Scanner: Efficient Video Analysis at Scale

ALEX POMS, Carnegie Mellon University, USA

WILL CRICHTON, PAT HANRAHAN, and KAYVON FATAHALIAN, Stanford University, USA

A growing number of visual computing applications depend on the analysis of large video collections. The challenge is that scaling applications to operate on these datasets requires efficient systems for pixel data access and parallel processing across large numbers of machines. Few programmers have the capability to operate efficiently at these scales, limiting the field's ability to explore new applications that leverage big video data. In response, we have created Scanner, a system for productive and efficient video analysis at scale. Scanner organizes video collections as tables in a data store optimized for sampling frames from compressed video, and executes pixel processing computations, expressed as dataflow graphs, on these frames. Scanner schedules video analysis applications expressed using these abstractions onto heterogeneous throughput computing hardware, such as multi-core CPUs, GPUs, and media processing ASICs, for high-throughput pixel processing. We demonstrate the productivity of Scanner by authoring a variety of video processing applications including the synthesis of stereo VR video streams from multi-camera rigs, markerless 3D human pose reconstruction from video, and data-mining big video datasets such as hundreds of feature-length films or over 70,000 hours of TV news. These applications achieve near-expert performance on a single machine and scale efficiently to hundreds of machines, enabling formerly long-running big video data analysis tasks to be carried out in minutes to hours.

CCS Concepts: • **Computing methodologies** → *Graphics systems and interfaces*; Image processing;

Additional Key Words and Phrases: large-scale video processing

ACM Reference Format:

Alex Poms, Will Crichton, Pat Hanrahan, and Kayvon Fatahalian. 2018. Scanner: Efficient Video Analysis at Scale. *ACM Trans. Graph.* 37, 4, Article 138 (August 2018), 13 pages. <https://doi.org/10.1145/3197517.3201394>

1 INTRODUCTION

The world is increasingly instrumented with sources of video: cameras are commonplace on people (smartphone cameras, GoPros), on vehicles (automotive cameras, drone videography), and in urban environments (traffic cameras, security cameras). Extracting value from these high-resolution video streams is a key research and commercial challenge, and a growing number of applications in fields like computer graphics, vision, robotics and basic science are based on analyzing large amounts of video.

The challenge is that scaling video analysis tasks to large video collections (thousands of hours of cable TV or YouTube clips, the output of a modern VR video capture rig) requires optimized systems

for managing pixel data as well as efficient, parallel processing on accelerated computing hardware (clusters of multi-core CPUs, GPUs, and ASICs). Unfortunately, very few programmers have the skill set to implement efficient software for processing large video datasets, inhibiting the field's ability to explore new applications that leverage this data. Inspired by the impact of data analytics frameworks such as MapReduce [Dean and Ghemawat 2004] and Spark [Zaharia et al. 2010], which facilitate rapid development of scalable big-data analytics applications, we have created *Scanner*, a system for productive and efficient *big video data* analysis.

Scanner provides integrated system support for two performance-critical aspects of video analysis: storing and accessing pixel data from large video collections, and executing expensive pixel-level operations in parallel on large numbers of video frames. Scanner addresses the first need by organizing video collections and derived raster data (depth maps, activation maps, flow fields, etc.) as tables in a data store whose implementation is optimized for compressed video. It addresses the second need by organizing pixel-analysis tasks as dataflow graphs that operate on sequences of frames sampled from tables. Scanner graphs support features useful for video processing, such as sparse sampling of video frames, access to temporal windows of frames, and state propagation across computations on successive frames. Scanner schedules these computations efficiently onto heterogeneous computing hardware such as multi-core CPUs, GPUs, and media processing ASICs.

We demonstrate that applications using Scanner for expensive, pixel-level video processing operations achieve near-expert performance when deployed on workstations with high-core count CPUs and multiple GPUs. The same applications also scale efficiently to hundreds of machines without source-level change. We report on experiences using Scanner to implement several large-scale video analysis applications including VR video processing, 3D human pose reconstruction from multi-viewpoint video, and data mining large video datasets of TV news. In these cases, Scanner enabled video analysis tasks that previously required days of processing (when implemented by researchers and data scientists using ad hoc solutions) to be carried out efficiently in hours to minutes.

2 CHALLENGES OF VIDEO ANALYSIS

Executing pixel-analysis pipelines (e.g., feature extraction, face/object detection, image similarity and alignment) on large image collections is the performance-critical component of many big visual data applications such as data-driven image manipulation and enhancement [Hays and Efros 2007; Kemelmacher-Shlizerman 2016], novel techniques for organizing and browsing photo collections [Sivic et al. 2008; Snavely et al. 2006], and exploratory data mining of the visual world [Chen et al. 2013; Doersch et al. 2012; Ginosar et al. 2017; Matzen et al. 2017; Zhu et al. 2014]. While these early applications analyzed collections of images, a growing class of applications now seek to manipulate large video datasets. To better understand

Authors' addresses: Alex Poms, Carnegie Mellon University, USA; Will Crichton; Pat Hanrahan; Kayvon Fatahalian, Stanford University, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2018/8-ART138 \$15.00

<https://doi.org/10.1145/3197517.3201394>

Fig. 1. We have implemented a set of video analysis applications in Scanner by expressing key pixel processing operations as dataflow graphs (Section 3.2). Each application contributes unique challenges to Scanner’s design, such as stateful processing, combining information across video streams, sparse frame access, and the need to process large numbers of video clips. Image credit, left to right: *The Rachel Maddow Show* © MSNBC 2015–2017, “Palace of Fine Arts Take 1” © Facebook 2017, “Run 5K” clip (top) and Figure 1 (bottom) from [Joshi et al. 2015], “160422_mafia2” scene from [Joo et al. 2016].

the challenges and requirements of these video analysis workloads, we selected a diverse set of video analysis applications to guide the design of Scanner. Fig. 1 summarizes the structure of these applications, which are implemented in Scanner and evaluated at scale in Section 5.2.

2.1 Workloads

Large-scale video data mining. Many applications now seek to perform labeling and data-mining of large video collections. Examples include autonomous vehicle development [Bojarski et al. 2016], surveillance, smart-city monitoring, and everyday egocentric video capture [Singh et al. 2016]. These computations require both traditional computer vision operations (optical flow, object tracking, etc.) and DNN inference (object detection, frame segmentation, activity recognition) to be executed on millions to billions of video frames. To keep costs manageable, it is common to sparsely sample frames from the video (e.g. every n -th frame, a list of frames likely to contain interesting objects). In Section 5.2.3 we report on experiences labeling and data mining two large video datasets: a dataset containing over 600 feature length films (106 million frames) and a dataset of 70,000 hours of TV news (12 billion frames, 20 TB).

360-degree stereo video generation for VR. Software for generating omnidirectional stereo (ODS) video, 360 degree stereo panoramas, provide a solution for authoring VR video. We ported the Surround 360 pipeline [Facebook 2017] for producing ODS video from 14 synchronized 2K video streams. This application involves *per-frame operations* (warping input frames to a spherical projection), *cross-video-stream operations* (depth estimation between frames from adjacent cameras), within-stream frame-to-frame dependencies (stateful temporal smoothing of computed flow fields), and the ability to output a final compressed high-resolution video stream. Surround 360 processing is computationally intense; it can take over twelve seconds to produce a single output frame on a

32-core server. The Jump VR Video processing pipeline has similar characteristics [Anderson et al. 2016].

Hyperlapse generation. Hyperlapses are stabilized timelapse videos synthesized from long videos captured with moving cameras. The challenge of generating a high-quality hyperlapse involves selecting source video frames that approximate a desired timelapse playback speed while minimizing apparent camera movement. We have implemented two variants of the frame-selection computation described by Joshi et al. [2015], which performs SIFT feature extraction and matching over *sliding windows* of frames from a video stream (temporal stencil computations).

3D human pose estimation. Recent computer vision advances make it possible to estimate temporally consistent human joint locations from dense multi-viewpoint video. This offers the promise of markerless human motion capture, even in high-occlusion scenarios, but comes at the cost of processing many video streams. For example, human motion capture sessions from the CMU Panoptic Dataset [Joo et al. 2015] feature 480 synchronized streams of 640×480 video (see visualization in Fig. 12). The dominant cost of a top-performing method for 3D pose reconstruction from these streams [Joo et al. 2016] involves evaluating a DNN on every frame of all streams to estimate 2D pose. The 2D poses are subsequently fused to obtain a 3D pose estimate.

2.2 Challenges

Scanner’s goal is to enable rapid development and scaling of applications such as those described above. This required a system with flexible abstractions to span a range of video analysis tasks, but also sufficiently constrained to allow efficient, highly-parallel implementations. Specifically, our experiences implementing the applications in Section 2.1 suggest that the size and temporal nature of video introduces several unique system requirements and challenges:

Organize and store compressed video. Managing tens of thousands of video clips, as well as per-frame raster products derived from its analysis (e.g., multiple resolutions of frames, flow fields, depth maps, feature maps, etc.), can be tedious and error prone without clear abstractions for organizing this data. The relational data model [Codd 1970] provides a natural representation for organizing video collections (e.g., a table per video, a row per video frame), however we are not aware of a modern database system optimized for managing, indexing, and *providing efficient frame-level access* to data stored compactly using video-specific compression (e.g., H.264). While some applications require video data to be maintained in a lossless form, in most cases it is not practical to store large video datasets as individual frames (even if frames are individually compressed). Video collections can fill TBs of storage even when encoded compactly using video-specific compression schemes. Ignoring inter-frame compression opportunities can increase storage footprint by an order or magnitude or more.

Support a flexible set of frame access patterns. Video compression schemes are designed for sequential frame access during video playback, however video analysis tasks exhibit a rich set of frame access patterns. While some applications access all video frames, others sample frames sparsely, select frame ranges, operate on sliding windows (e.g., optical flow, stabilization), or require joining frames from multiple videos (e.g., multi-view stereo). A system for video analysis must provide a rich set of streaming frame-level access patterns and implement these patterns efficiently on compressed video representations.

Support frame-to-frame (temporal) dependencies. Reasoning about a sequence of video frames as a whole (rather than considering individual frames in isolation) is fundamental to algorithms such as object tracking, optical flow, or activity recognition. Sequence-level reasoning is also key to achieving greater algorithmic efficiency when executing per-frame computations on a video stream since it is possible to exploit frame-to-frame coherence to accelerate analysis. Therefore, the system must permit video analysis computations to maintain state between processing of frames, but also constrain frame-to-frame dependencies to preserve opportunities for efficient data streaming and parallel execution.

Schedule pixel-processing pipelines (with black-box kernels) onto heterogeneous, parallel hardware. Authoring high-performance implementations of low-level image processing kernels (e.g., DNN evaluation, feature extraction, optical flow, object tracking) is difficult, so application developers typically construct analysis pipelines from pre-existing kernels provided by state-of-the-art performance libraries (e.g., cuDNN, OpenCV) or synthesized by high-performance DSLs (e.g., Halide [Ragan-Kelley et al. 2012]). Therefore, a video analysis system must assume responsibility for automatically scheduling these pipelines onto parallel, heterogeneous machines, and orchestrate efficient data movement between kernels. (The 3D human pose reconstruction pipeline presented in Section 5.2.1 involves computation on the CPU, GPU, and video decoding ASICs.) Although a single system for both kernel code generation and distributed execution provides opportunities for global optimization, it is not practical to

force applications to use a specific kernel code generation framework. For reasons of productivity and performance, Scanner should minimally constrain what 3rd-party kernels applications can use.

Scaling video analysis. Designing abstractions to address the above challenges is difficult because they must also permit an implementation which is able to scale from a workstation packed with GPUs underneath a researcher's desk to a cluster of thousands of machines, and from a dataset of a few 4K video streams to millions of 480p videos. Specifically, our examples require Scanner to scale in a number of ways:

- **Number of videos.** Scanner applications should scale to video datasets of arbitrary size (in our cases: millions or billions of frames), and consisting of both long videos (many feature length films or long-running vehicle capture sessions), or a large number of short video clips (e.g., millions of YouTube video clips).
- **Number of concurrent video streams.** We seek to handle applications that must process and combine a large number of video streams capturing a similar subject, scene, or event, such as VR video (14 streams) and 3D pose reconstruction (480 streams) discussed in Section 2.1. Scanner should accelerate computationally intensive pipelines to enable processing these streams at near-real time rates.
- **Number of throughput-computing cores.** Scanner applications should efficiently utilize throughput computing hardware (multi-core CPUs, multiple GPUs, media processing ASICs, and future DNN accelerators [Jouppi et al. 2017]) to achieve near-expert performance on a single machine, and also scale out to large numbers of compute-rich machines (thousands of CPUs or GPUs) with little-to-no source-level change.

We have designed Scanner to address these challenges. When our goals of productivity, scope, and performance conflict, we opted in favor of maintaining a scalable and performant system. This philosophy resulted in a number of clear *non-goals* for Scanner. For example, Scanner does not seek to aid with processing the results of pixel or feature-level analysis (image metadata, object labels, histograms, etc.). Post-processing these smaller derived data sets often involves a diverse set of algorithms that are well supported by existing data analysis frameworks. Also, Scanner does not seek to define its own programming language for authoring high-performance kernel functions. Many domain-specific programming frameworks exist for this purpose today and Scanner aims to inter-operate with and augment these best-in-class tools, not replicate their functionality.

3 SCANNER CONCEPTS

In this section we describe the primary abstractions used to construct Scanner applications. Scanner adopts two dataflow programming concepts familiar to users of existing data analytics frameworks and stream processing systems [Abadi et al. 2016; Chen et al. 2015; Dean and Ghemawat 2004; Zaharia et al. 2010], but extends and implements these concepts uniquely for the needs of efficient video processing.

Fig. 2. Scanner computation graphs (blue) operate on sequences of per-frame data extracted from data store tables (tan), and produce outputs that are stored as new tables (pink). This graph performs expensive face detection every 10th frame, and uses these detections to seed an object tracker run on each frame.

Videos as logical tables. Scanner represents video collections and the pixel-level products of video frame analysis (e.g., flow fields, depth maps, activations) as tables in a data store. Scanner’s data store features first-class support for video frame column types to facilitate key performance optimizations.

Video processing operations as dataflow graphs. Scanner structures video analysis tasks as dataflow graphs whose nodes produce and consume sequences of per-frame data. Scanner’s embodiment of the dataflow model includes operators useful for common video processing tasks such as sparse frame sampling, stenciled frame access, and stateful processing across frames.

We first provide an example of how Scanner’s abstractions are used to conduct a simple video analysis task, then describe the motivation and design of key system primitives in further detail.

3.1 Scanner Workflow

Fig. 2 illustrates a simple video analysis application (implemented using Scanner’s Python API) that annotates a video with bounding boxes for the faces in each frame.

First, the application ingests a collection of videos into the Scanner data store, shown in yellow. Logically, each video is represented by a table, with one row per video frame. In the example, ingest produces 100 tables, each with 18,000 rows, corresponding to 10-minute 30 FPS videos. The Scanner data store provides first-class support for table columns of *video frame* type, which facilitates compact storage and efficient frame-level access to compressed video data (Section 4.3). (See supplemental material for additional detail on how first-class video support enables Scanner’s storage formats to be optimized to specific access patterns without needing application-level change.)

Next, the application defines a five-stage *computation graph* that specifies what processing to perform on the video frames (code shaded in blue). Since accurate face-detection is costly, the application samples every 10th frame from the input video (*Stride*), down-samples the resulting frames (*Resize*), then evaluates a DNN to

detect faces in each downsampled frame to produce a per-frame list of bounding boxes (DNN). The 3 FPS (sparse-in-time) detections are then re-aligned (*Space*) with the original high-resolution, 30 FPS image sequence from the data store, and used to seed an object tracker (*Track*) that augments the original detections with additional detections produced by tracking on the original frames. The computation graph outputs a sequence of per-frame face bounding boxes that is stored as a new table with a column named *face_bboxes*.

A Scanner *job* specifies a computation graph to execute and the tables it consumes and produces. In this example, the application defines one job for each video (code shaded in pink). Scanner automatically schedules all jobs onto a target machine (potentially exploiting parallelism across jobs, frames in a job, and computations in the graph), resulting in the creation of new database tables (shown to the right in pink in Fig. 2). After using Scanner to perform the expensive pixel processing operations on video frames, an application typically exports results from Scanner, and uses existing data analysis frameworks to perform less performance-critical post-processing of the face bounding box locations.

3.2 Computation Graphs

Scanner applications express video processing tasks in the dataflow model by defining computation graphs. For consistency with [Abadi et al. 2016; Chen et al. 2015], we refer to graph nodes, which define stages of computation, as *operations*. Graph edges are *sequences* whose elements contain per-frame data communicated between operations. Figures 1 and 2 illustrate Scanner computation graphs for our example applications. These graphs range from simple pipelines defining stages of processing on a single video to complex DAGs of many operations on multiple input video streams.

3.2.1 Sequences. Scanner sequences are finite-length, 1D collections that are streamed element-by-element (or in small batches) to graph operations [Buck et al. 2004; Thies et al. 2002; Zaharia et al. 2010]. Each element in a length N sequence is associated with

Fig. 3. Scanner analyzes operation dependencies to reduce computation during graph execution. White boxes denote elements of a sequence, and are labeled with their corresponding sequence domain point. Black boxes denote the execution of a graph operation on an element. Grayed elements are not required to produce the graph’s required outputs and need not be computed.

a point in the $[0, N]$ domain. It is typical for sequence elements in Scanner applications to be video frames, or derived structures produced by graph operations, such as transformed images, flow fields, depth maps, or frame metadata (e.g., lists of per-frame object bounding boxes).

3.2.2 Graph Operations. A major challenge in Scanner’s design was selecting a set of graph operations that could be composed to express a rich set of video processing applications, but was sufficiently constrained to enable a streaming, data-parallel implementation. Scanner supports the following classes of graph operations, which are characterized by their input stream access patterns, and if state is propagated between invocations on consecutive stream elements.

Maps. Scanner operations may be mapped (Fig. 3-a) onto input sequences or onto multiple sequences of the same length (e.g., resizing a frame or evaluating a DNN to generate per-frame activations).

Sampling/spacing operations. Sampling and spacing operations (Figure 3-c,d) modify the length of sequences by selecting a subset of elements from the input sequence (sampling) or adding “fill” elements to it (spacing). Sampling operations enable computation on a sparse set of frames for computational efficiency or when specific frames must be selected for processing. For example, sampling every 30th row from a table representing a one-minute long, 30 FPS video (1800 frames) yields a length 60 sequence representing the video sampled at one frame per second. Spacing operations invert sampling and are used to align sequences representing data sampled at different frame rates. For example, in Fig. 2 a spacing operation was used to convert face detections computed at 3 FPS back into a 30 FPS stream. Both sampling and spacing operations can be defined by strides, ranges, or index lists.

Stencil operations. Stencil operations gain access to a window of elements from the input sequence defined by a constant-offset stencil. For example, the optical flow operation in Fig. 3-b requires elements i and $i + 1$ of the input sequence to generate output element i (the stencil is denoted by $S[0,1]$ next to the operation). Composing stencil and sampling operations yields a rich set of frame access patterns. For example, performing stride- N sampling prior to optical flow with stencil $(i, i + 1)$ yields flow vectors computed on a low

frame rate video sequence (Fig. 3-f), whereas sampling after the flow operation yields a sparse set of flow fields computed from differences between original video frames (Fig. 3-e).

Bounded State Operations. Video processing requires operations that maintain state from frame-to-frame, either because it is fundamental to the operation being performed (e.g., tracking) or as a compute optimization when there is little frame-to-frame change. However, if unconstrained, stateful processing would force serialization of graph execution. As a compromise, Scanner allows stateful operations, but limits the extent to which the processing of one sequence element can affect processing of later ones. Specifically, Scanner guarantees that prior to invoking an instance of a bounded state operation to generate output element i , the operation will have previously been invoked to produce *at least* the previous W elements of its output sequence. (The “warmup” value W is provided to Scanner by the stateful operation.) As a result, the operation is guaranteed that effects of processing element i will be visible when processing elements $(i + 1, \dots, i + W - 1)$ (Fig. 3-g: horizontal arrows). In Figs. 1 and 3, we denote the warmup size of bounded state operations (in elements) using the notation $W()$. An operation may have an infinite warmup, indicating that it must process input sequences serially (zero parallelism).

Warmup allows operations to benefit from element-to-element state propagation, while the bound on information flow provides Scanner flexibility to parallelize stateful operators at the cost of a small amount of redundant computation. For example, it is valid to execute a bounded state operation ($W = 2$) with a length-100 output sequence by producing output elements $[0, 50]$ on one machine independently from elements $[48, 99]$ on a second. Scanner automatically discards warmup elements 48 and 49 from the second worker (it does not include them in the output sequence), although effects of their processing may impact the value of subsequent elements (e.g., 50) generated by this worker. Bounded state operations use warmup to approximate the output of unbounded (fully serial) stateful execution when the influence of an operation’s effects is known to be localized in the video stream. For example, warmup of a few elements can be used to prime an object tracker prior to

producing required outputs, or to minimize temporal discontinuities in the outputs of a stateful operation at the boundary of two independently computed regions.

Slicing/unslicing operations. Slicing and unslicing operations insert and remove boundaries that affect stenciling and state propagation in a sequence. For example, slicing a video sequence at intervals according to shot boundaries would reset stencil operation access patterns and stateful processing to avoid information flow between shots (Fig. 3-h illustrates the use of slicing to partition a sequence into three independent slices). Unslicing removes these boundaries for all subsequent operations. Slicing and unslicing can be viewed as a constrained form of sequence nesting.

3.2.3 Computation Graph Limitations. Scanner’s design constrains the data flow expressible in computation graphs to permit two performance-critical graph scheduling optimizations: parallel graph execution and efficient graph scheduling in conditions of sparse sampling (Section 4). For similar reasons, Scanner currently disallows computation graphs with loops or operations that perform data-dependent filtering (discarding elements that do not pass a predicate) or amplification. Although Scanner operations are not provided mechanisms for dynamically modifying sequence length, sequence elements can be of tuple or list type (e.g., operations can produce variable length lists of face bounding boxes per frame).

3.3 Defining Graph Operations

Consistent with the goals from Section 2, Scanner does not provide mechanisms for defining the implementation of graph operations. With the exception of system-provided sampling, spacing, and slicing operations, Scanner operation definitions are implemented in 3rd party languages, externally compiled, and exposed to applications as Scanner graph operations using an operation definition API inspired by that of modern dataflow frameworks for machine learning [Abadi et al. 2016; Chen et al. 2015]. In the face detection example from Fig. 2, `Resize` is implemented in Halide, `DNN` by the Caffe library [Jia et al. 2014] in CUDA, and `Track` as multi-threaded C++. For bounded state operations, the allocation and management of mutable state carried across invocations is encapsulated entirely within the operation’s definition and is opaque to Scanner (e.g., internal object tracker state).

Although Scanner is oblivious to the details of an operation’s implementation, to facilitate efficient graph scheduling, all operations must declare their processing resource requirements (e.g., requires a GPU, requires N CPU cores) and data dependencies (warmup amount for stateful operations, stencil offsets for stencil operations) to Scanner. For efficiency, Scanner also supports operations that generate a batch of output elements (rather than a single element) per invocation (e.g., `DNN` inference on a batch of frames). We denote the batch size of operations as $B()$ in computation graph illustrations.

4 RUNTIME IMPLEMENTATION

Scanner jobs are executed by a high-performance runtime that provides applications high-throughput access to video frames and efficiently schedules computation graphs onto a parallel machine. While aspects of Scanner’s implementation constitute intelligent

Fig. 4. Left: Scanner creates multiple computation graph instances to process sequence elements in parallel. Here, three instances of the pose estimation graph (from Fig. 1-f) are distributed to single-GPU (left) and dual-GPU (right) machines. Instances of I/O and video decode stages that deliver data to and from application-defined graphs are shown in gray. Right: Scanner streams data through an execution graph at different bulk granularities to maximize data movement throughput and keep memory footprint low.

application of parallel systems design principles, the challenges of efficiently accessing compressed video data and executing compositions of sampling, stenciling, and bounded state graph operations led to unique implementation choices detailed here.

4.1 Graph Scheduling and Parallelization

The Scanner scheduler is responsible for efficiently distributing Scanner jobs onto the parallel processing resources within a machine and across large clusters of machines. Scanner implements data-parallel execution in the presence of stateful kernels by spawning multiple *instances* of the computation graph. In each instance, bounded state graph operations can maintain mutable state buffers, and all graph operations can preallocate a unique copy of read-only buffers (e.g., `DNN` weights, lookup tables). Scanner determines the maximum number of instances that can be created per machine by querying graph operations for their resource requirements, then maximizes parallelism without oversubscribing the machine. Fig. 4-left depicts a heterogeneous cluster of two machines, each containing an eight-core CPU and at least one GPU (worker 1 contains a single GPU, and worker 2 has two GPUs). To map the three-stage pose estimation pipeline (Fig. 1-f), which contains graph operations that require GPU execution and one operation that requires four CPU cores, onto this cluster, Scanner creates one computation graph instance on worker 1 and two instances of the pipeline on worker 2.

Scanner computation graphs can be statically analyzed to determine each sequence element’s dependencies before graph execution. This allows Scanner to partition the elements of a job’s output sequence into smaller *work packets* without violating graph operation dependencies. Work packets are then distributed to computation

graph instances, enabling parallelization within a single video and better load balancing (evaluated in Section 5.1.4). In addition to parallel work distribution, the Scanner runtime provides fault tolerance by automatically reassigning and restarting individual work packets (not entire jobs) assigned to failed workers. Scanner also distributes work to new worker machines that are added to a cluster while a job is running (supporting elasticity).

Scanner implements many common throughput-computing optimizations to sustain high-performance graph execution on machines with many cores and multiple GPUs. These include bulk transfer of sequence data between the data store and video decoders (particularly important in high latency cloud storage scenarios), bulk-granularity time-multiplexing of graph operations onto available machine compute resources, pipelining of CPU-GPU data transfers and data store I/O with graph operation execution (Fig. 4-right), and using custom GPU memory pools to reduce driver entry-point contention in multi-GPU environments.

In addition to processing work packets in parallel using multiple graph instances (data parallelism), Scanner also parallelizes computation within each graph instance by executing operations simultaneously on different CPU cores and GPU devices (pipeline parallelism). Scanner’s current implementation does not distribute the execution of a single graph instance across different machines. (We have not yet encountered applications that benefit from this functionality.) Multi-field elements are provided to operations in struct-of-arrays format to enable SIMD processing by batched operations without additional data shuffling. The granularities of bulk I/O (I/O packet size) and parallel work distribution (work packet size) are system parameters that can be tuned manually by a Scanner application developer to maximize performance, although auto-tuning solutions are possible. We evaluate the benefit of each of these key runtime optimizations in Section 5.1.3.

4.2 Unneeded Element Elimination

Scanner’s sequences are logically dense, however when a computation graph contains sampling operations, only a sparse set of intermediate sequence elements must be computed to generate a job’s required outputs. Since dependencies during graph execution do not depend on the values of sequence elements, Scanner determines which elements are required upfront through per-element graph dependency analysis. Interval analysis methods used to analyze stencil dependencies in image processing systems [Ragan-Kelley et al. 2013] are of little value when required graph outputs span the entire output domain, but are sparse (for example generating every N -th frame of an output sequence yields interval bounds that span the entire domain of all upstream sequences). Instead, given the set of output sequence points a job must produce, Scanner analyzes computation graph dependencies to determine the *exact set* of required points for all graph sequences. During graph execution, Scanner sparsely computes only the necessary sequence points. During dependency analysis, a bounded state operation with warmup size W is treated like a stencil operation with the footprint $(i - W, \dots, i - 1, i)$.

Fig. 3 illustrates the results of per-element dependency analysis for various example computation graphs. Gray boxes indicate sequence elements that are not required to compute the requested

Frames in work packet: 130, 134, 192, 320, 321,...

Fig. 5. The Scanner data store maintains an index of keyframe locations for video frame columns. The index is used to reduce I/O and video decode work when accessing a sparse set of frames from the video.

computation graph output elements and do not need to be computed by Scanner. Performing per-element dependency analysis to identify and eliminate unnecessary computation is unusual in a throughput-oriented system. However, Scanner graph operations typically involve expensive processing at the scale of entire frames, so the overhead of computing exact per-element liveness is negligible compared to the cost of invoking graph operations on elements that are not needed for the final job result.

To avoid the storage overhead of fully materializing lists of required sequence domain points, Scanner performs dependency analysis incrementally (at work packet granularity) as graph computation proceeds. Scanner also coalesces input sequence elements into dense batches to retain the efficiency of batch processing even when dependency analysis yields execution that is sparse.

For all stateless graph operations, sparse execution is a system implementation detail that does not influence the output of a Scanner application. It is valid, but inefficient, for Scanner to generate all sequence elements, even if they are never consumed. However, since prior invocations of a bounded state operation may impact future output, the values output by a bounded state operation may depend on which elements the Scanner runtime chooses to produce. (Different work distributions or conservative dependency analysis could yield different operation output.) However, Scanner applications are robust to this behavior since bounded state operations by definition are required to produce “acceptable” output provided their warmup condition is met.

4.3 Accessing Compressed Video Frames

Scanner presents the abstraction that videos are tables of individual frames, but internally stores video frame columns as compressed H.264 byte streams [Marpe et al. 2006] to minimize footprint and to reduce I/O. For example, the footprint of the 12 billion frame TVNEWS dataset (used in Section 5.2.3) is 20 TB when stored as H.264 byte streams, but exceeds 6 PB when expanded to an uncompressed N-D array of 24-bit pixels.

The cost of supporting compressed video storage in a system that must also support sparse frame-level data access is two-fold. First the byte stream must be decoded on the fly prior to graph execution. Second, video decode involves inherently sequential computation since most frames are encoded as deltas on data in prior frames. Therefore, to materialize a requested video frame, a decoder must locate the preceding “keyframe” (the last self-contained frame in the bytestream) then decode all frames up to the requested frame.

To accelerate access and decode of individual frames, the Scanner data store maintains an index of the byte stream offsets of keyframes in video columns, similar to indices maintained by video container formats to support scrubbing [ISO/IEC 2015]. The data store uses this index to minimize the amount of I/O and decode performed when servicing a sparse set of frame requests. For example, consider the sequence of elements in Fig. 5. To process this sequence, Scanner loads bytes from storage beginning from the keyframe preceding frame 130 (at byte offset 4,840). Decoding begins at this point, and continues until frame 192. Then, decoder state is reset to keyframe 310, and the process continues. When frames must be decoded but are not required by graph execution (e.g., frames 131-133, 135-191), Scanner skips decoder post-processing (extracting frames from the decoder, performing format conversion, etc.).

Scanner's data store implements a number of additional optimizations to maximize throughput, such as avoiding unnecessary reset of video decoder state when multiple required frames fall between two keyframes and time multiplexing decoders at bulk granularity to avoid unnecessary state resets when jobs draw video data from multiple tables. When available, Scanner also leverages ASIC hardware decode capabilities to accelerate video decode. For example, use of GPU-resident video decoding hardware frees programmable resources to execute other graph operations and also allows compressed video data to be communicated over the CPU-GPU bus.

5 EVALUATION

The goal of Scanner is to create a system that is sufficiently expressive to enable a rich set of video processing applications while also maintaining high performance. We evaluated Scanner's performance in terms of the efficiency of video frame access, efficiency in scheduling computation graphs onto a single machine, and scalability of applications to large numbers of CPUs and GPUs and very large video datasets. We evaluated Scanner's utility and expressiveness by implementing the video analysis workloads from Section 2.1 and deploying them at scale.

5.1 Performance

5.1.1 Video Decode Throughput. One of Scanner's goals is to provide applications with high-throughput access to compressed video frames, even when requested access patterns are sparse. We evaluated Scanner's H.264 decode performance against an OpenCV baseline under a varying set of frame access patterns drawn from our workloads:

- **STRIDE-1.** All video frames
- **STRIDE-24.** Every 24th frame.
- **GATHER.** A random list of frames that sparsely samples the video (0.25% of the video).
- **RANGE.** Blocks of 2,000 consecutive frames, each spread out by 20,000 frames.
- **KEYFRAME.** Only the keyframes from the video.

Figure 6 presents Scanner's decode throughput under these access patterns on a 2.2 hour, 202,525 frame, 1920×1080 H.264 encoded video (average keyframe distance of 104 frames) on a machine with two 8-core Intel Xeon E5-2620-v4 CPUs and one NVIDIA Titan Xp GPU. The throughput is normalized to a baseline implementation

Fig. 6. When executing a single graph instance, Scanner's sparse video decode optimizations improve throughput compared to OpenCV baselines on both the CPU and GPU. Scanner further improves CPU decode throughput by using multiple graph instances to more efficiently utilize all CPU cores.

which makes use of the OpenCV C++ API for video decode on the CPU and GPU (absolute throughput numbers in FPS are also given). The CPU version of this baseline delivers single-machine throughput that is similar to prior work on systems for large-scale video processing [Yang and Wu 2015]. For the CPU and GPU, we include results for a single graph instance to isolate the effect of sparse video decode optimizations. For the CPU, we also evaluate multiple graph instances to exploit Scanner's ability to decode different parts of the stream in parallel (we evaluate multiple graph instances on multiple GPUs in Section 5.1.3).

In all cases, Scanner's throughput matches or exceeds that of the baselines'. For a single graph instance, Scanner realizes higher throughput than the baselines when frame access is sparse (as much as 17× on the GPU). This speedup comes from Scanner avoiding post-decode processing of frames which must be decoded but that are not needed for graph execution (Section 4.3). Scanner uses the machine's 16 CPU cores more efficiently when executing multiple graph instances (Multi-instance on Fig. 6) since multiple instances of the decoder run in parallel (in addition to the parallelization available in H.264 decode which the baseline also exploits).

Even though Scanner's throughput can be higher than that of the CPU and GPU OpenCV baselines' in sparse access scenarios, overall throughput (FPS) of sparse access is fundamentally lower. If an application is flexible in which frames it can sample, such as accessing only a video's keyframes (KEYFRAME), it is possible to obtain higher throughput compared to other sparse access patterns (STRIDE-24 or GATHER), particularly when decoding on the GPU.

5.1.2 Scheduling Graphs with Optimized Kernels. In conjunction with video frame access, Scanner is also responsible for scheduling computation graphs of optimized kernels to machines with CPUs and GPUs. To test this, we chose three highly optimized kernels drawn from the applications in Section 2.1 and compared their native performance (when invoked from C++ and using OpenCV for video decode as in Section 5.1.1) to Scanner implementations using a single compute graph instance. These are:

- **HIST.** Compute and store the pixel color histogram for all frames (video decode bound) (OpenCV's `cv::calcHist` on the CPU and `cv::cuda::histEven` on the GPU).
- **FLOW.** Compute optical flow for all frames using a 2-frame stencil (OpenCV's CPU and GPU `FarnebackOpticalFlow` routines).

Fig. 7. Scanner executes graphs implemented using well-optimized kernels with nearly no overhead, matching or exceeding baseline implementations on both the CPU and GPU. Better orchestration of the compute graph produces modest improvements in HIST and DNN.

Fig. 8. Scanner computation graphs can be scaled to machines featuring multi-core CPUs and multiple GPUs without code modification. On the CPU, Scanner improves utilization of the 16-cores by parallelizing across multiple frames. On the GPU, Scanner achieves near linear speedup (at least 3.7 $\times$) when moving from one to four GPUs..

- **DNN.** Downsample and transform an input frame, then evaluate the Inception-v1 image DNN [Szegedy et al. 2015] for all frames. Image transformation is performed in Halide and DNN evaluation is performed using Caffe [Jia et al. 2014].

Figure 7 presents the throughput of CPU and GPU versions of the Scanner implementations of the microbenchmarks (using the libraries given above) normalized to their native implementations. We use the same multi-core CPU + single GPU machine from Section 5.1.1. In all cases, the Scanner implementations execute the kernels without incurring significant overhead, nearly matching or exceeding the native implementations. The Scanner implementations of HIST on the CPU and DNN on the GPU achieve a modest improvement in throughput due to better orchestration of the computation graph (pipelining of video decode, data transfers, and kernel execution).

5.1.3 Single Machine Scalability. It is common for high-end workstations and modern servers to be packed densely with multiple GPUs and CPUs. We evaluated Scanner’s scalability on multi-core CPU and multi-GPU platforms by running the microbenchmarks from Section 5.1.2 on a server with the same CPU but now with four Titan Xp GPUs. Figure 8 compares the microbenchmarks using *multiple graph instances* against their single graph instance counterparts from Section 5.1.2. Since OpenCV’s HIST and FLOW are not parallelized on the CPU, Scanner benefits from parallelization across video frames, providing a 5.1 and 12.5 $\times$ speedup respectively. Although the Caffe library is internally parallelized, Scanner still benefits from processing multiple frames simultaneously for DNN.

The GPU benchmarks realize near linear scaling (at least 3.7 $\times$) from one to four GPUs. The Scanner benchmarks realize these throughput improvements *without requiring modification to the*

Fig. 9. Scanner’s runtime optimizations result in a 5 to 19 $\times$ speedup of three microbenchmarks on a four-GPU machine. Each microbenchmark benefits differently from the optimizations, but the combination of all optimizations produces the best performance.

Scanner application. Achieving good multi-GPU scaling required the runtime optimizations discussed in Section 4. Figure 9 depicts a factor analysis of these optimizations for the three pipelines used in the four GPU scalability evaluation. The baseline configuration is Scanner with all optimizations disabled. Each data point adds one of the optimizations mentioned in Section 4:

- (1) Using multiple GPUs
- (2) Pipelining CPU-GPU computations and data-transfer
- (3) GPU HW ASIC decode
- (4) GPU memory pool
- (5) Increased work packet size
- (6) Batching input elements to kernels

Even when executing the simple computation graphs of HIST, FLOW, and DNN benchmarks, achieving multi-GPU scalability required combining several key optimizations. For example, HIST is decode bound, and benefits most from **GPU Memory Pool** because eliminating per-video frame memory allocations enables the GPU hardware video decoders (enabled by **GPU Decode**) to operate at high throughput. In the case of DNN, speedups from **Batching** are only possible after enabling a **Work Packet Size** that is greater than the batch size.

5.1.4 Multi-Machine and Dataset Scalability. The true benefit of Scanner is the ability to scale video processing applications to large numbers of machines and to very large video datasets. To evaluate Scanner’s scalability, we executed two benchmarks, the HIST computation graph from Section 5.1.2, and POSE, the OpenPose human pose estimation benchmark [Cao et al. 2016] which is central to several larger applications in Section 5.2, at scale on Google Compute Engine (GCE). We perform CPU scaling experiments on instances with 32 vCPUs (the unit of CPU hardware allocation on GCE, usually one hyper-thread), and GPU scaling experiments on instances with 16 vCPUs and two NVIDIA K80 GPUs. Since the POSE benchmark does not support CPU execution, we only evaluate it in GPU scaling experiments.

Single Video Scaling. One use of scaling to large machines is to deliver video processing results back to the user rapidly. (e.g., for quick preview or analysis.) Figure 10 shows Scanner executing HIST and POSE on a single 2.2 hour feature-length film on a cluster of 2,400 cores and a cluster of 75 GPUs. Executing HIST on this video took 4.3 minutes on a single machine (32 vCPUs) and nearly 15 minutes

Fig. 10. Scanner reduces the latency of analyzing a single video by using hundreds of GPUs and thousands of CPU cores. Scaling out reduces processing times from multiple minutes to seconds.

Fig. 11. Scanner applications efficiently scale to hundreds of GPUs and thousands of CPU cores when processing large datasets. Speedup is nearly linear until stragglers cause reduced scaling at high machine counts.

on a single GPU. These times were reduced to 20 and 31 seconds respectively when parallelizing this computation to large CPU and GPU clusters. Scaling POSE to the large GPU cluster reduced pose estimation processing time from 55 minutes (1 GPU) to two minutes (75 GPUs).

Large Dataset Scalability. Scanner facilitates scaling to large video datasets that would be impractical to process without the use of large numbers of machines. Figure 11 shows the speedup achieved running the H1ST and POSE benchmarks on datasets used by the video data mining applications in Section 5.2.3: CINEMA, a collection of 657 feature length films (107 million frames, 2.3 TB), and TVNEWS, a collection of short clips (approximately 10 seconds each) from 60K TV news videos (these shots total 86 million frames). Scanner scales linearly up to 3000 vCPUs and 150 GPUs while continuing to scale near linearly up to 250 GPUs. Speedups are sublinear at higher machine counts since a single slow machine (straggler) can delay job completion. Techniques for mitigating the effect of stragglers are well-studied and could be implemented by a future version of Scanner [Ananthanarayanan et al. 2013].

5.2 Application Experiences

We have used Scanner to scale a range of video processing applications (Section 2.1), enabling us to use many machines to obtain results faster, and to scale computations to much larger video datasets than previously practical. Each application presented a unique combination of frame access patterns, usage of Scanner computation graph features, and computational demands.

5.2.1 Video-Based 3D Pose Reconstruction. The video-based 3D pose reconstruction algorithm by Joo et al. [2016] requires efficient scheduling of compute graphs with both CPU and GPU operations to fully utilize machines packed densely with GPUs. The algorithm

involves evaluating a DNN on *every frame* of the 480 video streams in the Panoptic Studio (Figure 1-f). (Per-frame results from each video are then fused to estimate a per-frame 3D pose as in Figure 12, 3D Pose). An optimized implementation of the per-frame algorithm took 16.1 hours to process a 40-second sequence of captured video on a single Titan Xp GPU (frames 13,500 to 14,500 of the “160422_mafia2” scene from the CMU Panoptic Dataset). A version of this algorithm was previously parallelized onto four Titan Xp GPUs, reducing processing time to seven hours [Cao et al. 2016]. Using the exact same kernels, the Scanner implementation reduces runtime on the same 4-GPU machine to 2.6 hours due to more efficient graph scheduling (better pipelining and data transfer optimizations as discussed in Section 5.1.3).

Using Scanner, it was also simple to further accelerate the application using a large cluster of multi-GPU machines in the cloud. The same Scanner application scheduled onto 200 K80 GPUs (25 8-GPU machines on GCE) completed processing of the same video sequence in only 25 minutes. Dramatically reducing pose reconstruction time to minutes stands to enable researchers to capture longer and richer social interactions using emerging video-based capture infrastructure such as the Panoptic Studio.

5.2.2 Hyperlapse Generation. The real-time hyperlapse algorithm of [Joshi et al. 2015], which computes stabilized timelapses, makes use of computations that stencil over temporal windows. The computational bottleneck in the hyperlapse algorithm is feature extraction from the input images and pairwise feature matching between neighboring images. We implemented those portions of the algorithm as kernels in Scanner (Figure 1-d) using a GPU kernel to extract SIFT features from each frame and a second GPU kernel with a stencil window of size w to perform feature matching. Scanner’s stenciling mechanism simplified the implementation of the feature matching kernel (the runtime handles storing intermediate video frames and results) and made the pipeline easy to extend. For example, Joshi et al. [2015] suggest a performance optimization that approximates the reconstruction cost between two frames as the sum of successive costs, falling back to the full windowed feature matching when necessary. The corresponding Scanner pipeline (Figure 1-e) reduces the matching kernel’s stencil size to $[0, 1]$ to capture the adjacent reconstruction costs and adds a new kernel CondMatch which stencils over both the derived matching costs and the original features, conditionally determining if it is necessary to perform the full windowed feature matching.

5.2.3 Visual Data Mining at Scale. We have also used Scanner as the compute engine for two big video data mining research projects, requiring sparse sampling of videos, bounded state, and fault tolerance when scaling to hundreds of machines. The first involves visual analysis of a corpus of 657 feature length films (7.7 million frames, 2.3 TB). For example, Scanner applications are used to detect shot boundaries (via histogram differences, Figure 1-a), produce film summaries via montage (as in Figure 12-middle, with the Scanner pipeline in Figure 1-b), and detect faces. The second is a large-scale analysis of video from over three years of US TV news (FOX, MSNBC, and CNN), which includes over 70,000 hours of video (20 TB, 12 billion frames, six petapixels). In this project Scanner is being used

Fig. 12. Surround 360: Scanner’s port of Surround 360 fuses 14 video streams into a panoramic video for VR display. 3D Pose: Views of a social scene by 72 of the 480 cameras in the CMU Panoptic Studio (Joo et al. [2016]). Scanner performs pose estimation on all 480 camera streams which are then fused into 3D poses (shown projected onto a single view) Cinematography: A montage of one frame from each shot in *Star Wars* and *Mean Girls* computed using Scanner pipelines (Figure 1-a and -b). TV News: Scanner was used to calculate screen time given to people in 70,000 hours of TV News. Here we show instances of Rachel Maddow, a popular news host. Image credit, left to right: “Palace of Fine Arts Take 1”, © Facebook 2017; top image from [Joo et al. 2016] Figure 1, © Hanbyul Joo; *Star Wars: Episode IV - A New Hope*, © Lucasfilm Ltd. 1977; *Mean Girls*, © Paramount Pictures 2004; *The Rachel Maddow Show* © MSNBC 2015-2017.

Fig. 13. Single-machine scalability of the Surround 360 pipeline implemented in Scanner vs. the open source implementation. The Scanner implementation utilizes a single-machine well, scaling better to higher core counts.

to perform large-scale data mining tasks to discover trends in media bias and culture. These tasks involve visual analyses on video frames such as classifying news into shots, identifying the gender and identity of persons on screen, estimating screen time of various individuals, and understanding the movement of anchors on screen via pose estimation. Use of Scanner to manage and process billions of video frames was essential.

The large size of the feature length films and the TV news dataset stress-tested Scanner’s ability to scale. For example, to estimate the screen time allotted to male-presenting versus female-presenting individuals, we used Scanner to compute color histograms on every frame of the dataset (to detect shot boundaries), and then sparsely computed face bounding boxes and embeddings on a single frame per shot. To execute these tasks, we used a GCE cluster of 100 64-vCPU preemptible machines, relying on Scanner’s fault tolerance mechanism to handle preemption. The size of the dataset also required the use of cloud storage for both the videos and the derived metadata. Each computation took less than a day to complete and Scanner maintained 90%+ utilization of the 6,400 vCPUs throughout each run.

5.2.4 VR Video Stitching. We ported the Facebook Surround 360 Open Edition VR video stitching pipeline to Scanner [Facebook 2017]. The application requires simultaneously accessing 14 input video streams, scheduling up to 44 computation graph operations on a large number of CPU cores, employing kernels with temporal

dependencies (the *Flow* kernel is configured as a bounded state operation since it depends on the output of previous frames), and compressing output video frames to produce the final stereo panorama output (Figure 1-c). Given Scanner’s current scheduler implementation, we found it most efficient to execute each *Warp*, *Flow*, *Synth* block (the kernels surrounded by the blue box in Figure 1-c) as a separate job in Scanner and then feed each of those job’s outputs into the *Concat* stages using a second bulk launch. The Scanner implementation uses the same kernels as Facebook’s reference implementation.

In contrast to the reference Surround 360 implementation, which is parallelized across the 14 input video streams (but outputs frames serially), our Scanner implementation is also parallelized across segments of output frames, making use of bounded state operations with warmup of size 10 to maintain temporal coherence across segments of the video. Figure 13 plots the relative speedup of the reference and Scanner Surround 360 implementations on a machine with 32 CPUs (64 hyper-threaded). The Scanner implementation scales more efficiently on the large machine (5.3 seconds per frame versus 13.3 seconds per frame for the reference) due to the change in parallelization strategy. It is also faster due to pipelining (overlapping data movement and compute) and decreased IO since the Scanner implementation performs compression of the large output frames on the fly before writing out to disk.

We ran the Scanner version of Surround 360 implementation on a one minute sequence (28 GB, 25k total frames) over eight machines with 32 vCPU cores each (256 cores total) on Google Compute Engine and achieved a rate of 1.5 FPS. As was the case with our other applications, we were able to scale Surround 360 without any changes to the Scanner application.

6 RELATED WORK

Scanner contributes a unique integration of data-flow programming abstractions and systems implementation components that meet the productivity and performance needs of video analysis applications. However, many *individual components* of Scanner’s design were

influenced by prior systems for big data processing, databases, and machine learning.

Distributed data analytics frameworks. Frameworks such as MapReduce [Dean and Ghemawat 2004] and Spark [Zaharia et al. 2010] enable concise and productive expression of data analytics applications using data parallel operations on large collections. While these platforms handle the “scale-out” scheduling challenges of distributed computing (e.g. work distribution and fault tolerance), as identified in Section 2.2, they require new primitives and significant changes to their internal implementation to meet a broad set of video analysis needs. For example, while it is possible to use Spark to process video, prior implementations [Yang and Wu 2015] do not implement intra-video parallelism (precluding single-video speedups), do not target heterogeneous machines, and do not implement the video decode optimizations shown to provide significant benefits in Section 5.1.1. Scanner features such as bounded state operations (needed for intra-video parallelization in applications like VR video synthesis) and unneeded element elimination (needed for efficient sparse sampling common in data mining, Sec. 5.2.3) do not yet exist in popular distributed data-parallel systems.

Also, as we demonstrate in Fig. 9, Scanner execution graphs require a high-performance, heterogeneous (CPU, GPU, ASIC) runtime to be executed efficiently. While recent efforts have exposed popular GPU-accelerated machine learning libraries [Caf 2016; DataBricks 2016] to Spark applications, the Spark runtime, including its task scheduling, resource management, and data partitioning decisions, operates with no knowledge of the heterogeneous capabilities of the platform. Extending Spark to schedule tasks onto high-throughput accelerated computing platforms is known to require significant runtime redesign and extensions to application-visible abstractions (e.g., ability for kernels to specify resource requirements and data layouts, and to maintain local state) [Bordawekar 2016; Rosen and Xin 2015]. We hope that the design and implementation of Scanner influences ongoing development of the Spark runtime to better support video processing applications and accelerated computing.

Distributed machine learning frameworks. Modern machine learning frameworks [Abadi et al. 2016; Chen et al. 2015; Microsoft 2017] adopt a general dataflow programming model suitable for distributing GPU-accelerated training and inference pipelines across clusters of machines. While it may be possible to implement Scanner’s functionality as a library built upon these frameworks, doing so would require implementing new operations, runtime support for media accelerators, and integration with a pixel storage system providing the desired relational model and efficient video access—in other words, reimplementing most of Scanner itself. We elected to implement Scanner from the ground up as a lightweight runtime for simplicity and to achieve high performance.

Databases for raster and array data. Scanner models image and video collections as relations, and echoes the design of SparkSQL [Armbrust et al. 2015] in that row selection and joins on relations are used to define distributed datasets streamed to processing pipelines. Like relational Geographic Information Systems (GIS) ([PostGIS Project 2016]) or science/engineering-oriented Array Databases (ADBMS) such as SciDB [Cudre-Mauroux et al. 2009]

or RasDaMan [Baumann et al. 1998] which extend traditional database management systems with raster or multi-dimensional array types, Scanner natively supports image and video column types for efficiency. While GIS and ADBMS are optimized for operations such as range queries that extract pixel regions from high-resolution images, Scanner’s storage layer is designed to efficiently decompress and sample sequences of frames for delivery to computation graphs. As stated in Section 2.2, in contrast to array database designs, we intentionally avoided creating a new language for processing pixels in-database (e.g., SciDB’s Array Functional Language or RasDaMan’s RASCAL [Rasdaman.org 2015]). Instead we chose to support efficient delivery of video frame data to execution graphs with operations written in existing, well-understood languages like CUDA, C++, or Halide [Ragan-Kelley et al. 2012].

7 DISCUSSION

As large video collections become increasingly pervasive, and algorithms for interpreting their contents improve in capability, there will be an increasing number of applications that require efficient video analysis at scale. We view Scanner as an initial step towards establishing efficient parallel computing infrastructure to support these emerging applications. Future work should address higher-level challenges such as the design of query languages for visual data mining (what is SQL for video?), the cost of per-frame image analysis for the case of video (e.g., exploiting temporal coherence to accelerate DNN evaluation on a video stream), and integration of large-scale computation, visualization, and human effort to more rapidly label and annotate large video datasets [Ratner et al. 2018].

While the current version of Scanner achieves high efficiency, it requires the application developer to choose target compute platforms (CPU vs. GPU), video storage data formats, and key scheduling granularities (e.g., task size). It would be interesting to consider the extent to which these decisions could be made automatically for the developer as an application runs. Also, simple extensions of Scanner could expand system scope to provide high-throughput delivery of sampled video frames in model training scenarios (not just model inference) and to deliver regions of video frames rather than full frames (e.g., to support iteration over scene objects rather than video frames).

Most importantly, we are encouraged that Scanner has already proven to be useful. Our collaborations with video data analysts, film cinematographers, human pose reconstruction experts, and computer vision researchers show Scanner has enabled these researchers to iterate on big video datasets much faster than before, or attempt analyses that were simply not feasible given their level of parallel systems experience and existing tools. We hope that Scanner will enable many more researchers, scientists, and data analysts to explore new applications based on large-scale video analysis.

8 ACKNOWLEDGMENTS

This work was supported by the NSF (IIS-1422767, IIS-1539069), the Intel Science and Technology Center for Visual Cloud Computing, a Google Faculty Fellowship, and the Brown Institute for Media Innovation. TV News datasets were provided by the Internet Archive.

REFERENCES

2016. CaffeOnSpark. Github web site: <https://github.com/yahoo/CaffeOnSpark>. (2016).
- Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. USENIX Association, GA, 265–283.
- Ganesh Ananthanarayanan, Ali Ghodsi, Scott Shenker, and Ion Stoica. 2013. Effective Straggler Mitigation: Attack of the Clones. In *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*. USENIX, Lombard, IL, 185–198. <https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/ananthanarayanan>
- Robert Anderson, David Gallup, Jonathan T. Barron, Janne Kontkanen, Noah Snavely, Carlos Hernández, Sameer Agarwal, and Steven M. Seitz. 2016. Jump: Virtual Reality Video. *ACM Trans. Graph.* 35, 6, Article 198 (Nov. 2016), 13 pages.
- Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, and Matei Zaharia. 2015. Spark SQL: Relational Data Processing in Spark. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*. ACM, New York, NY, USA, 1383–1394.
- P. Baumann, A. Dehmel, P. Furtado, R. Ritsch, and N. Widmann. 1998. The Multidimensional Database System RasDaMan. *SIGMOD Rec.* 27, 2 (June 1998), 575–577.
- Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Praseen Goyal, Lawrence D Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, et al. 2016. End to End Learning for Self-Driving Cars. *arXiv preprint arXiv:1604.07316* (2016).
- Rajesh Bordawekar. 2016. Accelerating Spark workloads using GPUs. <https://www.oreilly.com/learning/accelerating-spark-workloads-using-gpus>. O'Reilly Media, Inc (2016).
- Ian Buck, Tim Foley, Daniel Horn, Jeremy Sugerman, Kayvon Fatahalian, Mike Houston, and Pat Hanrahan. 2004. Brook for GPUs: Stream Computing on Graphics Hardware. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 777–786.
- Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2016. Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. *arXiv preprint arXiv:1611.08050* (2016).
- Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems. *arXiv:1512.01274* (2015). [arXiv:eprint arXiv:1512.01274](https://arxiv.org/abs/1512.01274)
- X. Chen, A. Shrivastava, and A. Gupta. 2013. NEIL: Extracting Visual Knowledge from Web Data. In *2013 IEEE International Conference on Computer Vision*. 1409–1416.
- E. F. Codd. 1970. A Relational Model of Data for Large Shared Data Banks. *Commun. ACM* 13, 6 (June 1970), 377–387.
- P. Cudre-Mauroux, H. Kimura, K.-T. Lim, J. Rogers, R. Simakov, E. Soroush, P. Velikhov, D. L. Wang, M. Balazinska, J. Becla, D. DeWitt, B. Heath, D. Maier, S. Madden, J. Patel, M. Stonebraker, and S. Zdonik. 2009. A Demonstration of SciDB: A Science-oriented DBMS. *Proc. VLDB Endow.* 2, 2 (Aug. 2009), 1534–1537.
- DataBricks. 2016. TensorFrames. Github web site: <https://github.com/databricks/tensorframes>. (2016).
- Jeffrey Dean and Sanjay Ghemawat. 2004. MapReduce: Simplified Data Processing on Large Clusters. In *Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation - Volume 6 (OSDI'04)*. USENIX Association, Berkeley, CA, USA, 10–10.
- Carl Doersch, Saurabh Singh, Abhinav Gupta, Josef Sivic, and Alexei A. Efros. 2012. What Makes Paris Look Like Paris? *ACM Trans. Graph.* 31, 4, Article 101 (July 2012), 9 pages.
- Inc. Facebook. 2017. Facebook Surround 360. Web site: <https://facebook360.fb.com/facebook-surround-360/>. (2017).
- S. Ginosar, K. Rakelly, S. M. Sachs, B. Yin, C. Lee, P. Krahenbuhl, and A. A. Efros. 2017. A Century of Portraits: A Visual Historical Record of American High School Yearbooks. *IEEE Transactions on Computational Imaging* PP, 99 (2017).
- James Hays and Alexei A. Efros. 2007. Scene completion using millions of photographs. *ACM Trans. Graph.* 26, 3, Article 4 (July 2007). <https://doi.org/10.1145/1276377.1276382>
- ISO/IEC 2015. *ISO/IEC 14496-12:2015: Coding of audio-visual objects – Part 12: ISO base media file format*. Standard. International Organization for Standardization, Geneva, CH.
- Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional Architecture for Fast Feature Embedding. *arXiv preprint arXiv:1408.5093* (2014).
- H. Joo, H. Liu, L. Tan, L. Gui, B. Nabbe, I. Matthews, T. Kanade, S. Nobuhara, and Y. Sheikh. 2015. Panoptic Studio: A Massively Multiview System for Social Motion Capture. In *2015 IEEE International Conference on Computer Vision (ICCV)*. 3334–3342.
- Hanbyul Joo, Tomas Simon, Xulong Li, Hao Liu, Lei Tan, Lin Gui, Sean Banerjee, Timothy Godisart, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. 2016. Panoptic Studio: A Massively Multiview System for Social Interaction Capture. (2016). [arXiv:arXiv:1612.03153](https://arxiv.org/abs/1612.03153)
- Neel Joshi, Wolf Kienle, Mike Toelle, Matt Uyttendaele, and Michael F Cohen. 2015. Real-time hyperlapse creation via optimal frame selection. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 63.
- Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. 2017. In-Datcenter Performance Analysis of a Tensor Processing Unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA '17)*. ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/3079856.3080246>
- Ira Kemelmacher-Shlizerman. 2016. Transfiguring Portraits. *ACM Trans. Graph.* 35, 4, Article 94 (July 2016), 8 pages.
- D. Marpe, T. Wiegand, and G. J. Sullivan. 2006. The H.264/MPEG4 advanced video coding standard and its applications. *IEEE Communications Magazine* 44, 8 (Aug 2006), 134–143. <https://doi.org/10.1109/MCOM.2006.1678121>
- Kevin Matzen, Kavita Bala, and Noah Snavely. 2017. StreetStyle: Exploring world-wide clothing styles from millions of photos. (2017). <https://arxiv.org/abs/1706.01869>
- Microsoft. 2017. The Microsoft Cognitive Toolkit. Web site: <https://www.microsoft.com/en-us/cognitive-toolkit/>. (2017).
- PostGIS Project 2016. *PostGIS 2.3.2dev Manual*. PostGIS Project.
- Jonathan Ragan-Kelley, Andrew Adams, Sylvain Paris, Marc Levoy, Saman Amarasinghe, and Frédo Durand. 2012. Decoupling Algorithms from Schedules for Easy Optimization of Image Processing Pipelines. *ACM Trans. Graph.* 31, 4, Article 32 (July 2012), 12 pages.
- Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. *SIGPLAN Not.* 48, 6 (June 2013), 519–530.
- Rasdaman.org 2015. *Rasdaman Version 9.2 Query Language Guide*. Rasdaman.org.
- Alexander Ratner, Stephen H. Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2018. Snorkel: Rapid Training Data Creation with Weak Supervision. *Proc. VLDB Endow.* (to appear) 12, 1 (2018).
- Josh Rosen and Reynold Xin. 2015. Project Tungsten: Bringing Apache Spark Closer to Bare Metal. Databricks Engineering Blog: <https://databricks.com/blog/2015/04/28/project-tungsten-bringing-spark-closer-to-bare-metal.html>. (2015).
- Krishna Kumar Singh, Kayvon Fatahalian, and Alexei Efros. 2016. KrishnaCam: Using a longitudinal, single-person, egocentric dataset for scene understanding tasks. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- Josef Sivic, Biliiana Kaneva, Antonio Torralba, Shai Avidan, and William T. Freeman. 2008. Creating and Exploring a Large Photorealistic Virtual Space. In *First IEEE Workshop on Internet Vision*.
- Noah Snavely, Steven M. Seitz, and Richard Szeliski. 2006. Photo Tourism: Exploring Photo Collections in 3D. *ACM Trans. Graph.* 25, 3 (July 2006), 835–846.
- C. Szegedy, Wei Liu, Yangqing Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. 2015. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1–9.
- William Thies, Michal Karczmarek, and Saman P. Amarasinghe. 2002. StreamIt: A Language for Streaming Applications. In *Proceedings of the 11th International Conference on Compiler Construction (CC '02)*. Springer-Verlag, London, UK, UK, 179–196.
- S. Yang and B. Wu. 2015. Large Scale Video Data Analysis Based on Spark. In *2015 International Conference on Cloud Computing and Big Data (CCBD)*. 209–212.
- Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: Cluster Computing with Working Sets. In *Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing (HotCloud'10)*. USENIX Association, Berkeley, CA, USA, 10–10.
- Jun-Yan Zhu, Yong Jae Lee, and Alexei A. Efros. 2014. AverageExplorer: Interactive Exploration and Alignment of Visual Data Collections. *ACM Trans. Graph.* 33, 4, Article 160 (July 2014), 11 pages.