

# Scalable Text Processing with MapReduce



**Jimmy Lin**

The iSchool

University of Maryland

Saturday, July 19, 2008

2008 New Frontiers in Computing Technology



This work is licensed under a Creative Commons Attribution-Noncommercial-Share Alike 3.0 United States  
See <http://creativecommons.org/licenses/by-nc-sa/3.0/us/> for details

# Introduction

- In text processing, we've seen the emergence and dominance of:
  - Empirical techniques
  - Data-driven research
- Must scale up to larger datasets, or else:
  - Uninteresting conclusions on “toy” datasets
  - Ad hoc workarounds (e.g., approximations)
  - Unreasonably long experimental turnaround
- How do we *practically* scale up?
  - Managing concurrency is difficult
  - Clusters are expensive



# How much data?

- Wayback machine has ~2 PB (2006)
- Google processes 20 PB a day (2008)
- “all words ever spoken by human beings” ~ 5 EB
- CERN’s LHC will generate 15 PB a year (2008)
- NOAA has ~1 PB climate data (2007)



**640K** ought to be enough for anybody.



# What to do with more data?

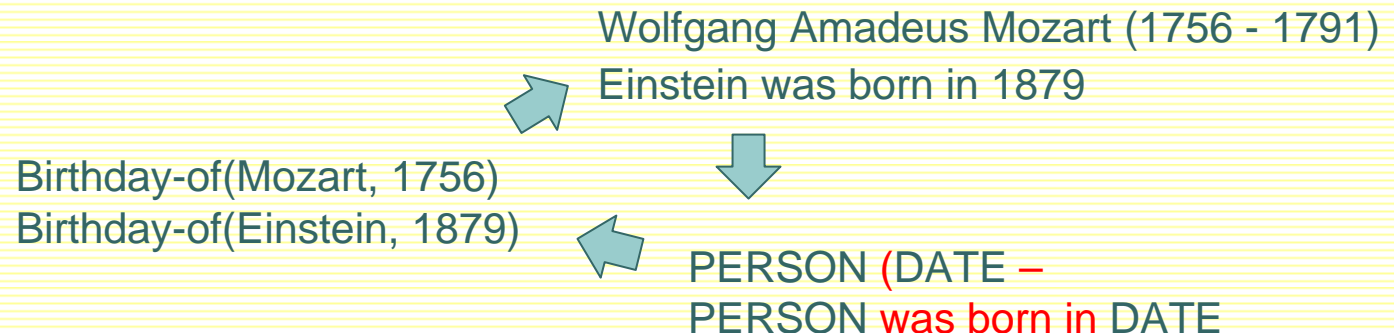
- Answering factoid questions

- Pattern matching on the Web
- Works amazingly well

Who shot Abraham Lincoln? → **X** shot Abraham Lincoln

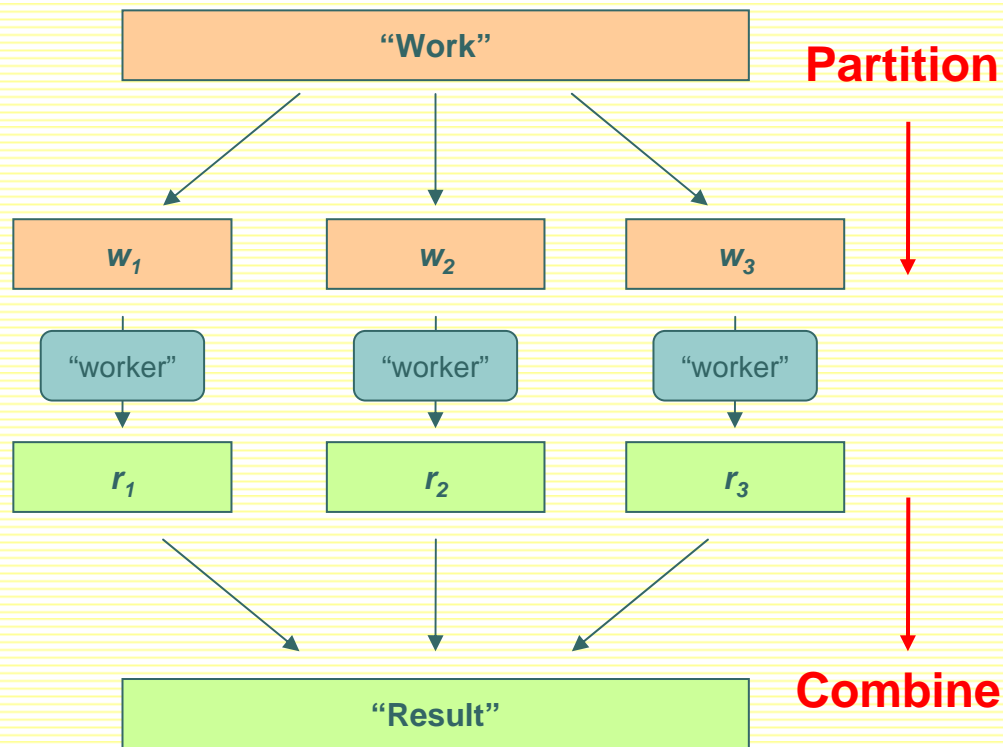
- Learning relations

- Start with seed instances
- Search for patterns on the Web
- Using patterns to find more instances



# Scaling Up: Present Solution


- Divide and conquer
- Throw more machines at it



# It's a bit more complex...

## Fundamental issues

scheduling, data distribution, synchronization, inter-process communication, robustness, fault tolerance, ...




## Architectural issues


Flynn's taxonomy (SIMD, MIMD, etc.), network typology, bisection bandwidth  
UMA vs. NUMA, cache coherence

## Different programming models

### Message Passing



### Shared Memory



## Different programming constructs

mutexes, conditional variables, barriers, ...  
masters/slaves, producers/consumers, work queues, ...

## Common problems

livelock, deadlock, data starvation, priority inversion...  
dining philosophers, sleeping barbers, cigarette smokers, ...

**The reality: programmer shoulders the burden of managing concurrency...**



# MapReduce in a Nutshell

- What's different?
  - Runtime transparently handles system-level issues
  - Programmer focuses on solving the problem
- General problem structure:
  - Iterate over a large number of records
  - **Map** Extract something of interest from each
  - Shuffle and sort intermediate results
  - Aggregate intermediate results **Reduce**
  - Generate final output
- MapReduce provides a functional abstraction:
  - Programmer supplies “Mapper” and “Reducer”
  - Runtime automatically handles everything else!



# MapReduce Runtime

- Handles scheduling
  - Assigns workers to map and reduce tasks
- Handles data distribution
  - Gets map workers to the data
- Handles synchronization
  - Shuffles intermediate key-value pairs to reduce workers
- Handles faults
  - Detects worker failures and restarts
- Everything happens on top of distributed FS
  - GFS = Google File System





# From MapReduce to Hadoop

- Google's proprietary MapReduce implementation is in C++
- Hadoop is an open-source MapReduce reimplementation in Java (lead by Yahoo)
  - HDFS is a reimplementation of GFS
  - Growing number of associated open source projects...



# “Cloud Computing” Initiative

- Google/IBM’s Academic Cloud Computing Initiative (October 2007)
  - Initial pilot institutions: Washington, Berkeley, CMU, MIT, UMD
- IBM provides UMD a Hadoop cluster
  - 20 machines (40 processors)
  - Couple of TB storage
  - Associated infrastructure support
- Maryland does good work with the cluster!
  - Use it to tackle open research problems
  - Use it in the classroom



# Statistical Machine Translation

**Chris Dyer** (Ph.D. student, Linguistics)

**Aaron Cordova** (undergraduate, Computer Science)

**Alex Mont** (undergraduate, Computer Science)

- Conceptually simple:  
(translation from foreign  $f$  into English  $e$ )

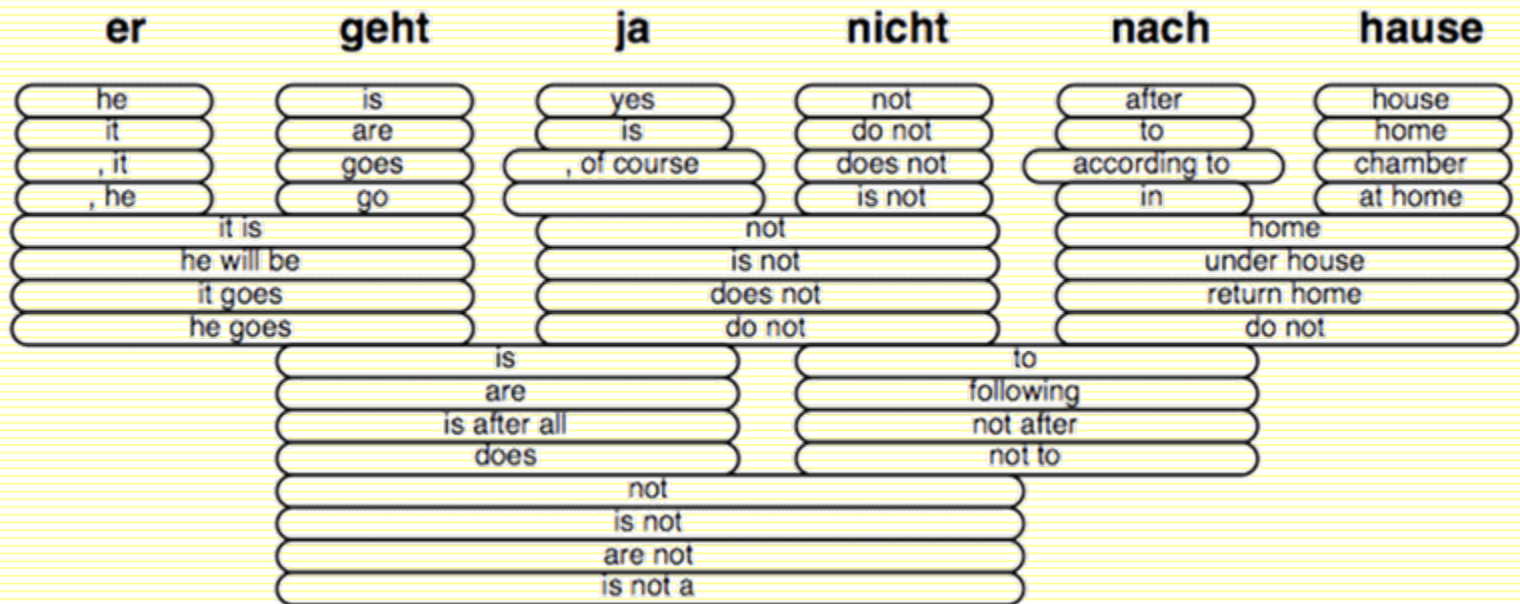
$$\hat{e} = \arg \max_e P(e | f)$$

$$\hat{e} = \arg \max_e P(f | e)P(e)$$

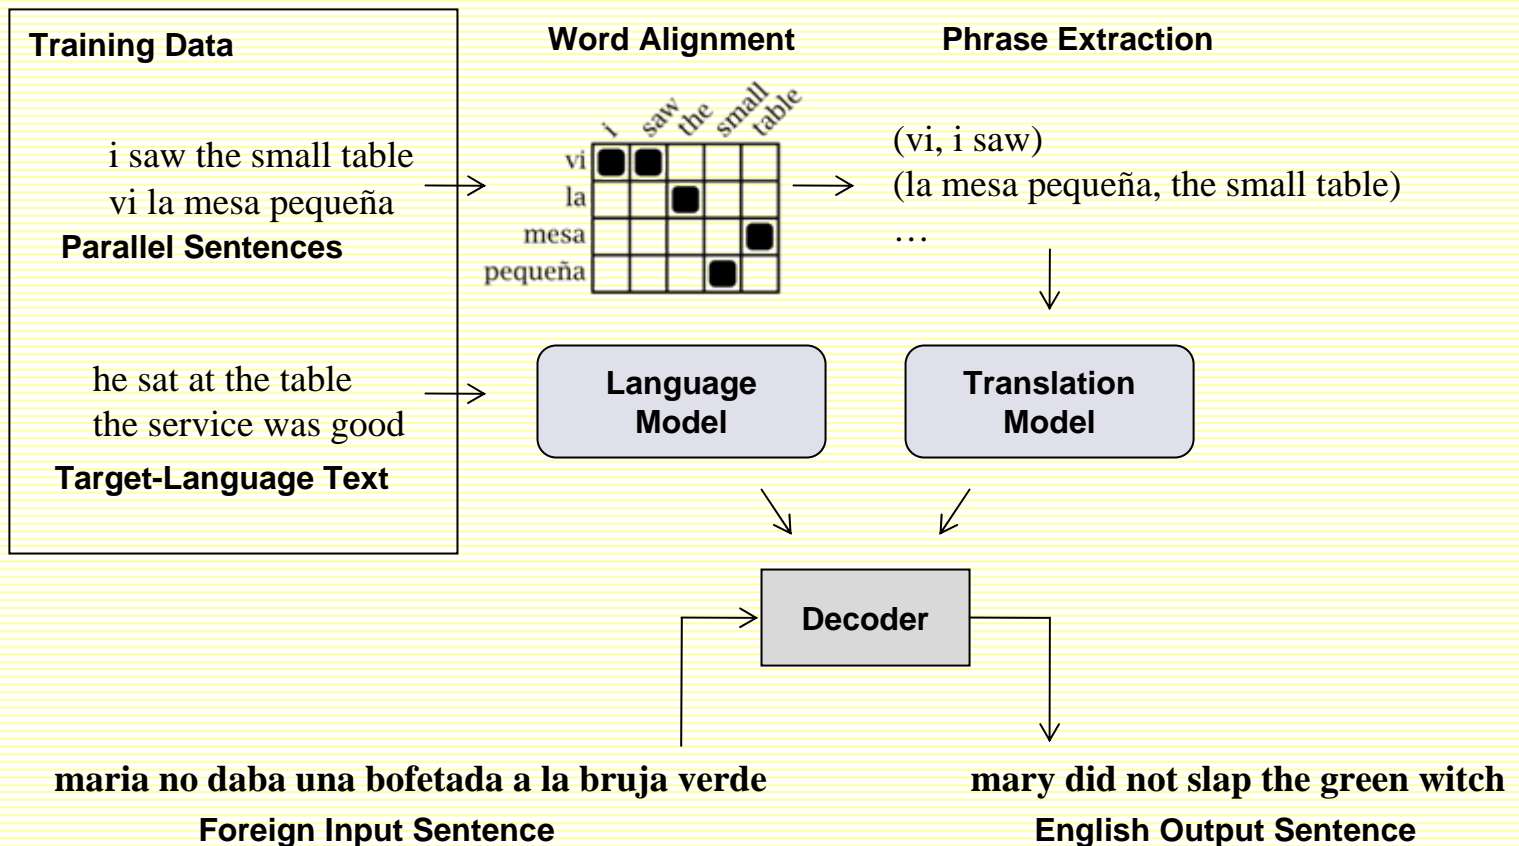
- Difficult in practice!
- Phrase-Based Machine Translation (PBMT) :
  - Break up source sentence into little pieces (phrases)
  - Translate each phrase individually



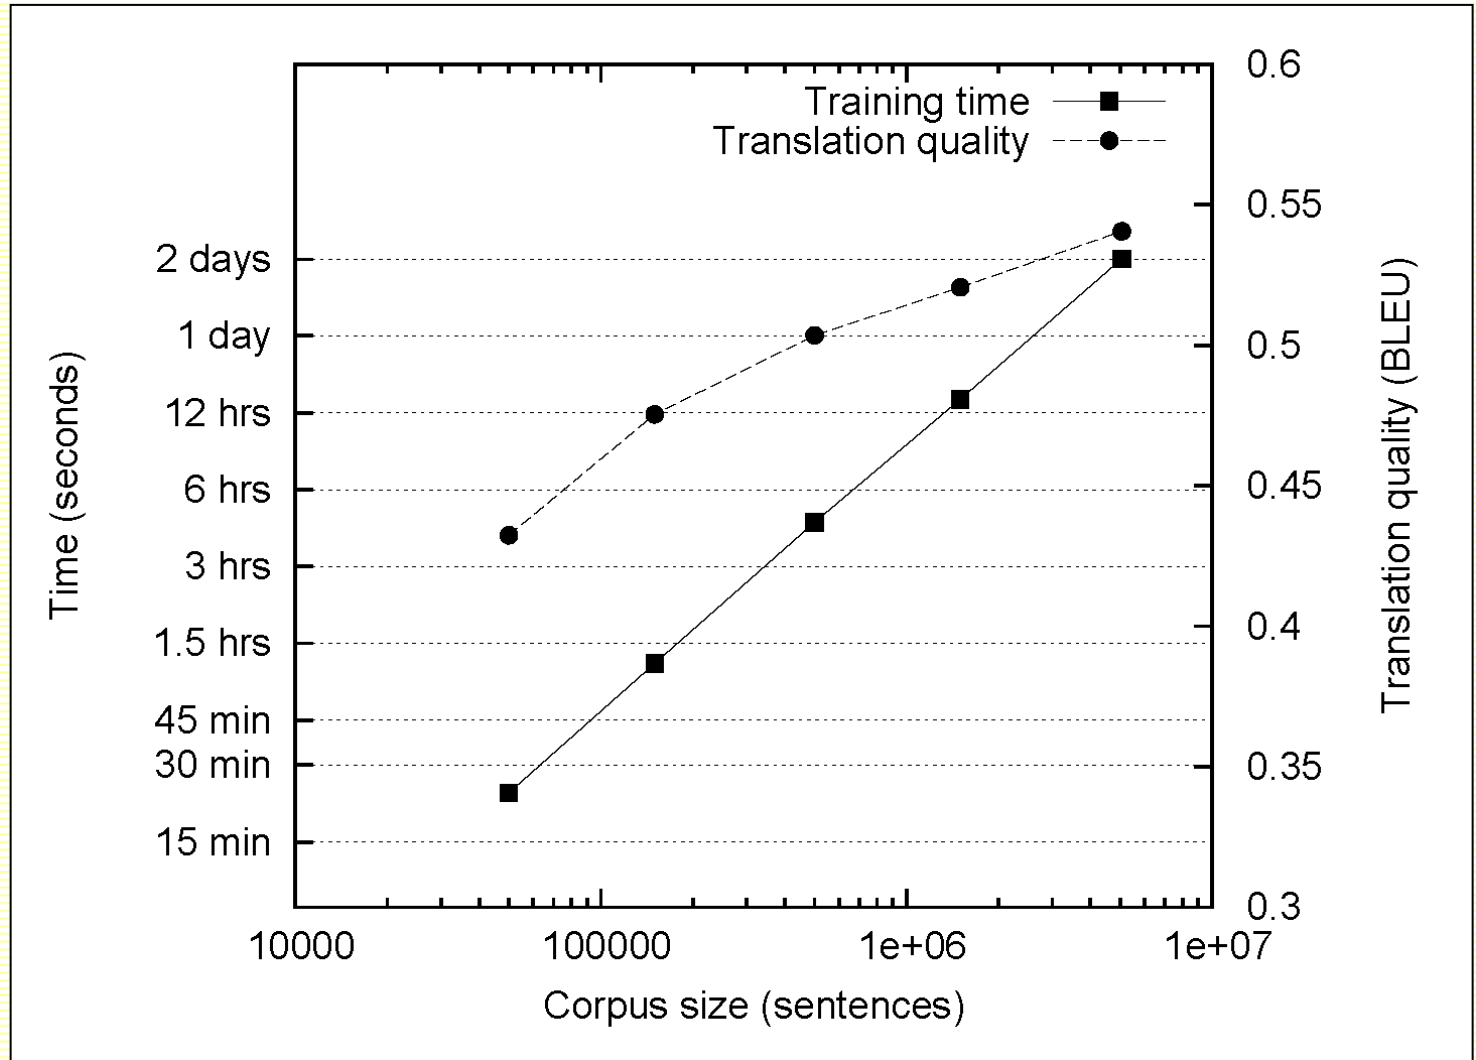
# Phrasal Decomposition



# MT Architecture

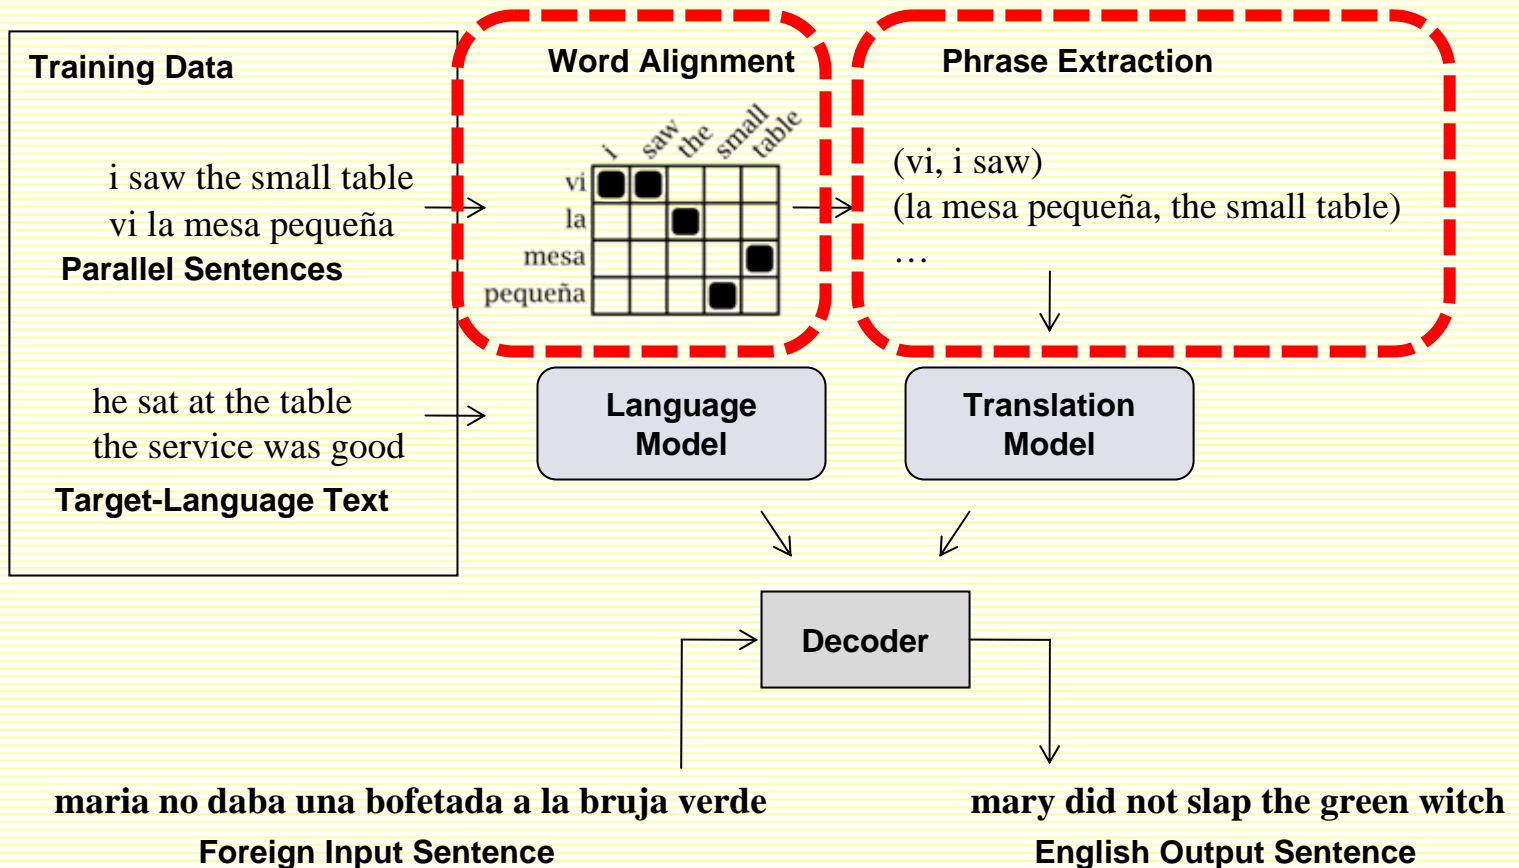


# The Data Bottleneck

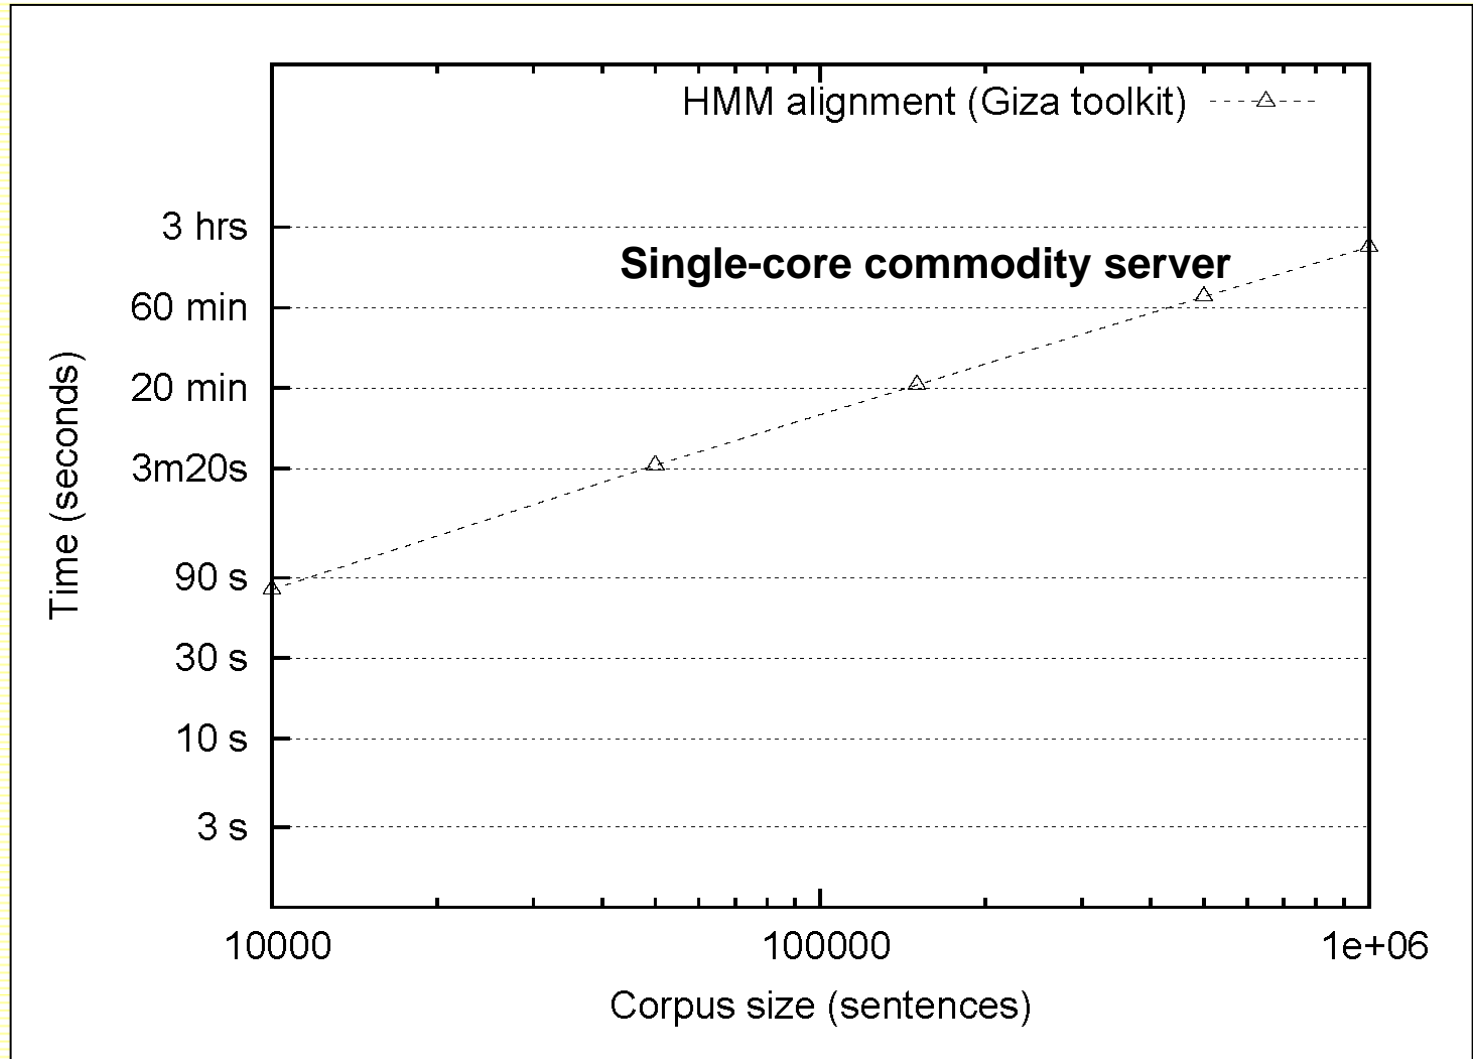


# MT Architecture

We've built MapReduce Implementations of these two components!




# HMM Alignment: Giza






# HMM Alignment: MapReduce



# HMM Alignment: MapReduce



# What's the point?

- The hypothetical, optimally-parallelized version doesn't exist!
- MapReduce occupies a sweet spot in the design space for a large class of problems:
  - **Fast**... in terms of running time
  - **Easy**... in terms of programming effort
  - **Cheap**... in terms of hardware costs

Chris Dyer, Aaron Cordova, Alex Mont, and Jimmy Lin. **Fast, Easy, and Cheap: Construction of Statistical Machine Translation Models with MapReduce.** Proceedings of the Third Workshop on Statistical Machine Translation at ACL 2008, June 2008, Columbus, Ohio.



# Beyond MapReduce

- Hadoop and HDFS provides a good start
- Everybody can play:
  - Different applications  
(e.g., from stat MT to biological sequence alignment)
  - Extensions to the programming model  
(e.g. MapReduceMerge)
  - Different hardware substrates  
(e.g., MapReduce on multicore, CELL, and GPU's)
- Development of a vibrant community
  - Academic-Industrial collaborations are the key
  - Government involvement: e.g., NSF's Cluster Exploratory (CLuE)
- Education plays a critical role!



# Acknowledgements

- **Google:** Christophe Bisciglia, et al.
- **IBM:** Dennis Quan, Eugene Hung, et al.
- Thirteen bright students from UMD:



**Chris Dyer**  
(Linguistics Ph.D.)



**Tamer Elsayed**  
(CS Ph.D.)

**Alex Mont**  
(CS ugrad)



**Greg Jablonski**  
(iSchool masters)



**Aaron Cordova**  
(CS ugrad)



**Alan Jackoway**  
(CS ugrad)

**Chang Hu**  
(CS Ph.D.)



**Denis Filimonov**  
(Linguistics Ph.D.)




**Punit Mehta**  
(iSchool masters)



**Hua Wei**  
(Geography Ph.D.)

**George Caragea**  
(CS Ph.D.)



**Mike Schatz**  
(CS Ph.D.)



**Christiam Camacho**  
(iSchool masters)

