark 2.3 (Implementation of rewinding). We work in the RAM model and perform reversible computation. One simple way of implementing reversible computation (e.g. [Ben73]) is to write down the change to memory cell in every step. The forward computation time only slows down by a constant factor and the backward (rewind) computation time equals the forward computation time. In practice for specific problems there may be faster ways to implement rewinding.

## 3 Reduction - Worst case to worst case

We next dedicate to prove Theorem 1.1, which translates the worst case guarantee from the incremental model to the deletions-look-ahead model. The major difference with the amortized reduction is that one cannot re-order/re-insert a large block of elements at once. A natural idea is to prepare the re-order/re-insertion in advance and split the cost. This brings new challenges as (unknown) future insertions/deletions interleave with the preparation step and one does not know the exact set of elements beforehand. To resolve it, we maintain multiple threads, and each thread further divides the preparation step into epochs of geometrically decreasing size.

The high-level idea is presented in Algorithm 2 with implementation details deferred to the proof of Lemma 3.2 Algorithm 2 maintains $m+1$ threads and $m+1$ buckets $B_{0}, \ldots, B_{m}$. During the execution of the algorithm, all existing elements are distributed over $B_{0}, \ldots, B_{m}$, and ideally, the $i$-th bucket $B_{i}$ should be of size $O\left(2^{i}\right)$ and $B_{0} \cup \cdots \cup B_{i}$ should contain the youngest $\Omega\left(2^{i}\right)$ elements. This guarantees the insertion/deletion of an element can be resolved in $O\left(\Gamma_{u}\right)$ time since one only needs to re-insert elements in $B_{0}$.

The crucial part is to maintain the ordered buckets $B_{0}, \ldots, B_{m}$, for which Algorithm 2 maintains $m$ threads; the $i$-th thread $(i \in[m])$ prepares the re-order/re-insertion ahead of $2^{i}$ updates. Precisely, the $i$-th thread re-starts every $2^{i+1}$ updates and operates over the upcoming $2^{i}$ updates (Line 12 - 28). It first rewinds the computation status to $B_{m} \ldots B_{i}$ (Line 144), which is prepared by the $(i+k(\tau))$-th thread and then re-inserts elements of $B_{0} \cup \cdots \cup B_{i-1}$ (Line 16- 24). Concretely, the re-insertion procedure is further divided into $i$ epochs, where epoch $j(j \in[i-1])$ lasts for $2^{j}$ updates and epoch 0 lasts for 2 updates. Let $t(i, \tau, j):=2^{i} \tau+\sum_{r=j+1}^{i-1} 2^{r}$ denote the end of epoch $j+1$ in the $(\tau / 2+1)$-th outer-for-loop-iteration of the $i$-th thread. During epoch $j$, the $i$-th thread leaves alone the youngest $2^{j+2}$ elements at the beginning of epoch $j$ to $B^{(i)}$ and Insert the remaining elements $B_{j}^{(i)}$ over the upcoming $2^{j}$ updates (i.e. $\left[t(i, \tau, j)+1: t(i, \tau, j)+2^{j}\right]$-th update, see Line (18). Meanwhile, the set $B^{(i)}$ is updated and elements are added and removed (Line 20). Finally, at the end of $t$-th update $(t \in[T]$ and $k(t) \geq 1)$, Algorithm 2 resets $B_{j}$ to $B_{j}^{(k(t))}$ for every $j \in[0: k(t)-1]\left(\right.$ Line 266) and this step takes $O(\log T)$ time as we only change the pointer of $B_{j}$.

Recall $E_{t, r}$ denotes the youngest $r$ elements and $E_{t}$ denotes all elements at the end of $t$-th update. We use $B_{t, j}$ and $B_{t, j}^{(i)}$ to denote the status of $B_{j}$ and $B_{j}^{(i)}$ at the end of $t$-th update. We first formalize the main invariant for Algorithm 2.

Lemma 3.1. For any thread $i \in[m]$, outer-for-loop-iteration $\tau \in\left\{0,2, \ldots, T / 2^{i}\right\}$, at the end of the $\left(2^{i} \tau\right)$-th update, we have

```
Algorithm 2 Reduction - Worst case to worst case
\(\triangleright\) Variables with superscript \({ }^{(i)}\) internal to Thread \((i)\)
\(\triangleright \operatorname{Thread}(i)\) only uses information from bigger threads (aka \(\operatorname{Thread}(j)\) for \(j>i\) )
\(\triangleright\) In particular, Thread \((i)\) 's Insert and rewind do not affect the state seen by bigger threads
\(\triangleright\) The output of the algorithm is maintained by Thread \((0)\)
    Initialize \(B_{i} \leftarrow \emptyset\) and run \(\operatorname{Thread}(i)(i \in[0: m]) \quad \triangleright m=\left\lceil\log _{2} T\right\rceil\)
    procedure Thread (0)
    for \(t=1,2, \ldots, T\) do
        Add/remove the element in \(B_{0} \quad \triangleright t\)-th update (deletions guaranteed to be from \(B_{0}\) )
        Rewind the computation to \(B_{m} \ldots B_{1} \quad \triangleright\) State prepared by Thread \((k(t-1))\)
        Insert \(B_{0}\)
    end for
end procedure
procedure Thread \((i)\)
    \(\triangleright i \in[1: m]\)
    for \(\tau=0,2,4, \ldots,\left\lfloor T / 2^{i}\right\rfloor\) do \(\quad \triangleright\) Restart every \(2^{i+1}\) updates
            \(B^{(i)} \leftarrow B_{0} \cup \cdots \cup B_{i-1}\)
            Rewind the computation to \(B_{m} \ldots B_{i} \quad \triangleright\) State prepared by Thread \((i+k(\tau))\)
            for \(j=i-1, i-2, \ldots, 1\) do \(\quad \triangleright j\)-th epoch, amortized over \(2^{j}\) updates
                \(B_{j}^{(i)} \leftarrow B^{(i)}\left[2^{j+2}+1:\right] \quad \triangleright\) Oldest \(\Theta\left(2^{j}\right)\) elements in \(B^{(i)}\)
                Insert \(B_{j}^{(i)}\)
                \(B^{(i)} \leftarrow B^{(i)} \backslash B_{j}^{(i)}\)
                Add/remove elements in \(B^{(i)} \quad \triangleright\) Updates \(\left[t(i, \tau, j)+1: t(i, \tau, j)+2^{j}\right]\)
            end for
            Add/remove elements of \(B^{(i)}\) in the remaining 2 updates \(\triangleright\) Epoch 0
            \(B_{0}^{(i)} \leftarrow B^{(i)}\)
            Insert \(B_{0}^{(i)}\)
            \(B_{j} \leftarrow B_{j}^{(i)}(\forall j \in[0: i-1])\)
            Do nothing for \(2^{i}\) updates
    end for
end procedure
```

- $E_{2^{i} \tau, 2^{j+2}} \subseteq B_{2^{i} \tau, 0} \cup \cdots \cup B_{2^{i} \tau, j}$,
- $\left|B_{2^{i} \tau, j}\right| \leq 3 \cdot 2^{j+1}$
holds for any $j \in[i-1]$. For $j=0$, we have $\left|B_{2^{i} \tau, 0}\right| \leq 12$ and $E_{2^{i} \tau, 4} \subseteq B_{2^{i} \tau, 0}$.
Proof. We prove the first bullet by an induction on $i$ (in the reverse order). The base case of $i=m$ holds trivially as all buckets are empty at the beginning.

Suppose the induction holds up to the ( $i+1$ )-th thread. For any outer-for-loop-iteration $\tau \in$ $\left\{2, \ldots, T / 2^{i}\right\}$, at the end of $2^{i} \tau$-th update, the buckets $B_{0}, \ldots, B_{i-1}$ are reset by the $(i+k(\tau))$-th thread, hence, it suffices to prove $E_{2^{i} \tau, 2^{j+2}} \subseteq B_{2^{i} \tau, 0}^{(i+k(\tau))} \cup \cdots \cup B_{2^{i} \tau, j}^{(i+k(\tau))}(\forall j \in[i-1])$. By the inductive hypothesis of the $(i+k(\tau))$-th thread, we know that

$$
E_{2^{i} \tau-2^{i+k(\tau)}, 2^{i+k(\tau)+1}} \subseteq B_{2^{i} \tau-2^{i+k(\tau)}}^{(i+k(\tau))}=B_{2^{i} \tau-2^{i+k(\tau)}, 0} \cup \cdots \cup B_{2^{i} \tau-2^{i+k(\tau), i+k(\tau)-1}},
$$

that is, the youngest $2^{i+k(\tau)+1}$ are contained in $B^{(i+k(\tau))}$ initially. We prove the desired claim by contradiction and assume for some $j \in[0: i-1]$, there exists an element $e$ such that $e \in E_{2^{i} \tau, 2^{j+2}}$ but $e \notin B_{2^{i} \tau, 0}^{(i+k(\tau))} \cup \cdots \cup B_{2^{i} \tau, j}^{(i+k(\tau))}$. This can only happen if (1) the element $e$ is inserted before epoch $j+1$; and (2) it is removed from $B^{(i+k(\tau))}$ at some epoch $\gamma \geq j+1$. The reason for (1) is that elements inserted on/after epoch $j+1$ would ultimately be included in $B_{2^{i} \tau, 0}^{(i+k(\tau))} \cup \cdots \cup B_{2^{i} \tau, j}^{(i+k(\tau))}$; the reason for (2) is similar.

Since the element $e$ is removed from $B^{(i+k(\tau))}$ at epoch $\gamma$, we have that

$$
e \notin E_{t\left(i+k(\tau), \tau / 2^{k(\tau)}-1, \gamma\right)} \backslash E_{t\left(i+k(\tau), \tau / 2^{k(\tau)}-1, \gamma\right), 2^{\gamma+2}} .
$$

There are at most $2+\sum_{r=1}^{\gamma} 2^{r}=2^{\gamma+1}$ deletions since epoch $\gamma$, the rank of $e$ can be improved to at most $2^{\gamma+2}-2^{\gamma+1}=2^{\gamma+1}$, hence we have $e \notin E_{2^{i} \tau} \backslash E_{2^{i} \tau, 2^{\gamma+1}}$ and therefore $e \notin E_{2^{i} \tau} \backslash E_{2^{i} \tau, 2^{j+2}}$ ( $\gamma \geq j+1$ ), this contradicts with the assumption.

For the second bullet, for any $i \in[m]$ and $\tau \in\left\{2, \ldots, T / 2^{i}\right\}$, consider the $(i+k(\tau))$-th thread. After executing Line 19 in the $(j+1)$-th epoch of Algorithm 2, we have that

$$
\left|B_{t\left(i+k(\tau), \tau / 2^{k(\tau)}-2^{j+1}, j\right)}^{(i+k(\tau))}\right| \leq 2^{j+3}
$$

For the rest of the epoch, there can be at most $2^{j+1}$ insertions, hence:

$$
\left|B_{t\left(i+k(\tau), \tau / 2^{k(\tau)}, j\right)}^{(i+k(\tau))}\right| \leq 2^{j+3}+2^{j+1}
$$

Finally, after executing Line 17 in the $j$-th epoch, we have:

$$
\left|B_{t\left(i+k(\tau), \tau / 2^{k(\tau)}, j\right), j}^{(i+k(\tau))}\right| \leq 2^{j+3}+2^{j+1}-2^{j+2}=3 \cdot 2^{j+1}
$$

We have proved the first claim for $j \in[i-1]$, the case of $j=0$ follows similarly.
We next bound the worst case update time.
Lemma 3.2. The update time per operation is at most $O\left(\Gamma_{u} \cdot \log ^{2}(T)\right)$.

Proof. By Lemma 3.1, the size of $B_{0}$ is $O(1)$ and it contains the youngest 2 elements, hence the rewinding and Insert step (Line 7) can be performed in $O\left(\Gamma_{u}\right)$ time per update. The major overhead comes from maintaining $m$ threads, and we bound the runtime of each thread separately.

For any thread $i$ and outer-for-loop-iteration $\tau \in\left\{0,2, \ldots, T / 2^{i}\right\}$, due to Line 17 of Algorithm 2. we have

$$
\left|B_{t(i, \tau, j-1), j}^{(i)}\right| \leq\left|B_{t(i, \tau, j)}^{(i)}\right| \leq 2^{j+3}+2^{j+1} \quad \forall j \in[i-2]
$$

and by Lemma 3.1,

$$
\left|B_{t(i, \tau, i-2), i-1}^{(i)}\right| \leq\left|B_{2^{\tau} i, 0} \cup \cdots \cup B_{2^{\tau} i, i-1}\right| \leq \sum_{r=0}^{i-1} 3 \cdot 2^{r+1} \leq 3 \cdot 2^{i+1}
$$

We analyse the update time step by step.
We first come to the rewinding step (Line 144). Unlike the amortized case, we cannot simply rewind by the reversible computation since we maintain multiple threads that need to access the state of the incremental algorithm with different sets of elements, in parallel. Instead, when we call Insert of each block $B_{j}^{(i)}$, we maintain a dictionary that records the location/value of changed memory cell. The construction of dictionary only incurs constant overhead. By doing this, during the execution of Algorithm 2, one can access any memory cell by looking up to at most $O(\log (T))$ dictionaries (note the lookup path is known to Algorithm (2) and find the last time it has been changed. Naively, looking up the memory updates in each dictionary takes $O\left(\log \left(\Gamma_{u} T\right)\right)$ time. This brings an $O\left(\left(\log (T) \log \left(\Gamma_{u} T\right)\right)\right)$ total overhead for every operation of Insert. Except for this, Line 14 essentially comes for free.

A more careful implementation leads to only $O(\log (T))$ overhead. We maintain an additional data structure, which links each memory cell of the incremental algorithm to $m+1$ lists, where the $i$-th list records the changes made by the $i$-th thread in chronological order. The maintenance of the data structure slows down the forward computation of InSERT by a constant factor. At the same time, in order to search the content of a memory cell, we only need to search through the lists (note again the look-up path is known), which takes $O(1)$ time per list and $O(\log (T))$ in total. Hence, it brings an $O((\log (T))$ total overhead for every operation of Insert.

Algorithm 2 updates $B^{(i)}$ and $B_{j}^{(i)}$ at the beginning of epoch $j$ (Lines 17 and 19). We do not rewrite, but instead, we copy $B^{(i)}$ and $B_{j}^{(i)}$ to new memory cells. Since both sets are of size $O\left(2^{j}\right)$, the copy operation can be done in the first $\frac{1}{4} \cdot 2^{j}$ updates during epoch $j$ and has $O(\log (T))$ cost per update using Binomial heap.

Algorithm 2 calls at most $O\left(2^{j}\right)$ times Insert during epoch $j$ (Line 18). Since elements in $B_{j}^{(i)}$ are known at the beginning so these operations can be averaged over the following $3 / 4 \cdot 2^{j}$ updates of epoch $j$ and take $O\left(\Gamma_{u} \log (T)\right)$ time per update.

The set $B^{(i)}$ receives new elements as well as removes old elements in epoch $j$ (Line 20). We buffer the changes in the first $\frac{1}{4} \cdot 2^{j}$ updates (as the "new" set $B^{(i)}$ is not yet ready) and add/remove elements during the following $3 / 4 \cdot 2^{j}$ updates. The size of $B^{(i)}$ is $O\left(2^{j}\right)$ during epoch $j$, so the update cost is $O(\log (T))$ per update.

Finally, we note (1) Lines 22-24 takes only $O\left(\Gamma_{u} \log (T)\right)$ time in total; (2) Line 13 can be done similarly to Line 17, (3) Line 26 resets $B_{j}(j \in[0: i-1])$ by changing the pointer, so it incurs only $O(\log T)$ cost.

Overall, Algorithm 2 has worst case update time $O\left(\Gamma_{u} \log (T)\right)$ per thread and $O\left(\Gamma_{u} \cdot \log ^{2}(T)\right)$ in total.

Proof of Theorem 1.1. The worst case guarantee has already been established in Lemma 3.2, it remains to prove the correctness of Algorithm 2.

By Lemma 3.1, the youngest 2 elements are always contained in $B_{0}$, hence insertions/deletions are operated correctly, i.e., the removal step (Line 5) indeed removes element in $B_{0}$. It remains to prove each thread operates normally, i.e., for any thread $i$ and outer-for-loop-iteration $\tau$, the removal operation would only remove elements in $B^{(i)}$ during epoch $j(j \in[i-1])$. It suffices to prove that $E_{t(i, \tau, j), 2^{j+2}} \in B_{t(i, \tau, j)}^{(i)}$. We prove by induction. This is true in epoch $i-1$ by Lemma 3.1. Suppose it holds to epoch $j+1$, i.e., $E_{t(i, \tau, j+1), 2^{j+3}} \in B_{t(i, \tau, j+1)}^{(i)}$, since there are at most $2^{j+1}$ deletions in epoch $j+1$, we have that $E_{t(i, \tau, j+1), 2^{j+2}} \in B_{t(i, \tau, j)}^{(i)}$. We complete the proof here.

## 4 Application

We provide a few applications of our reduction.

### 4.1 Submodular maximization

Dynamic submodular maximization In a submodular maximization problem, there is a ground set $N=[n]$ and a set function $f: N \rightarrow \mathbb{R}^{+}$. The function is said to be monotone if $f(A) \geq f(B)$ for any $B \subseteq A \subseteq N$ and it is said to be submodular if $f(A \cup\{u\})-f(A) \leq f(B \cup\{u\})-f(B)$ for any $B \subseteq A \subseteq N$ and element $u$. The task of submodular maximization under a cardinality constraint refers to $\max _{S \subseteq[n],|S|=k} f(S)$ for some parameter $1 \leq k \leq n$, and the task of submodular maximization under a matroid constraint $\mathcal{M}$ refers to $\max _{S \subseteq \mathcal{M}} f(S)$. Finally, in a dynamic submodular maximization problem, the ground set can be inserted and deleted, and the goal is to maintain a good solution set $S$.
[FLNF ${ }^{+} 22$ provides an 0.3178 -approximation algorithm (with a matroid constraint) under streaming setting, and one can adapt it to a dynamic algorithm with worst case update time under incremental setting.

Theorem 4.1 (Adapt from $\left[\overline{F L N F}^{+22}\right]$ ). For any $n, k>0$, under the incremental update, there exists a dynamic algorithm that maintains an 0.3178-approximate solution for monotone submodular maximization under a matroid constraint of rank $k$ and makes $\operatorname{poly}(k, \log n)$ queries per iteration.

The sliding window model is of interests to the community [ELVZ17]. The algorithm of ELVZ17] maintains an 1/2-approximation solution with polylogarithmic updates time for dynamic submodular maximization under cardinality constraints. Our reduction gives the first constant approximation algorithm for a matroid constraint.

Theorem 4.2 (Dynamic submodular maximization). For any $n, k>0$, there exists a dynamic algorithm that achieves 0.3178-approximation for the problem of submodular maximization under a matroid constraint using $\operatorname{poly}(k, \log n)$ queries per update under the sliding window model.

### 4.2 Depth first search (DFS) tree

Dynamic DFS Given an undirected graph $G=(V, E)$ with $|V|=n,|E|=m$, the task is to maintain a depth first search (DFS) tree under edge insertion/deletion. In the incremental model, BCCK16] obtains a dynamic algorithm with $O\left(n(\log n)^{3}\right)$ worst case update time, and it is improved to $O(n)$ by $\mathrm{CDW}^{+} 18$. While in the fully dynamic model, the current best known algorithm [BCCK16] has $\widetilde{O}(\sqrt{m n})$ update time.

Theorem 4.3 ([BCCK16, $\left.\left.\mathrm{CDW}^{+} 18\right]\right)$. Given a graph $G=(V, E)$, with $|V|=n,|E|=m$. There is a dynamic algorithm that maintains a DFS tree with $O(n)$ update time in the incremental model.

Using our reduction, one can immediately obtain
Theorem 4.4 (Dynamic DFS). Given a graph $G=(V, E)$, with $|V|=n,|E|=m$. There is a dynamic algorithm that maintains a DFS tree with $\widetilde{O}(n)$ worst case update time in the offline model.

### 4.3 All-Pair Shortest-Paths (APSP)

Dynamic APSP The APSP problem has been a central topic of graph algorithm. In a dynamic APSP problem, there is a undirected weighted graph $G=(V, E)(|V|=n,|E|=m)$ that subjects to edge insertion and deletion, The goal of the algorithm is to maintain an estimate $\delta(u, v)$ for every pair of node $u, v \in V$ that approximates the shortest path distance between $u$ and $v$. In the incremental setting, $\overline{\mathrm{CGH}^{+} 20}$ obtains an $O(1)$-approximate algorithm with $n^{o(1)}$ worst case update time.

Theorem 4.5 (Theorem 3.1 in $\left[\mathrm{CGH}^{+} 20\right]$ ). Let $G=(V, E)$ be a undirected weighted graph, there exists an incremental deterministic All-Pair Shortest-Paths algorithm that maintains $O(1)$ approximate shortest path in $n^{o(1)}$ worst case update time.

The offline model is of interest and it is already pointed out in [CGH $\left.{ }^{+} 20\right]$ that their data structure can be adapted to the offline model (but in a problem specific way). With Theorem 1.1 in hand, we can recover Theorem 4.8 of $\left[\mathrm{CGH}^{+} 20\right]$.

Theorem 4.6 (Dynamic APSP). Let $G=(V, E)$ be a undirected weighted graph, there exists a deterministic All-Pair Shortest-Paths algorithm that maintains $O(1)$ approximate shortest path in $n^{o(1)}$ worst case update time under the offline model.

## 5 Impossibility of general reduction

We prove both conditions (worst case guarantee and known deletion order) are indeed necessary to obtain a black box reduction.

Worst case guarantee is necessary A black box reduction is generally impossible if one only has amortized guarantee of incremental model. An example is the dynamic submodular maximization problem, where unconditional lower bound is known.

Theorem 5.1. Let $T \geq 1$ be the total number of updates. There exists a problem such that it is possible to have a dynamic algorithm with amortized update time $\Gamma_{u}$ in the incremental model, but any algorithm in the deletions-look-ahead model takes at $\Omega\left(\frac{T}{\log ^{4}(T)}\right) \cdot \Gamma_{u}$ amortized update time.

Proof. Let $[n]$ be the ground set and the total number of update be $T=O(n)$. By Theorem 1.3 of [CP22], there exists an algorithm with $O\left(\log (k / \epsilon) / \epsilon^{2}\right)$ amortized query complexity and maintains an $(1-1 / e-\epsilon)$-approximate solution for dynamic submodular maximization. While in the fully dynamic model (with known deletion order), by Theorem 1.2 of [CP22], no algorithm could maintain an 0.584-approximation with $o\left(n / k^{3}\right)$ amortized queries whenever $k=\Omega(\log n)$. Taking $k=C \log n$ for some constant $C>0$ exhibits a $\Omega\left(T / \log ^{4}(T)\right)$ separation.

Known deletion order is necessary If the deletion order is not known in advance, there exists a separation between the fully-dynamic and incremental model.

Theorem 5.2. Let $T \geq 1$ be the total number of updates. There exists a problem such that it is possible to have a dynamic algorithm with worst case update time $O(1)$ in the incremental model, but any algorithm in the fully dynamic model has amortized running time $\Omega\left(\frac{T}{\log (T)}\right)$.
Proof. Let $N=[n]$ be the ground set element and $T=2 n$ be the total number of updates. We first formalize the oracle model. For any subset of elements $A \subseteq N$, let $S(A) \in S$ be the transcript on $A$ and $Q(A) \in\{0,1\}$ be the answer to Query. There exists an oracle $O: S \times N \rightarrow S \times\{0,1\}$, it takes input of a transcript $S(A) \in S$ and an element $e \in N$, and returns the next transcript $S(A \cup\{e\}) \in S$ and the answer $Q(A \cup\{e\}) \in\{0,1\}$, i.e.,

$$
O(S(A), e)=(S(A \cup\{e\}), Q(A \cup\{e\})) .
$$

The oracle outputs empty when the input is invalid and we assume the oracle takes unit time.
It is clear that it takes only 1 oracle call per update in the incremental model. For the fully dynamic model, consider the update sequence of first inserting all elements in $N$ and delete them in random order. We prove $\Omega\left(n^{2} / \log n\right)$ oracle calls are necessary. For any $t \in[0: n / 100 \log n]$, let $N_{t}$ be the set of elements remain after deleting $4 t \log n$ elements and $N_{0}=N$. It suffices to prove the algorithm needs to make at least $\Omega(n)$ oracle queries between $N_{t}$ and $N_{t+1}$. Suppose the algorithm has the transcript $S\left(A_{1}\right), \ldots S\left(A_{\ell(t)}\right)$ at the beginning of the $(n+4 t \log n)$-th update, where $A_{i} \subseteq\left[N_{t}\right]$ and $\left|A_{i}\right| \geq n / 2(i \in[\ell(t)])$. We know that $\ell(t) \leq n^{2}$. After deleting the next $4 \log n$ elements (at random), the probability that $A_{i} \subseteq N_{t+1}$ is at most $1 / n^{4}$. Taking an union bound, with probability at least $1-1 / n^{2}$, none of the set satisfies $A_{i} \subseteq N_{t+1}$. Then we conclude that to get the transcript of $S\left(N_{t+1}\right)$, the algorithm needs at least $\left|N_{t+1}\right|-n / 2=\Omega(n)$ queries. We conclude the proof here.

## Acknowledgement

A.R. is grateful to Amir Abboud and Soheil Behnezhad for inspiring conversations. A.R. and B.P would like to thank Timothy Chan for explaining the interesting connections to related work from computational geometry.

## References

[ADK $\left.{ }^{+} 16\right]$ Ittai Abraham, David Durfee, Ioannis Koutis, Sebastian Krinninger, and Richard Peng. On fully dynamic graph sparsifiers. In Irit Dinur, editor, IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA, pages 335-344. IEEE Computer Society, 2016.
[AW14] Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014, pages 434-443. IEEE Computer Society, 2014.
[BCCK16] Surender Baswana, Shreejit Ray Chaudhury, Keerti Choudhary, and Shahbaz Khan. Dynamic dfs in undirected graphs: breaking the o (m) barrier. In Proceedings of the twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pages 730-739. SIAM, 2016.
$\left[\mathrm{BDM}^{+} 20\right]$ Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P Woodruff, and Samson Zhou. Near optimal linear algebra in the online and sliding window models. In 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), pages 517-528. IEEE, 2020.
[Ben73] Charles H Bennett. Logical reversibility of computation. IBM journal of Research and Development, 17(6):525-532, 1973.
[BKS12] Surender Baswana, Sumeet Khurana, and Soumojit Sarkar. Fully dynamic randomized algorithms for graph spanners. ACM Trans. Algorithms, 8(4):35:1-35:51, 2012.
[BS80] Jon Louis Bentley and James B Saxe. Decomposable searching problems i. static-todynamic transformation. Journal of Algorithms, 1(4):301-358, 1980.
[CDW ${ }^{+}$18] Lijie Chen, Ran Duan, Ruosong Wang, Hanrui Zhang, and Tianyi Zhang. An improved algorithm for incremental dfs tree in undirected graphs. In 16th Scandinavian Symposium and Workshops on Algorithm Theory, 2018.
[CGH ${ }^{+}$20] Li Chen, Gramoz Goranci, Monika Henzinger, Richard Peng, and Thatchaphol Saranurak. Fast dynamic cuts, distances and effective resistances via vertex sparsifiers. In 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), pages 1135-1146. IEEE, 2020.
[Cha12] Timothy M Chan. Three problems about dynamic convex hulls. International Journal of Computational Geometry and Applications, 22(4):341-364, 2012.
[CP22] Xi Chen and Binghui Peng. On the complexity of dynamic submodular maximization. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, pages 1685-1698, 2022.
[DGIM02] Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Maintaining stream statistics over sliding windows. SIAM journal on computing, 31(6):1794-1813, 2002.
[DS91] David Dobkin and Subhash Suri. Maintenance of geometric extrema. Journal of the ACM (JACM), 38(2):275-298, 1991.
[ELVZ17] Alessandro Epasto, Silvio Lattanzi, Sergei Vassilvitskii, and Morteza Zadimoghaddam. Submodular optimization over sliding windows. In Proceedings of the 26th International Conference on World Wide Web, pages 421-430, 2017.
[EMMZ22] Alessandro Epasto, Mohammad Mahdian, Vahab Mirrokni, and Peilin Zhong. Improved sliding window algorithms for clustering and coverage via bucketing-based sketches. In Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 3005-3042. SIAM, 2022.
[FLNF $\left.{ }^{+} 22\right]$ Moran Feldman, Paul Liu, Ashkan Norouzi-Fard, Ola Svensson, and Rico Zenklusen. Streaming submodular maximization under matroid constraints. In 49th International Colloquium on Automata, Languages, and Programming (ICALP 2022). Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
[HdLT01] Jacob Holm, Kristian de Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. J. ACM, 48(4):723-760, 2001.
[HHS22] Kathrin Hanauer, Monika Henzinger, and Christian Schulz. Recent advances in fully dynamic graph algorithms-a quick reference guide. ACM Journal of Experimental Algorithms (JEA), 2022.
[HK99] Monika Rauch Henzinger and Valerie King. Randomized fully dynamic graph algorithms with polylogarithmic time per operation. J. ACM, 46(4):502-516, 1999.
[HKNS15] Monika Henzinger, Sebastian Krinninger, Danupon Nanongkai, and Thatchaphol Saranurak. Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In Rocco A. Servedio and Ronitt Rubinfeld, editors, Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015, pages 21-30. ACM, 2015.
[HN21] Monika Henzinger and Alexander Noe. Random rank-based, hierarchical or trivial: Which dynamic graph algorithm performs best in practice? arXiv preprint arXiv:2108.04564, 2021.
[KMW98] Sanjeev Khanna, Rajeev Motwani, and Randall H Wilson. On certificates and lookahead in dynamic graph problems. Algorithmica, 21(4):377-394, 1998.
[KPP16] Tsvi Kopelowitz, Seth Pettie, and Ely Porat. Higher lower bounds from the 3sum conjecture. In Robert Krauthgamer, editor, Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016, pages 1272-1287. SIAM, 2016.
[LMNF ${ }^{+}$20] Silvio Lattanzi, Slobodan Mitrović, Ashkan Norouzi-Fard, Jakub M Tarnawski, and Morteza Zadimoghaddam. Fully dynamic algorithm for constrained submodular optimization. Advances in Neural Information Processing Systems, 33:12923-12933, 2020.
[OL81] Mark H Overmars and Jan van Leeuwen. Dynamization of decomposable searching problems yielding good worst-case bounds. In Theoretical Computer Science, pages 224-233. Springer, 1981.
[OvL81] Mark H Overmars and Jan van Leeuwen. Two general methods for dynamizing decomposable searching problems. Computing, 26(2):155-166, 1981.
[SM10] Piotr Sankowski and Marcin Mucha. Fast dynamic transitive closure with lookahead. Algorithmica, 56(2):180-197, 2010.
[vdBNS19] Jan van den Brand, Danupon Nanongkai, and Thatchaphol Saranurak. Dynamic matrix inverse: Improved algorithms and matching conditional lower bounds. In 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), pages 456-480. IEEE, 2019.
[WLFT18] Yanhao Wang, Yuchen Li, Ju Fan, and Kian-Lee Tan. Location-aware influence maximization over dynamic social streams. ACM Transactions on Information Systems (TOIS), 36(4):1-35, 2018.
[WZ22] David P Woodruff and Samson Zhou. Tight bounds for adversarially robust streams and sliding windows via difference estimators. In 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS), pages 1183-1196. IEEE, 2022.


[^0]:    *Supported by NSF CCF-1703925, IIS-1838154, CCF-2106429 and CCF-2107187.
    ${ }^{\dagger}$ Supported by NSF CCF-1954927, and a David and Lucile Packard Fellowship.

[^1]:    ${ }^{1}$ We believe that the update time can be improved at least to $O\left(\Gamma_{u} \log (T) \log \log (T)\right)$ at the expense of a more complicated algorithm.
    ${ }^{2}$ Equivalently, one can assume there exists an oracle that outputs the deletion order of existing ground set elements.

[^2]:    ${ }^{3}$ We thank Timothy Chan for pointing out this connection after we published the first version.

