Overview

1. Introduction
2. Our approach
3. Experiments
4. Summary

1. Introduction

- What is Display Advertising?
  – Advertisements appearing on websites
- Main purpose
  – Deliver general advertisements and brand messages

Example of Display AD

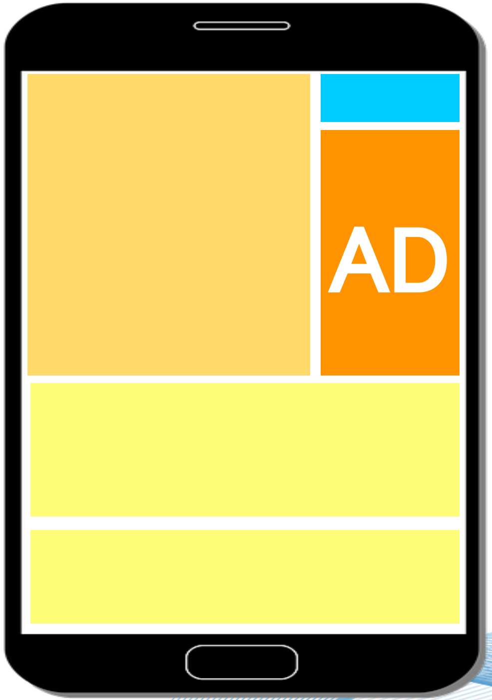

Behind the Scenes

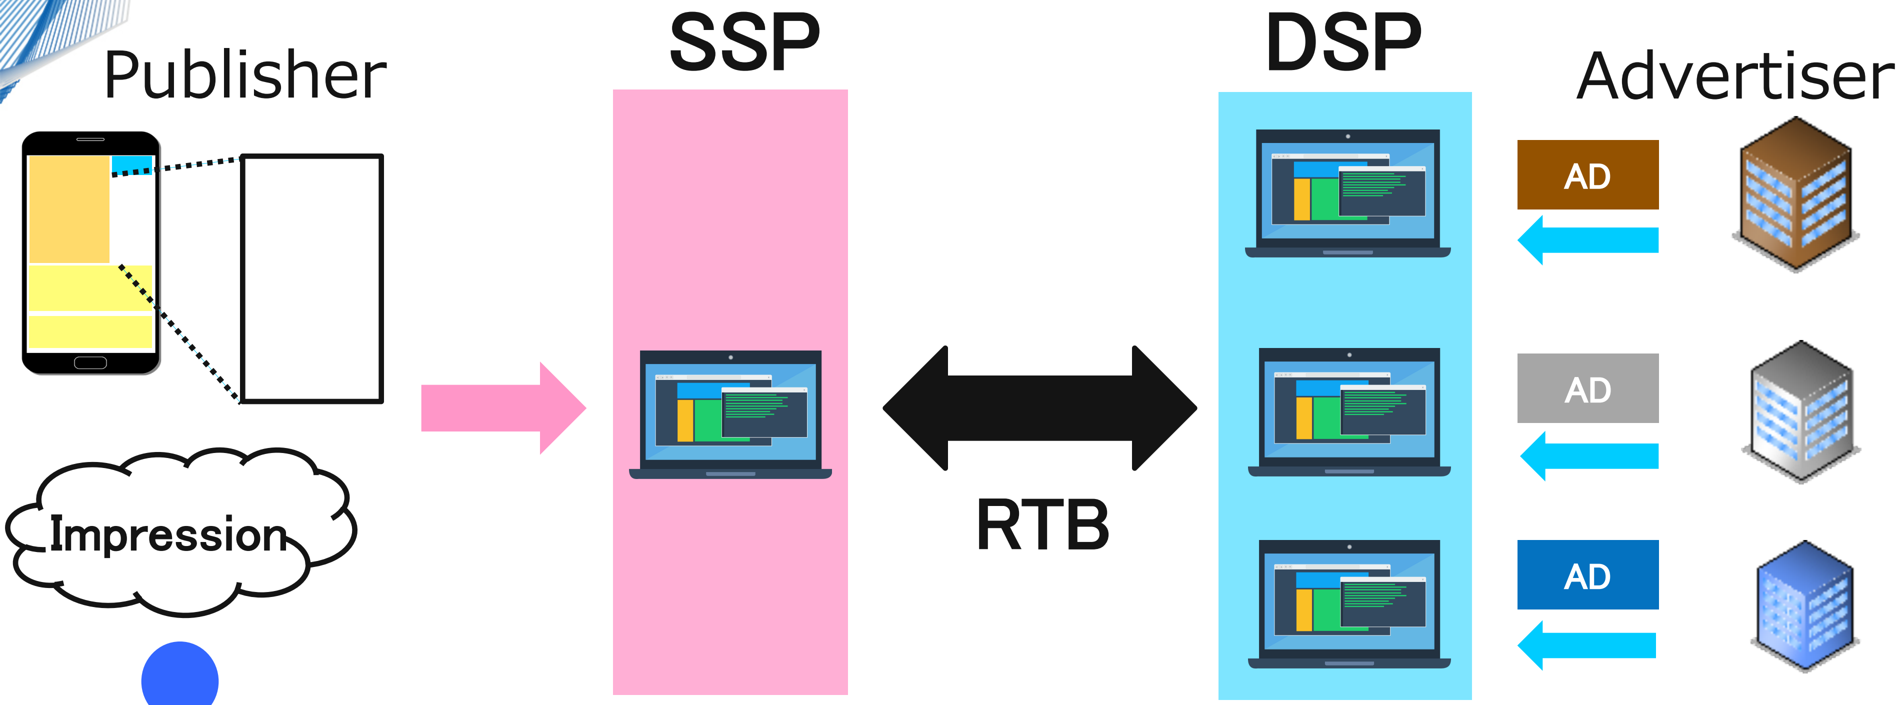

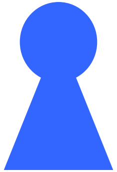

SSP: Supply-Side Platform
DSP: Demand-Side Platform
RTB: Real Time Bidding

Behind the Scenes

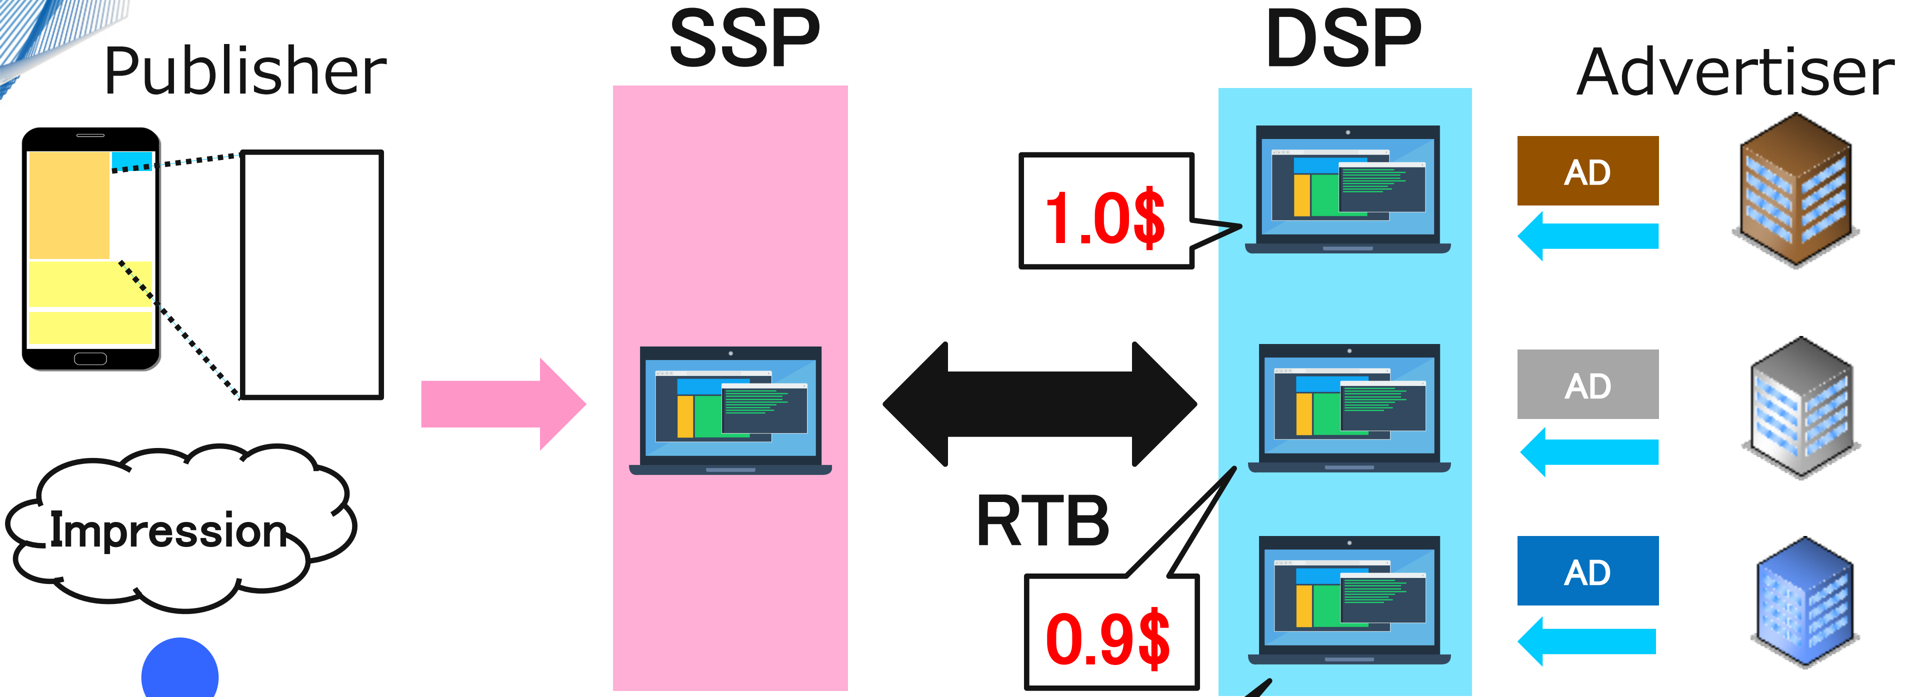

SSP: Supply-Side Platform
DSP: Demand-Side Platform
RTB: Real Time Bidding

Behind the Scenes

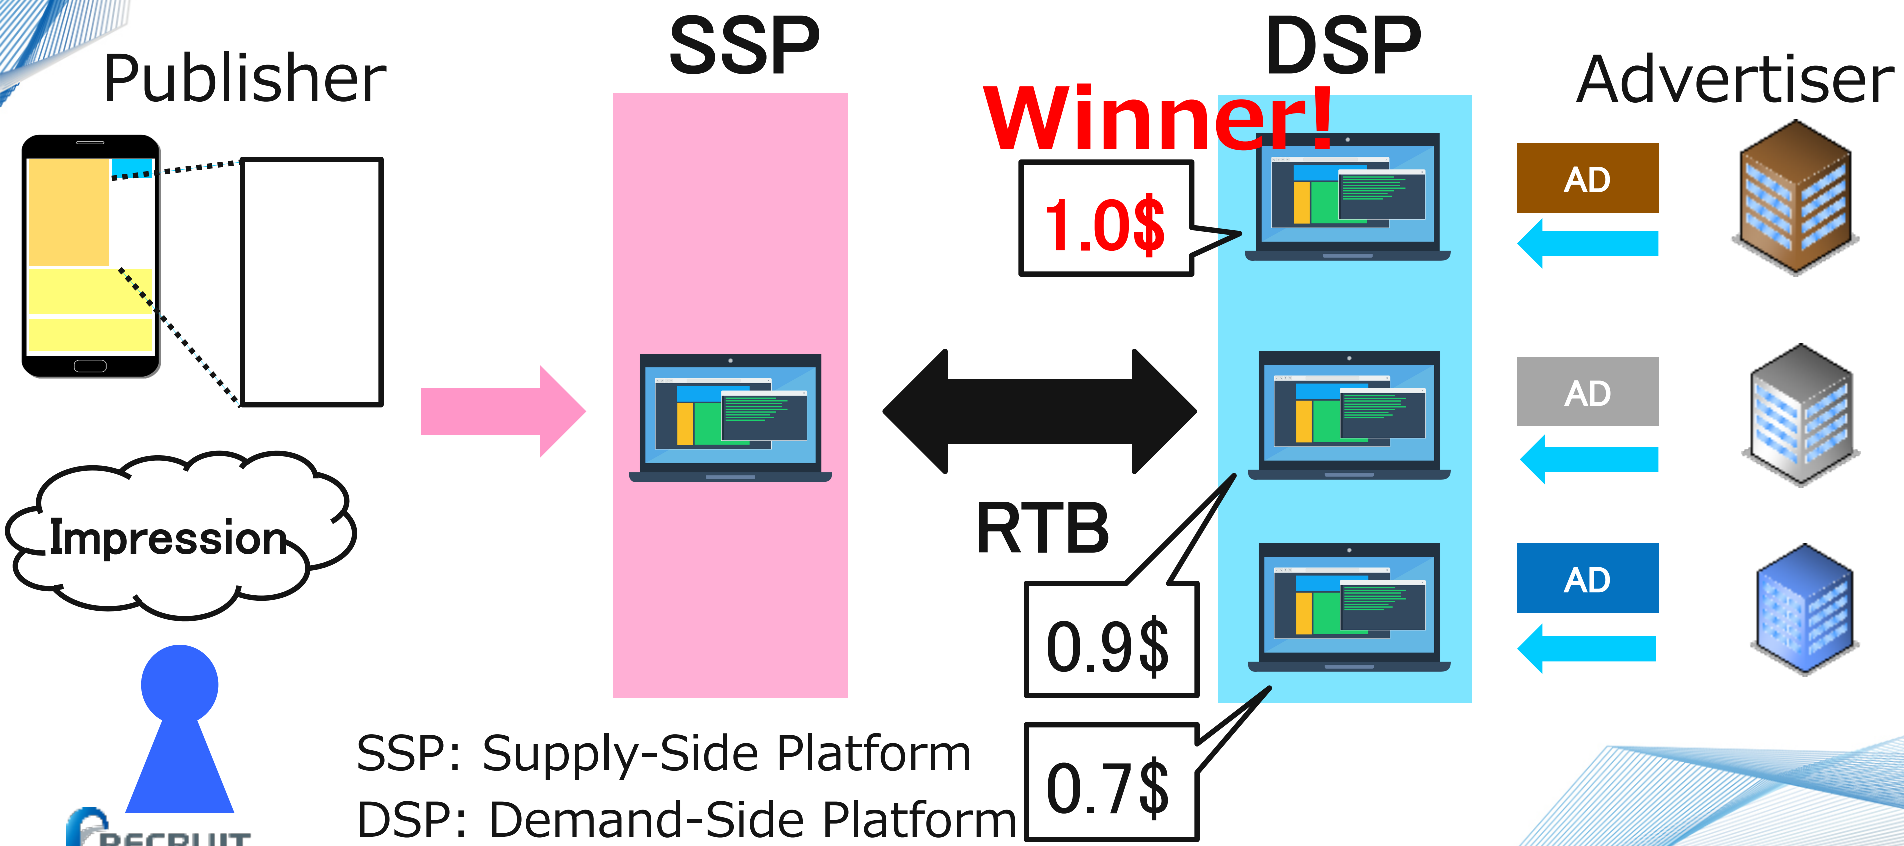

SSP: Supply-Side Platform
DSP: Demand-Side Platform
RTB: Real Time Bidding

Behind the Scenes

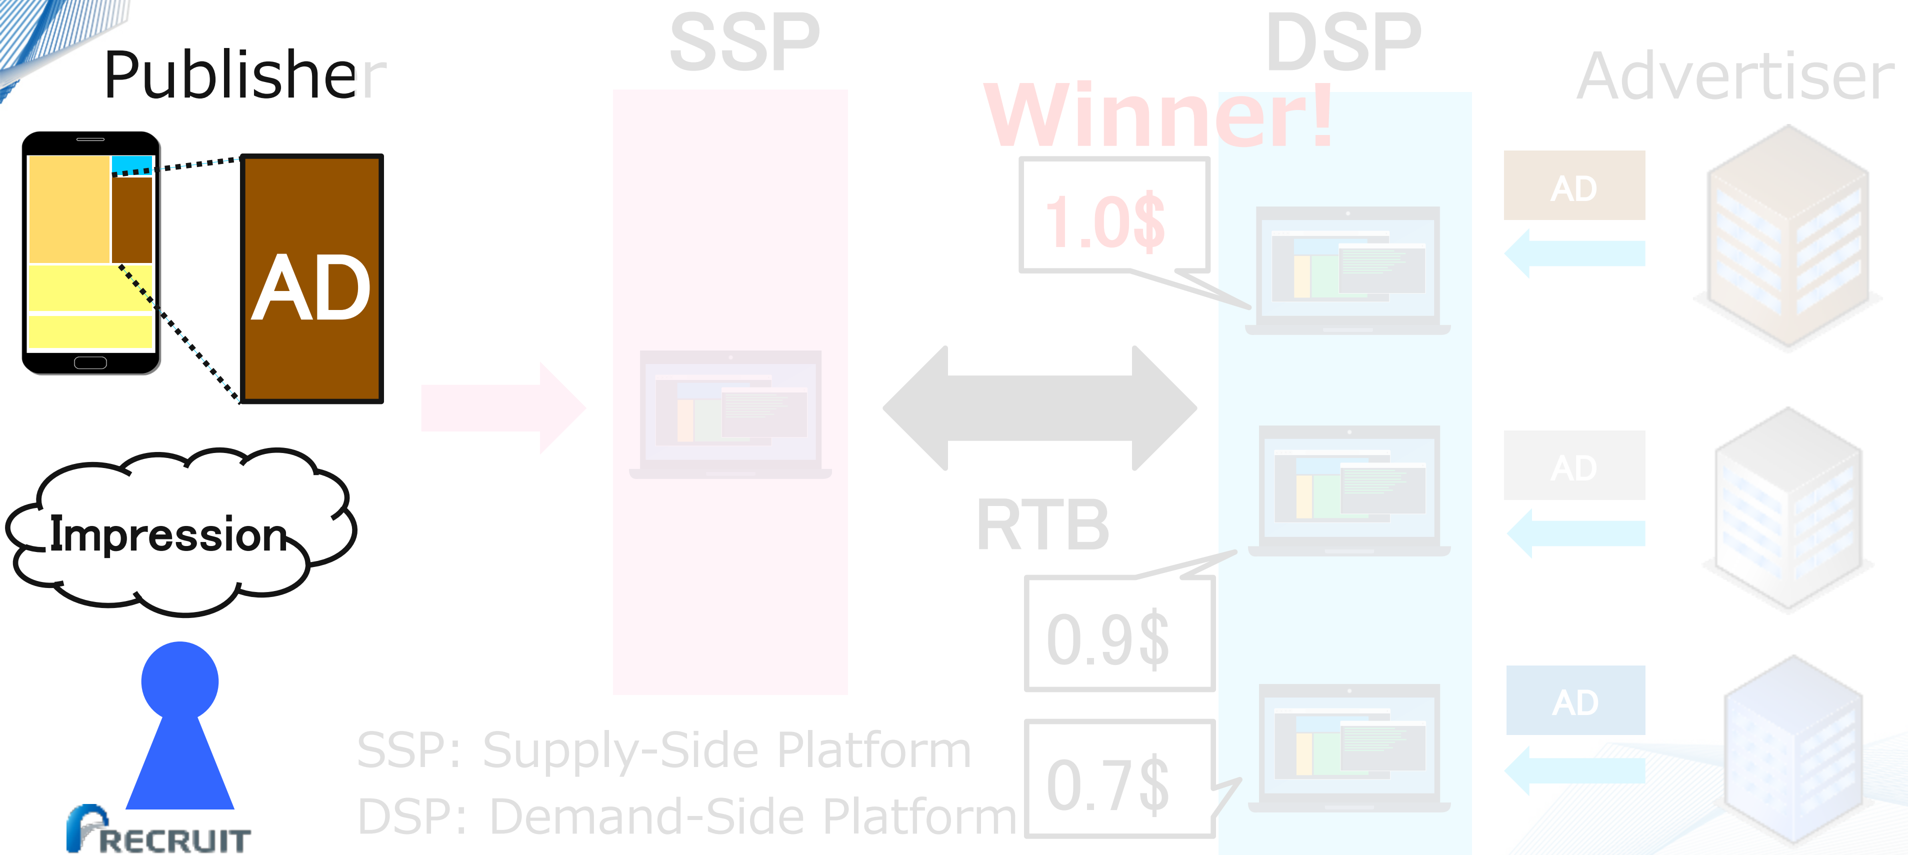

SSP: Supply-Side Platform
DSP: Demand-Side Platform
RTB: Real Time Bidding

Business Model

- Advertiser pays a publisher when the ad is clicked
- Performance Indicator
  - Click Through Rate (CTR)
  - Cost Per Impression (CPI)
  - Cost Per Click (CPC) ...etc

CTR Prediction with Machine-Learning

(Click-through-rate)

- Machine-Learning (ML) tech is often used for CTR prediction
- ML has succeeded in this field

Users

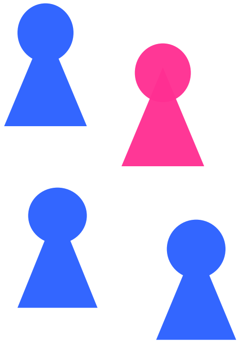

Matrix expression

Click	F1	F2	F3
1	M	01	2.13
0	F	07	2.12
0	F	23	4.2
?	F	99	1.2

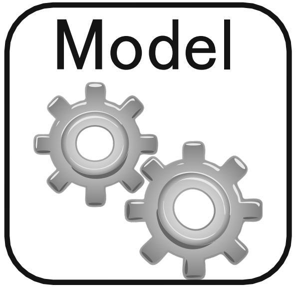

Click or Not Prediction

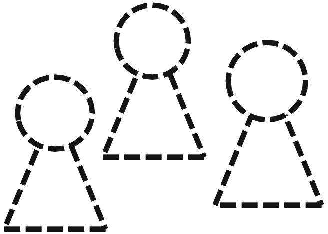

Recruit Communications

Budget Pacing

- Budget pacing is also important
- Control of budget pacing helps advertisers to…
  - Reach a wider range of audience
  - Avoid a premature campaign stop / overspending

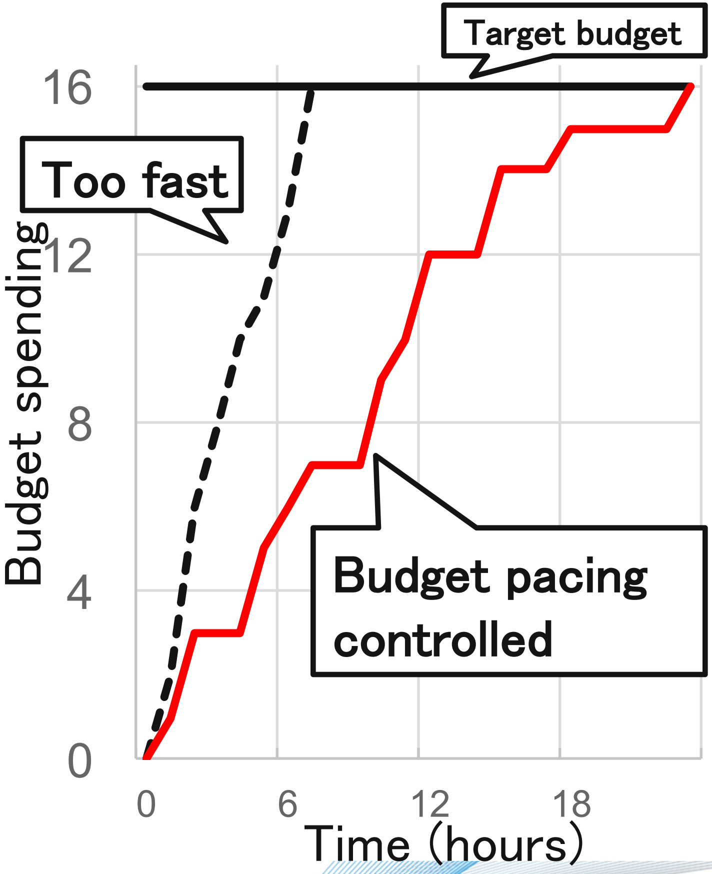

Budget Pacing

- Two lines of strategies
  - Bid modification
  - Probabilistic throttling

- We propose a different way based on a bipartite graph and inspired by financial theory

- Our research strengthens existing research rather than replacing it

Our Approach: Model and Formulation

- Display advertising is represented by the **edge-weighted bipartite matching problem**
- Edge weight = CTR
  (in our research)

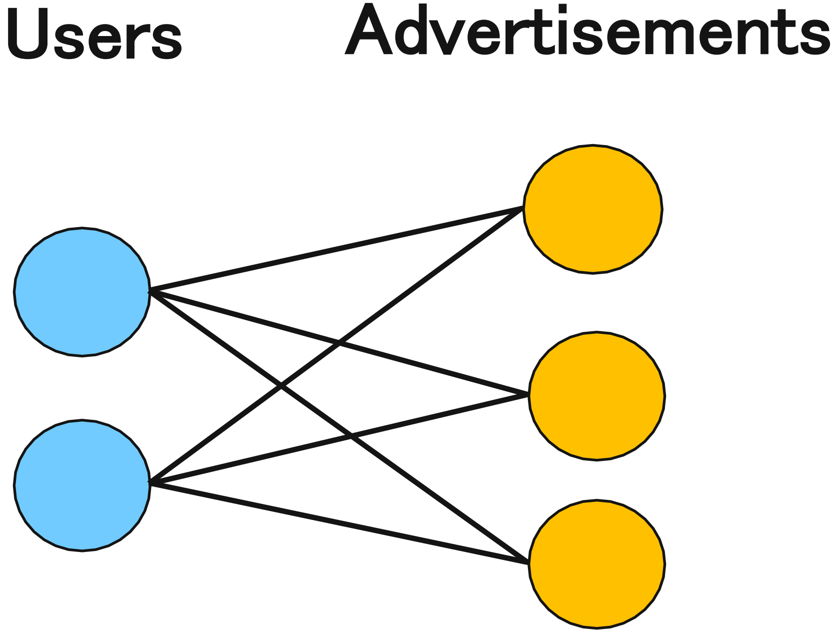

CTR: Click Through Rate

Our Approach: Model and Formulation

- Control the variation rate of CTR
- Find a matching with
  - High CTR
  - Low variation of CTR

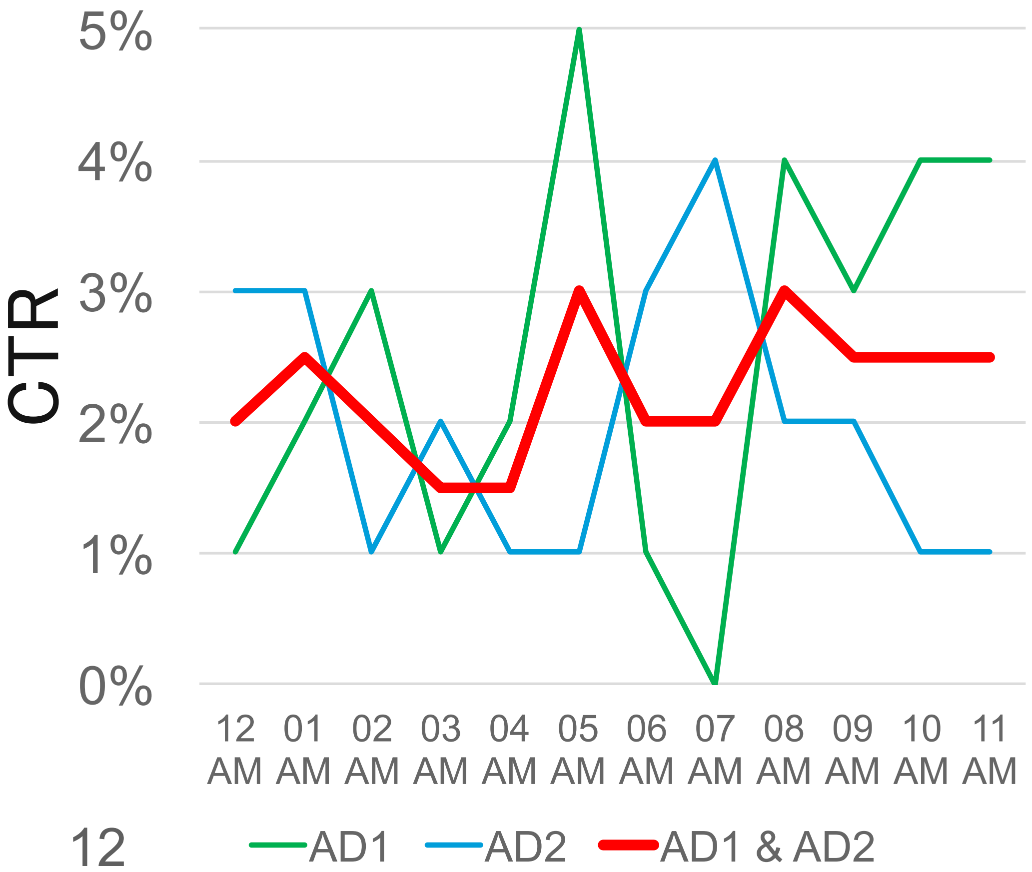

RECRUIT
Recruit Communications

Our Approach: Model and Formulation

- Objective Function as QUBO

$$\arg\max_{\boldsymbol{x}} \left\{ \boxed{\mathbb{E}\left[\boldsymbol{w}^{\mathrm{T}}\right]\boldsymbol{x}} - \alpha \boxed{\boldsymbol{x}^{\mathrm{T}} W \boldsymbol{x}} - \boxed{\beta\left(P\boldsymbol{x} - \mathbf{1}\right)^{2}} \right\}$$

Maximize CTR **Low variation** **Constraint**

N_a, N_c : Size of each vertice

$\boldsymbol{x} \in \{0,1\}^{N_a \times N_c}$: Decision variable

$\boldsymbol{w} : \Omega \to \mathbb{R}^{N_a \times N_c}$: Weights vector(CTR, CVR etc) for each edge

$W \in \mathbb{R}^{(N_a \times N_c)^2}$: Covariance matrix of $\boldsymbol{w}$

P : Matrix which expresses a restriction

$\mathbf{1}$: Vector with all elements as one

$\alpha, \beta \in \mathbb{R}$: Control parameters

RECRUIT
Recruit Communications

Experiments with Real World Data

- Setup
  - Data: our Display Ad campaign data
  - Combination: 14-advertisement campaign and 24 user (cluster)
  - CTR and its variation are estimated by historical data
    - Sample average and variance-covariance matrix

Experiments with Real World Data

- Use quantum annealing processor, D-Wave 2X
- Some techniques for optimization
  - Pruned edges with less impact
  - Reduce the solution space by clustering users

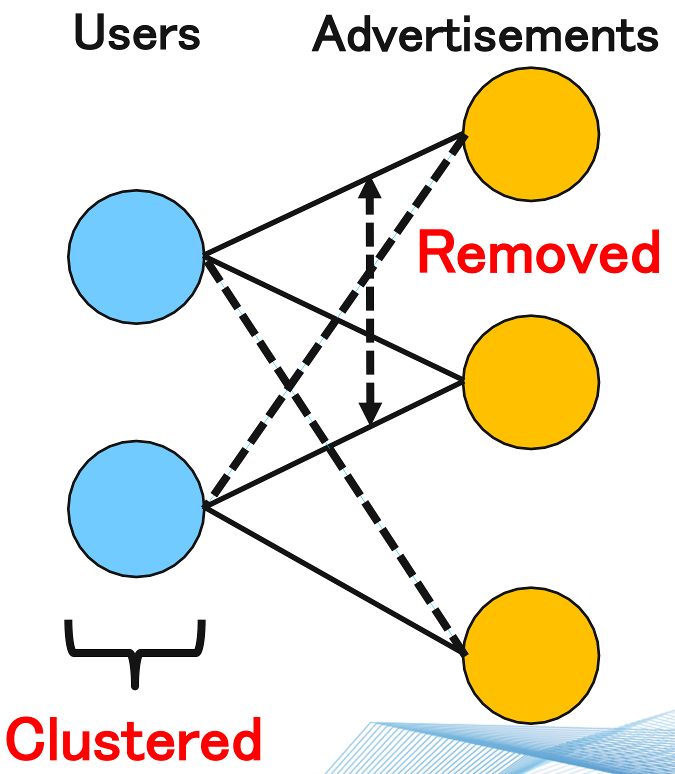

Experiments with Real World Data

- Correlation between edges (CTR correlation)

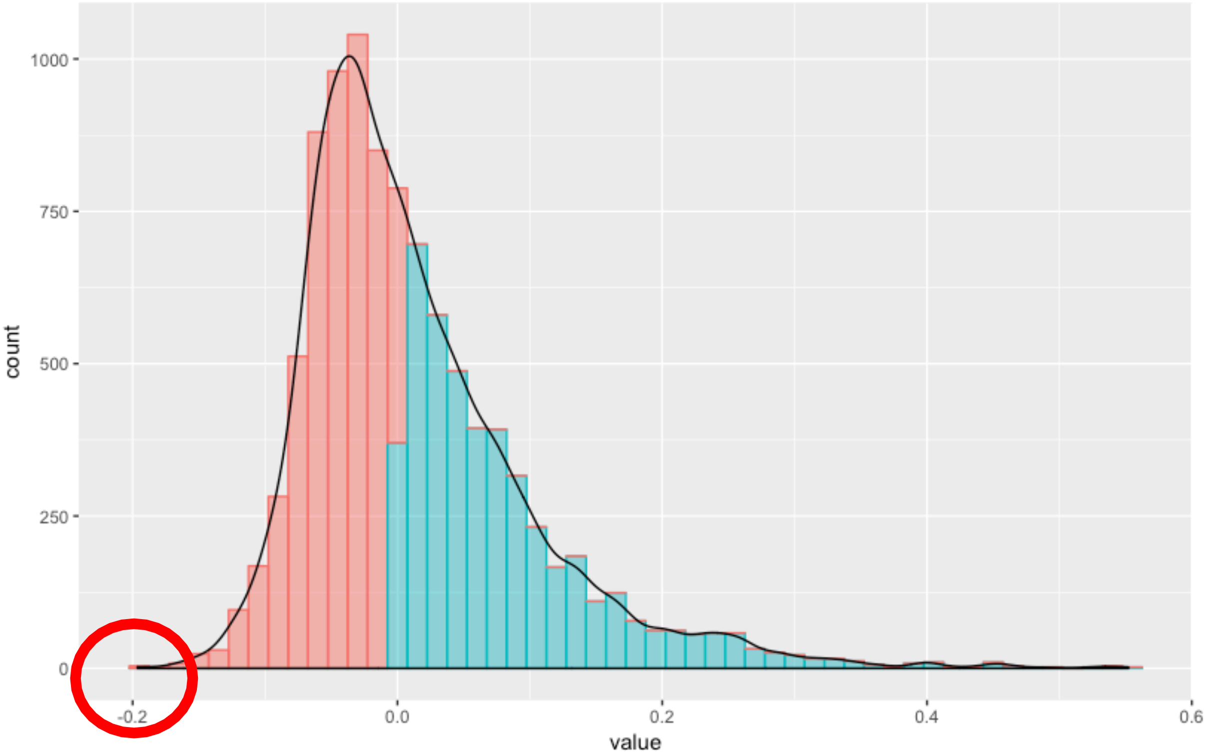

Recruit Communications

Performance Measure

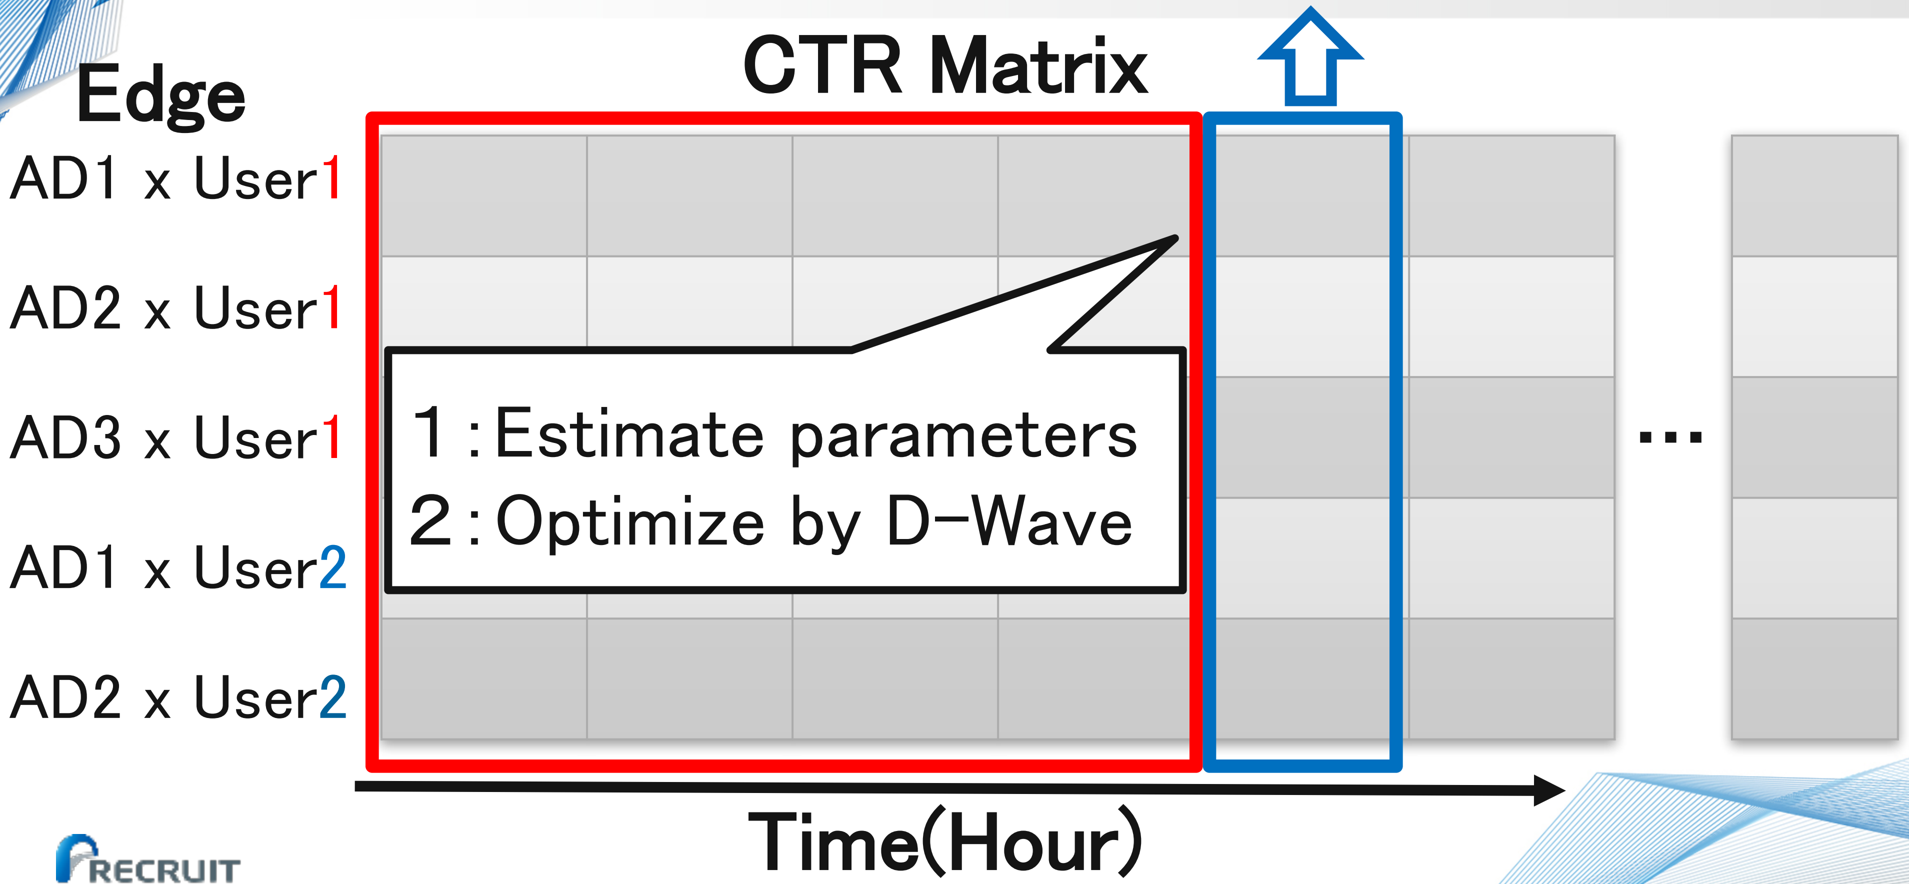

Performance Measure

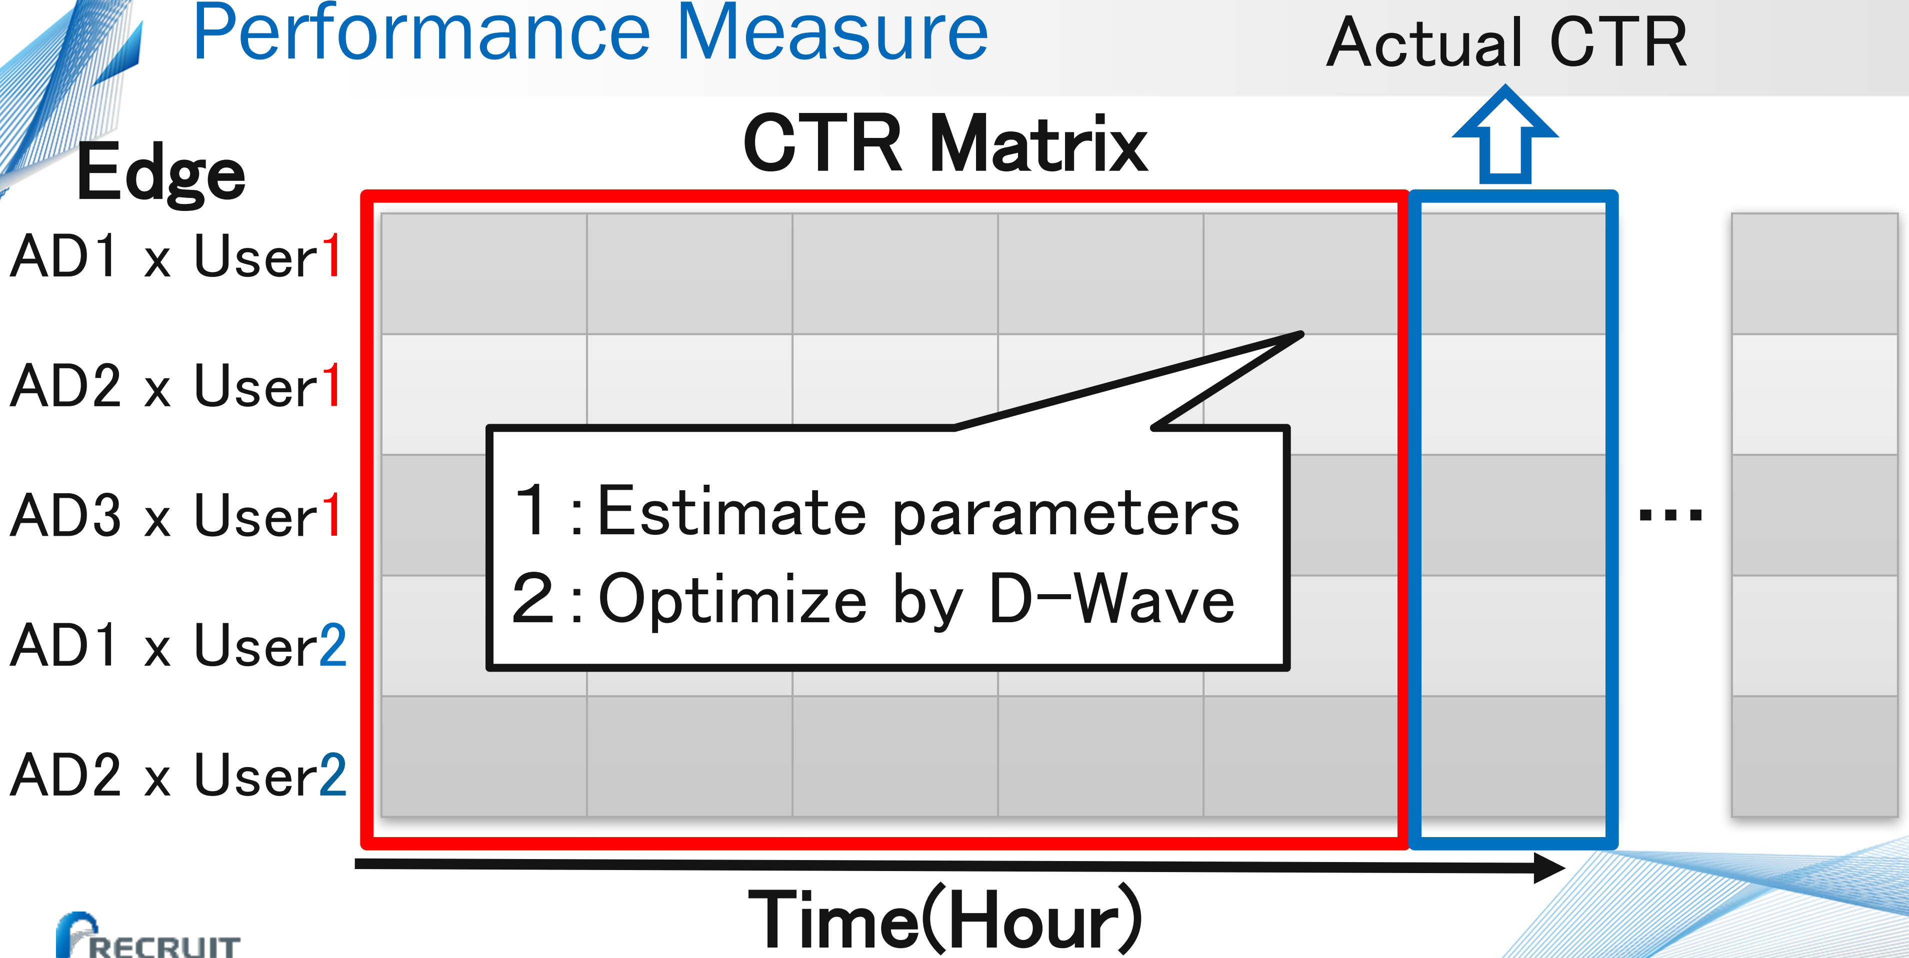

Experiments with Real World Data

- Hourly performance of each strategy
- Greedy method: Choose maximum CTR edge for each user

Experiments with Real World Data

- Quantum annealing(QA) finds a better solution than the greedy method
  - Almost same CTR level
  - Low variation of CTR

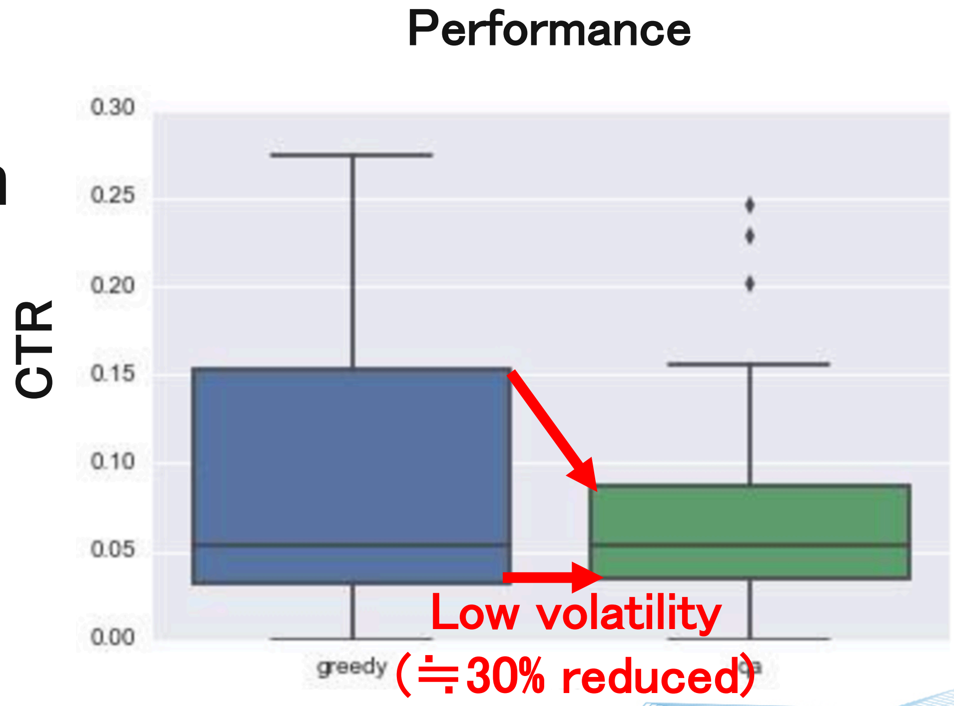

4. Summary

- Budget pacing is important for display advertising
- Formulate the problem as QUBO
- Use D-Wave 2X to solve budget pacing control optimization problem
- Quantum annealing finds a better solution than the greedy method.

Thank you for listening

RECRUIT
Recruit Communications

Memo

References

1. Budget pacing for targeted online advertisements at Linkedin
   - D. Agarwal, S. Ghosh, K. Wei, and S. You
   - Proceedings of the 20th ACM SIGKDD, 2014
2. Real time bid optimization with smooth budget delivery in online advertising
   - K.-C. Lee, A. Jalali, and A. Dasdan
   - The Seventh International Workshop on Data Mining for Online Advertising, 2013.
3. Online allocation of display ads with smooth delivery
   - A. Bhalgat, J. Feldman, and V. Mirrokni
   - Proceedings of the 18th ACM SIGKDD, 2012.

References

- •[Agarwal+ 2014] LASER: A scalable response prediction platform for online advertising, WSDM 2014 •
  [Ananthanarayanan+ 2013] Photon: Fault-tolerant and scalable joining of continuous data streams, SIGMOD 2013
  •[Chapelle+ 2014] Simple and scalable response prediction for display advertising, TIST 2014
  •[Graepel+ 2010] Web-scale bayesian click-through rate prediction for sponsored search advertising, ICML 2010
  •[He+ 2014] Practical lessons from predicting clicks on ads at Facebook, ADKDD 2014
  •[McMahan+ 2013] Ad click prediction : a view from the trenches, KDD 2013
  •[Yan+ 2014] Coupled group lasso for web-scale CTR prediction in display advertising, ICML 2014
  •[Yin+ 2014] Estimating ad group performance in sponsored search, WSDM 2014

HETEROGENEOUS QUANTUM COMPUTING FOR SATELLITE OPTIMIZATION

GIDEON BASS

BOOZ ALLEN HAMILTON

September 2017

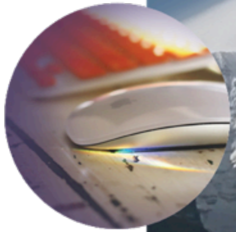

Agenda

- **Quantum Annealing in the field**
- **Problem Statement**
- **Results**
- **Conclusions**

Quantum Annealing has many real-world applications

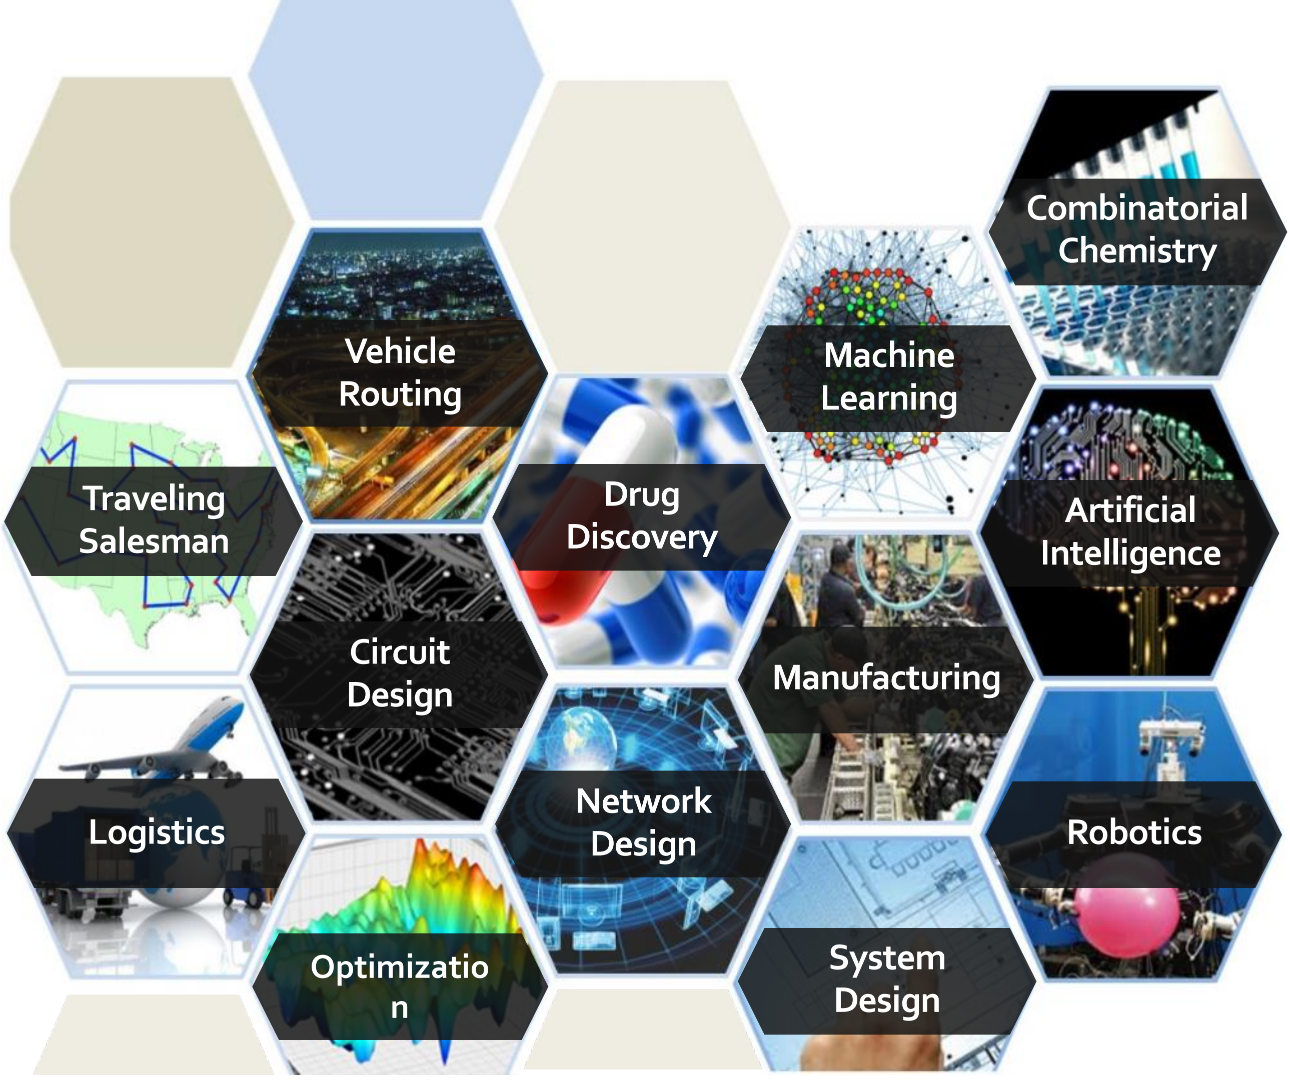

However most research has been theoretical

On the readiness of quantum optimization machines for industrial applications

Alejandro Perdomo-Ortiz,[1,2,3,*] Alexander Feldman,[4] Asier Ozaeta,[5] Sergei V. Isakov,[6] Zheng Zhu,[7] Bryan O'Gorman,[1,8,9] Helmut G. Katzgraber,[7,10,11] Alexander Diedrich,[12] Hartmut Neven,[13] Johan de Kleer,[4] Brad Lackey,[14,15,16] and Rupak Biswas[17]

[1]*Quantum Artificial Intelligence Lab., NASA Ames Research Center, Moffett Field, California 94035, USA*
[2]*USRA Research Institute for Advanced Computer Science (RIACS), Mountain View California 94043, USA*
[3]*Department of Computer Science, University College London, WC1E 6BT London, UK*
[4]*Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, California 94304, USA*
[5]*QC Ware Corp., 125 University Ave., Suite 260, Palo Alto, California 94301, USA*
[6]*Google Inc., 8002 Zurich, Switzerland*
[7]*Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843-4242, USA*
[8]*Berkeley Center for Quantum Information and Computation, Berkeley, California 94720 USA*
[9]*Department of Chemistry, University of California, Berkeley, California 94720 USA*
[10]*1QB Information Technologies (1QBit), Vancouver, British Columbia, Canada V6B 4W4*
[11]*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*
[12]*Fraunhofer IOSB-INA, Lemgo, Germany*
[13]*Google Inc., Venice, California 90291, USA*
[14]*Joint Center for Quantum Information and Computer Science, University of Maryland, College Park, Maryland 20742, USA*
[15]*Departments of Computer Science and Mathematics, University of Maryland, College Park, Maryland 20742, USA*
…al Security Agency, Ft. George G. Meade, Maryland 20755, USA
…NASA Ames Research Center, Moffett Field, California 94035, USA

(Dated: September 1, 2017)

…monstrate that quantum annealing and, in particular, quantum anneal- … the potential to outperform current classical optimization algorithms … he benchmarking of these devices has been controversial. Initially, however, these were quickly shown to be not well suited to detect …chmarking shifted to carefully crafted synthetic problems designed

Traffic flow optimization using a quantum annealer

Florian Neukart*[1], David Von Dollen[1], Gabriele Compostella[2], Christian Seidel[2], Sheir Yarkoni[3], and Bob Parney[3]

[1]Volkswagen Group of America, San Francisco, USA
[2]Volkswagen Data:Lab, Munich, Germany
[3]D-Wave Systems, Inc., Burnaby, Canada

Abstract

Quantum annealing algorithms belong to the class of meta-heuristic tools, applicable for solving binary optimization problems. Hardware implementations of quantum annealing, such as the quantum processing units (QPUs) produced by D-Wave Systems, have been subject to multiple analyses in research, with the aim of characterizing the technology's usefulness for optimization and sampling tasks. In this paper, we present a real-world application that uses quantum technologies. Specifically, we show how to map certain parts of a real-world traffic flow optimization problem to be suitable for quantum annealing. We show that time-critical optimization tasks, such as continuous redistribution of position data for cars in dense road networks, are suitable candidates for quantum computing. Due to the limited size and connectivity

Satellite Coverage
Quantum Optimization

Satellite Coverage Optimization

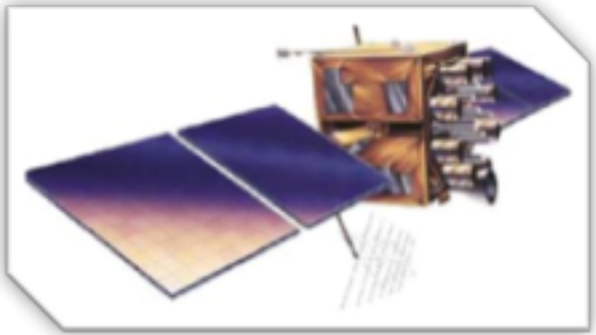

Summary: Group satellite together in such a way as to maximize coverage.

Data: For any possible grouping of satellites, a coverage percentage

Goal: Assign each of N satellites to k groups, such that total mean coverage is maximized

- Satellites change position and **require constant re-optimization**
- **Brute force solving is out of the question**; even trivial subsets of the satellites form too many combinations to check.
- Quantum technology offers a promise to **perform combinatorial optimization much faster**, while yielding better coverage outcomes.

The weighted k-clique problem

- This problem can be reformulated as a graph problem, called the **k-clique problem**

- Each potential group of satellites in a sub-constellation can be considered a node on a graph
  - Each node is given a weight equal to the coverage provided
  - If both sub-constellation use the same satellite, the nodes are unconnected
  - The goal is thus to find the k nodes with the highest total weight that are all mutually connected (a "clique")

- This problem can then be expressed as a QUBO, and sent to the quantum computer

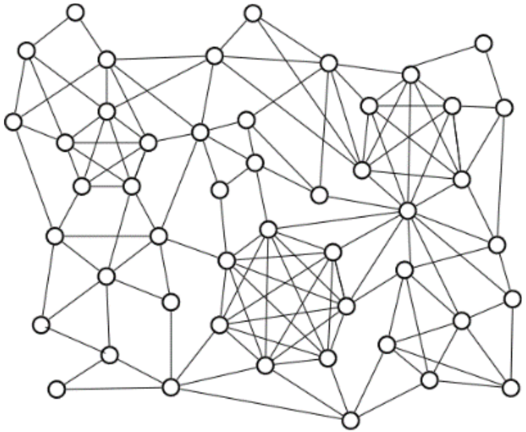

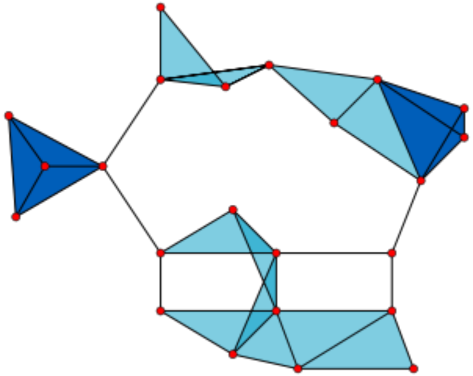

Designing the QUBO

Constraints:

1. Choose only nodes that are connected
2. Maximize the sum of coverages for each group chosen
3. Choose a number of qubits equal to the number of available satellites

Each (logical) qubit represents a potential grouping of satellites

Connections represent a grouping that is non-overlapping (does not use the same satellite in multiple groups)

$$H = \sum_{i<j} 2(w_i + w_j)$$

$$H = \sum_{i} -A w_i x_i$$

$$H = W\left(\sum_{i} x_i - 8\right)^2 = 64W - \sum_{i} 8W x_i + \sum_{i<j} x_i x_j$$

W is the qubit maximum weight

Quantum Hardware is rapidly maturing

- This satellite optimization problem is **a prime candidate for a quantum approach** when used in concert with classical computing resources.
- The application to satellites could be **the first major quantum success** when applied to a real-world full-scale problem.
- However, with current numbers, we would still need 10^4-10^5 qubits to fully embed this problem
- Thus, we created a heterogeneous approach that combines classical processing and quantum annealing

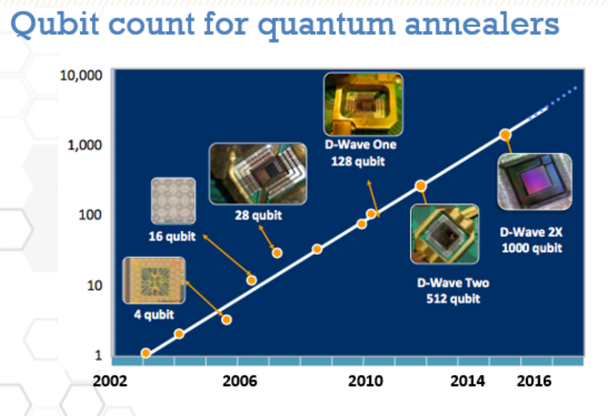

Heterogeneous techniques: TWO APPROACHES

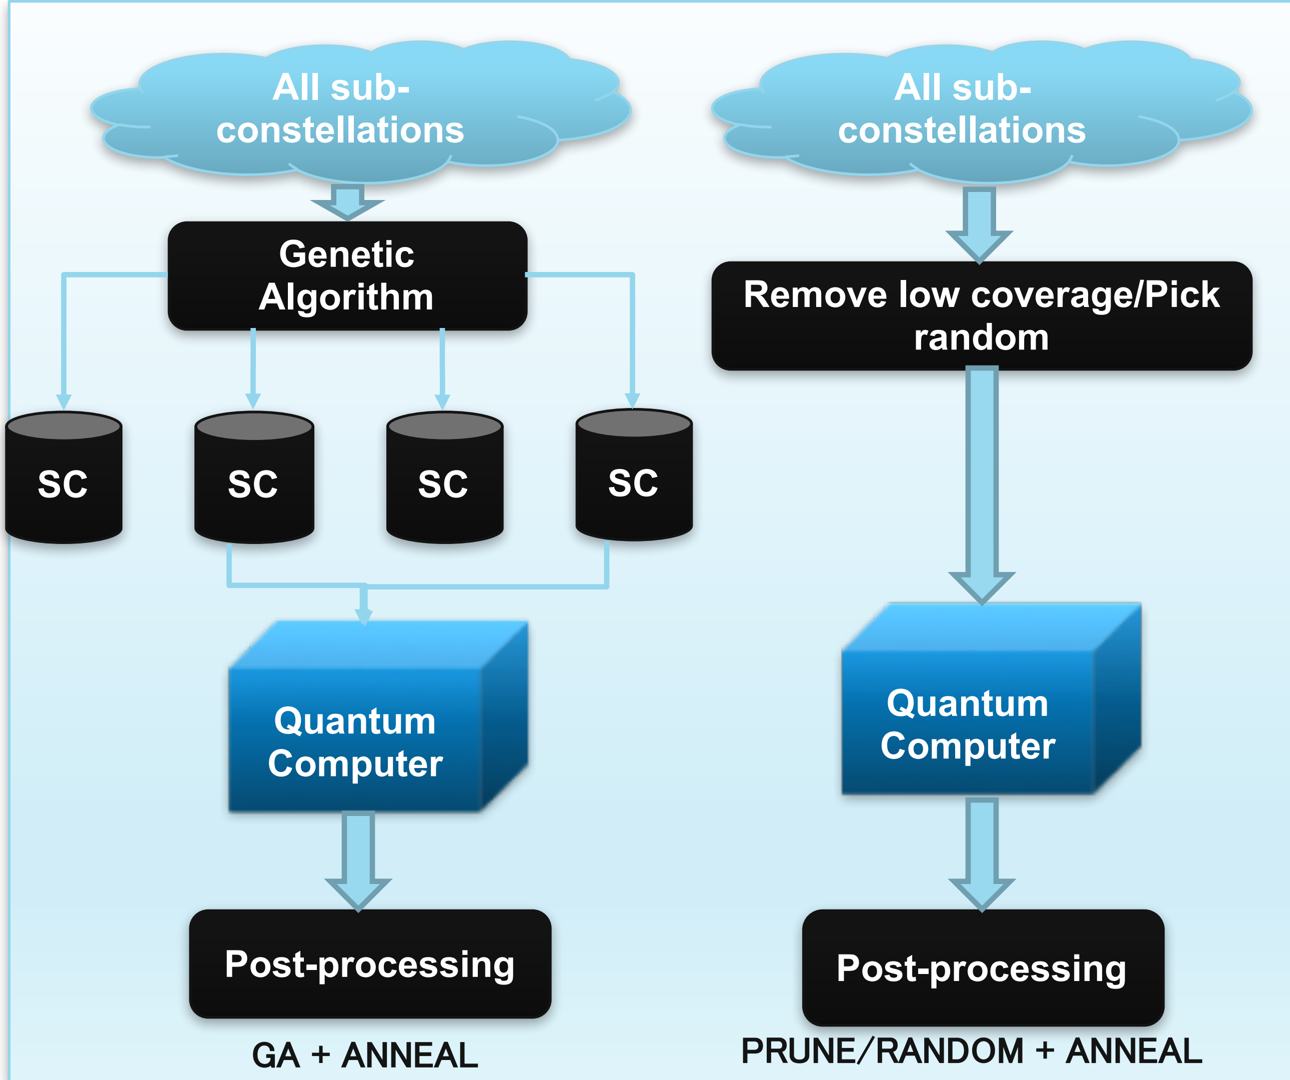

Heterogeneous Computing Models

Method	Pros	Cons
Classical Heuristics	Can provide fairly good results. Can be run on classical machine.	Cannot be run on current QA devices, no quantum speed-up, scaling uncertain
GA pre-processing	Searches full decision space, produces solid results	Middle of the road performance and speed, many parameters to tune
Prune and Anneal	Very good results in good time, most similar to existing technique	Does not explore full solution space, requires domain knowledge

Results Comparison: Quantum Simulator

- An 80% coverage(red) is the minimum acceptable average.
- The eight colored bars represent individual sets, black bar (and dotted line) is overall average
- Quantum approach is faster and finds a significantly better results

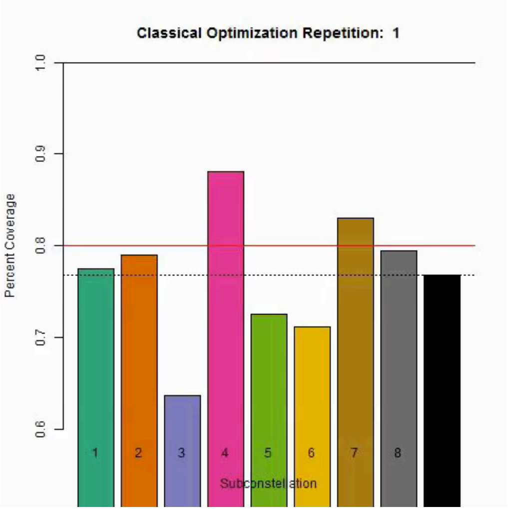

Purely Classical Genetic Algorithm

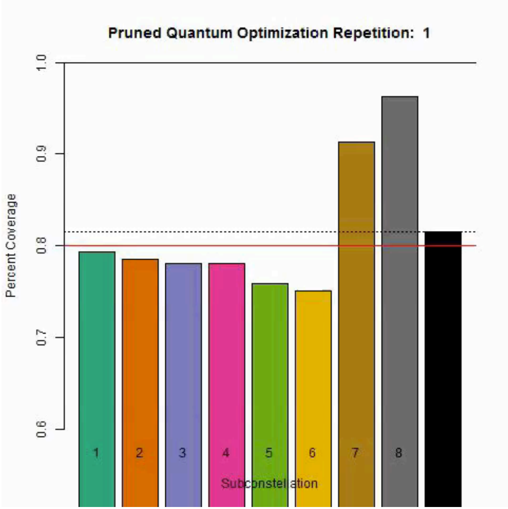

Simulated Quantum Prune and Anneal

Results Comparison: D-Wave

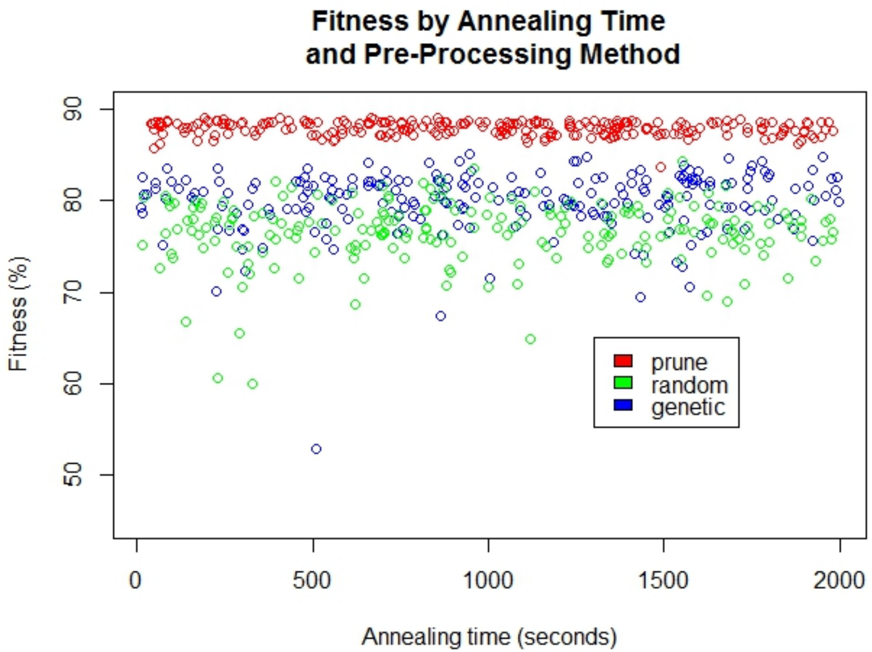

- Results are nearly constant with processing time
- Results are highly dependent on pre-processing method (color)
  - 80% is minimal acceptable
  - 90% is likely near the true maximum.

Results Comparison: D-Wave

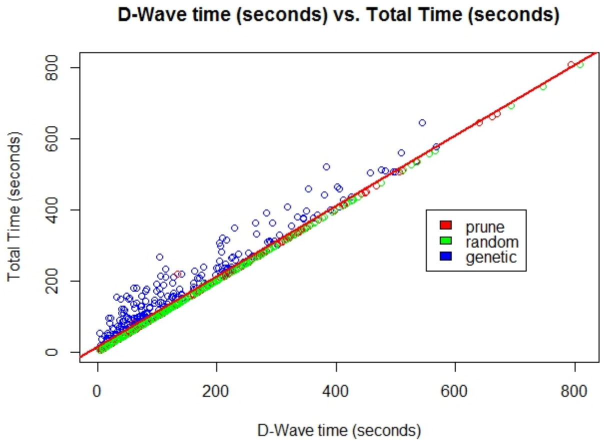

- D-Wave time makes up most of the time, GA adds a little more

- Including D-Wave's "Virtual Full Yield" does not significantly change performance while improving portability

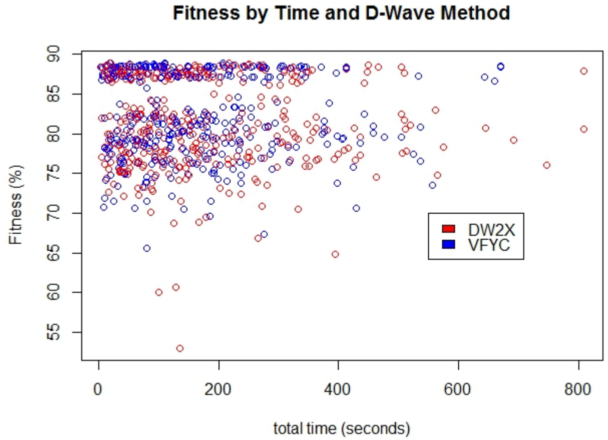

Summary

Method	Uses Domain-Knowledge	Time Needed	Performance
Prune + Anneal	✓	Very Little	90%
GA + Anneal	X	Some	80-85%
Random + Anneal	X	Very Little	75-80%

- The D-Wave functions best as a co-processor
- Performance is highly dependent on problem formulation, classical processing step
- Quantum portion does appear to provide significant improvement.

Conclusions

- As problems and datasets grow, modern computing systems have had to scale with them. **Quantum computing offers a totally new and potentially disruptive computing paradigm.**
- For problems like this satellite optimization problem, **heterogeneous quantum techniques will be required to solve the problem at larger scales.**
- Preliminary results on this problem using heterogeneous classical/quantum solutions **are very promising.**
- Exploratory studies in this area **have the potential to break new ground** as one of the first applications of quantum computing to a real-world problem

Thank You

GIDEON BASS

BOOZ ALLEN HAMILTON