A Poisson-Based Approximation Algorithm for Stochastic Bin Packing of Bernoulli Items

Tomasz Kanas, Krzysztof Rzadca

Scheduling in the cloud as Bin Packing

- Assignment
- of **tasks** (VMs/Containers)
- to **machines**
- **resource** (CPU, RAM) constraints
- **minimize** number of machines

Bin Packing:

$\rightarrow$ **item** of **size** s

$\rightarrow$ **bin** of **capacity** 1

$\rightarrow$ sum of sizes in a bin ≤ 1

$\rightarrow$ minimize #bins

Motivation

- Scheduling in the cloud $\implies$ Bin packing
- Unknown task resource requirements $\implies$ Random variable
- Resource usage changes in time $\implies$ Overcommitment

Stochastic Bin Packing

- Item size: random variable X_i
- Bounded probability of overload (SLO): $\mathbb{P}(\sum_{i \in B} X_i > 1) \leq \alpha$

Item Distribution

Central Limit Theorem $\implies$ Bin distribution tends to Gaussian.

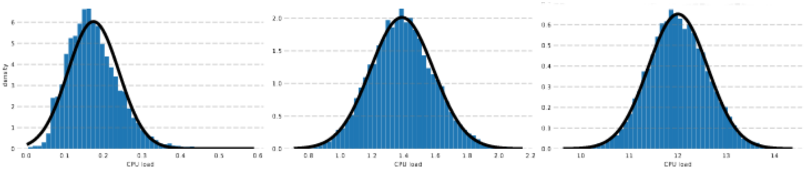

a) 10 Items b) 50 Items c) 500 Items

But the convergence is too slow in terms of the number of items! [Janus, Rzadca, 2017]

Bernoulli Distribution

- Realistic
- Not well studied
- Theoretically difficult

$$\mathbb{P}(X = t) = \begin{cases} p & t = s \\ 1 - p & t = 0 \end{cases}$$

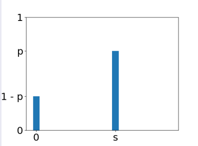

Distributions in Google Trace

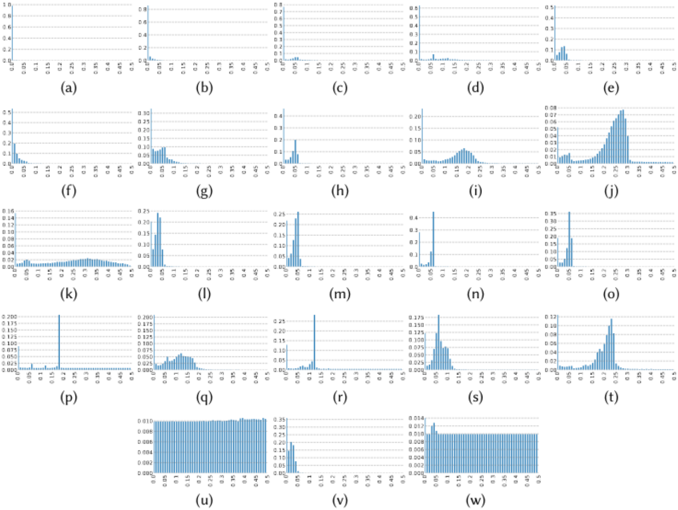

Centers of clusters of Google Trace task distributions.
[Janus, Rzadca, 2017]

Distributions in Google Trace

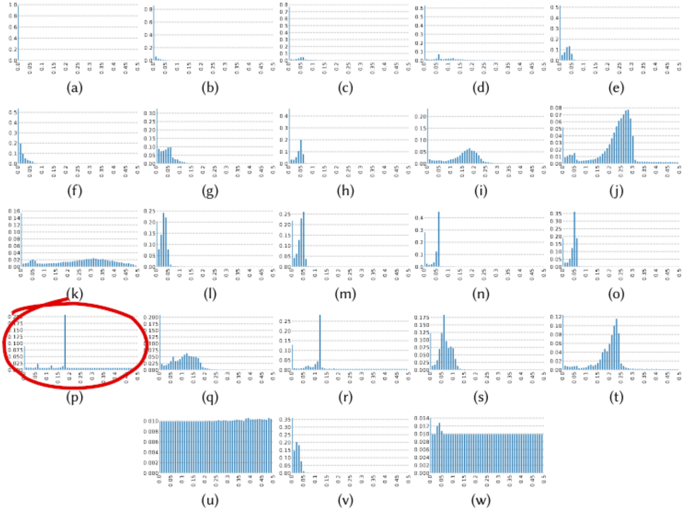

Centers of clusters of Google Trace task distributions.
[Janus, Rzadca, 2017]

Distributions in Google Trace

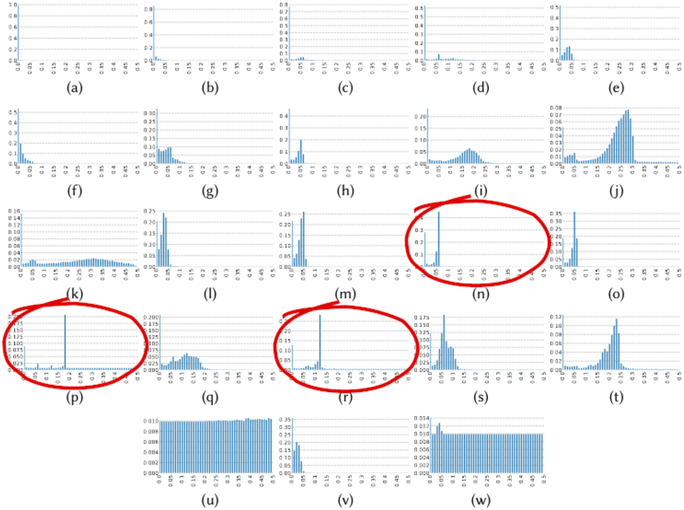

Centers of clusters of Google Trace task distributions.
[Janus, Rzadca, 2017]

Distributions in Google Trace

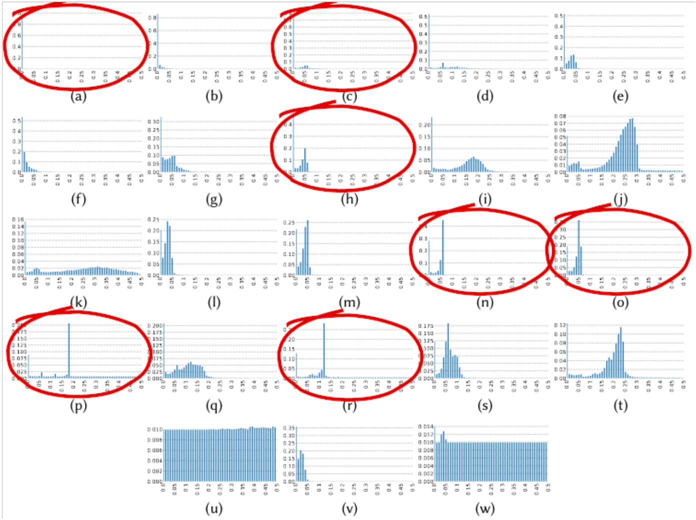

Centers of clusters of Google Trace task distributions.
[Janus, Rzadca, 2017]

Contribution

- **Refined Poisson Approximation Packing (RPAP)**: an online algorithm for bin packing of Bernoulli variables
- Closed-form approximation ratio:

$$2\frac{1+\alpha}{1-\alpha}\left(1+\frac{2}{Q^{-1}(2,1-\alpha)}\right)$$

- RPAPC: RPAP modification that beats First Fit in simulations on datasets with small items.

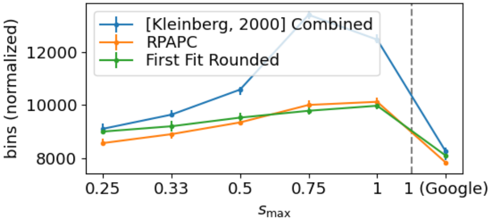

Algorithm idea

Item: size s with probability p
size 0 with probability $1 - p$

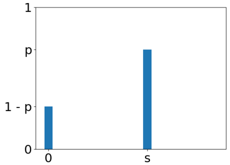

Types of items:

- Confident: p **large** → round to 1
  Pack by: **size** s
- Minor: s **small** → variance $(s^2 p)$ small
  Pack by: **expected value** sp
- Standard: More complicated...

Pack each type to **separate bins**.

Algorithm idea

Sum of independent Bernoulli → Poisson

- Round items up to Poisson variables.
- $\text{Poi}(a) + \text{Poi}(b) = \text{Poi}(a+b)$ $\implies$ classic Bin Packing!

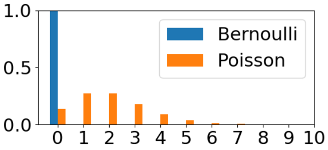

How to round up random variables?

Stochastic dominance:

$$A \geq_1 B \iff \forall_x \mathbb{P}(A \leq x) \leq \mathbb{P}(B \leq x)$$

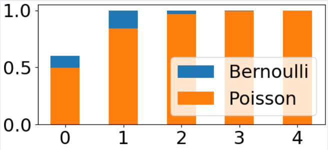

CDFs of Bernoulli and Poisson distributions

RPAP

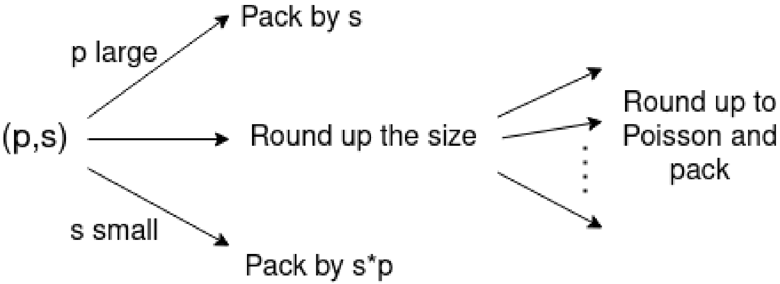
Pack by s
p large
(p,s)
Round up the size
s small
Pack by s*p
Round up to Poisson and pack

If the probability of overflow is small, then the expected value is also small!

Lemma 6.

- If $X_i \sim \mathrm{Ber}(p_i, s_i)$
- are independent,
- $s_i \leq 1$
- and their sum $B = \sum_i X_i$ satisfies $\mathbb{P}(B > 1) \leq \alpha$
- then $\mathbb{E}(B) \leq \frac{1+\alpha}{1-\alpha}$.

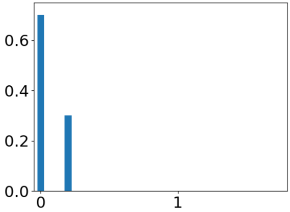

Approximation: Lower bound of optimum

If the probability of overflow is small, then the expected value is also small!

Lemma 6.

- If $X_i \sim \text{Ber}(p_i, s_i)$
- are independent,
- $s_i \leq 1$
- and their sum $B = \sum_i X_i$ satisfies $\mathbb{P}(B > 1) \leq \alpha$
- then $\mathbb{E}(B) \leq \frac{1+\alpha}{1-\alpha}$.

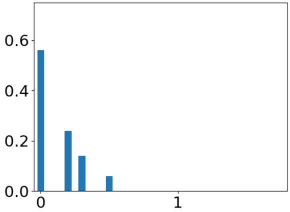

If the probability of overflow is small, then the expected value is also small!

Lemma 6.

- If $X_i \sim \text{Ber}(p_i, s_i)$
- are independent,
- $s_i \leq 1$
- and their sum $B = \sum_i X_i$ satisfies $\mathbb{P}(B > 1) \leq \alpha$
- then $\mathbb{E}(B) \leq \frac{1+\alpha}{1-\alpha}$.

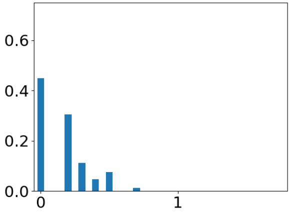

If the probability of overflow is small, then the expected value is also small!

Lemma 6.

- If $X_i \sim \text{Ber}(p_i, s_i)$
- are independent,
- $s_i \leq 1$
- and their sum $B = \sum_i X_i$ satisfies $\mathbb{P}(B > 1) \leq \alpha$
- then $\mathbb{E}(B) \leq \frac{1+\alpha}{1-\alpha}$.

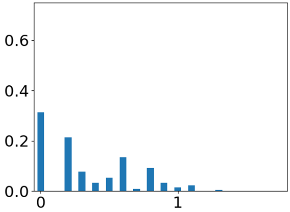

If the probability of overflow is small, then the expected value is also small!

Lemma 6.

- If $X_i \sim \mathrm{Ber}(p_i, s_i)$
- are independent,
- $s_i \leq 1$
- and their sum $B = \sum_i X_i$ satisfies $\mathbb{P}(B > 1) \leq \alpha$
- then $\mathbb{E}(B) \leq \frac{1+\alpha}{1-\alpha}$.

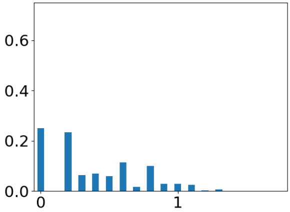

If the probability of overflow is small, then the expected value is also small!

Lemma 6.

- If $X_i \sim \mathrm{Ber}(p_i, s_i)$
- are independent,
- $s_i \leq 1$
- and their sum $B = \sum_i X_i$ satisfies $\mathbb{P}(B > 1) \leq \alpha$
- then $\mathbb{E}(B) \leq \frac{1+\alpha}{1-\alpha}$.

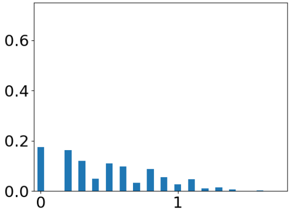

Simulation methodology

- Synthetic datasets
  - Sampled from Normal distribution
  - Different maximal item sizes (s_{max})
- Google trace derived dataset
  - Empirical task CPU usage distribution
  - Fitted Bernoulli distribution (by minimizing L_1 norm of CDF difference)
- Normalized results by average item expected value

Google trace dataset statistics:

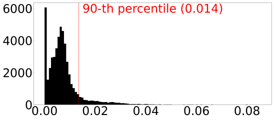

Distances from data

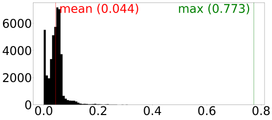

Sizes of items

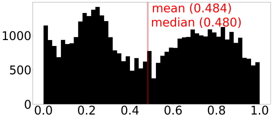

Probability of items

As suspected results are worse than the simple heuristic...

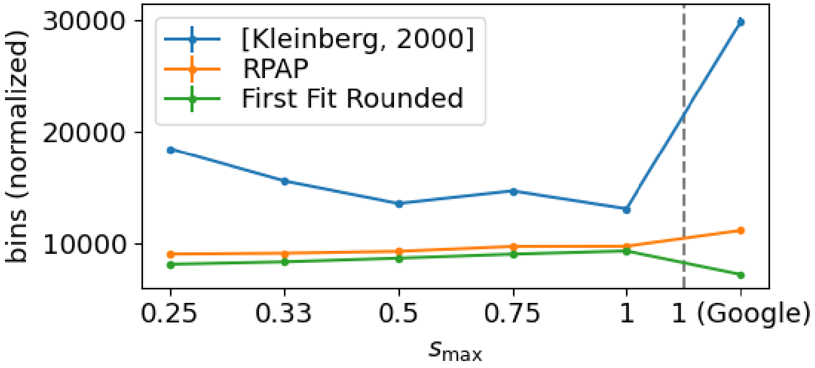

Combined approach

RPAP Combined (RPAPC):

- Try to pack item using **RPAP**
- If it *doesn't* fit into a bin
- Check if it would fit in **FF**

Notice: It will never open a bin if RPAP doesn't.

Approximation guarantees hold!

RPAPC slightly outperforms First Fit Rounded on datasets with small items

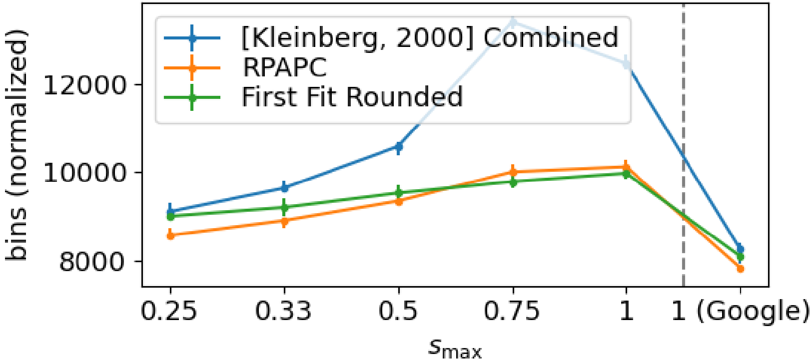

While using fewer bins than First Fit Rounded, RPAPC also achieves lower overflow probabilities

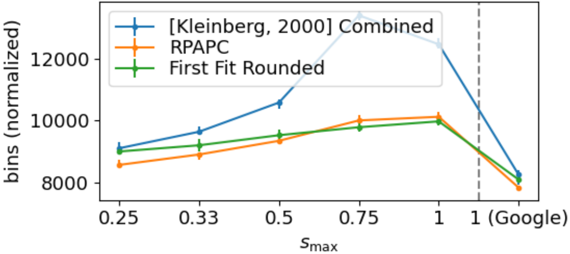

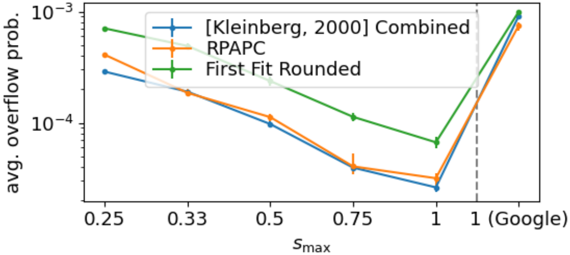

More experiments

$\alpha = 0.1$

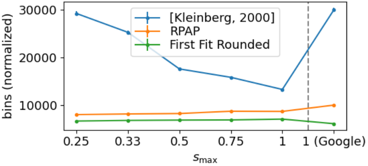

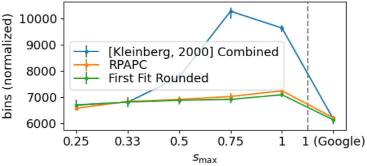

$\alpha = 0.01$

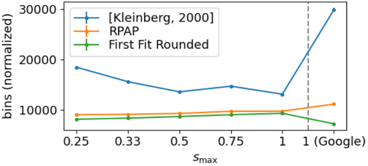

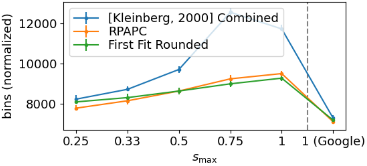

$\alpha = 0.001$

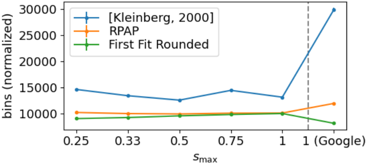

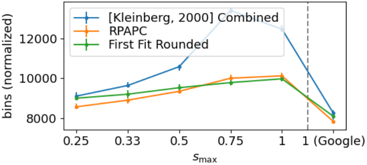

Key Takeaways

RPAP is an $2\frac{1+\alpha}{1-\alpha}\left(1+\frac{2}{Q^{-1}(2,1-\alpha)}\right)$ approximation algorithm for Stochastic Bin Packing with Bernoulli items.

- Probabilistic models for resource management are worth considering.
- Probabilistic models may lead to practical algorithms.
- Combining the theoretical approach with simple heuristics decreases risk and keeps the performance.

Thank you for your attention! Any questions?

Contact me: t.kanas@uw.edu.pl