# GPU-Efficient Recursive Filtering and Summed-Area Tables 

Diego Nehab ${ }^{1} \quad$ André Maximo ${ }^{1}$<br>${ }^{1}$ IMPA $\quad{ }^{2}$ Digitok

Rodolfo S. Lima ${ }^{2} \quad$ Hugues Hoppe ${ }^{3}$<br>${ }^{3}$ Microsoft Research


#### Abstract

Image processing operations like blurring, inverse convolution, and summed-area tables are often computed efficiently as a sequence of 1D recursive filters. While much research has explored parallel recursive filtering, prior techniques do not optimize across the entire filter sequence. Typically, a separate filter (or often a causal-anticausal filter pair) is required in each dimension. Computing these filter passes independently results in significant traffic to global memory, creating a bottleneck in GPU systems. We present a new algorithmic framework for parallel evaluation. It partitions the image into 2D blocks, with a small band of additional data buffered along each block perimeter. We show that these perimeter bands are sufficient to accumulate the effects of the successive filters. A remarkable result is that the image data is read only twice and written just once, independent of image size, and thus total memory bandwidth is reduced even compared to the traditional serial algorithm. We demonstrate significant speedups in GPU computation.




## 1 Introduction

Linear filtering (i.e. convolution) is commonly used to blur, sharpen, or downsample images. A direct implementation evaluating a filter of support $d$ on an $h \times w$-image has cost $O(h w d)$. For filters with a wide impulse response, the Fast Fourier Transform reduces the cost to $O(h w \log h w)$, regardless of filter support. Often, similar results can be obtained with a recursive filter, in which the computation reuses prior outputs, e.g. $y_{i}=x_{i}-\frac{1}{2} y_{i-1}$. Such feedback allows for an infinite impulse response (IIR), i.e. an effectively large filter support, at reduced cost $O(h w r)$, where the number $r$ of recursive feedbacks (a.k.a. the filter order) is small relative to $d$. Recursive filters are a key computational tool in several applications:

- Low-pass filtering. Filters like Gaussian kernels are well approximated by a pair of low-order causal and anticausal recursive filters [e.g. Deriche 1992; van Vliet et al. 1998].
- Inverse convolution. If an image $X$ is the result of convolving an image $V$ with a compactly supported filter $F$, i.e. $X=V * F$, the original image can be recovered as $V=X * F^{-1}$. Although the inverse filter $F^{-1}$ generally has infinite support, it can be expressed exactly as a sequence of low-order recursive filters.
- Summed-area tables. Such tables store the sum of all pixel values above and to the left of each pixel [Crow 1984]. They have many uses in graphics and vision. On a GPU, summed-area tables are typically computed with prefix sums over all columns then all rows of the image [Hensley et al. 2005]. Crucially, a prefix sum is a special case of a 1D first-order recursive filter.

These applications all have in common the fact that they invoke a sequence of recursive filters. First, a 2D operation is decomposed into separate 1D filters. Second, except for the case of summedarea tables, one usually desires a centered and well-shaped impulse response function, and this requires the combination of a causal and anticausal filter pair in each dimension.
As an example, consider the problem of finding the coefficient image $V$ such that reconstruction with the bicubic B -spline interpolates the pixel values of a given image $X$. (Such a preprocess is required in techniques for high-quality image filtering [Blu et al. 1999, 2001].) The coefficients and image are related by the convolution $X=V * \frac{1}{6}\left[\begin{array}{lll}1 & 4 & 1\end{array}\right] * \frac{1}{6}\left[\begin{array}{lll}1 & 4 & 1\end{array}\right]^{T}$. The inverse convolution can be decomposed into four successive 1D recursive filters:

$$
\begin{aligned}
y_{i, j} & =6 x_{i, j}+\alpha y_{i-1, j} \\
z_{i, j} & =-\alpha y_{i, j}+\alpha z_{i+1, j} \\
u_{i, j} & =6 z_{i, j}+\alpha u_{i, j-1} \\
v_{i, j} & =-\alpha u_{i, j}+\alpha v_{i, j+1}, \text { with } \alpha=\sqrt{3}-2
\end{aligned}
$$

(causal in $i$ )
(anticausal in $i$ )
(causal in $j$ )
(anticausal in $j$ )
Although the parallelization of such recursive filters is challenging due to the many feedback data dependencies, efficient algorithms have been developed as reviewed in section 2 .
However, an important element not considered in previous work is optimization over the entire sequence of recursive filters. If each filter is applied independently as is common, the entire image must be read from global memory, and then written back each time. Due to the low computational intensity of recursive filters, this memory traffic becomes a clear bottleneck. Because the data access pattern is different in successive filters (down, up, right, left), the different computation passes cannot be directly fused.

Contribution In this paper, we present a new algorithmic framework to reduce memory bandwidth by overlapping computation over the full sequence of recursive filters, across both multiple dimensions and causal-anticausal pairs. The resulting algorithm scales efficiently to images of arbitrary size.
Our approach partitions the image into 2D blocks of size $b \times b$. This blocking serves several purposes. Inter-block parallelism provides sufficient independent tasks to hide memory latency. The blocks fit in on-chip memory for fast intra-block computations. And, the square blocks enables efficient propagation of data in both dimensions.
We associate a set of narrow rectangular buffers with each block's perimeter. These buffers' widths are the same as the filter order $r$. Intuitively, the goal of these buffers is to measure the filter response of each block in isolation (i.e. given zero boundary conditions), as well as to efficiently propagate the boundary conditions across the blocks given these measured responses.
The surprising result is that we are able to accumulate, with a single block-parallel pass over the input image, all the necessary responses to the filter sequence (down, up, right, left), so as to efficiently initialize boundary conditions for a second block-parallel pass that produces the final solution. Thus, the source image is read only twice, and written just once. A careful presentation requires algebraic derivations, and is detailed in sections 4-5.
Our framework also benefits the computation of summed-area tables (section 6). Even though in that setting there are no anticausal filters, the reduction in bandwidth due solely to the overlapped processing of the two dimensions still provides performance gains.

## 2 Related work

Prefix sums and scans A prefix sum, $y_{i}=x_{i}+y_{i-1}$, is a simple case of a first-order recursive filter. A scan generalizes the recurrence using an arbitrary binary associative operator. Parallel prefix sums and scans are important building blocks for numerous algorithms [Iverson 1962; Stone 1971; Blelloch 1989; Sengupta et al. 2007]. Recent work has explored the efficient GPU implementation of scan operations using optimization strategies such as hierarchical reduction, bandwidth reduction through redundant computation, and serial evaluation within shared memory [Sengupta et al. 2007; Dotsenko et al. 2008; Merrill and Grimshaw 2009]. An optimized implementation comes with the CUDPP library [2011].

Recursive filtering Recursive filters generalize prefix sums by considering a weighted combination of prior outputs as indicated in equation (1). This generalization can in fact be implemented as a scan operation with appropriately redefined basic operators [Blelloch 1990]. For recursive filtering of images on the GPU, Ruijters and Thévenaz [2010] exploit parallelism across rows and columns. To expose parallelism within each row and each column, we build on the work of Sung and Mitra [1986, 1992]. The idea is to partition the input into blocks and decompose the computation over each block into two parts: one based only on the block data and assuming zero initial conditions, and the other based only on initial conditions and assuming zero block data (see section 4). We then overlap computation between successive filters passes. As noted by Sung and Mitra, there is also related work on pipelined circuit evaluation of recursive filters [e.g. Parhi and Messerschmitt 1989], but such approaches are more expensive on general-purpose computers.

Tridiagonal systems The forward- and back-substitution passes associated with solving banded linear systems by LU-decomposition are related to recursive filters, particularly for Toeplitz matrices, although the resulting $L$ and $U$ matrices have nonuniform entries at least near the domain boundaries. Forward- and back-substitution passes can be implemented with parallel scan operations [Blelloch 1990]. On GPUs, fast tridiagonal solvers process rows and columns independently with cyclic reduction [Hockney and Jesshope 1981], using either pixel-shader passes over global memory [Kass et al. 2006], or within shared memory using a single kernel [Zhang et al. 2010; Göddeke and Strzodka 2011]. The use of shared memory imposes a limit on the maximum image size. Recent work explores overcoming this limitation [Lamas-Rodrígues et al. 2011]. Our blocking strategy is able to overlap causal and anticausal computations to run efficiently on images of arbitrary size. We hope it can be applied it to general tridiagonal systems in future work.

Summed-area tables The typical approach for efficiently generating summed-area tables on the GPU is to compute parallel prefix sums over all columns, and then over all rows of the result. These prefix sums can be computed with recursive doubling using multiple passes of a pixel shader [Hensley et al. 2005]. As previously mentioned, recent prefix-sum algorithms reduce bandwidth by using GPU shared memory [Harris et al. 2008; Hensley 2010]. In fact, summed-area tables can be computed with little effort using the high-level parallel primitives of the CUDPP library and publicly available matrix transposition routines.

## 3 Problem definition

Causal recursive filters of order $r$ are characterized by a set of $r$ feedback coefficients $a_{k}$ in the following manner. Given a prologue vector $\boldsymbol{p} \in \mathbb{R}^{r}$ (i.e., the initial conditions) and an input vector $\boldsymbol{x} \in \mathbb{R}^{h}$, of any size $h$, the filter $F: \mathbb{R}^{r} \times \mathbb{R}^{h} \rightarrow \mathbb{R}^{h}$ produces the output vector $\boldsymbol{y}=F(\boldsymbol{p}, \boldsymbol{x})$ with the same size as the input $\boldsymbol{x}$ :

$$
\begin{equation*}
y_{i}=x_{i}-\sum_{k=1}^{r} a_{k} y_{i-k}, \quad i \in\{0, \ldots, h-1\} \tag{1}
\end{equation*}
$$

where on the right-hand side of (1) we set as initial conditions

$$
\begin{equation*}
y_{-k}=p_{r-k}, \quad k \in\{1, \ldots, r\} . \tag{2}
\end{equation*}
$$

The definition of an anticausal filter $R: \mathbb{R}^{h} \times \mathbb{R}^{r} \rightarrow \mathbb{R}^{h}$ of order $r$ is analogous. Given input vector $\boldsymbol{y} \in \mathbb{R}^{h}$ and epilogue vector $\boldsymbol{e} \in \mathbb{R}^{r}$, the output vector $\boldsymbol{z}=R(\boldsymbol{y}, \boldsymbol{e})$ is defined by

$$
\begin{equation*}
z_{i}=y_{i}-\sum_{k=1}^{r} a_{k}^{\prime} z_{i+k}, \quad i \in\{0, \ldots, h-1\}, \tag{3}
\end{equation*}
$$

where on the right-hand side of (3) we set as initial conditions

$$
\begin{equation*}
z_{h+k-1}=e_{k-1}, \quad k \in\{1, \ldots, r\} . \tag{4}
\end{equation*}
$$

Causal and anticausal filters frequently appear in pairs. In such cases, we are interested in computing

$$
\begin{equation*}
\boldsymbol{z}=R(\boldsymbol{y}, \boldsymbol{e})=R(F(\boldsymbol{p}, \boldsymbol{x}), \boldsymbol{e}) \tag{5}
\end{equation*}
$$

In image processing applications, it is often desirable to apply filters independently to each row and to each column of an image (i.e., in a separable way). To process all image columns, we extend the definition of $F$ to operate independently on all corresponding columns of an $r \times w$ prologue matrix $P$ and an $h \times w$ input matrix $X$. The extended filter is described by $F: \mathbb{R}^{r \times w} \times \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^{h \times w}$.
For row processing, we define a new extended filter $F^{\mathcal{T}}$ that operates independently on corresponding rows of an $h \times w$ input matrix $X$ and an $h \times r$ prologue matrix $P^{\prime}$. This transposed filter is described by $F^{\tau}: \mathbb{R}^{h \times r} \times \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^{h \times w}$. Anticausal filters can also be extended to operate independently over multiple columns, $R: \mathbb{R}^{h \times w} \times \mathbb{R}^{r \times w} \rightarrow \mathbb{R}^{h \times w}$, and over multiple rows, $R^{\mathcal{T}}: \mathbb{R}^{h \times w} \times \mathbb{R}^{h \times r} \rightarrow \mathbb{R}^{h \times w}$.

With these definitions, we are ready to formulate the problem of applying the full sequence of four recursive filters (down, up, right, left) to an image. As shown in figure 1, given an input matrix $X$ and the prologues $\left(P, P^{\prime}\right)$ and epilogues $\left(E, E^{\prime}\right)$ that surround its columns and rows, we wish to compute matrix $V$, where

$$
\begin{array}{ll}
Y=F(P, X), & Z=R(Y, E) \\
U=F^{\mathcal{T}}\left(P^{\prime}, Z\right), & V=R^{\mathcal{T}}\left(U, E^{\prime}\right) \tag{6}
\end{array}
$$

In particular, we want to take advantage of the massive parallelism offered by modern GPUs.


Figure 1: The problem is to compute for an image $X$ a sequence of four $1 D$ recursive filters (down, up, right, left), with boundary conditions given by prologue $P, P^{\prime}$ and epilogue $E, E^{\prime}$ matrices.

Design goals on GPUs Modern GPUs are composed of $p$ streaming multiprocessors (SMs), each containing $c$ computational cores. For example, our test hardware (an NVIDIA GTX 480) contains $p=15$ SMs, each with $c=32$ cores. Because the cores in each SM execute instructions in lockstep, divergence in execution paths should be minimized for maximum performance. SMs execute independently from each other, but typically run the same program, called a kernel, operating over distinct parts of the input.
Although bandwidth to global memory is higher than on a CPU, access must be coalesced for maximum performance. Ideally, all cores should access aligned, consecutive memory addresses. Furthermore, hiding memory latency requires a large number of independent tasks to be available for scheduling. Each SM contains a small amount of fast memory ( 48 KB in our case) shared among its cores. This memory is typically used as a managed cache to reduce the number of expensive accesses to global memory.
Mapping recursive filtering to GPU architectures presents several challenges. One is the dependency chain in the computation, which reduces the number of tasks that can be executed independently. The other problem is the high bandwidth-to-FLOP ratio, particularly in low-order filters. Without a large number of schedulable tasks to hide memory latency, processing cores become idle waiting for memory transactions to complete, and performance suffers.
Given these considerations, the main design goal of our work is to increase the amount of parallelism without increasing the total memory bandwidth required to complete the computation.

## 4 Prior parallelization strategies revisited

In large images, the independence of column and row processing offers a significant level of parallelism. With that in mind, here is the approach followed by Ruijters and Thévenaz [2010]:

## Algorithm RT

RT. 1 In parallel for each column in $X$, apply $F$ sequentially according to (1). Store $Y$.
RT. 2 In parallel for each column in $Y$, apply $R$ sequentially according to (3). Store $Z$.
RT. 3 In parallel for each row in $Z$, apply $F^{\mathcal{\tau}}$ sequentially according to (1). Store $U$.
RT. 4 In parallel for each row in $U$, apply $R^{\top}$ sequentially according to (3). Store $V$.
Algorithm RT uses bandwidth proportional to $8 h w$, and completes in $O\left(\frac{h w}{c p} 4 r\right)$ steps. The number of tasks that can be performed in parallel is limited by $h$ or $w$. In a GPU with 15 SMs and 32 cores per SM, we need at least 480 rows and columns to keep all cores busy. Much larger images are required for effective latency hiding.

Roadmap for improvements A common approach to improving performance is to restructure the computation so memory accesses are coalesced (during row as well as column processing). This is one of the reasons we employ a blocking strategy (section 4.1). We introduce inter-block parallelism by adapting the method of Sung and Mitra [1986] to modern GPU architectures (section 4.3). The inter-block parallelism also lets us fuse computations, thereby reducing traffic to global memory (section 4.4).
The fundamental contributions of our work begin in section 5.1, where we develop an overlapping strategy for merging causal and anticausal filtering passes. The idea culminates in section 5.2 , where we present an overlapped algorithm that merges row and column processing as well. Finally, in section 6 we specialize it to compute summed-area tables. The resulting reduction in memory bandwidth combined with the high degree of parallelism enable our overlapped algorithms to outperform prior approaches.

(a)

(b)

Figure 2: (a) 2D block notation showing a block and its boundary data from adjacent blocks. (b) 1D block notation, showing vectors as rows for layout reasons.

### 4.1 Block notation

We partition the image into blocks of $b \times b$ elements, where $b$ matches the number of threads executed in lockstep (the warp size) within an SM, i.e., $b=c=32$ throughout our work.
As shown in figure 2(a), we define operators to select a block and its boundary data in adjacent blocks. $B_{m, n}(X)$ identifies in matrix $X$ the $b \times b$ data block with index $(m, n) . P_{m-1, n}(X)$ extracts the $r \times b$ submatrix in the same shape and position as the column-prologue to block $B_{m, n}(X)$. Similarly, $E_{m+1, n}(X)$ denotes the $r \times b$ submatrix in the position of the column-epilogue to block $B_{m, n}(X)$. For row processing, transposed operators $P_{m, n-1}^{\tau}(X)$ and $E_{m, n+1}^{\tau}(X)$ are defined accordingly. ${ }^{1}$
It will often be useful to select prologue- and epilogue-shaped submatrices from a $b \times b$ block. To that end, we define a tail operation $T$ (for prologues) and a head operation $H$ (for epilogues). Transposed counterparts $T^{\mathcal{T}}$ and $H^{\mathcal{T}}$ are defined analogously:

$$
\begin{align*}
& P_{m, n}(X)=T\left(B_{m, n}(X)\right), \quad E_{m, n}(X)=H\left(B_{m, n}(X)\right) \\
& P_{m, n}^{\mathcal{T}}(X)=T^{\mathcal{T}}\left(B_{m, n}(X)\right), \quad E_{m, n}^{\mathcal{T}}(X)=H^{\mathcal{T}}\left(B_{m, n}(X)\right) \tag{7}
\end{align*}
$$

For simplicity, when describing unidimensional operations, we will drop the vertical block index $n$, as shown in figure 2(b).
Given these definitions, the blocked version of the problem formulation in (6) is to compute all output blocks $B_{m, n}(V)$ where:

$$
\begin{align*}
& B_{m, n}(Y)=F\left(P_{m-1, n}(Y), B_{m, n}(X)\right) \\
& B_{m, n}(Z)=R\left(B_{m, n}(Y), E_{m+1, n}(Z)\right)  \tag{8}\\
& B_{m, n}(U)=F^{\mathcal{T}}\left(P_{m, n-1}^{\tau}(U), B_{m, n}(Z)\right) \\
& B_{m, n}(V)=R^{\mathcal{T}}\left(B_{m, n}(U), E_{m, n+1}^{\mathcal{\tau}}(V)\right)
\end{align*}
$$

All algorithms described from now on employ blocking to guarantee coalesced memory access patterns. Typically, each stage loads an entire block of data to shared memory, performs computation in shared memory, and finally writes the data back to global memory.

### 4.2 Important properties

We first introduce key properties of recursive filtering that will significantly simplify the presentation. The first property of interest is superposition, and it comes from linearity. The effects of the input and the prologue (or epilogue) on the output can be computed independently and later added together:
(10)

$$
\begin{align*}
F(\boldsymbol{p}, \boldsymbol{x}) & =F(\mathbf{0}, \boldsymbol{x})+F(\boldsymbol{p}, \mathbf{0})  \tag{9}\\
R(\boldsymbol{y}, \boldsymbol{e}) & =R(\boldsymbol{y}, \mathbf{0})+R(\mathbf{0}, \boldsymbol{e})
\end{align*}
$$

[^0]Notice that the first term $F(\mathbf{0}, \boldsymbol{x})$ does not depend on the prologue, whereas the second term $F(\boldsymbol{p}, \mathbf{0})$ does not depend on the input. Similar remarks hold for $R$. We will use superposition extensively to split dependency chains in the computation.
Because filtering operations are linear, we can express them as matrix products. This will help reveal useful associative relations. For any $d$, let $I_{d}$ denote the identity matrix of size $d \times d$. Then,

$$
\begin{array}{ll}
F(\boldsymbol{p}, \mathbf{0})=F\left(I_{r}, \mathbf{0}\right) \boldsymbol{p}=A_{F P} \boldsymbol{p}, & A_{F P}=F\left(I_{r}, \mathbf{0}\right) \\
R(\mathbf{0}, \boldsymbol{e})=R\left(\mathbf{0}, I_{r}\right) \boldsymbol{e}=A_{R E} \boldsymbol{e}, & A_{R E}=R\left(\mathbf{0}, I_{r}\right), \\
F(\mathbf{0}, \boldsymbol{x})=F\left(\mathbf{0}, I_{b}\right) \boldsymbol{x}=A_{F B} \boldsymbol{x}, & A_{F B}=F\left(\mathbf{0}, I_{b}\right), \text { and }  \tag{11}\\
R(\boldsymbol{y}, \mathbf{0})=R\left(I_{b}, \mathbf{0}\right) \boldsymbol{y}=A_{R B} \boldsymbol{y}, & A_{R B}=R\left(I_{b}, \mathbf{0}\right)
\end{array}
$$

Here, $A_{F B}$ and $A_{R B}$ are $b \times b$ matrices, $A_{F P}$ is $b \times r$, and $A_{R E}$ is $b \times r$. Notice that whereas the columns of $A_{F B}$ and $A_{R B}$ are shifted impulse responses (which allows for efficient storage), the columns of $A_{F P}$ and $A_{R E}$ are generally not shifts of each other.
Row processing is analogous, except for transposition. This is the justification for our notation choice:

$$
\begin{array}{ll}
F^{\mathcal{T}}\left(\boldsymbol{p}^{T}, \mathbf{0}\right)=\boldsymbol{p}^{T}\left(A_{F P}\right)^{T}, & R^{\mathcal{T}}\left(\mathbf{0}, \boldsymbol{e}^{T}\right)=\boldsymbol{e}^{T}\left(A_{R E}\right)^{T}  \tag{12}\\
F^{\mathcal{T}}\left(\mathbf{0}, \boldsymbol{x}^{T}\right)=\boldsymbol{x}^{T}\left(A_{F B}\right)^{T}, & R^{\mathcal{T}}\left(\boldsymbol{y}^{T}, \mathbf{0}\right)=\boldsymbol{y}^{T}\left(A_{R B}\right)^{T}
\end{array}
$$

A trivial property of matrix product lets us write

$$
\begin{equation*}
T(A \boldsymbol{p})=T(A) \boldsymbol{p} \quad \text { and } \quad H(A \boldsymbol{e})=H(A) \boldsymbol{e} \tag{13}
\end{equation*}
$$

Finally, as shown in appendix A,

$$
\begin{equation*}
T\left(A_{F P}\right)=A_{F}^{b} \quad \text { and } \quad H\left(A_{R E}\right)=A_{R}^{b} \tag{14}
\end{equation*}
$$

where $A_{F}^{b}, A_{R}^{b}$ are precomputed $r \times r$-matrices that depend only on the feedback coefficients $a_{k}, a_{k}^{\prime}$ of filters $F$ and $R$, respectively.

### 4.3 Inter-block parallelism

We start the description in one dimension for simplicity. When attempting to perform block computations independently, we are faced with the difficulty that each output block

$$
\begin{equation*}
B_{m}(\boldsymbol{y})=F\left(P_{m-1}(\boldsymbol{y}), B_{m}(\boldsymbol{x})\right) \tag{15}
\end{equation*}
$$

depends on the prologue $P_{m-1}(\boldsymbol{y})$, i.e., on the tail end of the previous output block, $T\left(B_{m-1}(\boldsymbol{y})\right)$. This creates a dependency chain that we must work around.

By the superposition property (9), we can decompose (15) as

$$
\begin{equation*}
B_{m}(\boldsymbol{y})=F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right)+F\left(P_{m-1}(\boldsymbol{y}), \mathbf{0}\right) \tag{16}
\end{equation*}
$$

The first term can be computed independently for each block (interblock parallelism) because it assumes a zero prologue. We use symbol $\overline{\boldsymbol{y}}$ to denote this incomplete causal output:

$$
\begin{equation*}
B_{m, n}(\overline{\boldsymbol{y}})=F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right), \tag{17}
\end{equation*}
$$

where each block is filtered as if its prologue was null.
The second term of (16) is more challenging due to the dependency chain. Nevertheless, as soon as the dependency is satisfied, each element in a block can be computed independently (intra-block parallelism). To see this, apply property (11):

$$
\begin{align*}
F\left(P_{m-1}(\boldsymbol{y}), \mathbf{0}\right) & =F\left(I_{r}, \mathbf{0}\right) P_{m-1}(\boldsymbol{y})  \tag{18}\\
& =A_{F P} P_{m-1}(\boldsymbol{y}) \tag{19}
\end{align*}
$$

and recall that $A_{F P}$ depends only on the feedback coefficients.

Sung and Mitra [1986] perform unidimensional recursive filtering on a CRAY X-MP computer. Their algorithm works in two stages. First, each processor works on a separate block, computing $B_{m}(\overline{\boldsymbol{y}})$ sequentially. Second, the blocks are corrected sequentially, according to (16) and (19), with each block's elements updated in parallel across processors. Although this algorithm requires twice the bandwidth and computation of the sequential algorithm, the amount of available parallelism is increased by a factor $h / b$ (i.e., the number of blocks) in the first stage, and by a factor $b$ in the second stage.
Inspired by the bandwidth reduction strategy described by Dotsenko et al. [2008] in the context of scan primitives, we proceed differently. Notice that the dependency chain only involves prologues. Using properties (11), (13), and (14), we can make this apparent:

$$
\begin{align*}
P_{m}(\boldsymbol{y}) & =T\left(B_{m}(\boldsymbol{y})\right)=T\left(F\left(P_{m-1}(\boldsymbol{y}), B_{m}(\boldsymbol{x})\right)\right)  \tag{20}\\
& =T\left(F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right)\right)+T\left(F\left(P_{m-1}(\boldsymbol{y}), \mathbf{0}\right)\right)  \tag{21}\\
& =T\left(B_{m}(\overline{\boldsymbol{y}})\right)+T\left(F\left(I_{r}, \mathbf{0}\right) P_{m-1}(\boldsymbol{y})\right)  \tag{22}\\
& =P_{m}(\overline{\boldsymbol{y}})+T\left(A_{F P}\right) P_{m-1}(\boldsymbol{y})  \tag{23}\\
& =P_{m}(\overline{\boldsymbol{y}})+A_{F}^{b} P_{m-1}(\boldsymbol{y}) . \tag{24}
\end{align*}
$$

Therefore, in contrast to the work of Sung and Mitra [1986], we refrain from storing the entire blocks $B_{m}(\overline{\boldsymbol{y}})$. Instead, we store only the incomplete prologues $P_{m}(\overline{\boldsymbol{y}})$. Equation (24) then allows us to sequentially compute one complete prologue from the previous by means of multiplication by the $r \times r$ matrix $A_{F}^{b}$ and an $r$-vector addition with the incomplete $P_{m}(\overline{\boldsymbol{y}})$. The resulting algorithm is:

## Algorithm 1

1.1 In parallel for all $m$, compute and store each $P_{m}(\overline{\boldsymbol{y}})$.
1.2 Sequentially for each $m$, compute and store the $P_{m}(\boldsymbol{y})$ according to (24) and using the previously computed $P_{m}(\overline{\boldsymbol{y}})$.
1.3 In parallel for all $m$, compute and store output block $B_{m}(\boldsymbol{y})$ using (15) and the previously computed $P_{m-1}(\boldsymbol{y})$.
The algorithm does some redundant work by computing the recursive filter within each block in both stages 1.1 and 1.3. However, the reduction in memory bandwidth due to not writing the block in stage 1.1 provides a net benefit.
Although stage 1.2 is still sequential, it typically executes faster than the remaining stages. This is because we focus on small filter orders $r$ with block size $b$ large in comparison. This causes stage 1.2 to account for a small fraction of the total I/O and computation. Although we have not experienced a need for it, a hierarchical parallelization strategy is possible, as outlined in section 7.
The derivation for the anticausal counterpart to algorithm 1 follows closely along the same lines and is omitted for brevity. We include the anticausal versions of equations (15), (17), and (24) for completeness. Assuming $\boldsymbol{z}=R(\boldsymbol{y}, \boldsymbol{e})$ :

$$
\begin{align*}
& B_{m}(\boldsymbol{z})=R\left(B_{m}(\boldsymbol{y}), E_{m+1}(\boldsymbol{z})\right)  \tag{25}\\
& B_{m}(\overline{\boldsymbol{z}})=R\left(B_{m}(\boldsymbol{y}), \mathbf{0}\right)  \tag{26}\\
& E_{m}(\boldsymbol{z})=E_{m}(\overline{\boldsymbol{z}})+A_{R}^{b} E_{m+1}(\boldsymbol{z}) \tag{27}
\end{align*}
$$

Row operations are analogous and are also omitted.
When applying causal and anticausal filter pairs to process all columns and rows in an image, the four successive applications of algorithm 1 require $O\left(\frac{h w}{c p}\left(8+4 \frac{1}{b}\left(r^{2}+r\right)\right)\right.$ steps to complete. The $h w / b$ tasks that can be performed independently provide the scheduler with sufficient freedom to hide memory access latency.
The repeated passes over the input increase memory bandwidth to $(12+16 r / b) h w$, which is significantly more than the $8 h w$ required by algorithm RT. Fortunately, we can reduce bandwidth by employing the kernel fusion technique.

### 4.4 Kernel fusion

As it is known in the literature [Kirk and Hwu 2010], kernel fusion is an optimization that consists of directly using the output of one kernel as the input for the next without going through global memory. The fused kernel executes the code of both original kernels, keeping intermediate results in shared memory.

Using algorithm 1 for all filters, we fuse the last stage of $F$ with the first stage of $R$. We also fuse the last stage of $F^{\mathcal{T}}$ with the first stage of $R^{\mathcal{T}}$. Combined, the fusion of causal and anticausal processing lowers the required bandwidth to $(10+16 r / b) h w$. Row and column processing are also fused. The last stage of $R$ and the first stage of $F^{\mathcal{T}}$ are combined to further reduce bandwidth usage to $(9+16 r / b) h w$.

The resulting algorithm is presented below:

## Algorithm 2

2.1 In parallel for all $m$ and $n$, compute and store the $P_{m, n}(\bar{Y})$.
2.2 Sequentially for each $m$, but in parallel for each $n$, compute and store the $P_{m, n}(Y)$ according to (24) and using the previously computed $P_{m, n}(\bar{Y})$.
2.3 In parallel for all $m$ and $n$, compute and store $B_{m, n}(Y)$ using the previously computed $P_{m-1, n}(Y)$. Then immediately compute and store the $E_{m, n}(\bar{Z})$.
2.4 Sequentially for each $m$, but in parallel for each $n$, compute and store the $E_{m, n}(Z)$ according to (27) and using the previously computed $E_{m, n}(\bar{Z})$.
2.5 In parallel for all $m$ and $n$, compute and store $B_{m, n}(Z)$ using the previously computed $E_{\underline{m}+1, n}(Z)$. Then immediately compute and store the $P_{m, n}^{\mathcal{T}}(\bar{U})$.
2.6 Sequentially for each $n$, but in parallel for each $m$, compute and store the $P_{m, n}^{\mathcal{T}}(U)$.
2.7 In parallel for all $m$ and $n$, compute and store $B_{m, n}(U)$ using the previously computed $P_{m, n-1}^{\mathcal{T}}(U)$. Then compute and store the $E_{m, n}^{\mathcal{T}}(\bar{V})$.
2.8 Sequentially for each $n$, but in parallel for each $m$, compute and store the $E_{m, n}^{\mathcal{T}}(V)$.
2.9 In parallel for all $m$ and $n$, compute and store $B_{m, n}(V)$ using the previously computed $E_{m, n+1}^{\mathcal{T}}(V)$.

The combination of coalesced memory accesses, increased parallelism, and kernel fusion make algorithm 2 substantially faster than algorithm RT (see section 8).
We could aggressively reduce I/O even further by refraining from storing any of the intermediary results $Y, Z$, and $U$ at stages 2.3, 2.5 , and 2.7 , respectively. These can be recomputed from $X$ and the appropriate prologues/epilogues when needed. Bandwidth would decrease to $(6+22 r / b) h w$. (Notice that this is even less than the sequential algorithm RT!) Unfortunately, the repeated computation of $Y, Z$, and $U$ increases the total number of steps to $O\left(\frac{h w}{c p}\left(14 r+4 \frac{1}{b}\left(r^{2}+r\right)\right)\right.$, and offsets the gains from the bandwidth reduction. Even though the tradeoff is not advantageous for our test hardware, future GPU generations may tip the balance in the other direction.
Conceptually, buffers $P_{m, n}(\bar{Y}), E_{m, n}(\bar{Z}), P_{m, n}^{\mathcal{T}}(\bar{U}), E_{m, n}^{\mathcal{T}}(\bar{V})$ can be thought of as narrow bands around the perimeter of each block $B_{m, n}(X)$. However, these buffers are stored compactly in separate data arrays to maintain efficient memory access.

## 5 Overlapping

Our main contribution is the development of an overlapping technique which achieves further bandwidth reductions without any substantial increase in computational cost.

### 5.1 Causal-anticausal overlapping

One source of inefficiency in algorithm 2 is that we wait for the complete causal output block $B_{m, n}(Y)$ in stage 2.3 before we obtain the incomplete anticausal epilogues $E_{m, n}(\bar{Z})$. The important insight is that it is possible to work instead with twiceincomplete anticausal epilogues $E_{m, n}(\hat{Z})$, computed directly from the incomplete causal output block $B_{m, n}(\bar{Y})$. Even though these are obtained before $B_{m, n}(Y)$ are available, we are still able to compute the complete epilogues $E_{m, n}(Z)$ from them. We name this strategy for obtaining prologues $P_{m, n}(Y)$ and epilogues $E_{m, n}(Z)$ with one fewer pass over the input causal-anticausal overlapping.
The trick is to apply properties (9) and (10) to split the dependency chains of anticausal epilogues. Proceeding in one dimension for simplicity:

$$
\begin{align*}
E_{m}(\boldsymbol{z})= & H\left(R\left(B_{m}(\boldsymbol{y}), E_{m+1}(\boldsymbol{z})\right)\right)  \tag{28}\\
= & H\left(R\left(F\left(P_{m-1}(\boldsymbol{y}), B_{m}(\boldsymbol{x})\right), E_{m+1}(\boldsymbol{z})\right)\right)  \tag{29}\\
= & H\left(R\left(\mathbf{0}, E_{m+1}(\boldsymbol{z})\right)\right)  \tag{30}\\
& +H\left(R\left(F\left(P_{m-1}(\boldsymbol{y}), B_{m}(\boldsymbol{x})\right), \mathbf{0}\right)\right) \\
= & H\left(R\left(\mathbf{0}, E_{m+1}(\boldsymbol{z})\right)\right) \\
& +H\left(R\left(F\left(P_{m-1}(\boldsymbol{y}), \mathbf{0}\right), \mathbf{0}\right)\right)  \tag{31}\\
& +H\left(R\left(F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right), \mathbf{0}\right)\right)
\end{align*}
$$

We can simplify further using (11), (13) and (14) to reach

$$
\begin{align*}
E_{m}(\boldsymbol{z})= & H\left(R\left(\mathbf{0}, I_{r}\right) E_{m+1}(\boldsymbol{z})\right) \\
& +H\left(R\left(I_{b}, \mathbf{0}\right) F\left(I_{r}, \mathbf{0}\right) P_{m-1}(\boldsymbol{y})\right)  \tag{32}\\
& +H\left(R\left(F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right), \mathbf{0}\right)\right) \\
= & H\left(R\left(\mathbf{0}, I_{r}\right)\right) E_{m+1}(\boldsymbol{z}) \\
& +H\left(R\left(I_{b}, \mathbf{0}\right)\right) F\left(I_{r}, \mathbf{0}\right) P_{m-1}(\boldsymbol{y})  \tag{33}\\
& +H\left(R\left(F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right), \mathbf{0}\right)\right) \\
= & A_{R}^{b} E_{m+1}(\boldsymbol{z}) \\
& +\left(H\left(A_{R B}\right) A_{F P}\right) P_{m-1}(\boldsymbol{y})  \tag{34}\\
& +E_{m}(\hat{\boldsymbol{z}})
\end{align*}
$$

where the twice-incomplete $\hat{z}$ is such that

$$
\begin{align*}
B_{m}(\hat{\boldsymbol{z}}) & =R\left(F\left(\mathbf{0}, B_{m}(\boldsymbol{x})\right), \mathbf{0}\right)  \tag{35}\\
& =R(\overline{\boldsymbol{y}}, \mathbf{0}) \tag{36}
\end{align*}
$$

Notice that the $r \times r$ matrix $H\left(A_{R B}\right) A_{F P}$ used in (34) can be precomputed. Furthermore, each twice-incomplete epilogue $E_{m}(\hat{\boldsymbol{z}})$ depends only on the corresponding input block $B_{m}(\boldsymbol{x})$ and therefore they can all be computed in parallel already in the first pass. As a byproduct of that same pass, we can compute and store the $P_{m}(\overline{\boldsymbol{y}})$ that will be needed to obtain $P_{m}(\boldsymbol{y})$. With $P_{m}(\boldsymbol{y})$, we can compute all $E_{m}(\boldsymbol{z})$ sequentially with equation (34).

The resulting one-dimensional algorithm is as follows:

## Algorithm 3

3.1 In parallel for all $m$, compute and store $P_{m}(\overline{\boldsymbol{y}})$ and $E_{m}(\hat{\boldsymbol{z}})$.
3.2 Sequentially for each $m$, compute and store the $P_{m}(\boldsymbol{y})$ according to (24) and using the previously computed $P_{m}(\overline{\boldsymbol{y}})$.
3.3 Sequentially for each $m$, compute and store $E_{m}(\boldsymbol{z})$ according to (34) using the previously computed $P_{m-1}(\boldsymbol{y})$ and $E_{m}(\hat{\boldsymbol{z}})$.
3.4 In parallel for all $m$, compute each causal output block $B_{m}(\boldsymbol{y})$ using (15) and the previously computed $P_{m-1}(\boldsymbol{y})$. Then compute and store each anticausal output block $B_{m}(\boldsymbol{z})$ using (25) and the previously computed $E_{m+1}(\boldsymbol{z})$.

Using algorithm 3 for both row and column processing and fusing the two stages leads to:

## Algorithm 4

4.1 In parallel for all $m$ and $n$, compute and store the $P_{m, n}(\bar{Y})$ and $E_{m, n}(\hat{Z})$.
4.2 Sequentially for each $m$, but in parallel for each $n$, compute and store the $P_{m, n}(Y)$ according to (24) and using the previously computed $P_{m, n}(\bar{Y})$.
4.3 Sequentially for each $m$, but in parallel for each $n$, compute and store the $E_{m, n}(Z)$ according to (34) using the previously computed $P_{m-1, n}(Y)$ and $E_{m, n}(\hat{Z})$.
4.4 In parallel for all $m$ and $n$, compute $B_{m, n}(Y)$ using the previously computed $P_{m-1, n}(Y)$. Then compute and store the $B_{m, n}(Z)$ using the previously computed $E_{m+1, n}(Z)$. Finally, compute and store both $P_{m, n}^{\tau}(\bar{U})$ and $E_{m, n}^{\tau}(\hat{V})$.
4.5 Sequentially for each $n$, but in parallel for each $m$, compute and store the $P_{m, n}^{\tau}(U)$ from $P_{m, n}^{\tau}(\bar{U})$.
4.6 Sequentially for each $n$, but in parallel for each $m$, compute and store each $E_{m, n}^{\mathcal{T}}(V)$ using the previously computed $P_{m, n-1}^{\tau}(U)$ and $E_{m, n}^{\tau}(\hat{V})$.
4.7 In parallel for all $m$ and $n$, compute $B_{m, n}(V)$ using the previously computed $P_{m, n-1}^{\mathcal{T}}(V)$ and $B_{m, n}(Z)$. Then compute and store the $B_{m, n}(U)$ using the previously computed $E_{m, n+1}(U)$.
As expected, algorithm 4 is faster than algorithm 2 (see section 8 ). The advantage stems from a similar $O\left(\frac{h w}{c p}\left(8 r+6 \frac{1}{b}\left(r^{2}+r\right)\right)\right.$ number of steps combined with lower $(5+18 r / b) h w$ bandwidth.

### 5.2 Row-column causal-anticausal overlapping

There is still one source of inefficiency: we wait until the complete block $B_{m, n}(Z)$ is available in stage 4.4 before computing incomplete $P_{m, n}^{\tau}(\bar{U})$ and twice-incomplete $E_{m, n}^{\mathcal{T}}(\hat{V})$. Fortunately, we can overlap row and column computations and work with thriceincomplete transposed prologues and four-times-incomplete transposed epilogues obtained directly during stage 4.1. From these, we can compute the complete $P_{m, n}^{\top}(U)$ and $E_{m, n}^{\top}(V)$ without going over the input again.
The derivation of the procedure for completing the transposed prologues and epilogues is somewhat tedious, and can be found in full in appendix B. Below is the formula for completing thriceincomplete transposed prologues:

$$
\begin{align*}
P_{m, n}^{\tau}(U)= & P_{m, n-1}^{\tau}(U)\left(A_{F}^{b}\right)^{T} \\
& +A_{R E}\left(E_{m+1, n}(Z)\left(T\left(A_{F B}\right)\right)^{T}\right)  \tag{37}\\
& +\left(A_{R B} A_{F P}\right)\left(P_{m-1, n}(Y)\left(T\left(A_{F B}\right)\right)^{T}\right) \\
& +P_{m, n}^{\tau}(\check{U}),
\end{align*}
$$

where $b \times r$ matrix $A_{R E}$ as well as $b \times r$ matrices $A_{R B} A_{F P}$ and $T\left(A_{F B}\right)$ can all be precomputed. The thrice-incomplete $\check{U}$ satisfies

$$
\begin{equation*}
B_{m, n}(\check{U})=F^{\mathcal{T}}\left(\mathbf{0}, B_{m, n}(\hat{Z})\right) . \tag{38}
\end{equation*}
$$

To complete the four-times-incomplete transposed epilogues of $V$ :

$$
\begin{aligned}
E_{m, n}^{\tau}(V)= & E_{m, n+1}^{\tau}(V)\left(A_{R}^{b}\right)^{T} \\
& +P_{m, n-1}^{\tau}(U)\left(H\left(A_{R B}\right) A_{F P}\right)^{T} \\
& +A_{R E}\left(E_{m+1, n}(Z)\left(H\left(A_{R B}\right) A_{F B}\right)^{T}\right) \\
& +\left(A_{R B} A_{F P}\right)\left(P_{m-1, n}(Y)\left(H\left(A_{R B}\right) A_{F B}\right)^{T}\right) \\
& +E_{m, n}^{\tau}(\tilde{V})
\end{aligned}
$$

Here, the $r \times r$ matrix $H\left(A_{R B}\right) A_{F P}$ is the same as in (34), and the $b \times r$ matrix $A_{R B} A_{F P}$ is the same as in (37). The new $r \times b$ matrix $H\left(A_{R B}\right) A_{F B}$ is also precomputed. Finally, the four-timesincomplete $V$ is

$$
\begin{equation*}
B_{m, n}(\tilde{V})=R^{\top}\left(B_{m, n}(\check{U}), \mathbf{0}\right) \tag{40}
\end{equation*}
$$

The fully overlapped algorithm is:

## Algorithm 5

5.1 In parallel for all $m$ and $n$, compute and store each $P_{m, n}(\bar{Y})$, $E_{m, n}(\hat{Z}), P_{m, n}^{\mathcal{T}}(\check{U})$, and $E_{m, n}^{\mathcal{T}}(\tilde{V})$.
5.2 In parallel for all $n$, sequentially for each $m$, compute and store the $P_{m, n}(Y)$ according to (24) and using the previously computed $P_{m-1, n}(\bar{Y})$.
5.3 In parallel for all $n$, sequentially for each $m$, compute and store $E_{m, n}(Z)$ according to (34) and using the previously computed $P_{m-1, n}(Y)$ and $E_{m+1, n}(\hat{Z})$.
5.4 In parallel for all $m$, sequentially for each $n$, compute and store $P_{m, n}^{\top}(U)$ according to (37) and using the previously computed $P_{m, n}^{\top}(\check{U}), P_{m-1, n}(Y)$, and $E_{m+1, n}(Z)$.
5.5 In parallel for all $m$, sequentially for each $n$, compute and store $E_{m, n}^{\tau}\left(V_{\tilde{N}}\right)$ according to (39), using the previously computed $E_{m, n}^{\tau}(\tilde{V}), P_{m, n-1}^{\tau}(U), P_{m-1, n}(Y)$, and $E_{m+1, n}(Z)$.
5.6 In parallel for all $m$ and $n$, successively compute $B_{m, n}(Y)$, $B_{m, n}(Z), B_{m, n}(U)$, and $B_{m, n}(V)$ according to (8) and using the previously computed $P_{m-1, n}(Y), E_{m+1, n}(Z)$, $P_{m, n-1}^{\tau}(U)$, and $E_{m, n+1}^{\tau}(V)$. Store $B_{m, n}(V)$.
The number of steps required by the fully overlapped algorithm 5 is only $O\left(\frac{h w}{c p}\left(8 r+\frac{1}{b}\left(18 r^{2}+10 r\right)\right)\right)$. In our usage scenario, the filter order is much smaller than the block size. The step complexity is therefore about twice that of the sequential algorithm. On the other hand, algorithm 5 uses only $(3+22 r / b) h w$ bandwidth. This is less than half of what is needed by the sequential algorithm. In fact, considering that the lower bound on the required amount of traffic to global memory is $2 h w$ (assuming the entire image fits in shared memory), this is a remarkable result. It is not surprising that algorithm 5 is our fastest, at least in first-order filters (see section 8).

## 6 Overlapped summed-area tables

A summed-area table is obtained using prefix sums over columns and rows, and the prefix-sum filter $S$ is a special case of first-order causal recursive filter (with feedback weight $a_{1}=-1$ ). We can therefore directly apply the idea of overlapping to optimize the computation of summed-area tables.
In blocked form, the problem is to obtain output $V$ from input $X$ where the blocks satisfy the relations

$$
\begin{align*}
& B_{m, n}(V)=S^{\mathcal{T}}\left(P_{m, n-1}^{\tau}(V), B_{m, n}(Y)\right) \quad \text { and } \\
& B_{m, n}(Y)=S\left(P_{m-1, n}(Y), B_{m, n}(X)\right) . \tag{41}
\end{align*}
$$

With the framework we developed, overlapping the computation of $S$ and $S^{\mathcal{T}}$ is easy. In the first stage, we compute incomplete output blocks $B_{m, n}(\bar{Y})$ and $B_{m, n}(\hat{V})$ directly from the input:

$$
\begin{align*}
& B_{m, n}(\bar{Y})=S\left(\mathbf{0}, B_{m, n}(X)\right) \quad \text { and }  \tag{42}\\
& B_{m, n}(\hat{V})=S^{\mathcal{T}}\left(\mathbf{0}, B_{m, n}(\bar{Y})\right) \tag{43}
\end{align*}
$$

We store only the incomplete prologues $P_{m, n}(\bar{Y})$ and $P_{m, n}^{\tau}(\hat{V})$. Then we complete them using:

$$
\begin{align*}
& P_{m, n}(Y)=P_{m-1, n}(Y)+P_{m, n}(\bar{Y}) \quad \text { and }  \tag{44}\\
& P_{m, n}^{\mathcal{T}}(V)=P_{m, n-1}^{\mathcal{T}}(V)+s\left(P_{m-1, n}(Y)\right)+P_{m, n}^{\mathcal{T}}(\hat{V}) . \tag{45}
\end{align*}
$$



Figure 3: Overlapped summed-area table computation according to algorithm SAT. Stage S. 1 reads the input (in gray) then computes and stores incomplete prologues $P_{m, n}(\bar{Y})$ (in red) and $P_{m, n}^{\top}(\hat{V})$ (in blue). Stage $S .2$ completes prologues $P_{m, n}(Y)$ and computes scalars $s\left(P_{m-1, n}(Y)\right)$ (in yellow). Stage S. 3 completes prologues $P_{m, n}^{\tau, n}(V)$. Finally, stage $S .4$ reads the input and completed prologues, then computes and stores the final summed-area table.

Scalar $s\left(P_{m-1, n}(Y)\right)$ in (45) denotes the sum of all entries in vector $P_{m-1, n}(Y)$. The simplified notation means the scalar should be added to all entries of $P_{m, n}^{\top}(V)$. The resulting algorithm is depicted in figure 3 and described below:

## Algorithm SAT

S. 1 In parallel for all $m$ and $n$, compute and store the $P_{m, n}(\bar{Y})$ and $P_{m, n}^{\tau}(\hat{V})$ according to (42) and (43).
S. 2 Sequentially for each $m$, but in parallel for each $n$, compute and store the $P_{m, n}(Y)$ according to (44) and using the previously computed $P_{m, n}(\bar{Y})$. Compute and store $s\left(P_{m, n}(Y)\right)$.
S. 3 Sequentially for each $n$, but in parallel for each $m$, compute and store the $P_{m, n}^{\tau}(V)$ according to (45) using the previously computed $P_{m-1, n}(Y), P_{m, n}^{\top}(\hat{V})$ and $s\left(P_{m, n}(Y)\right)$.
S. 4 In parallel for all $m$ and $n$, compute $B_{m, n}(Y)$ then compute and store $B_{m, n}(V)$ according to (41) and using the previously computed $P_{m, n}(Y)$ and $P_{m, n}^{\tau}(V)$.
Our overlapped algorithm SAT uses only $\left(3+8 / b+2 / b^{2}\right) h w$ of bandwidth and requires $O\left(\frac{h w}{c p}(4+3 / b)\right)$ steps for completion. ${ }^{2}$ In contrast, using a stock implementation of prefix scan to process rows, followed by matrix transposition, followed by column processing (as in [Harris et al. 2008]) requires at least 8 hw bandwidth. Even our transposition-free, fused, bandwidth-saving variant of the multiwave approach of Hensley [2010] requires at least $(4+9 / b) h w$ bandwidth. As shown in section 8, the bandwidth reduction leads to the higher performance of the overlapped algorithm SAT.

## 7 Additional considerations

Feedforward in recursive filters A more general formulation for causal recursive filters would include an additional $a_{0}$ coefficient as well as a feedforward filter component with coefficients $b_{k}$ :

$$
\begin{equation*}
\boldsymbol{y}_{i}=\frac{1}{a_{0}}\left(\sum_{k=0}^{q} b_{k} \boldsymbol{x}_{i-k}-\sum_{k=1}^{r} a_{k} \boldsymbol{y}_{i-k}\right) . \tag{46}
\end{equation*}
$$

The added complication due to the feedforward coefficients is small compared to the basic challenge of the feedback coefficients, because the feedforward computation does not involve any dependency chains. One implementation approach is to simply pre-convolve the input with the feedforward component and then apply the appropriate feedback-only filter. To reduce memory bandwidth, it is best to fuse the feedforward computation with the parallelized feedback filter. This is relatively straightforward.

[^1]Flexibility with higher-order filters Higher-order causal or anticausal recursive filters can be implemented using different schemes: direct, cascaded, or in parallel. A direct implementation uses equations (1) or (3). Alternatively, we can decompose a higher-order filter into an equivalent set of first- and second-order filters, to be applied in series one after the other (i.e., cascaded) or independently and then added together (i.e., in parallel). See [Oppenheim and Schafer 2010, section 6.3] for an extensive discussion.
Although the different schemes are mathematically equivalent, in practice we may prefer one to the other due to bandwidth-to-FLOP ratio, implementation simplicity, opportunities for code reuse, and numerical precision. (Precision differences are less of a concern in floating-point [Oppenheim and Schafer 2010, section 6.9].)
Our overlapping strategy is compatible with all these schemes. The tools we provide allow the implementation of overlapped recursive filters of any order, in a way that maps well to modern GPUs. Overlapping can be used to directly implement a high-order filter, or alternatively to implement the first- and second-order filters used as building blocks for a cascaded or parallel realization.

Hierarchical reduction of prologues and epilogues In our algorithms, incomplete prologues and epilogues are completed sequentially (i.e., without inter-block parallelism). It is possible to perform such computations in a hierarchical fashion, thereby increasing parallelism and potentially speeding up the process. Completion equations such as (24) have the same structure as a recursive filter. This is clear when we consider the formulation in (52) of appendix A, where we process the input in groups of $r$ elements at a time. Taking stage 5.2 as an example, we could group $b$ prologues $P_{i b, n}(\bar{Y})$ to $P_{(i+1) b, n}(\bar{Y})$ into input blocks and recursively apply a variant of algorithm 1 to compute all $P_{m, n}(Y)$. Step 2 of this new computation is $b$ times smaller still, consisting of only $w / b^{2}$ prologues. We can repeat the recursion until we reach a problem size small enough to be solved sequentially, the results of which can be propagated back up to complete the original filtering operation. We skipped this optimization in our work due to diminishing returns. For example, the amount of time spent by algorithm 5 on stages 5.2-5.5 for a $4096^{2}$ image amounts to only about $17 \%$ of the total. This means that, even if we could run these stages instantaneously, performance would improve by only $20 \%$.

Recursive doubling for intra-block parallelism In our imageprocessing applications, the processing of each two-dimensional block exposes sufficient parallelism to occupy all cores in each SM. When processing one-dimensional input with algorithms 1 or 3, however, it may be necessary to increase intra-block parallelism.

Table 1: Properties of the presented algorithms, for row and column processing of an $h \times w$ image with causal and anticausal recursive filters of order $r$, assuming block-size $b$, and $p$ SMs with c cores each. For each algorithm, we show an estimate of the number of steps required, the maximum number of parallel independent threads, and the required memory bandwidth.

| Alg. | Step complexity | Max. \# of threads | Bandwidth |
| :---: | :---: | :---: | :---: |
| RT | $\frac{h w}{c p} 4 r$ | $h, w$ | $8 h w$ |
| 2 | $\frac{h w}{c p}\left(8 r+4 \frac{1}{b}\left(r^{2}+r\right)\right)$ | $\frac{1}{b} h w$ | $\left(9+16 \frac{r}{b}\right) h w$ |
| 4 | $\frac{h w}{c p}\left(8 r+6 \frac{1}{b}\left(r^{2}+r\right)\right)$ | $\frac{1}{b} h w$ | $\left(5+18 \frac{r}{b}\right) h w$ |
| 5 | $\frac{h w}{c p}\left(8 r+\frac{1}{b}\left(18 r^{2}+10 r\right)\right)$ | $\frac{1}{b} h w$ | $\left(3+22 \frac{r}{b}\right) h w$ |
| SAT | $\frac{h w}{c p}\left(8+\frac{3}{b}\right)$ | $\frac{1}{b} h w$ | $\left(3+\frac{8}{b}+\frac{2}{b^{2}}\right) h w$ |

Recursive doubling [Stone 1973] is a well known strategy for firstorder recursive filter parallelization we can use to perform intrablock computations. The idea maps well to GPU architectures, and is related to the tree-reduction optimization employed by efficient one-dimensional parallel scan algorithms [Sengupta et al. 2007; Dotsenko et al. 2008; Merrill and Grimshaw 2009]. Using a block size $b$ that matches the number of processing cores $c$, the idea is to break the computation into steps in which each entry is modified by a different core. Using recursive doubling, computation of $b$ elements completes in $O\left(\log _{2} b\right)$ steps.
The extension of recursive doubling to higher-order recursive filters has been described by Kooge and Stone [1973]. The key idea is to group input and output elements into $r$-vectors and consider equation (1) in the matrix form of (52) in appendix A. Since the algebraic structure of this form is the same as that of a first-order filter, the same recursive doubling structure can be reused.

## 8 Results

Table 1 summarizes the main characteristics of all the algorithms that we evaluated, in terms of number of required steps, the progression in parallelism, and the reduction in memory bandwidth.
Our test hardware consisted of an NVIDIA GTX 480 with 1.5 GB of RAM ( 480 CUDA cores, $p=15$ SMs, $c=32$ cores/SM). All algorithms were implemented in C for CUDA, under CUDA 4.0. All our experiments ran on single-channel 32-bit floating-point images. Image sizes ranged from $64^{2}$ to $4096^{2}$ pixels, in 64-pixel increments. Measurements were repeated 100 times to reduce variation. Note that small images are solved in sequence, not in batches that could be processed independently for added parallelism and performance.

First-order filters As an example of combined row-column causal-anticausal first-order filter, we solve the bicubic B-spline interpolation problem (see also figure 7(top)). Algorithm RT is the original implementation by Ruijters and Thévenaz [2010] (available on-line). Algorithm 2 adds blocking for memory coalescing, inter-block parallelism, and kernel fusion. (Algorithms 1 and 3 work on 1D input and are omitted from our tests.) Algorithm 4 employs overlapped causal-anticausal processing and fused row-column processing. Finally, algorithm 5 is fully overlapped. Performance numbers in figure 4(top) show the progression in throughput described throughout the text. As an example, the throughput of algorithm 5 solving the bicubic B-spline interpolation problem for $1024^{2}$ images is $4.7 \mathrm{GiP} / \mathrm{s}$ (gibi-pixels per second). Each image is transformed in just 0.21 ms or equivalently at more than 4800 fps . The algorithm appears to be compute-bound. Eliminating computation and keeping data movements, algorithm 5 attains $13 \mathrm{GiP} / \mathrm{s}$ ( $152 \mathrm{~GB} / \mathrm{s}$ ) on large images, whereas with computation it reaches $6 \mathrm{GiP} / \mathrm{s}(72 \mathrm{~GB} / \mathrm{s})$.


Figure 4: Throughput of the various algorithms for row-column causal-anticausal recursive filtering. (Top plot) First-order filter (e.g. bicubic B-spline interpolation). (Bottom plot) Second-order filter (e.g. biquintic $B$-spline interpolation).

Second-order filters The second-order, causal-anticausal, rowcolumn separable recursive filter used in our tests solves the biquintic B-spline interpolation problem. Figure 4(bottom) compares three alternative structures: $2 \times 5_{1}$ is a cascaded implementation using two fused fully-overlapped passes of first-order algorithm $5,4_{2}$ is a direct second-order implementation of algorithm 4, and $5_{2}$ is a direct fully-overlapped second-order implementation of algorithm 5. Our implementation of $4_{2}$ is the fastest, despite using more bandwidth than $5_{2}$. The higher complexity of second-order equations slows down stages 5.4 and 5.5 substantially. We believe this is an optimization issue that may be resolved with a future hardware, compiler, or implementation. Until then, the best alternative is to use the simpler and faster $4_{2}$. It runs at $3.1 \mathrm{GiP} / \mathrm{s}$ for $1024^{2}$ images, processing each image in less than 0.32 ms , or equivalently at more than 3200fps.
Precision A useful measure of numerical precision in the solution of a linear system such as the bicubic B-spline interpolation problem is the relative residual. Using random input images with entries in the interval $[0,1]$, the relative residual was less than $2 \times 10^{-7}$ for all algorithms and for all image sizes.
Summed-area tables Our overlapped summed-area table algorithm was compared with the algorithm of Harris et al. [2008] implemented in the CUDPP library [2011], and with the multi-wave method of Hensley [2010]. We also compare against a version of Hensley's method improved by two new optimizations: fusion of processing across rows and columns, and storage of just "carries" (e.g. $P_{m, n}(\bar{Y})$ ) between intermediate stages to reduce bandwidth. As expected, the results in figure 5 confirm that our specialized overlapped summed-area table algorithm outperforms the others.

Recursive Gaussian filters As mentioned in section 1, Gaussian filters of wide support are well approximated by recursive filters (see also figure 7(bottom)). To that end, we implemented the third-order


Figure 5: Throughput of summed-area table computation, using our overlapped algorithm, the multi-wave algorithm of Hensley [2010], a revisited version of it, and the CUDPP code of Harris et al. [2008].
approximation ( $N=3$ ) of van Vliet et al. [1998] with cascaded first- and second-order filters. The first-order component uses algorithm 5, and is fused with the second-order component, which uses algorithm $4_{2}$. Figure 6 compares performance of the overlapped recursive implementation with convolution in the frequency domain using the CUFFT library [2007], and with highly optimized direct convolution (separable in shared memory) [Podlozhnyuk 2007]. Because computation in the direct convolution is proportional to the number of filter taps $d$, we show results for filter kernels with $d=17,33,65$ taps. Although the direct approach is fastest for small filter windows, both the recursive filter and CUFFT are faster for larger windows. For reasonably sized images, overlapped recursive filtering is fastest. Moreover, its advantage over CUFFT should continue to increase on larger images due to its lower, linear-time complexity.

## 9 Conclusions

We describe an efficient algorithmic framework that reduces memory bandwidth over a sequence of recursive filters. It splits the input into blocks that are processed in parallel by modern GPU architectures, and overlaps the causal, anticausal, row, and column filter processing. The full sequence of filter passes requires reading the input image only twice, and writing only once. The reduced number of accesses to global memory leads to substantial performance gains. We demonstrate practicality in several scenarios: solving bicubic and biquintic B -spline interpolation problems, computing summed-area tables, and performing Gaussian filtering.
Future work Although we have focused on 2D processing, our framework of overlapping successive filters should extend to volumes (or even higher dimensions) using analogous derivations. However, one practical difficulty is the limited shared memory (e.g. 48k bytes) available in current GPUs. Making the volumetric blocks of size $b^{3}$ fit within shared memory would require significantly smaller block sizes, and this would decrease the efficiency of the algorithm.
Another area for future work is to adapt our algorithms to modern multicore CPUs (rather than GPUs). Although row and column parallelism is already sufficient for the fewer CPU threads, the memory bandwidth reduction may still be beneficial. In particular, our blocking structure may allow more accesses to be served by the faster L1 cache (or by the L2 cache in the case of volumetric processing). Whether this would lead to performance gains remains to be seen. The reduced bandwidth may also be beneficial on architectures like the Cell microprocessor in which memory transfers must be managed explicitly.
Finally, we would like to investigate whether the idea of overlapped computation can be used to optimize the solution of general narrow-band-diagonal linear systems.


Figure 6: Throughput of 2D Gaussian filtering, using our overlapped 3 rd-order recursive implementation, direct convolution with $d$ taps ( $\sigma=d / 6$ ), and the Fast Fourier Transform.

## 10 Acknowledgements

This work has been funded in part by a post-doctoral scholarship from CNPq and by an INST grant from FAPERJ.

## References

Blelloch, G. E. 1989. Scans as primitive parallel operations. IEEE Transactions on Computers, 38(11).
Blelloch, G. E. 1990. Prefix sums and their applications. Technical Report CMU-CS-90-190, Carnegie Mellon University.
Blu, T., Thénavaz, P., and Unser, M. 1999. Generalized interpolation: Higher quality at no additional cost. In IEEE International Conference on Image Processing, volume 3, 667671.

Blu, T., Thévenaz, P., and Unser, M. 2001. MOMS: Maximalorder interpolation of minimal support. IEEE Transactions on Image Processing, 10(7):1069-1080.
Crow, F. 1984. Summed-area tables for texture mapping. In Computer Graphics (Proceedings of ACM SIGGRAPH 90), volume 18, ACM, 207-212.
CUDPP library. 2011. URL http://code.google.com/p/ cudpp/.
CUFFT library. 2007. URL http://developer.nvidia.com/ cuda-toolkit. NVIDIA Corporation.
DERICHE, R. 1992. Recursively implementing the Gaussian and its derivatives. In Proceedings of the $2^{\text {nd }}$ Conference on Image Processing, 263-267.
Dotsenko, Y., Govindaraju, N. K., Sloan, P.-P., Boyd, C., and Manferdelli, J. 2008. Fast scan algorithms on graphics processors. In Proceedings of the $22^{\text {nd }}$ Annual International Conference on Supercomputing, 205-213.
Göddeke, D. and Strzodka, R. 2011. Cyclic reduction tridiagonal solvers on GPUs applied to mixed-precision multigrid. IEEE Transactions on Parallel and Distributed Systems, 22(1):22-32.
Harris, M., Sengupta, S., and Owens, J. D. 2008. Parallel prefix sum (scan) with CUDA. In GPU Gems 3, chapter 39.
Hensley, J. 2010. Advanced rendering technique with DirectX 11: High-quality depth of field. Gamefest 2010 talk.
Hensley, J., Scheuermann, T., Coombe, G., Singh, M., and Lastra, A. 2005. Fast summed-area table generation and its applications. Computer Graphics Forum, 24(3):547-555.
Hockney, R. W. and Jesshope, C. R. 1981. Parallel Computers: Architecture, Programming and Algorithms. Adam Hilger.


Figure 7: Effect of successive recursive filtering passes on a test image. (Top row) Cubic B-Spline interpolation filter, which acts as a high-pass. (Bottom row) Gaussian filter with $\sigma=5$, which acts as a low-pass. Insets show the combined filter impulse responses (not to scale). Our overlapped algorithm computes all four of these successive filtering passes using just two traversals of the input image.

IVERSON, K. E. 1962. A Programming Language. Wiley.
Kass, M., Lefohn, A., and Owens, J. D. 2006. Interactive depth of field using simulated diffusion on a GPU. Technical Report \#06-01, Pixar Animation Studios.

Kirk, D. B. and Hwu, W. W. 2010. Programming Massively Parallel Processors. Morgan Kaufmann.

Kooge, P. M. and Stone, H. S. 1973. A parallel algorithm for the efficient solution of a general class of recurrence equations. IEEE Transactions on Computers, C-22(8):786-793.
Lamas-Rodrígues, J., Heras, D. B., Bóo, M., and Argübllo, F. 2011. Tridiagonal solvers internal report. Technical report, University of Santiago de Compostela.

Merrill, D. and Grimshaw, A. 2009. Parallel scan for stream architectures. Technical Report CS2009-14, University of Virginia.

Oppenheim, A. V. and Schafer, R. W. 2010. Discrete-Time Signal Processing. Prentice Hall, $3^{\text {rd }}$ edition.

Parhi, K. and Messerschmitt, D. 1989. Pipeline interleaving and parallelism in recursive digital filters-Part II: Pipelined incremental block filtering. IEEE Transactions on Acoustics, Speech and Signal Processing, 37(7):1099-1117.
Podlozhnyuk, V. 2007. Image convolution with CUDA. NVIDIA whitepaper.
Ruijters, D. and Thévenaz, P. 2010. GPU prefilter for accurate cubic B-spline interpolation. The Computer Journal.
Sengupta, S., Harris, M., Zhang, Y., and Owens, J. D. 2007. Scan primitives for GPU computing. In Proceedings of Graphics Hardware, 97-106.

Stone, H. S. 1971. Parallel processing with the perfect shuffle. IEEE Transactions on Computers, C-20(2):153-161.

Stone, H. S. 1973. An efficient parallel algorithm for the solution
of a tridiagonal linear system of equations. Journal of the ACM, 20(1):27-38.

Sung, W. and Mitra, S. 1986. Efficient multi-processor implementation of recursive digital filters. In IEEE International Conference on Acoustics, Speech and Signal Processing, 257260.

Sung, W. and Mitra, S. 1992. Multiprocessor implementation of digital filtering algorithms using a parallel block processing method. IEEE Transactions on Parallel and Distributed Systems, 3(1):110-120.
van Vliet, L. J., Young, I. T., and Verbeek, P. W. 1998. Recursive Gaussian derivative filters. In Proceedings of the $14^{\text {th }}$ International Conference on Pattern Recognition, 509-514 (v. 1).

Zhang, Y., Cohen, J., and Owens, J. D. 2010. Fast tridiagonal solvers on the GPU. In Proceedings of the $15^{\text {th }}$ ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming.

## A Derivation of $T\left(A_{F P}\right)=A_{F}^{b}$

In this appendix we derive the $r \times r$-matrix $T\left(A_{F P}\right)$ that is used to transfer a prologue vector across a zero-valued block in the recursive filter computation of (24).

When dealing with a higher-order recursive filter $F$, it is convenient to group input and output elements into $r$-vectors, so that $F$ is expressed in a simplified matrix form. To that end, let

$$
\dot{\boldsymbol{y}}_{i}=\left(\begin{array}{c}
y_{i-r}  \tag{47}\\
\vdots \\
y_{i-2} \\
y_{i-1}
\end{array}\right), \dot{\boldsymbol{x}}_{i}=\left(\begin{array}{c}
x_{i-r} \\
\vdots \\
x_{i-2} \\
x_{i-1}
\end{array}\right) \text { and } \ddot{\boldsymbol{x}}_{i}=\left(\begin{array}{c}
0 \\
\vdots \\
0 \\
x_{i-1}
\end{array}\right) .
$$

It is easy to verify that if

$$
\begin{gather*}
\dot{\boldsymbol{y}}_{0}=\boldsymbol{p}, \text { then } \\
\dot{\boldsymbol{y}}_{i}=\ddot{\boldsymbol{x}}_{i}+A_{F} \dot{\boldsymbol{y}}_{i-1},  \tag{48}\\
A_{F}=\left(\begin{array}{ccccc}
0 & 1 & \cdots & 0 & 0 \\
\vdots & \vdots & \ddots & \vdots & \vdots \\
0 & 0 & \cdots & 1 & 0 \\
0 & 0 & \cdots & 0 & 1 \\
-a_{r} & -a_{r-1} & \cdots & -a_{2} & -a_{1}
\end{array}\right) . \tag{49}
\end{gather*}
$$

A nicer expression involves $\dot{\boldsymbol{x}}_{i}$ rather than $\ddot{\boldsymbol{x}}_{i}$. Iterating (48):

$$
\begin{align*}
\dot{\boldsymbol{y}}_{i} & =\ddot{\boldsymbol{x}}_{i}+A_{F} \dot{\boldsymbol{y}}_{i-1}=\ddot{\boldsymbol{x}}_{i}+A_{F} \ddot{\boldsymbol{x}}_{i-1}+A_{F}^{2} \dot{\boldsymbol{y}}_{i-2}=\cdots  \tag{50}\\
& =\left(\sum_{k=0}^{r-1} A_{F}^{k} \ddot{\boldsymbol{x}}_{i-k}\right)+A_{F}^{r} \dot{\boldsymbol{y}}_{i-r} \\
& =\bar{A}_{F} \dot{\boldsymbol{x}}_{i}+A_{F}^{r} \dot{\boldsymbol{y}}_{i-r} .
\end{align*}
$$

Column $k$ of matrix $\bar{A}_{F}$ is simply the last column of $A_{F}^{r-k-1}$, for $k \in\{0,1, \ldots, r-1\}$. Under this formulation, $r$ new output elements are produced and $r$ elements of input are consumed. For example, given a second-order filter:

$$
A_{F}=\left(\begin{array}{cc}
0 & 1  \tag{53}\\
-a_{2} & -a_{1}
\end{array}\right),
$$

$$
\bar{A}_{F}=\left(\begin{array}{cc}
1 & 0  \tag{54}\\
-a_{1} & 1
\end{array}\right), \quad \text { and } \quad A_{F}^{2}=\left(\begin{array}{cc}
-a_{2} & -a_{1} \\
a_{1} a_{2} & a_{1}^{2}-a_{2}
\end{array}\right)
$$

Equivalent derivations appear in the works of Kooge and Stone [1973] and Blelloch [1990].
Anticausal filters can also be expressed in terms of $r$-vectors in a form similar to (52), with corresponding matrices $\bar{A}_{R}$ and $A_{R}$.
Recall that
(55) $T(F(\boldsymbol{p}, \mathbf{0}))=T\left(F\left(I_{r}, \mathbf{0}\right) \boldsymbol{p}\right)=T\left(A_{F P} \boldsymbol{p}\right)=T\left(A_{F P}\right) \boldsymbol{p}$.

When $\boldsymbol{x}_{i}$ is null, equations (48) reduce to

$$
\begin{equation*}
\dot{\boldsymbol{y}}_{0}=\boldsymbol{p}, \quad \quad \dot{\boldsymbol{y}}_{i}=A_{F} \dot{\boldsymbol{y}}_{i-1}=A_{F}^{i} \boldsymbol{p} \tag{56}
\end{equation*}
$$

Since the last $r$ output elements of the $b$-wide block are given by

$$
\begin{equation*}
T(F(\boldsymbol{p}, \mathbf{0}))=\dot{\boldsymbol{y}}_{b}, \tag{57}
\end{equation*}
$$

we have

$$
\begin{equation*}
T\left(A_{F P}\right) \boldsymbol{p}=A_{F}^{b} \boldsymbol{p} \quad \Rightarrow \quad T\left(A_{F P}\right)=A_{F}^{b} . \tag{58}
\end{equation*}
$$

## B Derivation of full overlapping

Consider the transposed prologues of $U$. From the definition and repeated application of properties (9) and (10):

$$
\begin{align*}
P_{m, n}^{\tau}(U)= & T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(P_{m, n-1}^{\tau}(U), \mathbf{0}\right)\right) \\
& +T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(\mathbf{0}, E_{m+1, n}(Z)\right)\right)\right) \\
& +T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(P_{m-1, n}(Y), \mathbf{0}\right), \mathbf{0}\right)\right)\right)  \tag{59}\\
& +T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(\mathbf{0}, B_{m, n}(X)\right), \mathbf{0}\right)\right)\right)
\end{align*}
$$

Now applying (11), (13) and (14):

$$
\begin{aligned}
P_{m, n}^{\mathcal{T}}(U)= & P_{m, n-1}^{\mathcal{T}}(U) T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(I_{r}, \mathbf{0}\right)\right) \\
& +R\left(\mathbf{0}, I_{r}\right) E_{m+1, n}(Z) T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, I_{b}\right)\right) \\
& +R\left(I_{b}, \mathbf{0}\right) F\left(I_{r}, \mathbf{0}\right) P_{m-1, n}(Y) T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, I_{b}\right)\right) \\
& +T^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(\mathbf{0}, B_{m, n}(X)\right), \mathbf{0}\right)\right)\right) \\
= & P_{m, n-1}^{\mathcal{T}}(U)\left(A_{F}^{b}\right)^{T} \\
& +A_{R E}\left(E_{m+1, n}(Z)\left(T\left(A_{F B}\right)\right)^{T}\right) \\
& +\left(A_{R B} A_{F P}\right)\left(P_{m-1, n}(Y)\left(T\left(A_{F B}\right)\right)^{T}\right) \\
& +P_{m, n}^{\mathcal{T}}(\check{U}),
\end{aligned}
$$

where $\check{U}$ is such that

$$
\begin{align*}
B_{m, n}(\check{U}) & =F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(\mathbf{0}, B_{m, n}(X)\right), \mathbf{0}\right)\right)  \tag{62}\\
& =F^{\mathcal{T}}\left(\mathbf{0}, B_{m, n}(\hat{Z})\right) \tag{63}
\end{align*}
$$

In (61) we have exploited matrix-product associativity to select the best computation ordering.
Following a similar derivation for the transposed epilogues of $V$ :

$$
\begin{aligned}
& E_{m, n}^{\mathcal{T}}(V)=H^{\mathcal{T}}\left(R^{\mathcal{T}}\left(\mathbf{0}, E_{m, n+1}^{\mathcal{\tau}}(V)\right)\right) \\
& +H^{\mathcal{T}}\left(R^{\tau}\left(F^{\mathcal{T}}\left(P_{m, n-1}^{\tau}(U), \mathbf{0}\right), \mathbf{0}\right)\right) \\
& +H^{\mathcal{T}}\left(R^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(\mathbf{0}, E_{m+1, n}(Z)\right)\right), \mathbf{0}\right)\right) \\
& +H^{\mathcal{T}}\left(R^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(P_{m-1, n}(Y), \mathbf{0}\right), \mathbf{0}\right)\right), \mathbf{0}\right)\right) \\
& +H^{\mathcal{T}}\left(R^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(\mathbf{0}, B_{m, n}(X)\right), \mathbf{0}\right)\right), \mathbf{0}\right)\right) \\
& =E_{m, n+1}^{\mathcal{T}}(V)\left(A_{R}^{b}\right)^{T} \\
& +P_{m, n-1}^{\tau}(U)\left(H\left(A_{R B}\right) A_{F P}\right)^{T} \\
& +A_{R E}\left(E_{m+1, n}(Z)\left(H\left(A_{R B}\right) A_{F B}\right)^{T}\right) \\
& +\left(A_{R B} A_{F P}\right)\left(P_{m-1, n}(Y)\left(H\left(A_{R B}\right) A_{F B}\right)^{T}\right) \\
& +E_{m, n}^{\tau}(\tilde{V}),
\end{aligned}
$$

where

$$
\begin{align*}
B_{m, n}(\tilde{V}) & =R^{\mathcal{T}}\left(F^{\mathcal{T}}\left(\mathbf{0}, R\left(F\left(\mathbf{0}, B_{m, n}(X)\right), \mathbf{0}\right)\right), \mathbf{0}\right)  \tag{66}\\
& =R^{\mathcal{T}}\left(B_{m, n}(\check{U}), \mathbf{0}\right) \tag{67}
\end{align*}
$$


[^0]:    ${ }^{1}$ Note that $P^{\top}(\cdot)$ is not the transpose of $P(\cdot)$ (but has the same shape).

[^1]:    ${ }^{2}$ Actually, $O\left(\frac{h w}{c p}\left(4+(b+1) / b^{2}+2 / b\right)\right)$, but $(b+1) / b \approx 1$.

