

# Parallel Dynamic Programming: Breaking the Parallel Block with PiCloud

Computation in California, 2013

Bo Cowgill

UC Berkeley

July 24, 2013

## Agenda:

- ▶ Brief intro.
- ▶ Tradeoffs in parallel computing.
- ▶ Sample problem: Dynamic Programming Parallelization
- ▶ Picloud engineering details
- ▶ Using PiCloud to break through parallel block.

## About me:

- ▶ BA, Stanford 2003 (Public Policy) + roughly a minor in CS.
- ▶ Google 2003-2009
- ▶ MS Statistics, Stanford 2008-2009
- ▶ UC Berkeley, PhD 2009-Present
- ▶ Google (again): Present.

## Research interests: Micro theory and empirics

- ▶ Organizations and labor
  - ▶ Hiring and referral networks.
  - ▶ “Do employee stock options incent”?
  - ▶ Prediction markets (R&R, ReStud)
- ▶ Online commerce and markets
  - ▶ Television ads and online search
  - ▶ Ad auctions, targeting and Publisher Revenue
  - ▶ SAT Cheating via Search Engines
- ▶ Papers, etc: <http://bocowgill.com>

Latest project (July 3, 2013):



## Tradeoffs in parallel computing:

- ▶ Language/platform flexibility.
- ▶ Latency and communication between cores.
- ▶ Threads spawning threads.
- ▶ Quality of scheduler and job management:
  - ▶ Load balancing.
  - ▶ Blocking and chunking.
  - ▶ Dependencies between jobs.
  - ▶ Returning “partial” results.
- ▶ Communication between threads.
- ▶ Access and marginal cost.

Dynamic programming w/ value function iteration:

- ▶  $N$  exogenous states (“ $\theta$ s”)
- ▶ Exogenous states transition via a transition matrix  $P$ .
- ▶  $K$  endogenous states (“ $x$ es”)
- ▶  $T$  iterations of the value function.

Simple, sequential/looping algorithm:

- ▶ For each of  $N$   $\theta$ s ...
- ▶ For each  $K$  of tomorrow's  $x$ s ...
- ▶ For each  $K$  of today's  $x$ s ...
- ▶ Calculate the value of choosing that  $x$
- ▶ Figure out best  $k$  after calculating them all.



## PiCloud:

- ▶ <http://picloud.com>
- ▶ “A supercomputer at your fingertips.”
- ▶ Simplified cloud computing for scientists and researchers.
- ▶ Uses Amazon’s cloud, but can plug into others.
- ▶ Google has comparable functionality internally.
- ▶ Signup today. No grant application or credit card needed.

## PiCloud:

- ▶ \$0.05 per CPU hour.
- ▶ 20 free hours per month (\$1).
- ▶ Development on your machine, deployment on theirs.
- ▶ Language of your choice (Extra goodies for Python)

Three key functions:

- ▶ `picloud exec [program]`
- ▶ `picloud result [jobid(s)]`
- ▶ `picloud join [jobid(s)]`

Additional functionality:

- ▶ Mapping
- ▶ Dependencies

Discount for people in this room:

- ▶ \$20 of free hours (400 computing hours).
- ▶ Email [support@picloud.com](mailto:support@picloud.com) and cc [bo.cowgill@gmail.com](mailto:bo.cowgill@gmail.com).
- ▶ State your account ID and say that you are with “Kenn Judd’s 2013 Computational Economics Workshop at Stanford.”

Implementation details of algorithm.