

# L1-MAGIC BENCHMARKS AND POTENTIAL IMPROVEMENTS

Stephen Conover  
September 2, 2015

With Thanks to Chris Turnes and Adam Charles

# INTRODUCTION

1. Improvements to Optimization Library L1-Magic
2. A Few Techniques
  1. Don't waste memory or allocate unnecessarily
  2. Avoid big operations when small ops work
  3. Do smart block operations when you can to exploit sparsity and structure

# L1-MAGIC

1. **P<sub>1</sub>: Basis Pursuit (L1 with Equality Constraints)**  
 $\min \|x\|_1$  subject to  $Ax = b$  (**P<sub>1</sub>**)
2. **P<sub>A</sub>: Decode**  
 $\min_x \|y - Ax\|_1$
3. **P<sub>2</sub>: L1 with Quadratic Constraints**  
 $\min \|x\|_1$  subject to  $\|Ax - b\|_2 \leq \epsilon$
4. **P<sub>D</sub>: Dantzig Selector**  
 $\min \|x\|_1$  subject to  $\|A^*(Ax - b)\|_\infty \leq \gamma$
5. **TV1: Total Variation with Equality Constraints**  
 $\min TV(x)$  subject to  $Ax = b$
6. **TV2: Total Variation with Quadratic Constraints**  
 $\min TV(x)$  subject to  $\|Ax - b\|_2 \leq \epsilon$
7. **TV<sub>D</sub>: Dantzig TV**

# GOAL

- Improve Performance
  - Reduce Execution Time
  - Reduce Memory Requirements
- Improvement Philosophy
  - No new fundamentals
  - Avoid MEX
  - CPU Now, GPU Later

# POTENTIAL IMPROVEMENTS

1. Memory Allocation and Sparse Data Structures
2. Woodbury Identity
3. Block Matrix for Sparse + Dense Matrices

# PROBLEM P1: BASIS PURSUIT

Definition:  $\min \|x\|_1$  subject to  $Ax = b$

Improvement: Memory Usage with Sparse Storage

# PROFILER RESULTS

## Lines where the most time was spent

Line Number	Code	Calls	Total Time	% Time
<a href="#">139</a>	<code>H11p = A*(sparse(diag(1./sigx)...</code>	10	0.066 s	73.3%
<a href="#">148</a>	<code>Adx = A*dx;</code>	10	0.003 s	3.3%
<a href="#">141</a>	<code>[dv,hcond] = linsolve(H11p, w1...</code>	10	0.003 s	3.3%
<a href="#">138</a>	<code>w1p = -(w3 - A*(w1./sigx - w2....</code>	10	0.002 s	2.2%
<a href="#">87</a>	<code>u = (0.95)*abs(x0) + (0.10)*ma...</code>	1	0.002 s	2.2%

# AMDAHL'S LAW

- The program can be divided into fraction of the code we can improve,  $P$ , and the fraction of code we cannot,  $(1 - P)$ , to calculate the overall speedup of the program,  $S(N)$ , resulting from a factor  $N$  speedup of  $P$



$$S(N) = \frac{1}{(1 - P_t) + \frac{P_t}{N}}$$

- $P$  may be defined in terms of  $P_t$ , the fraction of total execution time, rather than lines of code
- If  $P_t \cong 1$  and corresponds to a small number of lines of code, then then rate of return can be significant



# CONCENTRATION AND SPEEDUP

$$S(N) = \frac{1}{(1 - P_t) + \frac{P_t}{N}}$$



$P_t$	SPEEDUP		
	Max	N=100	N=50
25%	1.3	1.3	1.3
50%	2	2	2
90%	10	9	8
95%	20	17	14
98%	50	34	25
99%	100	50	34

# CODE

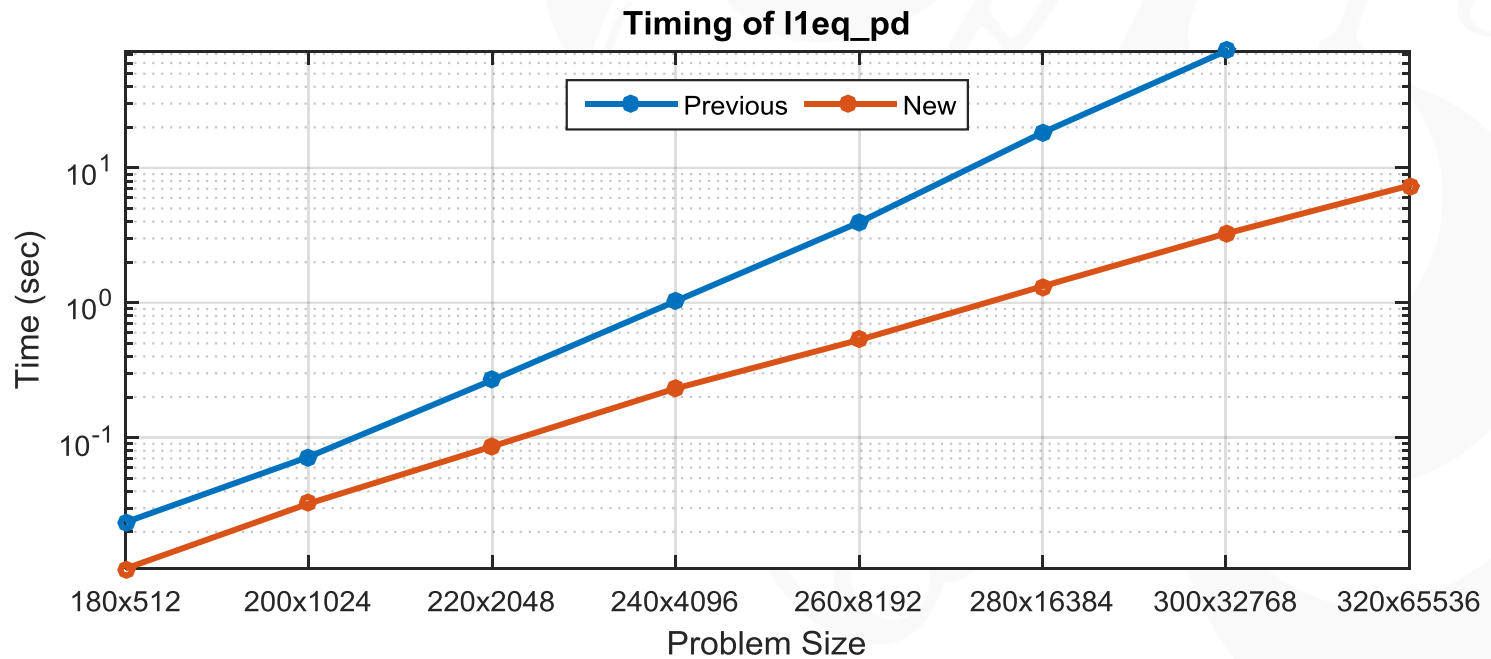
## 1. Original

```
1. w1p = -(w3 - A*(w1./sigx -  
    w2.*sig2./(sigx.*sig1)));  
2. H11p = A*(sparse(diag(1./sigx))*A');  
3. opts.POSDEF = true; opts.SYM = true;  
    [dv,hcond] = linsolve(H11p, w1p, opts);
```

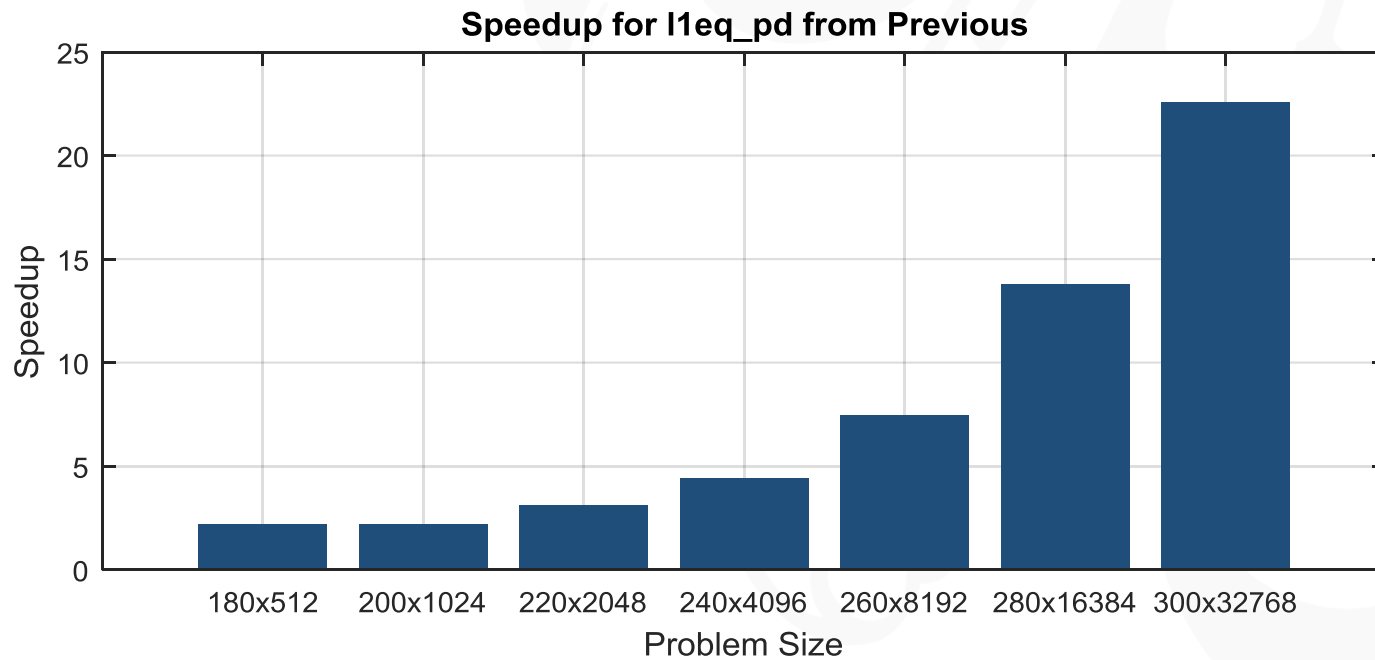
## 2. Memory Allocation Improvement

```
1. w1p = -(w3 - A*(w1./(sigx) -  
    w2.*sig2./(sigx.*sig1)));  
2. SS = (sparse((1:N), (1:N), (1./(sigx))', N,N));  
3. H11p = A*(SS*A');  
4. opts.POSDEF = true; opts.SYM = true;  
    [dv,hcond] = linsolve(H11p, w1p, opts);
```


# RESULTS PART I




# RESULTS II



# RESULTS III



# RESULTS IV



# PROBLEM P2: QUADRATIC CONSTRAINTS

Definition:  $\min \|x\|_1$  subject to  $\|Ax - b\|_2 \leq \epsilon$

Improvement: Sparse Memory & Low Rank Updates

# THE CODE

```
N = length(sigx);  
SX = (sparse((1:N), (1:N), ((sigx))', N,N));  
H11p = SX - (1/fe)*AtA + (1/fe)^2*(atr*atr');  
opts.POSDEF = true; opts.SYM = true;  
[dx,hcond] = linsolve(H11p, w1p, opts);
```

$$dx = \left( SX - \frac{1}{fe} A^T A + \left( \frac{1}{fe^2} A^T r r^T A \right) \right)^{-1} w1p$$

## Two Problems

1. Code is slower than it could be
2. Has stability issues



# LOW RANK UPDATES

- Woodbury Matrix Identity

$$(S + UCV)^{-1} = S^{-1} - S^{-1}U(C^{-1} + VS^{-1}U)^{-1}VS^{-1}$$

- Sherman-Morrison Formula

$$(M - uv^T)^{-1} = M^{-1} - \frac{M^{-1}uv^T M^{-1}}{1 + v^T M^{-1}u}$$

# REARRANGEMENT

1.  $dx = H^{-1} * w1p = \left( SX - \frac{1}{f_e} A^T A + \left( \frac{1}{f_e^2} A^T r r^T A \right) \right)^{-1} w1p$
2.  $H = S_x - \frac{1}{f_e} A^T A + \frac{1}{f_e^2} A^T r r^T A = S_x + \frac{1}{f_e} A^T \left( -I + \frac{1}{f_e} r r^T \right) A$
3. Apply to the Woodbury Matrix Identity with  
 $S = S_x = \text{diag}(\text{sig}_x) \quad U = \frac{1}{f_e} A^T \quad V = A \quad C = -I + \frac{1}{f_e} r r^T$
4.  $H^{-1} = S^{-1} - S^{-1} U (C^{-1} + V S^{-1} U)^{-1} V S^{-1}$   
 $= S_x^{-1} - \frac{1}{f_e} S_x^{-1} A^T \left( C^{-1} + \frac{1}{f_e} A S_x^{-1} A^T \right)^{-1} A S_x^{-1}$

# MORE REARRANGEMENT

Use the Sherman-Morrison Formula:

$$C^{-1} = (M + uv^T)^{-1} = M^{-1} - \frac{M^{-1}uv^T M^{-1}}{1 + v^T M^{-1}u}$$

and set  $M = M^{-1} = -I$  and  $u = v = \frac{1}{\sqrt{f_e}}\mathbf{r}$

to get

$$C^{-1} = \frac{\mathbf{r}\mathbf{r}^T}{\mathbf{r}^T\mathbf{r} - f_e} - I$$

# FINALLY

Altogether, this gives us

$$H^{-1}y = S_x^{-1}y - \frac{1}{f_e} S_x^{-1} A^T (G)^{-1} A S_x^{-1} y$$

Where


$$G = \frac{r r^T}{r^T r - f_e} - I + \frac{1}{f_e} A S_x^{-1} A^T$$

And so


```
B = bsxfun(@times, A', sqrt(1./sigx));
```

```
dx = wlp./sigx - 1/fe*A'*(-linsolve(-  
(r*r'/(r'*r-fe)-eye(k) + (1/fe)*(B'*B)),  
A*(wlp./sigx)))./sigx;
```

# RESULTS



# RESULTS




# PROBLEM TV: TOTAL VARIATION

Definition:  $\min TV(x)$  subject to  $Ax = b$

Improvement: Block Matrix Inverse

# THE SETUP

- The Newton Step Involves both the operator  $A$  and the TV operator



$$A = \begin{bmatrix} H & \sigma A \\ \sigma A' & 0 \end{bmatrix}$$



# THE CODE

```
H11p = Dh'*sparse(diag(-1./ft + sigb.*Dhx.^2))*Dh + ...
Dv'*sparse(diag(-1./ft + sigb.*Dvx.^2))*Dv + ...
Dh'*sparse(diag(sigb.*Dhx.*Dvx))*Dv + ...
Dv'*sparse(diag(sigb.*Dhx.*Dvx))*Dh;
afac = max(diag(H11p));
Hp = full([H11p afac*A'; afac*A zeros(K)]);
opts.SYM = true;
[dxv, hcond] = linsolve(Hp, wp, opts);
```

# THE CODE WITH SPARSE

```
S1 = sparse((1:N), (1:N), -1./ft + sigb.*Dhx.^2, N,N);
S2 = sparse((1:N), (1:N), -1./ft + sigb.*Dvx.^2, N,N);
S3 = sparse((1:N), (1:N), sigb.*Dhx.*Dvx, N,N);
S4 = sparse((1:N), (1:N), sigb.*Dhx.*Dvx, N,N);
H11p = Dh'*S1*Dh + ...
      Dv'*S2*Dv + ...
      Dh'*S3*Dv + ...
      Dv'*S4*Dh;
afac = max(diag(H11p));
Hp = full([H11p afac*A'; afac*A zeros(K)]);
opts.SYM = true;
[dxv, hcond] = linsolve(Hp, wp, opts);
```

# RESULTS COMING

- Never call full
- Do Block Inverse on H
- But H is poorly conditioned!
- Calculate  $\mathbf{C} = (\mathbf{B}' * \mathbf{A}) * \mathbf{B}$ ; and do a full SVD on this  $k \times k$  matrix and determine how many extra rows and columns H needs.
- Tentatively Up to x10 Faster but highly depended on  $K \ll N$ .
- Results Pending

# ARROWHEAD SOLVER

```
1. function [blockwiseSolution,
   conditionWarning] =
   arrowHeadSolver(A, B, w)
2. k0 = size(A,1);
3. C = (B'*A)*B;
   [U, Sc] = svd(full(C)); % grab
   the singular vectors, too
   Sc = diag(Sc); % make them a
   vector
   U = fliplr(U); % put worst
   singular vectors first
   Sc = Sc / sum(Sc);
   Sc = cumsum(Sc);
   del = size(C,1) - nnz(Sc < (1-1e-
   4));
   % now, solve a different syste
4. w((k0+1):end) = U'*w((k0+1):end);

5. % Amod = [ A      B*U ]
   %         [ U'*B'   0   ]
   Bmod = B*U;
   A = [A Bmod(:, 1:del); Bmod(:,
   1:del)'], zeros(del)];
   B = [Bmod(:, (del+1):end); zeros(del,
   size(B, 2)-del)];
6. k = k0 + del;
   w1 = w(1:k);
   w2 = w((k+1):end);
7. [L,D,P,Sx] = ldl(A);
8. % S*P*(L*D*L')^(-1)*P'*S*y
   om = (Sx*P)*(L'\((L*D)\((P'*Sx)*[w1,
   B])));
   mu = (B'*om(:,1) - w2);
   nu = (-B'*om(:,2:end)) \ mu;
9. blockwiseSolution = [om*[1; nu]; -nu];
   % re-calibrate the solution:
10. blockwiseSolution((k0+1):end,:) =
   U*blockwiseSolution((k0+1):end,);
```

# L1-MAGIC COMPARISON

How Does L1-Magic Compare with  
Alternatives FISTA and NESTA?

# FISTA NOTES

- FISTA Had Memory Issues  $> 2^{16}$  size signal because it calculates  $G=A'A$  for use in backtracking, some stopping criterion

```
Creating measurment matrix...
```

```
Done.
```

```
Warning: Requested 65536x65536 (32.0GB) array exceeds maximum  
array size preference. Creation of arrays greater than this  
limit may take a long time and cause  
MATLAB to become unresponsive. See array size limit or  
preference panel for more information.
```

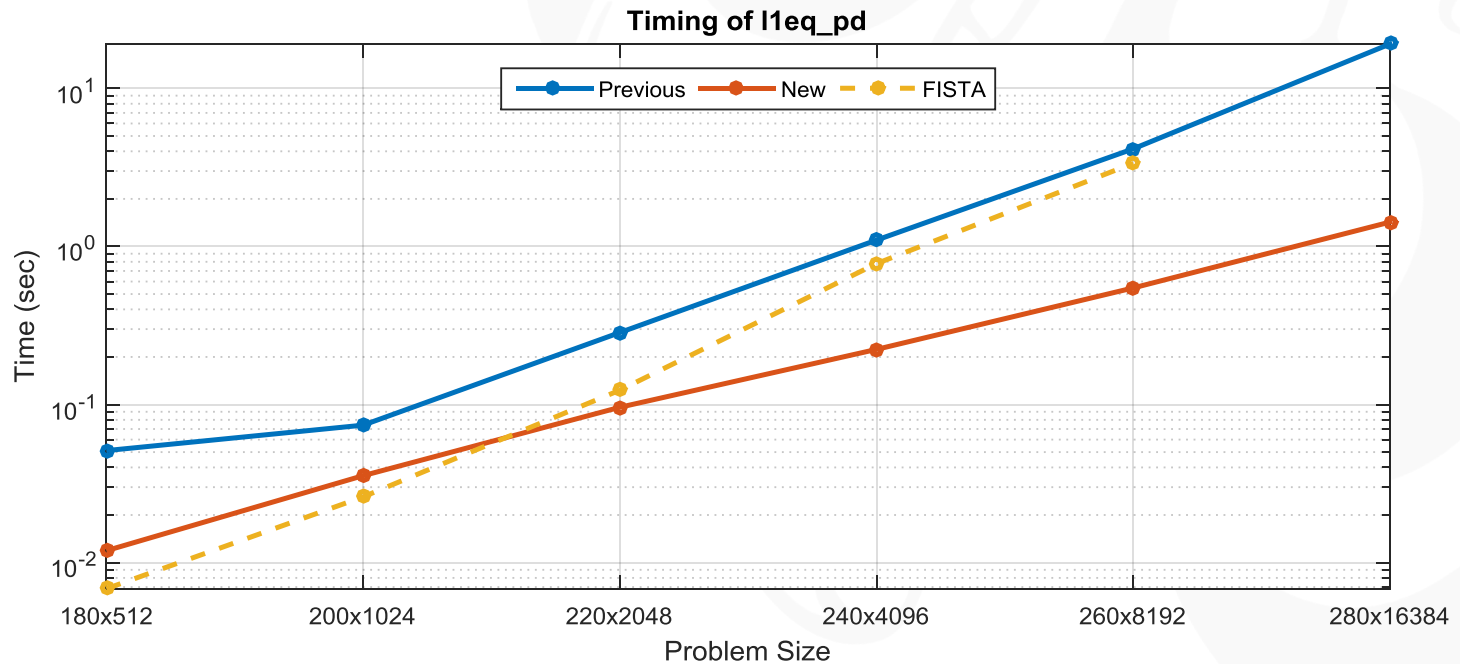
- Mixed Performance – Fixed by not calculating  $G=A'A$

```
% G=A'*A; temp1 = G*yk - c ;
```


```
temp = A'*(A*yk) - c ; % gradient of f at yk
```

- Some convergence issues so may need to go back and tweek configuration for fair comparison

# FISTA @ SMALL N

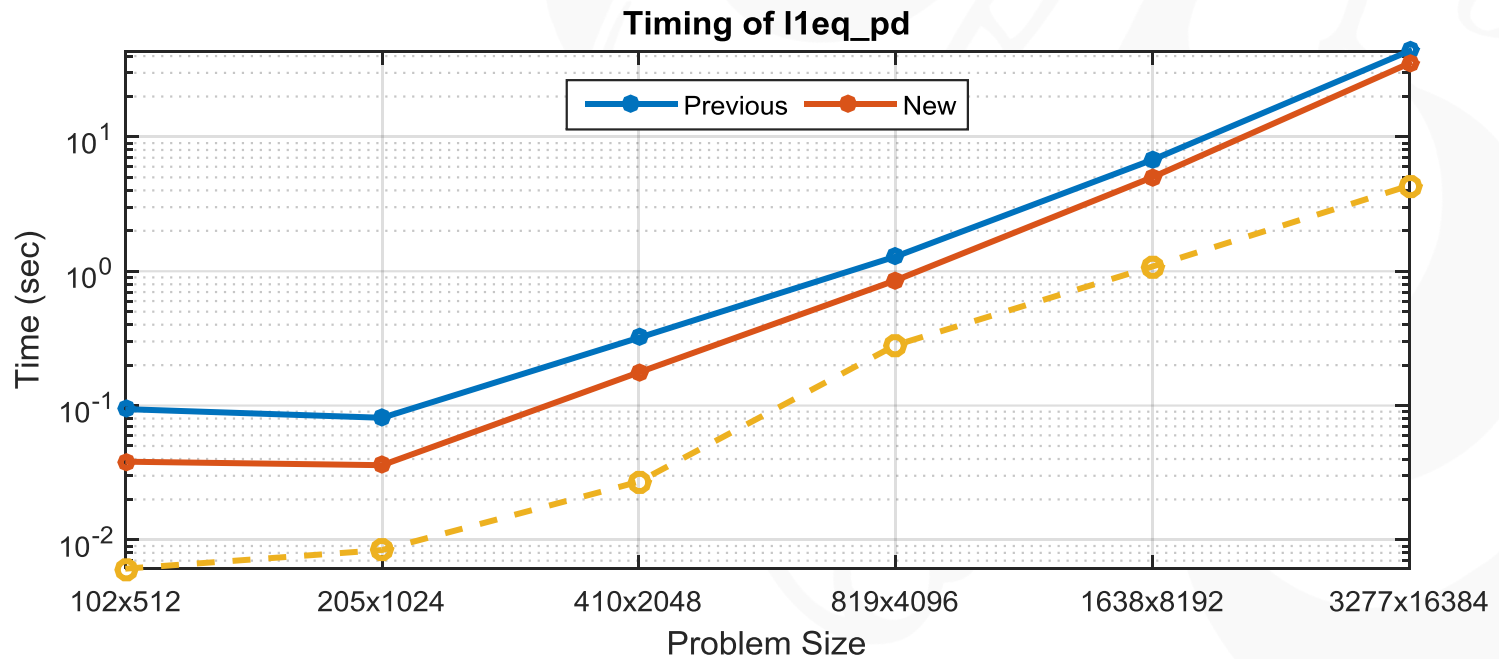


# FISTA @ $K=N/5$

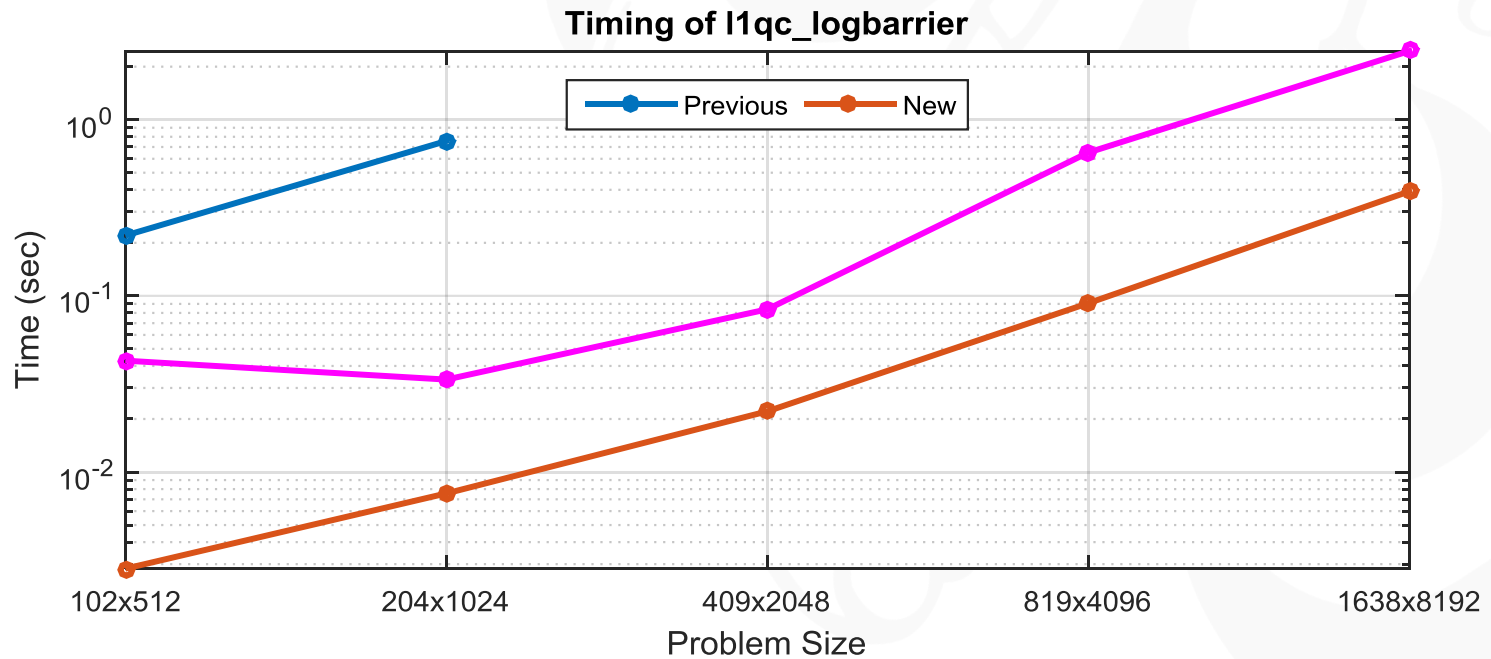




# FISTA @ $\kappa = N/5$ WITH IMPROVEMENT



# PRELIMINARY NESTA RESULTS



# CONCLUSION

- **Memory**
  - Memory Handling is important
  - If you don't run out of memory, you'll eventually run out of improvement
- **Low Rank Updates with Woodbury Identity**
  - Better Stability –  $k \times k$  instead of  $N \times N$  updates
  - Lower Memory Usage
  - Dependent on  $K \ll N$
- **Block Matrix**
  - Results Pending ... but we don't just have to call `full(...)` and `linsolve(...)` blindly for best results and order of magnitude speed increases

END.



# REFERENCES


1. S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
2. S. Conover. “L<sub>1</sub>-Magic Benchmarks and Potential Improvements.” Pending Report. September 2015.
3. FISTA: [eecs.berkeley.edu/~yang/software/l1benchmark/](http://eecs.berkeley.edu/~yang/software/l1benchmark/)
4. NESTA: [statweb.stanford.edu/~candes/nesta/](http://statweb.stanford.edu/~candes/nesta/)
5. J. Romberg. L1-Magic.  
Website: [l1-magic.org](http://l1-magic.org)  
Mirror: [statweb.stanford.edu/~candes/l1magic/](http://statweb.stanford.edu/~candes/l1magic/)

# MACHINE: ICanHaswell

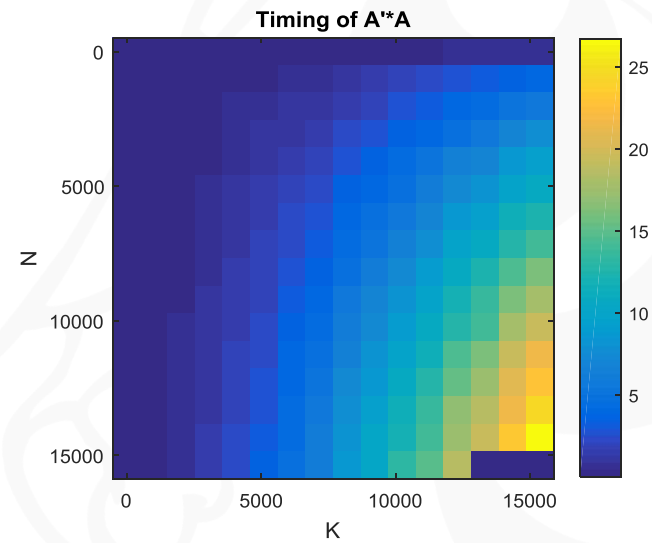
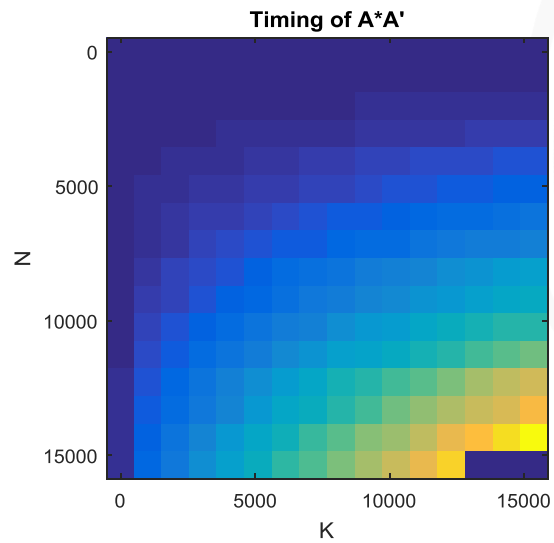
All Benchmarks Performed on the ICanHaswell development machine.

Field	Value
Name	ICanHaswell
OS	Windows 7 Pro 6.1, 7601
Motherboard	Asus
Processor	Intel Core i7-4770K @ 3.5Hz (8CPUs)
RAM	16GB
MATLAB	8.5.0.197613 (R2015a)
GPU	NVIDIA GeForce GTX 750, 4GB

# L1EQ\_PD(...) MEMORY SPIKES




# MATRIX-MATRIX MULTIPLY I





# MATRIX-MATRIX MULTIPLY II

Timing of  $A \cdot A'$



Timing of  $A \cdot A$

