## Probabilistic Graphical Models

## Approximate Inference: Parallel MCMC



Eric Xìng
Lecture XX, April 23rd, 2014

## Recap of MCMC

- Markov Chain Monte Carlo methods use adaptive proposals $\mathrm{Q}\left(\mathrm{x}^{\prime} \mid x\right)$ to sample from the true distribution $\mathrm{P}(\mathrm{x})$
- Metropolis-Hastings allows you to specify any proposal $\mathrm{Q}\left(\mathrm{x}^{\prime} \mid \mathrm{x}\right)$
- But choosing a good $\mathrm{Q}\left(\mathrm{x}^{\prime} \mid \mathrm{x}\right)$ requires care
- Gibbs sampling sets the proposal $Q\left(x^{\prime} \mid x\right)$ to the conditional distribution $P\left(x^{\prime} \mid x\right)$
- Acceptance rate always 1 !


## Parallel MCMC for Large Scales

- Datasets and models can be very large
- Millions to billions of data points
- Millions to billions of random variables
- Compute time measured in CPU-years
- Need GBs to TBs of memory
- E.x. Yahoo web graph has $\sim 1.4$ billion nodes and 6.6 billion edges
- Imagine doing a Markov Random Field on that network
- Without parallelism, we cannot use large datasets and models!
- Today: how to use multiple CPUs and machines in MCMC


## Taking Multiple Chains

- Proper use of MCMC actually requires parallelism
- To determine convergence, you need to take multiple MCMC chains
- Chains are independent, so you can run one chain per CPU
- Once converged, you can combine samples from all chains



## Taking Multiple Chains

- Taking multiple chains doesn't solve all issues, though
- If burn-in is long, then all chains will take a long time to converge!
- We need a way to take each sample faster...



## Idea: Run Gibbs Sampling in Parallel?



- Recall the alarm network
- Initialize all variables at $t=0$ to False
- Idea: parallel Gibbs sample all variables at step t conditioned on t-1


## Naïve Parallel Gibbs Sampling



- Sampling $\mathrm{P}(\mathrm{B} \mid \mathrm{A}, \mathrm{E})$ at $\mathrm{t}=1$ : Using Bayes Rule,

$$
P(B \mid A, E) \propto P(A \mid B, E) P(B)
$$

- $(\mathrm{A}, \mathrm{E})=(\mathrm{F}, \mathrm{F})$, so we compute the following, and sample $\mathrm{B}=\mathrm{F}$

$$
\begin{aligned}
& P(B=T \mid A=F, E=F) \propto(0.06)(0.01)=0.0006 \\
& P(B=F \mid A=F, E=F) \propto(0.999)(0.999)=0.9980
\end{aligned}
$$

## Naïve Parallel Gibbs Sampling



- Sampling $\mathrm{P}(\mathrm{E} \mid \mathrm{A}, \mathrm{B})$ : Using Bayes Rule,

$$
P(E \mid A, B) \propto P(A \mid B, E) P(E)
$$

- $(A, B)=(F, F)$, so we compute the following, and sample $\mathrm{E}=\mathrm{T}$

$$
\begin{aligned}
& P(E=T \mid A=F, B=F) \propto(0.71)(0.02)=0.0142 \\
& P(E=F \mid A=F, B=F) \propto(0.999)(0.998)=0.9970
\end{aligned}
$$

## Naïve Parallel Gibbs Sampling



- Notice the difference
- Normal Gibbs sampling: compute $\mathrm{P}(\mathrm{E} \mid \mathrm{A}, \mathrm{B})$ based on $\mathrm{B}_{\mathrm{t}=1}, \mathrm{~A}_{\mathrm{t}=0}$
- Naïve Parallel GS: compute $P(E \mid A, B)$ based on $B_{t=0}, A_{t=0}$
- At step $t$, always condition on $\mathrm{t}-1$ instead of most recently sampled value


## Naïve Parallel Gibbs Sampling



- Sampling $\mathrm{P}(\mathrm{A} \mid \mathrm{B}, \mathrm{E}, \mathrm{J}, \mathrm{M})$ : Using Bayes Rule,

$$
P(A \mid B, E, J, M) \propto P(J \mid A) P(M \mid A) P(A \mid B, E)
$$

- $\quad(B, E, J, M)=(F, F, F, F)$, so we compute the following, and sample $A=F$

$$
\begin{aligned}
& P(A=T \mid B=F, E=F, J=F, M=F) \propto(0.1)(0.3)(0.001)=0.00003 \\
& P(A=F \mid B=F, E=F, J=F, M=F) \propto(0.95)(0.99)(0.999)=0.9396
\end{aligned}
$$

## Naïve Parallel Gibbs Sampling



- Sampling $P(J \mid A)$ : No need to apply Bayes Rule
- $\mathrm{A}=\mathrm{F}$, so we compute the following, and sample $\mathrm{J}=\mathrm{T}$

$$
\begin{aligned}
& P(J=T \mid A=F) \propto 0.05 \\
& P(J=F \mid A=F) \propto 0.95
\end{aligned}
$$

## Naïve Parallel Gibbs Sampling



- Sampling $P(M \mid A)$ : No need to apply Bayes Rule
- $A=F$, so we compute the following, and sample $\mathrm{M}=\mathrm{F}$

$$
\begin{aligned}
& P(M=T \mid A=F) \propto 0.01 \\
& P(M=F \mid A=F) \propto 0.99
\end{aligned}
$$

## Naïve Parallel Gibbs Sampling



- We just finished sampling variables $t=1$
- Why is the sampling parallelizable?
- We only conditioned on variable state at $\mathrm{t}=0$, which is known in advance!
- We can sample B,E,A,J,M on separate processors, without having to send information between processors


## Naïve Parallel Gibbs Sampling

- In practice, works very well for some graphical models
- E.g. collapsed Gibbs Sampling for LDA

$$
P\left(z_{i}=j \mid \mathbf{z}_{-i}, \mathbf{w}\right) \propto \frac{n_{-i, j}^{\left(w_{i}\right)}+\beta}{n_{-i, j}^{(\cdot)}+W \beta} \frac{n_{-i, j}^{\left(d_{i}\right)}+\alpha}{n_{-i, .}^{\left(d_{i}\right)}+T \alpha}
$$

- Just assign different $z_{i}^{\prime}$ s to different processors or machines
- But there's a problem...


## Where Naïve Parallel GS Fails

- Naïve Parallel GS may not converge to the stationary distribution
- Consider the following Bayes Net:

- Essentially an XOR relation between $(A, B)$ and $(A, C)$
- Joint distribution $\mathrm{P}(\mathrm{A}, \mathrm{B}, \mathrm{C})$ has only 8 states, so we can compute the stationary distribution. It is dominated by 2 equally-probable states:
- $(A, B, C)=(T, F, T)$ and $(A, B, C)=(F, T, F)$


## Where Naïve Parallel GS Fails



- Let's initialize $(A, B, C)=(F, F, F)$ and see what happens when we naively Gibbs sample in parallel...


## Where Naïve Parallel GS Fails



- Sampling $\mathrm{P}(\mathrm{A} \mid \mathrm{B}, \mathrm{C})$ :

$$
P(A \mid B, C) \propto P(B \mid A) P(C \mid A)
$$

- $(\mathrm{B}, \mathrm{C})=(\mathrm{F}, \mathrm{F})$ so we sample $\mathrm{A}=\mathrm{T}$

$$
\begin{gathered}
P(A=T \mid B=F, C=F) \propto(0.999)(0.999) \approx 1 \\
P(A=F \mid B=F, C=F) \propto(0.001)(0.001) \approx 0 \\
\text { 〇Eric xing @ сми, 2005-2014 }
\end{gathered}
$$

## Where Naïve Parallel GS Fails



| $\mathbf{t}$ | $\mathbf{A}$ | $\mathbf{B}$ | $\mathbf{C}$ |
| :--- | :--- | :--- | :--- |
| 0 | F | F | F |
| 1 | T | T |  |
| 2 |  |  |  |
| 3 |  |  |  |
| 4 |  |  |  |

- Sampling $P(B \mid A)$ : No need to apply Bayes Rule
- $\mathrm{A}=\mathrm{F}$ so we sample $\mathrm{B}=\mathrm{T}$

$$
\begin{gathered}
P(B=T \mid A=F) \propto(0.999) \approx 1 \\
P(B=F \mid A=F) \propto(0.001) \approx 0 \\
\text { ©Eric Xing @ Cmu, 2005-2014 }
\end{gathered}
$$

## Where Naïve Parallel GS Fails



- Sampling $P(C \mid A)$ : No need to apply Bayes Rule
- $\mathrm{A}=\mathrm{F}$ so we sample $\mathrm{C}=\mathrm{T}$

$$
\begin{gathered}
P(C=T \mid A=F) \propto(0.999) \approx 1 \\
P(C=F \mid A=F) \propto(0.001) \approx 0 \\
\text { © Enic xing @ CMU, 2005-2014 }
\end{gathered}
$$

## Where Naïve Parallel GS Fails



- Easy to see that at $t=2$, we will get $(A, B, C)=(F, F, F)$


## Where Naïve Parallel GS Fails



- Easy to see that at $t=2$, we will get $(A, B, C)=(F, F, F)$
- At $t=3,(A, B, C)=(T, T, T)$


## Where Naïve Parallel GS Fails



- Easy to see that at $t=2$, we will get $(A, B, C)=(F, F, F)$
- At $t=3,(A, B, C)=(T, T, T)$
- At $t=4,(A, B, C)=(F, F, F)$


## Where Naïve Parallel GS Fails



- Easy to see that at $t=2$, we will get $(A, B, C)=(F, F, F)$
- At $t=3,(A, B, C)=(T, T, T)$
- $A t=4,(A, B, C)=(F, F, F)$
- Can you see the problem?


## Where Naïve Parallel GS Fails

- We know the stationary distribution is [(F,T,F), (T,F,T)]
- But naïve parallel GS gets stuck in [(T,T,T), (F,F,F)]

- Naïve parallel GS performs poorly on near-discrete distributions
- What is the correct way to Gibbs sample in parallel?


## Correct Parallel Gibbs Sampling

- Recall that in MRFs, we Gibbs sample by sampling from $P(x \mid$ $\mathrm{MB}(\mathrm{x})$ ), the conditional distribution of x given its Markov Blanket MB(x)
- For MRFs, the Markov Blanket of $x$ is just its neighbors
- In the MRF below, the red node's Markov Blanket consists of the blue nodes

© Eric Xing @ CMU, 2005-2014


## Correct Parallel Gibbs Sampling

- Observe that we can correctly Gibbs sample the two green nodes simultaneously
- Neither node is part of the other's Markov Blanket, so their conditional distributions do not depend on each other
- Sampling one of the green nodes doesn't change the conditional distribution of the other node!

© Eric Xing @ CMU, 2005-2014


## Correct Parallel Gibbs Sampling

- How do we generalize this idea to the whole graph?
- Find subsets of nodes, such that all nodes in a given subset are not in each other's Markov Blankets, and the subsets cover the whole graph
- The subsets should be as large as possible
- Because we can Gibbs sample all nodes in a subset at the same time
- At the same time, we want as few subsets as possible
- The Markov Blankets of different subsets overlap, so they cannot be sampled at the same time. We must process the subsets sequentially.

© Eric Xing @ CMU, 2005-2014


## Correct Parallel Gibbs Sampling

- We can find these covering subsets with k-coloring algorithms (Gonzales et al., 2011)
- A k-coloring algorithm colors a graph using k colors, such that:
- Every node gets one color
- No edge has two nodes of the same color
- Trees always admit a 2-coloring (e.g. below)
- Assign one color to some node, and alternate colors as you move away

© Eric Xing @ CMU, 2005-2014


## Correct Parallel Gibbs Sampling

- Bipartite graphs are always 2-colorable
- Color each side of the bipartite graph with opposite colors
- e.x. Latent Dirichlet Allocation model is bipartite
- However, not all graphs have $k$-colorings for all $k \geq 2$
- In the worst case, a graph with $n$ nodes can require $n$ colors
- The full clique is one such graph
- Determining if a graph is $k$-colorable for $k>2$ is NP-complete
- In practice, we employ heuristics to find k-colorings
- Instead of using k-colorings, why not just Gibbs sample all variables at the same time?
- The Markov Chain may become non-ergodic, and is no longer guaranteed to converge to the stationary distribution!


## Online ParalleI MCMC

- In "online" algorithms, we need to process new data points one-at-a-time
- Moreover, we have to "forget" older data points because memory is finite
- For such applications to be viable, we can only afford constant time work per new data point
- Otherwise we will reach a point where new data can no longer be processed in a reasonable amount of time
- We also want the algorithm to be parallel for scaling up
- What MCMC techniques can we use to make an online parallel algorithm?


## Sequential Monte Carlo

- SMC is a generalization of Particle Filters
- Recall that PFs incrementally sample $P\left(X_{t} \mid Y_{1: t}\right)$, where the $X s$ are latent r.v.s and the $Y$ s are observations under a state-space model
- SMC does not assume the GM is a state-space model, or has any particular structure at all
- Suppose we have n r.v.s $\mathrm{x}_{1}, \ldots, \mathrm{x}_{\mathrm{n}}$
- SMC first draws samples from the marginal distribution $P\left(x_{1}\right)$, then $\mathrm{P}\left(\mathrm{x}_{1: 2}\right)$, and so on until $\mathrm{P}\left(\mathrm{x}_{1: n}\right)$
- Key idea: Construct proposals such that we sample from $P\left(x_{1: k+1}\right)$ in constant time, given samples from $\mathrm{P}\left(\mathrm{x}_{1: \mathrm{k}}\right)$
- Like other MCMC algorithms, we only require that we can evaluate $P^{\prime}\left(x_{1: n}\right)=a P\left(x_{1: n}\right)$ for some unknown a


## Sequential Importance Sampling

- SIS is the foundation of Sequential Monte Carlo
- It allows new variables to be sampled in constant time, without resampling older variables
- SIS uses proposal distributions with the following structure:

$$
\begin{aligned}
q_{n}\left(x_{1: n}\right) & =q_{n-1}\left(x_{1: n-1}\right) q_{n}\left(x_{n} \mid x_{1: n-1}\right) \\
& =q_{1}\left(x_{1}\right) \prod_{k=2}^{n} q_{k}\left(x_{k} \mid x_{1: k-1}\right)
\end{aligned}
$$

- Notice we can propose $x_{k+1}$ if we've already drawn $x_{1: k}$, without having to redraw $\mathrm{X}_{1 \text { :k }}$


## Sequential Importance Sampling

- In normalized importance sampling, recall how the sample weights $w^{i}$ are defined:

$$
\begin{gathered}
\langle f(x)\rangle_{p}=\sum_{i} f\left(x^{i}\right) w^{i} \\
\text { where } \quad w^{i}=\frac{r^{i}}{\sum_{r^{\prime}}} \text { and } r^{i}=\frac{P^{\prime}\left(x^{i}\right)}{Q\left(x^{i}\right)}
\end{gathered}
$$

- In SIS, the unnormalized weights $r$ can be rewritten as a telescoping product:

$$
\begin{aligned}
r\left(x_{1: n}\right) & =\frac{P_{n}^{\prime}\left(x_{1: n}\right)}{q_{n}\left(x_{1: n}\right)} \\
& =\frac{P_{n-1}^{\prime}\left(x_{1: n-1}\right)}{q_{n-1}\left(x_{1: n-1}\right)} \frac{P_{n}^{\prime}\left(x_{1: n}\right)}{P_{n-1}^{\prime}\left(x_{1: n-1}\right) q_{n}\left(x_{n} \mid x_{1: n-1}\right)} \quad \text { where } \quad \alpha_{n}\left(x_{1: n}\right)=\frac{P_{n}^{\prime}\left(x_{1: n}\right)}{P_{n-1}^{\prime}\left(x_{1: n-1}\right) q_{n}\left(x_{n} \mid x_{1: n-1}\right)} \\
& =r_{n-1}\left(x_{1: n-1}\right) \alpha_{n}\left(x_{1: n}\right) \\
& =r_{1}\left(x_{1}\right) \prod_{k=2}^{n} \alpha_{k}\left(x_{1: k}\right)
\end{aligned}
$$

## Sequential Importance Sampling

$$
r\left(x_{1: n}\right)=r_{1}\left(x_{1}\right) \prod_{k=2}^{n} \alpha_{k}\left(x_{1: k}\right) \quad \text { where } \quad \alpha_{n}\left(x_{1: n}\right)=\frac{P_{n}^{\prime}\left(x_{1: n}\right)}{P_{n-1}^{\prime}\left(x_{1: n-1}\right) q_{n}\left(x_{n} \mid x_{1: n-1}\right)}
$$

- This means the unnormalized weights $r$ can be computed incrementally
- Compute $\alpha_{n}$ and use it to update $r\left(x_{1: n-1}\right)$ to $r\left(x_{1: n}\right)$
- NB: For this update to be constant time, we also require $\mathrm{P}_{\mathrm{n}}\left(\mathrm{x}_{1: n}\right)$ to be computable from $P_{n-1}^{\prime}\left(x_{1: n-1}\right)$ in constant time
- We remember the unnormalized weights $r$ at each iteration, and compute the normalized weights $w$ as needed from $r$
- Thus, we can sample x AND compute the normalized weights w using constant time per new variable $x_{n}$
- So SIS meets the requirements for an online inference algorithm!
- Even better, the samples don't depend on each other
- Assign one CPU core per sample to make the SIS algorithm paralle!!


## Sequential Importance Sampling

- SIS algorithm:
- At time $\mathrm{n}=1$
- Parallel draw samples $x_{1}{ }_{1} \sim q_{1}\left(x_{1}\right)$
- Parallel compute unnormalized weights $r_{1}^{i}=P_{1}^{\prime}\left(x_{1}^{i}\right) / q_{1}\left(x_{1}^{i}\right)$
- Compute normalized weights $w_{1}{ }_{1}$ by normalizing $r_{1}{ }_{1}$
- Although this step is sequential, it takes almost no time to perform
- At time $\mathrm{n} \geq 2$
- Parallel draw samples $x_{n}^{i} \sim q_{n}\left(x_{n} \mid x_{1: n-1}^{i}\right)$
- Parallel compute unnorm. wgts. $\quad r_{n}^{i}=r_{n-1}^{i} \alpha_{n}\left(x_{1: n}^{i}\right)=r_{n-1}^{i} \frac{P_{n}^{\prime}\left(x_{1: n}^{i}\right)}{P_{n-1}^{\prime}\left(x_{1: n-1}^{i}\right) q_{n}\left(x_{n}^{i} \mid x_{1: n-1}^{i}\right)}$
- Compute normalized weights $w_{n}^{i}$ by normalizing $r_{n}^{i}$
- Although this step is sequential, it takes almost no time to perform


## Sequential Importance Sampling

- But we are not done yet!
- Unfortunately, SIS suffers from a severe drawback: the variance of the samples increases exponentially with $n$ !
- See eq (31) of Doucet's SMC tutorial for an example
- Resampling at each iteration will decrease the sample variance!
- Similar to weighted resampling from the first MC lecture!


## Multinomial Resampling

- Suppose we have $m$ samples $x^{1}, \ldots, x^{m}$ with corresponding importance weights $w^{1}, \ldots, w^{m}$
- Construct a categorical distribution from these samples:
- This distribution has m categories (choices)
- The probability of drawing category $k$ is $w^{k}$
- Drawing category k gets us $\mathrm{x}^{\mathrm{k}}$
- To resample, just draw N times from this distribution
- Note that N can be greater/less than m !
- For more advanced strategies such as systematic and residual resampling, refer to page 13 of Doucet's SMC tutorial


## Why Resample?

- Apart from decreasing variance, there are other reasons...
- Resampling removes samples $x^{k}$ with low weights $w^{k}$
- Low-weight samples come from low-probability regions of $P(x)$
- We want to focus computation on high-probability regions of $\mathrm{P}(\mathrm{x})$
- Notice that each sample gets an equal amount of computation, regardless of its weight $w_{k}$
- Resampling ensures that more computation is spent on samples $x_{k}$ that come from high-probability regions of $P(x)$
- Resampling prevents a small number of samples $x_{k}$ from dominating the empirical distribution
- Resampling resets all weights $w_{k}$ to $1 / \mathrm{N}$
- This prevents sample weights $w_{k}$ from growing until they reach 1


## Sequential Monte Carlo

- The SMC algorithm is just SIS with resampling:
- At time $\mathrm{n}=1$
- Parallel draw samples $x_{1}{ }_{1} \sim q_{1}\left(x_{1}\right)$
- Parallel compute unnormalized weights $r_{1}^{i}=P_{1}^{\prime}\left(x_{1}^{i}\right) / q_{1}\left(x_{1}^{i}\right)$
- Compute normalized weights $w_{1}{ }_{1}$ by normalizing $r_{1}{ }_{1}$
- Parallel resample $w_{1}^{i}, x_{1}^{i}$ into $N$ equally-weighted particles $x_{1}^{i}$
- At time $\mathrm{n} \geq 2$
- Parallel draw samples $x_{n}^{i} \sim q_{n}\left(x_{n} \mid x_{1: n-1}^{i}\right)$
- Parallel compute unnorm. wgts. $\quad r_{n}^{i}=r_{n-1}^{i} \alpha_{n}\left(x_{1: n}^{i}\right)=r_{n-1}^{i} \frac{P_{n}^{\prime}\left(x_{1: n}^{i}\right)}{P_{n-1}^{\prime}\left(x_{1: n-1}^{i}\right) q_{n}\left(x_{n}^{i} \mid x_{1: n-1}^{i}\right)}$
- Compute normalized weights $w_{n}^{i}$ by normalizing $r_{n}^{i}$
- Parallel resample $w_{n}^{i}, x_{1: n}^{i}$ into $N$ equally-weighted particles $x_{1: n}^{i}$


## Summary

- Parallel Gibbs sampling
- Naïve strategy: sample all variables at the same time
- Correct strategy: perform graph colorings and sample same-colored nodes in parallel
- Sequential Monte Carlo
- Uses incremental proposal distributions
- Provides a framework for designing online, parallel MCMC algorithms


## Parallel Inference for Bayesian Nonparametric

- Dirichlet Process Mixture Model (recap)
- Inference schemes (recap)
- Parallel inference schemes
- Results


# Finite Mixture Model:- Restaurant Perspective 

People sit on the table with the most preferred dish/color


## Finite Mixture Model:- Restaurant Perspective

- Table:
- Cluster
- People:
- Items to be clustered
- Parameters:
- Dish/color on each table
- Center of each cluster
- Hidden Variable:
- Assignment of people to each table


# Finite Mixture Model:- Restaurant Perspective 

People sit on the table with the most preferred dish/color


Which clustering algorithm will it lead to?

# Finite Mixture Model:- Restaurant Perspective 

People sit on the table with the most preferred dish/color


Which clustering algorithm will it lead to?

Hard Kmeans

## Finite Mixture Model:- Restaurant Perspective

People sit on the table proportional to appreciation of dish/color


Which clustering algorithm will it lead to?

## Finite Mixture Model:- Restaurant Perspective

People sit on the table proportional to appreciation of dish/color


Which clustering algorithm will it lead to?
Soft Kmeans

## Soft Kmeans Generative Model

$$
\begin{gathered}
\text { for } k=1, \ldots K \\
\eta_{k} \sim H \\
\text { for } i=1, \ldots N \\
Z_{i} \sim U(1, K) \\
X_{i} \sim f\left(\eta_{z i}\right)
\end{gathered}
$$



# Finite Mixture Model:- Restaurant Perspective 

People sit on the table proportional to appreciation of dish/color and number of people sitting on the table


Which clustering algorithm will it lead to?
Dirichlet Distribution
Mixture Model

## Finite MM Generative Model



# Finite Mixture Model:- Restaurant Perspective 

People sit on the table proportional to appreciation of dish/color and number of people sitting on the table


Which clustering algorithm will it lead to?
Dirichlet Distribution
Mixture Model

## Infinite Mixture Model:Restaurant Perspective

People sit on the table proportional to appreciation of dish/color and number of people sitting on the table


## Infinite Mixture Model:Restaurant Perspective

People sit on the table proportional to appreciation of dish/color and


## Turning the definition

Proportional to selecting a table

Dish on the table


## Stick Breaking Construction

Step 1:-Take a stick of unit length


## Stick Breaking Construction

Proportional to selecting a table

Dish on the table


## Graphical Model Representation



## Inference

- Gibbs Sampling:-
- Sample each of the variable given the rest.
- Variables to sample are table proportion $\mathrm{V}_{\mathrm{k}}$, table assignment to each customer ( $Z$ ) and dish at each table $\eta$



## Inference

- Gibbs Sampling:-
- Sample each of the variable given the rest.
- Variables to sample are table proportion $\mathrm{V}_{\mathrm{k}}$, table assignment to each customer ( $Z$ ) and dish at each table $\eta$
- Parallel inference: Easy



## Inference

- Gibbs Sampling:-
- Sample each of the variable given the rest.
- Variables to sample are table proportion $\mathrm{V}_{\mathrm{k}}$, table assignment to each customer ( $Z$ ) and dish at each table $\eta$
- Parallel inference: Easy
- Poor mixing



## Inference

- Collapsed Gibbs Sampler:-
- Integrate out $\mathrm{V}_{\mathrm{k}}$ and $\eta_{\mathrm{k}}$
- Leads to better mixing
- Parallel inference: Hard



## Inference

- Collapsed Gibbs suffer from large computational cost

Running Example:
10 million data points to be clustered.


## Inference

- Variational Inference
- Approximate the posterior with a distribution belonging to a more manageable family of distribution
- Parallel inference: Easy
- Search within a restricted class of models, looses the expressiveness
- Typically less accuracy than MCMC methods

© Eric Xing @ CMU, 2005-2014


## Inference

- Sequential Monte Carlo Method:-
- Keep a pools of particles, approximate the distribution using weighted combination of the pool
- Parallel inference: Easy
- High variance for naïve implementation, needs resampling (MCMC )

© Eric Xing @ CMU, 2005-2014


## Parallel MCMC

- Naïve
- Run collapsed sampler on individual core
- Combine the result approximately !!



## Parallel MCMC

- Naïve
- Run collapsed sampler on individual core
- Combine the result approximately !!
- How
- Why should two newly discovered clustered in two different processor be the same?

© Eric Xing @ CMU, 2005-2014


## Parallel MCMC

- Idea: Dirichlet Mixture of Dirichlet processes are Dirichlet processes
- Skeptic
(proof coming)


## Parallel MCMC

- Idea: Dirichlet Mixture of Dirichlet processes are Dirichlet processes


Restaurant 1

$$
D_{j} \sim \operatorname{DP}\left(\frac{\alpha}{P}, H\right), \quad j=1, \ldots, P
$$



Restaurant $P$

## Parallel MCMC

- Idea: Dirichlet Mixture of Dirichlet processes are Dirichlet processes



## Parallel MCMC

- Idea: Dirichlet Mixture of Dirichlet processes are Dirichlet processes



## Auxiliary Variable Model For DP

- The generative process is as follows :-

$$
\begin{aligned}
D_{j} & \sim \operatorname{DP}\left(\frac{\alpha}{P}, H\right), \quad j=1, \ldots, P \\
\phi & \sim \text { Dirichlet }\left(\frac{\alpha}{P}, \ldots, \frac{\alpha}{P}\right) \\
\pi_{i} & \sim \phi \\
\theta_{i} & \sim D_{\pi_{i}} \\
x_{i} & \sim f\left(\theta_{i}\right), \quad i=1, \ldots, N .
\end{aligned}
$$

## Proof

- If $G \sim D P\left(\alpha, G_{0}\right)$ and $\theta_{1} \sim G$ Then posterior distribution is given by:

$$
\theta_{n+1} \mid \theta_{1}, \ldots, \theta_{n} \sim \sum_{l=1}^{n} \frac{1}{n+\alpha} \delta_{\theta_{l}}+\frac{\alpha}{n+\alpha} G_{0}
$$

- If $D_{j} \sim \operatorname{DP}\left(\alpha / P, G_{0}\right), \phi \sim \operatorname{Dir}\left(\frac{\alpha}{P}, \ldots, \frac{\alpha}{P}\right), \pi_{i} \sim \phi \quad$ and $\theta_{i} \sim D_{\phi_{i}}$, Then

$$
\begin{aligned}
\theta_{n+1} \mid \theta_{1}, \ldots, \theta_{n} \sim & \sum_{j=1}^{P} P\left(\pi_{n+1}=j \mid \pi_{1}, \ldots, \pi_{n}\right) \\
& P\left(\theta_{n+1} \mid \pi_{n+1}=j, \pi_{1}, \ldots, \pi_{n}, \theta_{1}, \ldots \theta_{n}, G_{0}\right) \\
= & \sum_{j} \frac{n_{j}+\alpha / P}{n-1+\alpha} \\
& \left\{\sum_{l=1}^{n} \frac{1}{n_{j}+\alpha / P} \delta_{\theta_{l}} \delta_{\pi_{l}=j}+\frac{\alpha / P}{n_{j}+\alpha / P} G_{0}\right\} \\
= & \sum_{l=1}^{n} \frac{1}{n+\alpha} \delta_{\theta_{l}}+\frac{\alpha}{n+\alpha} G_{0}
\end{aligned}
$$

## Inference

- Conditioned on the Restaurant allocation data are distributed according to P independent Dirichlet process
- Perform local collapsed gibbs sampling on the independent DPs
- For the global parameters perform MH
- Select a cluster 'c' and a processor 'p'
- Propose: move 'c' to 'p'
- Acceptance ratio depends on cluster size
- Can pass the indices of the cluster item.
- Can be done asynchronously without affecting the performance.


## Result



## Extension to HDP


© Eric Xing @ CMU, 2005-2014

## Take home message

- Naïve parallel inference scheme does not always work
- Utilize structure of the problem: Conditional independence
- Exact parallel inference or bound on error

