

DCAF: A Dynamic Computation Allocation Framework for Online Serving System

Biye Jiang, Pengye Zhang, Rihan Chen*

Binding Dai, Xinchun Luo, Yin Yang, Guan Wang, Guorui Zhou, Xiaoqiang Zhu, Kun Gai
Alibaba Group

ABSTRACT

Modern large-scale systems such as recommender system and online advertising system are built upon computation-intensive infrastructure. The typical objective in these applications is to maximize user experience, revenue and other business targets under a limited computation resource. Usually, the online serving system follows a multi-stage cascade architecture, which consists of several stages including retrieval, pre-ranking, ranking, etc. These stages usually allocate resource manually with specific computing power budgets, which requires the serving configuration to adapt accordingly. As a result, the existing system easily falls into suboptimal solutions with respect to maximizing the total revenue. The limitation is due to the fact that, although the value of requests varies greatly, online serving system still allocate equal computing resource among them.

In this paper, we introduce a novel idea that online serving system could treat each request differently and allocate "personalized" computation resource based on its value. We formulate this resource allocation problem as a knapsack problem and propose a Dynamic Computation Allocation Framework (DCAF). Under some general assumptions, DCAF can theoretically guarantee that the system can maximize the total revenue within given computation budget. DCAF brings significant improvement and has been deployed in the display advertising system of Taobao for serving the main traffic. With DCAF, we are able to maintain the same business performance with 20% computation resource reduction.

KEYWORDS

Dynamic Computation Allocation, Online Serving System

ACM Reference Format:

Biye Jiang, Pengye Zhang, Rihan Chen and Binding Dai, Xinchun Luo, Yin Yang, Guan Wang, Guorui Zhou, Xiaoqiang Zhu, Kun Gai. 2020. DCAF: A Dynamic Computation Allocation Framework for Online Serving System. In *2nd Workshop on Deep Learning Practice for High-Dimensional Sparse Data with KDD 2020, Aug 24, 2020, San Diego, CA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

*These three authors contributed equally, corresponding author: Guorui Zhou <guorui.xgr@alibaba-inc.com>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DLP-KDD '20, Aug 24, 2020, San Diego, CA

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Modern large-scale systems such as recommender system and online advertising are built upon computation-intensive infrastructure [7] [21] [20]. With the popularity of e-commerce shopping, e-commerce platform such as Taobao, the world's leading e-commerce platforms, are now enjoying a huge boom in traffic [4], e.g. user requests at Taobao are increasing year by year. As a result, the system load is under great pressures [19]. Moreover, request fluctuation also gives critical challenge to online serving system. For example, the Taobao recommendation system always bears many spikes of requests during the period of Alibaba's Singles Day sales.




Figure 1: Illustration of our cascaded display advertising system. Each request will be served through these modules sequentially. Considering the limitation of computation resource and latency for online serving system, the fixed quota of candidate advertisements, denoted by N for each module, is usually pre-defined manually by experience.

To address the above challenges, the prevailing practices for online engine are: 1) decomposing the cascade system [14] into multiple modules and manually allocating a fixed quota for each module by experience, as shown in Figure 1; 2) designing many computation downgrade plans in case the sudden traffic spikes arrive and manually executing these plans when needed.

These non-automated strategies are often lack of flexibility and require human interventions. Furthermore, most of these practices often impact on all requests once they are executed and ignore the fact that the value of requests varies greatly. Obviously it is a straightforward but better strategy to allocate the computation

resource by biasing towards the requests that are more valuable than others, for maximizing the total revenue.

Considering the shortcomings of existing works, we aim at building a dynamic allocation framework that can allocate the computation budget flexibly among requests. Moreover, this framework should also take into account the stability of the online serving system which are frequently challenged by request boom and spike. Specifically, we formulate the problem as a knapsack problem, of which objective is to maximize the total revenue under a computation budget constraint. We propose a dynamic allocation framework named DCAF which could consider both computation budget allocation and stability of online serving system simultaneously and automatically.

Our main contributions are summarized as follow:

- We break through the stereotypes in most cascade systems where each individual module is limited by a fixed computation budget separately. We introduce a brand-new idea that computation budget can be allocated w.r.t the value of requests in a "personalized" manner.
- We propose a dynamic allocation framework DCAF which could guarantee in theory that the total revenue can be maximized under a computation budget constraint. Moreover, we provide an powerful control mechanism that could always keep the online serving system stable when encountering the sudden spike of requests.
- DCAF has been deployed in the display advertising system of Taobao, bringing a notable improvement. To be specific, system with DCAF maintains the same business performance with 20% Graphics Processing Unit (GPU) resource reduction for online ranking system. Meanwhile, it greatly boosts the online engine's stability.
- By defining the new paradigm, DCAF lays the cornerstone for jointly optimizing the cascade system among different modules and raising the ceiling height for performance of online serving system further.

2 RELATED WORK

Quite a lot of research have been focusing on improving the serving performance. Park et al. [15] describes practice of serving deep learning models in Facebook. Clipper [10] is a general-purpose low-latency prediction serving system. Both of them use latency, accuracy, and throughput as the optimization target of the system. They also mentioned techniques like batching, caching, hyper-parameters tuning, model selection, computation kernel optimization to improve overall serving performance. Also, many research and system use model compression [12], mix-precision inference, quantization [9, 11], kernel fusion [6], model distillation [13, 19] to accelerate deep neural net inference.

Traditional work usually focus on improving the performance of individual blocks, and the overall serving performance across all possible queries. Some new systems have been designed to take query diversity into consideration and provide dynamic planning. DQBarge [8] is a proactive system using monitoring data to make data quality tradeoffs. RobinHood [2] provides tail Latency aware caching to dynamically allocate cache resources. Zhang et al. [18]

takes user heterogeneity into account to improve quality of experience on the web. Those systems provide inspiring insight into our design, but existing systems did not provide solutions for computation resource reduction and comprehensive study of personalized planning algorithms.

3 FORMULATION

We formulate the dynamic computation allocation problem as a knapsack problem which is aimed at maximizing the total revenue under the computation budget constraint. We assume that there are N requests $\{i = 1, \dots, N\}$ requesting the e-commerce platform within a time period. For each request, M actions $\{j = 1, \dots, M\}$ can be taken. We define Q_{ij} and q_j as the expected gain for request i that is assigned action j and the cost for action j respectively. C represents the total computation budget constraint within a time period. For instance, in the display advertising system deployed in e-commerce, q_j can be defined as items (ads) quota that request the online engine to evaluate, which positively correlate with system load in usual. For Q_{ij} , the Effective Cost Per Mile (eCPM) conditioned on action j can be an appropriate choice. A few potential choices are given below:

- There are a bunches of CTR models which are different in their complexities and providing online services simultaneously. For each request i , the platform can choose an appropriate CTR model for it.
- In a cascade system, it is not necessary for all requests to go through every module. For example, the platform let request i skip *Pre-Ranking* stage and go to *Ranking* stage directly.

x_{ij} is the indicator that request i is assigned action j . For each request, there is one and only one action j can be taken, in other words, x_i is an one-hot vector.

Following the definitions above, for each request, our target is to maximize the total revenue under computation budget by assigning each request i with appropriate action j . Formally,

$$\begin{aligned} \max \quad & \sum_{ij} x_{ij} Q_{ij} \\ \text{s.t.} \quad & \sum_{ij} x_{ij} q_j \leq C \\ & \sum_j x_{ij} \leq 1 \\ & x_{ij} \in \{0, 1\} \end{aligned} \quad (1)$$

where we assume that each individual request has its "personalized" value, thus should be treated differently. Besides, request expected gain is correlated with action j which will be automatically taken by the platform in order to maximize the objective under the constraint. In this paper, we mainly focus on proving the effectiveness of DCAF's framework. However, in real case, we are faced with several challenges which are beyond the scope of this paper. We simply list them as below for considering in the future:

- The dynamic allocation problem [2] are usually coupled with real-time request and system status. As the online traffic and system status are both varying with time, we should consider the knapsack problem to be real-time, e.g. real-time computation budget.

- Q_{ij} are unknown, thus needs to be estimated. Q_{ij} prediction is vital to maximize the objective, which means real-time and efficient approaches are required to estimate the value. Besides, to avoid increasing the system's burden, it is essential for us to consider light-weighted methods.

4 METHODOLOGY

4.1 Global Optimal Solution and Proof

To solve the problem, we first construct the Lagrangian from the formulation above,

$$\begin{aligned} L &= - \sum_{ij} x_{ij} Q_{ij} + \lambda (\sum_{ij} x_{ij} q_j - C) + \sum_i (\mu_i (\sum_j x_{ij} - 1)) \\ &= \sum_{ij} x_{ij} (-Q_{ij} + \lambda q_j + \mu_i) - \lambda C - \sum_i \mu_i \\ &\quad \text{s.t. } \lambda \geq 0 \\ &\quad \mu_i \geq 0 \\ &\quad x_{ij} \geq 0 \end{aligned} \quad (2)$$

where we relax the discrete constraint for the indicator x_{ij} , we could show that the relaxation does no harm to the optimal solution. From the primal, the dual function [3] is

$$\max_{\lambda, \mu} \min_{x_{ij}} (\sum_{ij} x_{ij} (-Q_{ij} + \lambda q_j + \mu_i) - \lambda C - \sum_i \mu_i) \quad (3)$$

With $x_{ij} \geq 0$ ($x_{ij} \leq 1$ is implicitly described in the Lagrangian), the linear function is bounded below only when $-Q_{ij} + \lambda q_j + \mu_i \geq 0$. And only when $-Q_{ij} + \lambda q_j + \mu_i = 0$, the inequality $x_{ij} > 0$ could hold which means $x_{ij} = 1$ in our case (remember that the x_i is an one-hot vector). Formally,

$$\begin{aligned} &\max_{\lambda, \mu} (-\lambda C - \sum_i \mu_i) \\ &\text{s.t. } -Q_{ij} + \lambda q_j + \mu_i \geq 0 \\ &\quad \lambda \geq 0 \\ &\quad \mu_i \geq 0 \\ &\quad x_{ij} \geq 0 \end{aligned} \quad (4)$$

As the dual objective is negatively correlated with μ , the global optimal solution for μ would be

$$\mu_i = \max_j (Q_{ij} - \lambda q_j) \quad (5)$$

Hence, the global optimal solution to x_{ij} that indicate which action j could be assigned to request i is

$$j = \arg \max_j (Q_{ij} - \lambda q_j) \quad (6)$$

From Slater's theorem [17], it can be easily shown that the Strong Duality holds in our case, which means that this solution is also the global optimal solution to the primal problem.

4.2 Parameter Estimation

4.2.1 Lagrange Multiplier.

The analytical form of Lagrange multiplier cannot be easily, or even possibly derived in our case. And meanwhile, the exact global optimal solution in arbitrary case is computationally prohibitive. However, under some general assumptions, simple bisection search

could guarantee that the global optimal λ could be obtained. Without loss of generality, we reset the indices of action space by following the ascending order of q_j 's magnitude.

ASSUMPTION 4.1. Q_{ij} is monotonically increasing with j .

ASSUMPTION 4.2. Q_{ij}/q_j is monotonically decreasing with j .

LEMMA 1. Suppose Assumptions (4.1) and (4.2) hold, for each i , $Q_{i_{j_1}/q_{j_1}} \geq Q_{i_{j_2}/q_{j_2}}$ will hold if $\lambda_1 \geq \lambda_2$, where j_1 and j_2 are the actions that maximize the objective under λ_1 and λ_2 respectively.

PROOF. As Equation (5) and $\mu_i \geq 0$, the inequality $Q_{ij} - \lambda q_j \geq 0$ holds. Equally, $Q_{ij}/q_j \geq \lambda$ holds. Suppose $Q_{i_{j_1}/q_{j_1}} < Q_{i_{j_2}/q_{j_2}}$, we have $Q_{i_{j_2}/q_{j_2}} > Q_{i_{j_1}/q_{j_1}} \geq \lambda_1 \geq \lambda_2$. However, we could always find j_2^* such that $Q_{ij_2^*} \geq Q_{ij_2}$ and $q_{j_2^*} > q_{j_2}$ where $Q_{i_{j_1}/q_{j_1}} \geq Q_{i_{j_2^*}/q_{j_2^*}} \geq \lambda_1 \geq \lambda_2$ such that $Q_{ij_2^*} \geq Q_{ij_2}$ by following the Assumptions (4.1) and (4.2). In order words, j_2 is not the action that maximize the objective. Therefore, we have $Q_{i_{j_1}/q_{j_1}} \geq Q_{i_{j_2}/q_{j_2}}$. \square

LEMMA 2. Suppose Assumptions (4.1) and (4.2) could be satisfied, both $\max \sum_{ij} x_{ij} Q_{ij}$ and its corresponding $\sum_{ij} x_{ij} q_j$ are monotonically decreasing with λ .

PROOF. With λ increasing, Q_{ij}/q_j is also increasing monotonically by Lemma (1). Moreover, by Assumptions (4.1) and (4.2), we conclude that both $\max \sum_{ij} x_{ij} Q_{ij}$ and its corresponding $\sum_{ij} x_{ij} q_j$ are monotonically decreasing with λ . \square

THEOREM 1. Suppose Lemma (2) holds, the global optimal Lagrange Multiplier λ could be obtained by finding a solution that make $\sum_{ij} x_{ij} q_j = C$ hold through bisection search.

PROOF. By Lemma (2), this proof is almost trivial. We denote the Lagrange Multiplier that makes $\sum_{ij} x_{ij} q_j = C$ hold as λ^* . Obviously, the increase of λ^* will result in computation overload and the decrease of λ^* will inevitably reduce $\max \sum_{ij} x_{ij} Q_{ij}$ due to the monotonicity in Lemma (2). Hence, λ^* is the global optimal solution to the constrained maximization problem. Besides, the bisection search must work in this case which is also guaranteed by the monotonicity. \square

Assumption (4.1) usually holds because the gain is directly proportional to the cost in general, e.g. more sophisticated models usually bring better online performance. For Assumption (4.2), it follows the law of diminishing marginal utility [16], which is an economic phenomenon and reasonable in our constrained dynamic allocation case.

The algorithm for searching Lagrange Multiplier λ is described in Algorithm 1. In general, we implement the bisection search over a pre-defined interval to find out the global optimal solution for λ . Suppose $\min_j \sum_j q_j \leq C \leq \max_j \sum_j q_j$ (o.w there is no need for dynamic allocation), it can be easily shown that λ locates in the interval $[0, \min_{ij} (Q_{ij}/q_j)]$. Then we get the global optimal λ through bisection search of which target is the solution of $\sum_{ij} x_{ij} q_j = C$.

For more general cases, more sophisticated method other than bisection search, e.g. reinforcement learning, will be conducted to explore the solution space and find out the global optimal λ .

Algorithm 1 Calculate Lagrange Multiplier

1: **Input:** Q_{ij} , q_j , C , interval $[0, \min_{ij}(\frac{Q_{ij}}{q_j})]$ and tolerance ϵ
2: **Output:** Global optimal solution of Lagrange Multiplier λ
3: Set $\lambda_l = 0$, $\lambda_r = \min_{ij}(\frac{Q_{ij}}{q_j})$, $gap = +\infty$
4: **while** ($gap > \epsilon$):
5: $\lambda_m = \lambda_l + \frac{\lambda_r - \lambda_l}{2}$
6: Choose action j_m^* by
 $\{j : \arg \max_j (Q_{ij} - \lambda_m q_j), Q_{ij} - \lambda_m q_j \geq 0\}$
7: Calculate the $\sum_i q_{j_i^*}$ denoted by C_m
8: $gap = |C_m - C|$
9: **if** $gap \leq \epsilon$:
10: **return** λ_m
11: **else if** $C_m \leq C$:
12: $\lambda_l = \lambda_m$
13: **else**:
14: $\lambda_r = \lambda_m$
15: **end while**
16: **Return** the global optimal λ_m which satisfies $|\sum_i q_{j_i^*} - C| \leq \epsilon$.

4.2.2 Request Expected Gain.

In e-commerce, the expected gain is usually defined as online performance metric e.g. Effective Cost Per Mile (eCPM), which could directly indicate each individual request value with regard to the platform. Four categories of feature are mainly used: User Profile, User Behavior, Context and System status. It is worth noticing that our features are quite different from typical Click-through rate (CTR) model:

- Specific target ad feature isn't provided because we estimate the CTR conditioned on actions.
- System status is included where we intend to establish the connection between system and actions.
- The context feature consists of the inference results from previous modules in order to re-utilize the request information efficiently.

5 ARCHITECTURE

In general, DCAF is comprised of online decision maker and offline estimator:

- The online modules make the final decision based on personalized request value and system status.
- The offline modules leverage the logs to calculate the Lagrange Multiplier λ and train an estimator for the request expected value conditioned on actions by historical data.

5.1 Online Decision Maker

5.1.1 Information Collection and Monitoring.

This module monitors provide timely information about the system current status which includes GPU-util, CPU-util, runtime, failure rate, and etc. The acquired information enables the framework to dynamically allocate the computation resource without exceeding the budget by limiting the action space.

5.1.2 Request Value Estimation.

This module estimates the request's Q_{ij} based on the features provided in information collection module. Notably, to avoid growing

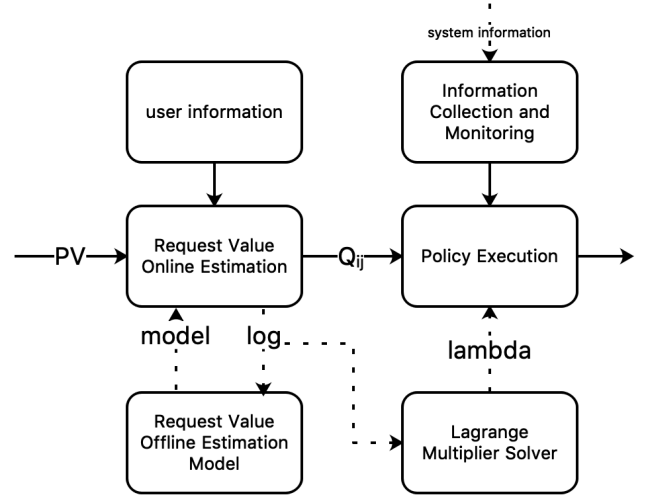


Figure 2: Illusion of the system of DCAF. *Request Value Online Estimation* module will score each request conditioned on action j through online features of which estimator is trained offline. *Policy Execution* module mainly takes charge of executing the final action j for each request based on the system status collected by *Information Collection and Monitoring* module, λ calculated offline and Q_{ij} obtained from previous module.

the system load, the online estimator need to be light-weighted, which necessitates the balance between efficiency and accuracy. One possible solution is that the estimation of Q_{ij} should re-utilize the request context features adequately, e.g. high-level features generated by other models in different modules.

5.1.3 Policy Execution.

Basically, this module assigns the best action j to request i by Equation (6). Moreover, for the stability of online system, we put forward a concept called *MaxPower* which is an upper bound for q_j to which each request must subject. DCAF sets a limit on the *MaxPower* in order to strongly control the online engine. The *MaxPower* is automatically controlled by system's runtime and failure rate through control loop feedback mechanism, e.g. Proportional Integral Derivative (PID) [1]. The introduction of *MaxPower* guarantees that the system can adjust itself and remain stable automatically and timely when encountering sudden request spikes.

According to the formulation of PID, $u(t)$ and $e(t)$ are the control action and system unstability at time step t . For $e(t)$, we define it as the weighted sum of average runtime and fail rates over a time interval which are denoted by rt and fr respectively. k_p , k_i and k_d are the corresponding tuned weights for proportional, integral and derivative control. θ means a tuned scale factor for the weighted sum between rt and fr . Formally,

$$u(t) = k_p e(t) + k_i \sum_{n=1}^t e(n) + k_d (e(t) - e(t-1)) \quad (7)$$

Algorithm 2 PID Control for MaxPower

```

1: Input:  $k_p, k_i, k_d, MaxPower$ 
2: Output:  $MaxPower$ 
3: while (true):
4:   Obtain  $rt$  and  $fr$  from Information Collection and Monitoring
5:    $e(t) = rt + \theta fr$ 
6:    $u(t) = k_p e(t) + k_i \sum_{n=1}^t e(t) + k_d (e(t) - e(t-1))$ 
7:   Update  $MaxPower$  with  $u(t)$ 
8: end while

```

5.2 Offline Estimator

5.2.1 Lagrange Multiplier Solver.

As mentioned above, we could get the global optimal solution of the Lagrange Multipliers by a simple bisection search method. In real case, we take logs as an offline request pool to search a best candidate Lagrange Multiplier λ . Formally,

- Sample N records from the logs with Q_{ij}, q_j and computation cost C , e.g. the total amount of advertisements that request the CTR model within a time interval.
- Adjust the computation cost C by the current system status in order to keep the dynamic allocation problem under constraint in time. For example, we denote regular QPS by QPS_r and current QPS by QPS_c . Then the adjusted computation cost \hat{C} is $C \times QPS_r / QPS_c$, which could keep the N records under the current computation constraint.
- Search the best candidate Lagrange Multiplier λ by Algorithm (1)

It's worth noting that we actually assume the distribution of the request pool is the same as online requests, which could probably introduce the bias for estimating Lagrange Multiplier. However, in practice, we could partly remove the bias by updating the λ frequently.

5.2.2 Expected Gain Estimator.

In our settings, for each request, Q_{ij} is associated with eCPM under different action j which is the common choice for performance metric in the field of online display advertising. Further, we build a CTR model to estimate the CTR, because the eCPM could be decomposed into $ctr \times bid$ where the bids are usually provided by advertisers directly. It is notable that the CTR model is conditioned on actions in our case, where it is essential to evaluate each request gain under different actions. And this estimator is updated routinely and provides real-time inference in *Policy Execution* module.

6 EXPERIMENTS

6.1 Offline Experiments

For validating the framework's correctness and effectiveness, we extensively analyse the real-world logs collected from the display advertising system of Taobao and conduct offline experiments on it. As mentioned above, it is common practice for most systems to ignore the differences in value of requests and execute same procedure on each request. Therefore, we set equally sharing the computation budget among different requests as the **baseline** strategy. As shown in Figure 1, we simulate the performance of DCAF in Ranking stage by offline logs. In advance, we make it clear that all data has been rescaled to avoid breaches of sensitive commercial

data. We conduct our offline and online experiments in Taobao's display advertisement system where we spend the GPU resource automatically through DCAF. In detail, we instantiate the dynamic allocation problem as follow:

- Action j controls the number of advertisements that need to be evaluated by online CTR model in Ranking stage.
- q_j represents the advertisement's quota for requesting the online CTR model.
- Q_{ij} is the sum of top-k ad's eCPM for request i conditioned on action j in Ranking stage which is equivalent to online performance closely. And Q_{ij} is estimated in the experiment.
- C stands for the total number of advertisements that are requesting online CTR model in a period of time within the serving capacity.
- **Baseline:** The original system, which allocates the same computation resource to different requests. With the baseline strategy, system scores the same number of advertisements in Ranking stage for each request.




Figure 3: Global optima under different λ candidates. In Figure 3, x-axis stands for λ 's candidate; left y-axis represents $\sum_{ij} x_{ij} Q_{ij}$; right-axis denotes the corresponding cost. The red shadow area corresponds to the exceeding part of $\max \sum_{ij} x_{ij} Q_{ij}$ beyond the baseline. And yellow shadow area is the reduction of $\sum_{ij} x_{ij} q_j$ under these λ 's compared with the baseline. Random strategy is also shown in Figure 3 for comparison with DCAF.

Global optima under different λ candidates. In DCAF, the Lagrange Multiplier λ works by imposing constraint on the computation budget. Figure 3 shows the relation among λ 's magnitude, $\max \sum_{ij} x_{ij} Q_{ij}$ and its corresponding $\sum_{ij} x_{ij} q_j$ under fixed budget constraint. Clearly, λ could monotonically impact on both $\max \sum_{ij} x_{ij} Q_{ij}$ and its corresponding $\sum_{ij} x_{ij} q_j$. And it shows that the DCAF outperforms the baseline when λ locates in an appropriate interval. As demonstrated by the two dotted lines, in comparison with the baseline, DCAF can achieve both higher performance with same computation budget (eCPM+3.7% in total) and same performance with much less computation budget (computation budget -49.0% in total). Moreover, the result of random strategy is displayed for comparison with DCAF. Random strategy is the strategy that

we randomly cut down the advertisements quota for each request for saving computing cost by 49%. With the same reduction of computation resource, DCAF's performance outmatches the random strategy's to a large extent.




Figure 4: Comparison of DCAF with the original system on computation cost. In Figure 4, x-axis denotes the $\sum_{ij} x_{ij} Q_{ij}$; y-axis represents the $\sum_{ij} x_{ij} q_j$. For points on the two lines with same x-coordinate, Figure 4 shows that DCAF always perform as well as the baseline by much less computation resource.

Comparison of DCAF with the original system on computation cost. Figure 4 shows that DCAF consistently accomplish same performance as the baseline and save the cost by a huge margin. Furthermore, DCAF plays much more important role in more resource-constrained systems.

Total eCPM and its cost over different actions. As shown by the distributions in Figure 5, we could see that DCAF treats each request differently by taking different action j . And $\sum_{ij} Q_{ij} / \sum_{ij} q_j$ is decreasing with action j 's which empirically show that the relation between expected gain and its corresponding cost follows the law of diminishing marginal utility in total.

6.2 Online Experiments

DCAF is deployed in Alibaba display advertising system since 2020. From 2020-05-20 to 2020-05-30, we conduct online A/B testing experiment to validate the effectiveness of DCAF. The settings of online experiments are almost identical to offline experiments. Action j controls the number of advertisements for requesting the CTR model in Ranking stage. And we use a simple linear model to estimate the Q_{ij} . The original system without DCAF is set as baseline. The online performances are evaluated by Revenue Per Mille (RPM), an common metric in online display advertising. The DCAF is deployed between Pre-Ranking stage and Ranking stage which

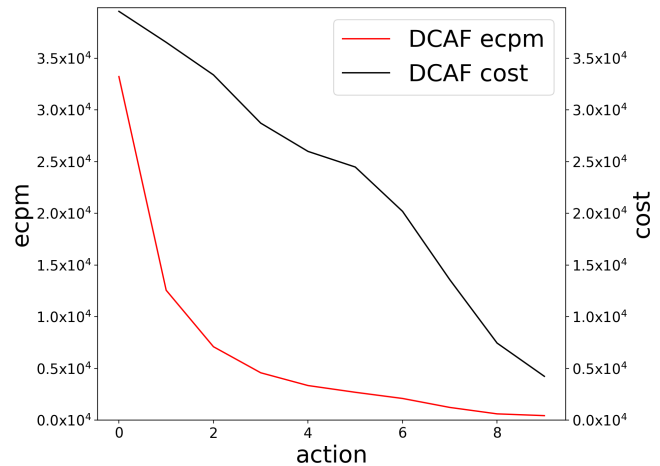


Figure 5: Total eCPM and its cost over different actions. In this figure, x-axis stands for λ candidates and left y-axis represents $\sum_{ij} x_{ij} Q_{ij}$ conditioned on action j ; right-axis denotes the corresponding cost. For each action j , we sum over Q_{ij} for requests that are assigned action j by DCAF.

is aimed at dynamically allocating the GPU resource consumed by Ranking's CTR model. Table 1 shows that DCAF could bring improvement while using the same computation cost. Considering the massive daily traffic of Taobao, we deploy DCAF to reduce the computation cost while not hurting the revenue of the ads system. The results are illustrated in Table 2, and DCAF reduces the computation cost with respect to the total amount of advertisements requesting CTR model by 25% and total utilities of GPU resource by 20%. It should be noticed that, in online system, the Q_{ij} is estimated by a simple linear model which may be not sufficiently complex to fully capture data distribution. Thus the improvement of DCAF in online system is less than the results of offline experiments. This simple method enables us to demonstrate the effectiveness of the overall framework which is our main concern in this paper. In the future, we will dedicate more efforts in modeling Q_{ij} . Figure 6 shows the performance of DCAF under the pressures of online traffic in extreme case e.g. Alibaba's Singles Day sales. By the control mechanism of *MaxPower*, the online serving system can react to the sudden rising of traffic quickly. It makes the system back to normal status by consistently keeping the fail rate and runtime at a low level. It is worth noticing that the control mechanism of *MaxPower* is superior to human interventions in the scenario that the large traffic arrives suddenly and human interventions inevitably delay.

Table 1: Results with Same Computation Budget

	CTR	RPM
Baseline	+0.00%	+0.00%
DCAF	+0.91%	+0.42%

Table 2: Results with Same Revenue

	CTR	RPM	Computation Cost	GPU-utils
Baseline	+0.00%	+0.00%	-0.00%	-0.00%
DCAF	-0.57%	+0.24%	-25%	-20%

7 CONCLUSION

In this paper, we propose a novel dynamic computation allocation framework (DCAF), which can break pre-defined quota constraints within different modules in existing cascade system. By deploying DCAF online, we empirically show that DCAF consistently maintains the same performance with great computation resource reduction in online advertising system, and meanwhile, keeps the system stable when facing the sudden spike of requests. Specifically, we formulate the dynamic computation allocation problem as a knapsack problem. Then we theoretically prove that the total revenue can be maximized under a computation budget constraint by properly allocating resource according to the value of individual request. Besides, under some general assumptions, the global optimal Lagrange Multiplier λ can also be obtained which finally completes the constrained optimization problem in theory. Moreover, we put forward a concept called *MaxPower* which is controlled by a designed control loop feedback mechanism in real-time. Through *MaxPower* which imposes constraints on the range of action candidates, the system could be controlled powerfully and automatically.

8 FUTURE WORK

DCAF is able to dynamically allocate resources according to the value of each request. But we also noticed that DCAF may discriminate among users. While the allocated computation budgets varying with users, DCAF may leave a impression that it would aggravate the unfairness phenomenon of system further.

Fairness has attracted more and more concerns in the fields of recommendation system and online display advertising. In our opinion, the unfair problem stems from that all the approaches to model users are data-driven. Meanwhile, most of systems create a data feedback loop that a system is trained and evaluated on the data impressed to users [5]. We think the fairness of recommender system and ads system is important and needs to be paid more attention to. In the future, we will analyse the long-term effect for fairness of DCAF extensively and include the consideration of it in DCAF carefully.

Besides, DCAF is still in the early stage of development, where modules in the cascade system are considered independently and the action j is defined as the number of advertisements quota to be evaluated in our experiments. Obviously, DCAF could work with diverse actions, such as models with different calculation complexity. Meanwhile, instead of maximizing the total revenue in particular module, DCAF will achieve the global optima in the view of the whole cascade system in the future. Moreover, in the subsequent stages, we will endow DCAF with the abilities of quick adaptation and fast reactions. These abilities will enable DCAF to exert its full effect in any scenario immediately.




Figure 6: The effect of *MaxPower* mechanism. In this experiments, we manually change the traffic of system at time 158 and the requests per second increase 8-fold. Figure 6a shows the trend of *MaxPower* over time and Figure 6b shows the trend of fail rate over time. As shown in Figure 6, the *MaxPower* takes effect immediately when the QPS is rising suddenly which makes the fail rate keep at a lower level. At the same time, the base strategy fails to serve some requests, because it does not change the computing strategy while the computation power of system is insufficient.

9 ACKNOWLEDGMENT

We thanks Zhenzhong Shen, Chi Zhang for helping us on deep neural net inference optimization and conducting the dynamic resource allocation experiments.

REFERENCES

- [1] Kiam Heong Ang, Gregory Chong, and Yun Li. 2005. PID control system analysis, design, and technology. *IEEE transactions on control systems technology* 13, 4 (2005), 559–576.
- [2] Daniel S Berger, Benjamin Berg, Timothy Zhu, Siddhartha Sen, and Mor Harchol-Balter. 2018. RobinHood: Tail latency aware caching–dynamic reallocation from cache-rich to cache-poor. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 195–212.
- [3] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. 2004. *Convex optimization*. Cambridge university press.
- [4] Valeria Cardellini, Michele Colajanni, and Philip S Yu. 1999. Dynamic load balancing on web-server systems. *IEEE Internet computing* 3, 3 (1999), 28–39.
- [5] Allison JB Chaney, Brandon M Stewart, and Barbara E Engelhardt. 2018. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems*. 224–232.
- [6] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. 2018. {TVM}: An automated end-to-end optimizing compiler for deep learning. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 578–594.
- [7] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.
- [8] Michael Chow, Mosharaf Chowdhury, Kaushik Veeraraghavan, Christian Cachin, Michael Cafarella, Wonho Kim, Jason Flinn, Marko Vukolić, Sonia Margulis, Inigo Goiri, et al. 2016. Dqbarge: Improving data-quality tradeoffs in large-scale internet services. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. 771–786.
- [9] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*. 3123–3131.
- [10] Daniel Crankshaw, Xin Wang, Guilio Zhou, Michael J Franklin, Joseph E Gonzalez, and Ion Stoica. 2017. Clipper: A low-latency online prediction serving system. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*. 613–627.
- [11] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. 2015. Deep learning with limited numerical precision. In *International Conference on Machine Learning*. 1737–1746.
- [12] Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149* (2015).
- [13] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
- [14] Shichen Liu, Fei Xiao, Wenwu Ou, and Luo Si. 2017. Cascade ranking for operational e-commerce search. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1557–1565.
- [15] Jongsoo Park, Maxim Naumov, Protonu Basu, Summer Deng, Aravind Kalaiah, Daya Khudia, James Law, Parth Malani, Andrey Malevich, Satish Nadathur, et al. 2018. Deep learning inference in facebook data centers: Characterization, performance optimizations and hardware implications. *arXiv preprint arXiv:1811.09886* (2018).
- [16] Anthony Scott. 1955. The fishery: the objectives of sole ownership. *Journal of political Economy* 63, 2 (1955), 116–124.
- [17] L Slater Morton. 1950. Lagrange Multipliers Revisited. *CCDP Mathematics* 403 (1950).
- [18] Xu Zhang, Siddhartha Sen, Daniar Kurniawan, Haryadi Gunawi, and Junchen Jiang. 2019. E2E: embracing user heterogeneity to improve quality of experience on the web. In *Proceedings of the ACM Special Interest Group on Data Communication*. 289–302.
- [19] Guorui Zhou, Ying Fan, Runpeng Cui, Weijie Bian, Xiaoqiang Zhu, and Kun Gai. 2018. Rocket launching: A universal and efficient framework for training well-performing light net. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [20] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 5941–5948.
- [21] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1059–1068.