


Michael J. Anderson, Stephan J. Ritter, James Demmel

## Robust PCA

- Decompose a matrix into a **low rank** component and a **sparse** component:

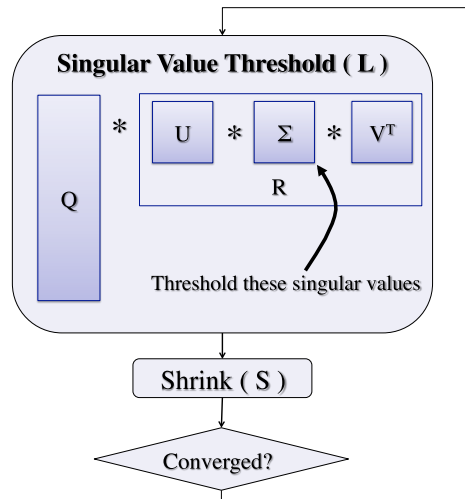


- M, L and S are **videos**
- Each column is a frame
- Dimension roughly  $100,000 \times N$  for N frames

Emmanuel Candes, 2009

## Algorithm


$$\begin{aligned} &\text{minimize } \|L\|_F + \lambda \|S\|_1 \\ &\text{subject to } L + S = M \end{aligned}$$



- Challenge: QR decomposition of a tall-skinny matrix**

## GPU QR

- Householder QR
- for  $(i = 1..n)$   
Form  $u$  from  $A^{(i)}$   
SGEMV  
SGER  
Global Barrier
- Had to make custom SGEMV for tall-skinny matrix




Implementation	Time	Speedup
LAPACK QR (SGEQR)	0.71 sec.	-
GPU QR	.072 sec.	9.9x

## Matrix Dimensions Determine the Algorithms

- Matrix is generally not wide enough to provide enough parallelism in the 'n' direction. Must also divide in the 'm' direction and use atomic add when necessary.
- Best results came from minimizing the number of parallel reductions in our SGEMV
- Global barrier could potentially be avoided, but synchronization costs were a second-order concern for this size matrix

## SV Threshold Optimization



Implementation	Time	Speedup
LAPACK without optimization	1.54 sec	-
LAPACK with optimization	0.81 sec	1.9x
GPU with optimization	0.093 sec	17x

## Robust PCA Speedup

Can process 10 seconds of video in 11 seconds

Went from 28x **slower** than real-time to near real-time