Multidimensional Included and Excluded Sums



Helen Xu



Sean Fraser

ACDA June 21, 2021





Charles E. Leiserson



Tensor Region Sums

Several scientific computing applications involve reducing **many (potentially overlapping) regions** of a tensor to a single value for each region, using a binary associative operator \oplus .

Tensor Region Sums

Several scientific computing applications involve reducing many (potentially overlapping) regions of a tensor to a single value for each region, using a binary associative operator \oplus .

Inclusion: The summed-area table (SAT) method preprocesses an image to answer queries about the sum in rectangular subregions of a tensor [C84].



Tensor Region Sums

Several scientific computing applications involve reducing many (potentially overlapping) regions of a tensor to a single value for each region, using a binary associative operator \oplus .

Inclusion: The summed-area table (SAT) method preprocesses an image to answer queries about the sum in rectangular subregions of a tensor [C84].

Exclusion: The essence of the **fast multipole method** (FMM) is a reduction of a subregion's elements, excluding elements too close [BG97, CRW93, D00].

Summed-area Table





\oplus

Fast Multipole Method



Excluded-sums Problem

The excluded-sums problem [DDELP05] underlies applications that require reducing regions of a tensor to a single value using \bigoplus .

Excluded-sums Problem

The excluded-sums problem [DDELP05] underlies applications that require reducing regions of a tensor to a single value using \bigoplus .

In 2D, it takes as input an $n_1 \times n_2$ matrix A and "box size" $\mathbf{k} = (k_1, k_2)$ where $k_1 \leq n_1, k_2 \leq n_2$.

 n_1

 $N = n_1 n_2$



Excluded-sums Problem

The excluded-sums problem [DDELP05] underlies applications that require reducing regions of a tensor to a single value using \bigoplus .

In 2D, it takes as input an $n_1 \times n_2$ matrix A and "box size" $\mathbf{k} = (k_1, k_2)$ where $k_1 \le n_1, k_2 \le n_2$.

The problem involves reducing the excluded region outside of every $k\mbox{-box}$ in the matrix.



Included-sums Problem

problem.

In 2D, the included sum at coordinate (x_1, x_2) involves reducing (accumulating with \oplus) all elements in the k-box cornered at (x_1, x_2) .



- The included-sums problem takes the same input as the excluded-sums



Included-sums Problem

problem.

In 2D, the included sum at coordinate (x_1, x_2) involves reducing (accumulating with \oplus) all elements in the k-box cornered at (x_1, x_2) .



- The included-sums problem takes the same input as the excluded-sums



Included-sums Problem

problem.

In 2D, the included sum at coordinate (x_1, x_2) involves reducing (accumulating with \oplus) all elements in the k-box cornered at (x_1, x_2) .

Can be computed straightforwardly with four nested loops in $\Theta(n_1n_2k_1k_2)$ time.

 n_1

 $N = n_1 n_2$

- The included-sums problem takes the same input as the excluded-sums



Inclusion and Exclusion Example



Inclusion



Inclusion and Exclusion Example





Inclusion and Exclusion Example



We present an example with addition for ease of understanding, but in general an algorithm for these problems should work with general operators.



Included and Excluded Sums With and Without Operator Inverse



Included and Excluded Sums With and Without Operator Inverse



FMM's functions, which may exhibit singularities [DDELP05].

This approach fails for operators without inverse such as max, or the

Included and Excluded Sums With and Without Operator Inverse



FMM's functions, which may exhibit singularities [DDELP05].

We refine the included- and excluded-sums problems into weak and strong version does not.

- This approach fails for operators without inverse such as max, or the
- strong versions. The weak version requires an operator inverse, while the

State of the Art for Included and **Excluded Sums**

Given a d-dimensional tensor with N elements:

Weak/ Problem Algorithm

Summed-area table Included W

/Strong	Time	Space		
eak 😐	$\Theta(2^d N)$	$\Theta(N)$		

State of the Art for Included and **Excluded Sums**

Given a d-dimensional tensor with N elements:

Algorithm	Problem	Weak/Strong	Time	Space
Summed-area table	Included	Weak 😐	$\Theta(2^d N)$	$\Theta(N)$
Corners(c)	Excluded	Strong 💛	$\Omega(2^d N)$	$\Theta(cN)$

Our Contributions

Given a d-dimensional tensor with N elements: Weak/ Problem Algorithm

Summed-area table Included We

> Corners(c) Excluded Str

This work

Bidirectional Included Str box-sum (BDBS)

/Strong	Time	Space
eak 😐	$\Theta(2^d N)$	$\Theta(N)$
ong 💗	$\Omega(2^d N)$	$\Theta(cN)$
ong 💓	$\Theta(dN)$	$\Theta(N)$

Our Contributions

Given a d -dimensional tensor with N elements:						
Algorithm	Problem	Weak/Strong	Time	Space		
Summed-area table	Included	Weak 😐	$\Theta(2^d N)$	$\Theta(N)$		
Corners(c)	Excluded	Strong 💛	$\Omega(2^d N)$	$\Theta(cN)$		
This work						
Bidirectional box-sum (BDBS)	Included	Strong 💗	$\Theta(dN)$	$\Theta(N)$		
Box complement	Excluded	Strong 💛	$\Theta(dN)$	$\Theta(N)$		

Strong Excluded Sums in 3D





We will start with the bidirectional box-sum algorithm (BDBS) in one dimension then show how to extend the technique to higher dimensions.

Given a list A of length N and a (scalar) box size k, output a list A' of included sums. k = 4

Position	1	2	3	4	5	6	7	8
Input A	2	5	3	1	6	3	9	0

We will start with the bidirectional box-sum algorithm (BDBS) in one dimension then show how to extend the technique to higher dimensions.

Given a list A of length N and a (scalar) box size k, output a list A' of included sums. *k* = 4 1 2 3 4 5 6 Position 8 5 3 6 3 9 1 0 Input $\blacktriangleright A$ 2 10 11 6 9 18 18 **Prefix** 7 2 18 12 4 9 Suffix 9 0 Target A' 11 15 13 19 18 12 9 0

Compute intermediate prefix and suffix arrays with N/k prefixes and suffixes of length k each.

We will start with the bidirectional box-sum algorithm (BDBS) in one dimension then show how to extend the technique to higher dimensions.

Given a list A of length N and a (scalar) box size k, output a list A' of included sums. *k* = 4 1 2 3 4 5 6 Position 8 5 3 6 3 9 1 0 Input $\blacktriangleright A$ 2 10 11 6 9 18 18 **Prefix** 7 2 18 12 9 4 Suffix 9 0 Target A' 11 15 13 19 18 12 9 0

Compute intermediate prefix and suffix arrays with N/k prefixes and suffixes of length k each.

We will start with the bidirectional box-sum algorithm (BDBS) in one

Given a list A of length N and a (scalar) box size k, output a list A' of included sums. *k* = 4 1 2 3 4 5 6 Position 8 7 5 3 9 6 3 Input $\blacktriangleright A$ 0 2 10 11 6 18 18 9 Prefix 18 12 9 Suffix 4 9 0 11 15 13^{+,*} 19 18 12 9 Target > 0

dimension then show how to extend the technique to higher dimensions.





Multidimensional Bidirectional Box Sum

The BDBS technique extends into arbitrary dimensions by performing the prefixes and suffixes along each dimension in turn.

Bidirectional box sum along first dimension



Multidimensional Bidirectional Box Sum

The BDBS technique extends into arbitrary dimensions by performing the prefixes and suffixes along each dimension in turn.

Bidirectional box sum along first dimension



Bidirectional box sum along second dimension on result



Multidimensional Bidirectional Box Sum

The BDBS technique extends into arbitrary dimensions by performing the prefixes and suffixes along each dimension in turn.

Bidirectional box sum along first dimension



Given a *d*-dimensional tensor with *N* elements, BDBS solves the strong included-sums problem in $\Theta(dN)$ time and $\Theta(N)$ space.

Bidirectional box sum along second dimension on result



Formulating the Excluded Sum as the Box Complement





Given a *d*-dimensional tensor and a "box size", we will first sketch how to decompose the excluded region for each point into 2d disjoint regions.

Formulating the Excluded Sum as the Box Complement

At a high level, the "*i*-complement" of a box such that there is some the coordinates are "in range" for all dimensions $m \in [i + 1, d]$.





- Given a d-dimensional tensor and a "box size", we will first sketch how to decompose the excluded region for each point into 2d disjoint regions.
- coordinate in dimension $j \in [1,i]$ that is "out of range" in dimension j, and

Formulating the Excluded Sum as the Box Complement

At a high level, the "*i*-complement" of a box such that there is some the coordinates are "in range" for all dimensions $m \in [i + 1, d]$.



Dimension 2

- Given a d-dimensional tensor and a "box size", we will first sketch how to decompose the excluded region for each point into 2d disjoint regions.
- coordinate in dimension $j \in [1,i]$ that is "out of range" in dimension j, and



Box Complement Algorithm for Strong Excluded Sums

The box-complement algorithm uses dimension reduction to compute the "*i*-complement" for all i = 1, ..., d.

The **BDBS** algorithm for included sums is a major subroutine in the boxcomplement algorithm to sum up elements "in the range."



each row

each row

Box Complement Algorithm for Strong Excluded Sums

The box-complement algorithm uses **dimension reduction** to compute the "*i*-complement" for all i = 1, ..., d.

The **BDBS algorithm** for included sums is a major subroutine in the boxcomplement algorithm to sum up elements "in the range."



each row

each row

Suffix







3





Each dimension-reduction step takes $\Theta(N)$ time and reuses the same temporaries, for a total of $\Theta(dN)$ time and $\Theta(N)$ space.

Conclusion

This work introduces the box-complement algorithm for the excluded-sums problem, which improves the running time of the state-of-the-art corners algorithm from $\Omega(2^d N)$ to $\Theta(dN)$.

Conclusion

This work introduces the box-complement algorithm for the excluded-sums problem, which improves the running time of the state-of-the-art corners algorithm from $\Omega(2^d N)$ to $\Theta(dN)$.

Since floating-point subtraction is less numerically stable than addition, our next step is to study the **numerical accuracy** of the different algorithms.

Conclusion

This work introduces the box-complement algorithm for the excluded-sums problem, which improves the running time of the state-of-the-art corners algorithm from $\Omega(2^d N)$ to $\Theta(dN)$.

Since floating-point subtraction is less numerically stable than addition, our next step is to study the **numerical accuracy** of the different algorithms.

Future work includes incorporating the BDBS and box-complement algorithm into the FMM.

