

Testing fine-grained parallelism for the ADMM on a factor-graph

Ning Hao AmirReza Oghbae Mohammad Rostami Nate Derbinsky José Bento

1. Problem

There is a shortage of not-problem-specific general-purpose optimization tools that can automatically exploit GPU parallelism.

2. Focus

The Alternating Direction Method of Multipliers (ADMM).

- deals with non-smooth functions;
- better bounds for global convergence rates and variants of the ADMM than for other first order methods such as Nesterov's and gradient descent [Lessard et al. 14; França and Bento 15];
- convergence guarantees for convex problems;
- breaks problems into a series of parallel computations.

3. Questions

Can we use the ADMM to automatically exploit parallelism in general optimization without having to write problem specific parallel code?

Do we get good speedup on a GPU?

Is there an advantage on using a GPU vs other parallel frameworks?

Our numerical experiments with three different domains say

YES

4. "Typical" ADMM

$$\begin{aligned} & \text{minimize } f(w_1) + g(w_2) \\ & \text{subject to } Aw_1 + Bw_2 = c \end{aligned}$$

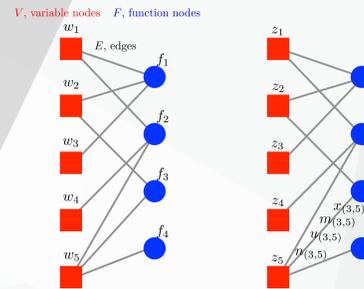
```

while !stopping criteria do
  x ← arg min_s f(s) + (ρ/2) ||As + Bz - c + u||2
  z ← arg min_r g(r) + (ρ/2) ||Ax + Br - c + u||2
  u ← u + Ax + Bz - c
end while
    
```

It does not expose parallelism immediately although for specific problem it does become visible.

5. ADMM on factor-graph

$$\begin{aligned} \min_w & f_1(w_1, w_2, w_3) + f_2(w_1, w_4, w_5) \\ & + f_3(w_2, w_5) + f_4(w_5) \end{aligned}$$



```

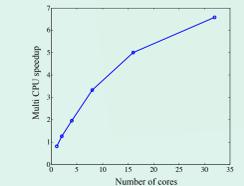
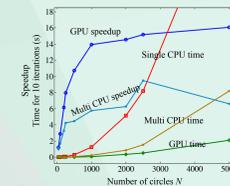
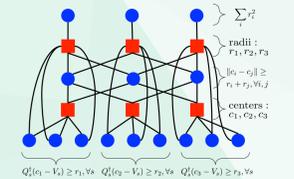
while !stopping criteria do
  for a ∈ F do
    x(a,δa) ← Proxfa,ρ(a,b)(n(a,b))
  end for
  for (a,b) ∈ E do
    m(a,b) ← x(a,b) + u(a,b)
  end for
  for b ∈ V do
    z_b ← (∑a∈δb ρ(a,b)m(a,b)) / (∑a∈δb ρ(a,b))
  end for
  for (a,b) ∈ E do
    u(a,b) ← u(a,b) + α(a,b)(x(a,b) - z_b)
  end for
  for (a,b) ∈ E do
    n(a,b) ← z_b - u(a,b)
  end for
end while
    
```

6. Testing fine-grained par.

GPU = TeslaK40; CPU = 2.8GHz AMD;
Multi CPU-cores = up to 32 × 2.8GHz AMD.
Software: <https://github.com/parADMM/>

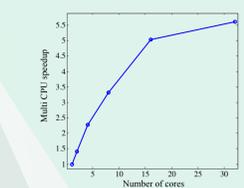
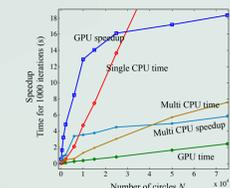
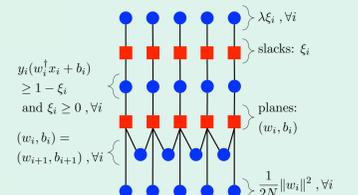
- ❖ Combinatorial optimization: What is the best way to pack N circles inside a triangle to maximize covered area?

$$\begin{aligned} & \text{minimize } \sum_{i=1}^N r_i^2 \\ & \text{subject to} \\ & \|c_i - c_j\| \geq r_i + r_j \quad \forall i, j \\ & Q_i^T(c_i - V_s) \geq r_i \quad \forall s, i \end{aligned}$$



- ❖ Machine learning: What is the best separating hyper plane for N data points labeled +1 or -1?

$$\begin{aligned} & \text{minimize } \sum_{i=1}^N \frac{1}{2N} \|w_i\|^2 + \lambda \xi_i \\ & \text{subject to} \\ & w_i = w_{i+1}, b_i = b_{i+1} \quad \forall i \\ & y_i(w_i^T x_i + b_i) \geq 1 - \xi_i \\ & \xi_i \geq 0 \quad \forall i \end{aligned}$$



- ❖ Optimal control: What is the best sequence of K inputs that filters perturbations?

$$\begin{aligned} & \text{minimize } \sum_{t=0}^{K-1} q^T(t) Q_t q(t) \\ & \quad + u^T(t) R_t u(t) \\ & \text{subject to} \\ & q(t+1) - q(t) = Aq(t) + Bu(t) \quad \forall t \\ & q(0) = q_0 \end{aligned}$$

