

Angler: Dark Pool Resource Allocation

James Choncholas, Ketan Bhardwaj, Vlad Kolesnikov, Ada Gavrilovska

Edge Infrastructure Platforms Are Growing

- Cloud Edge Offerings
 - GCP partners with Verizon, T-Mobile
 - Azure partners with Verizon, AT&T
 - AWS partnerships with Verizon, T-Mobile
- Programmable CDNs
 - Cloudflare
 - Akamai
 - Fastly
- ISPs
 - Cox Edge
 - Comcast partners with Vapor IO
- Edge Native Providers
 - fly.io
 - Equinix Metal
- Edge as On-Prem
 - Dell
 - Oxide
 - Google Distributed Cloud Edge
 - Red Hat
 - ...



EQUINIX

COMCAST



The Best Provider Has...

Low Network Latency

Affected by geography and connectivity.

Sufficient Capacity







PoPs are small and load is dynamic.

Competitive Pricing

Across full application footprint.

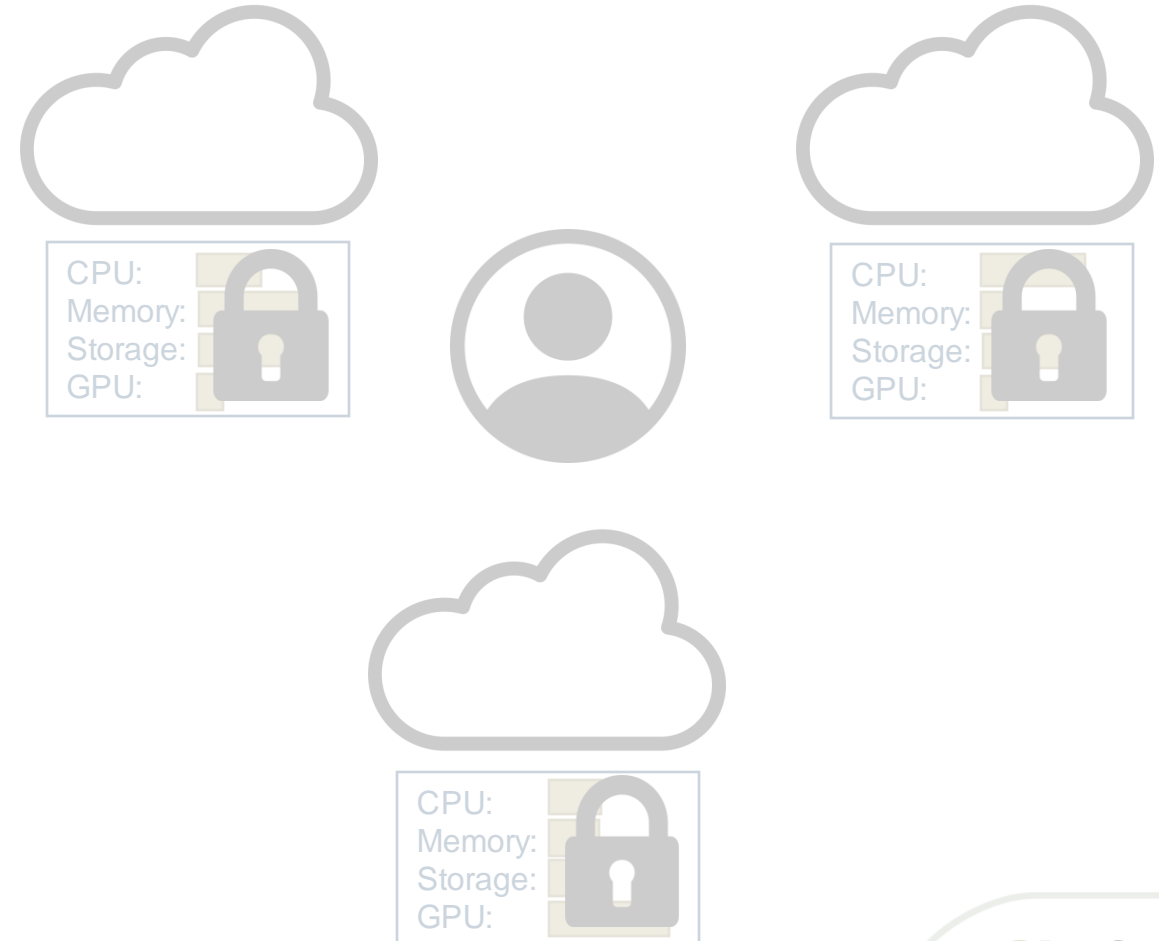
Distributed Resource Allocation

- Early Days
 - Eucalyptus ('09)
 - OpenNebula ('11)
 - Nimbus ('12)
 - RESERVOIR ('11)
- Both Cloud and Edge use LP solvers.
 - TetriSched
 - Edge Federation
 - ENTS (SEC'22)
- Requires providers share:
 - Number of machines.
 - Current Utilization.
 - Other infrastructure details.
- These are **trade secrets**.




Distributed Resource Allocation

How to allocate resources across providers without sharing confidential information?



Angler: Dark Pool Resource Allocation


- Supports resource allocation across multiple edge providers.
 - Manages the complexity of finding nearby providers.
 - Integrates with Kubernetes.
- Protects privacy of providers and requestors.
 - Contributions to the pool are secret.
 - Requests from the pool are secret.
- **Angler** leverages AGMPC.
 - Authenticated Garbling Multi-Party Computation.
 - Naively applying MPC to resource allocation is too slow.
- **Angler** has only 2x overhead vs. non-private baseline.



Angler Design

Challenges:

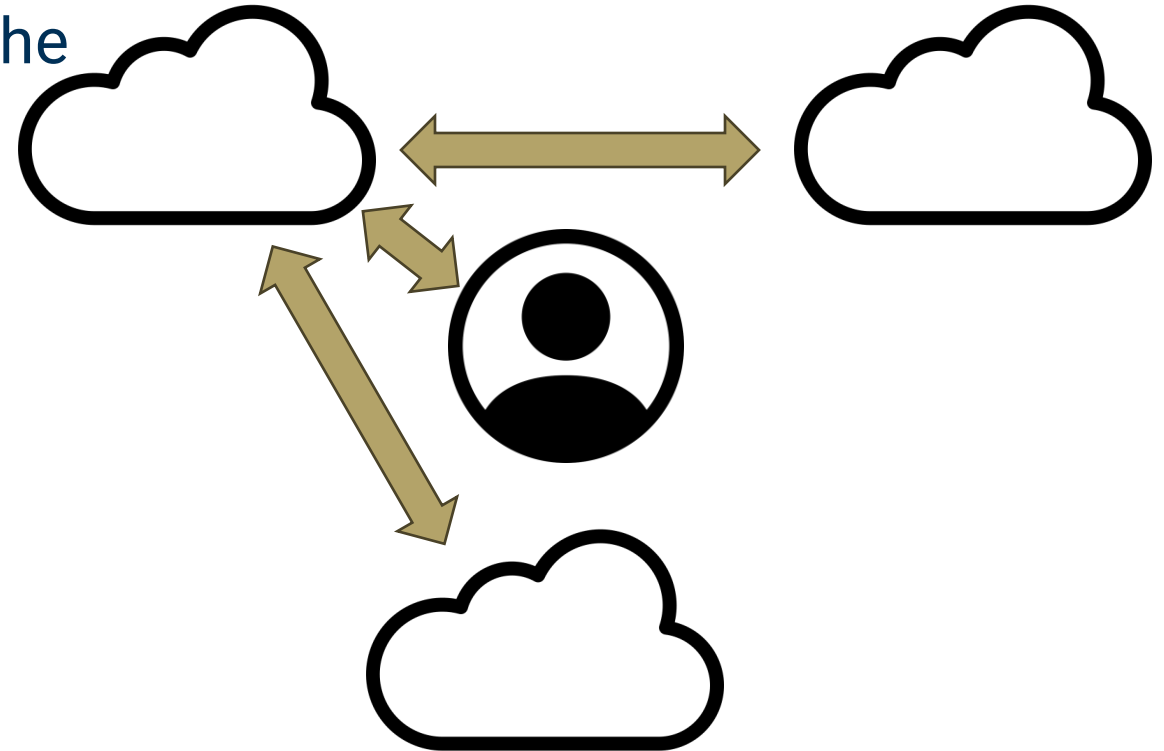
- AGMPC has polynomial complexity in the number of participants.
 - Reduce participant set through localized DHT-based discovery.



Angler Design

Challenges:

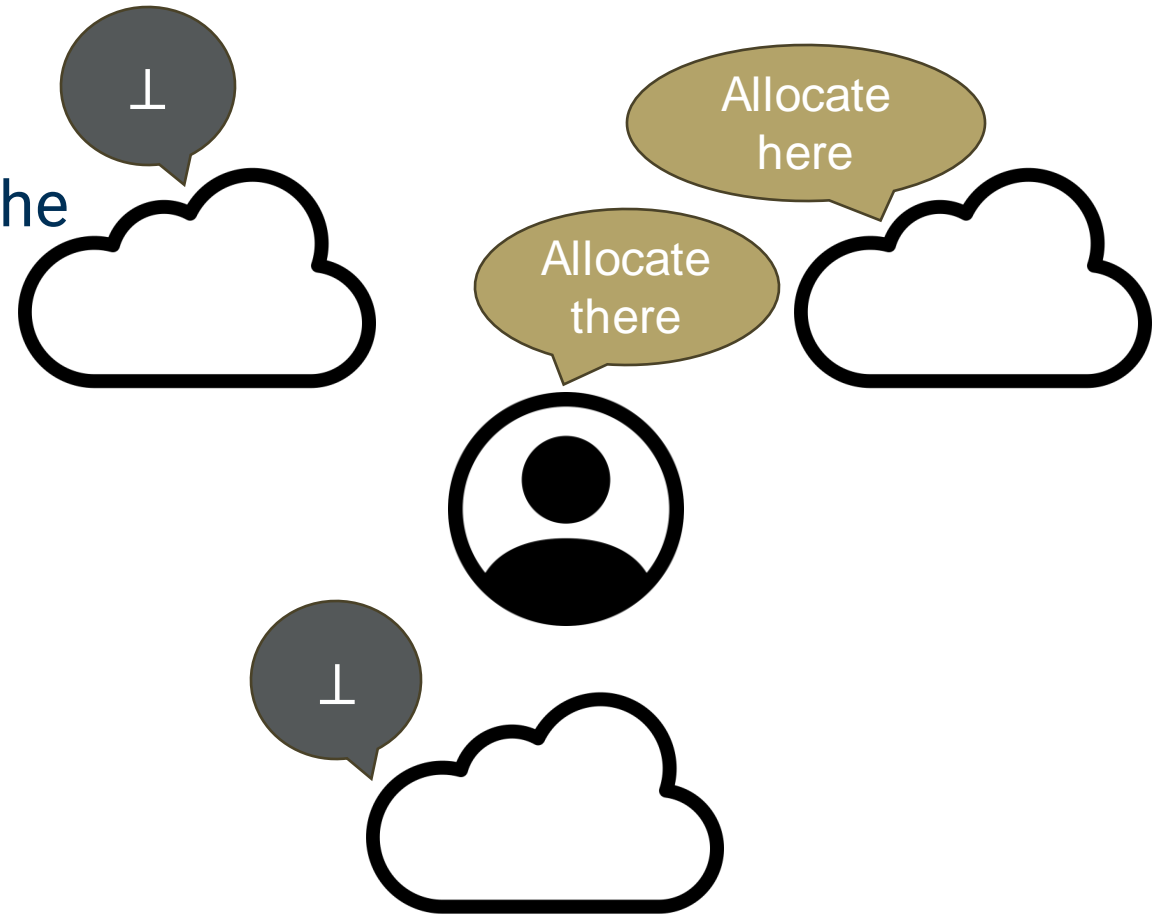
- AGMPC has polynomial complexity in the number of participants.
 - Reduce participant set through localized DHT-based discovery.
- AGMPC is network constrained.
 - Optimize network stack for AGMPC. (parallelize initialization and BBR)



Angler Design

Challenges:





- AGMPC has polynomial complexity in the number of participants.
 - Reduce participant set through localized DHT-based discovery.
- AGMPC is network constrained.
 - Optimize network stack for AGMPC. (parallelize initialization and BBR)
- All parties learn placement.
 - Output delivery on a "need-to-know" basis.
 - Overlap output delivery with provisioning.



Angler Design

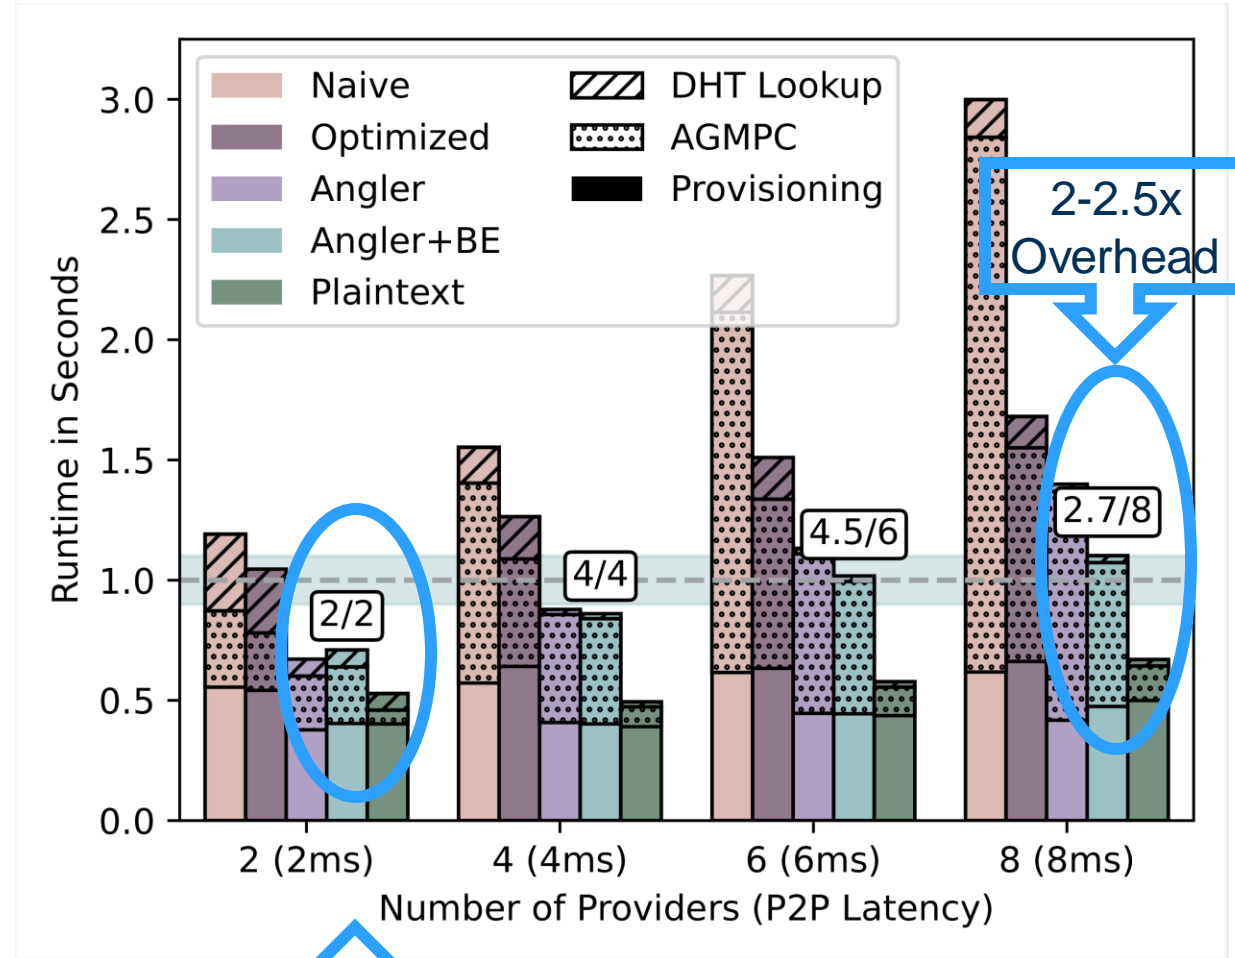
Challenges:

- AGMPC has polynomial complexity in the number of participants.
 - Reduce participant set through localized DHT-based discovery.
- AGMPC is network constrained.
 - Optimize network stack for AGMPC. (parallelize initialization and BBR)
- All parties learn placement.
 - Output delivery on a "need-to-know" basis.
 - Overlap output delivery with provisioning.
- Still too slow?
 - Best-effort semantics.




End-to-End Runtime Composite

- TCP congestion control (BBR).
- Parallelize network initialization.
- Targeted output delivery.
- Overlap output delivery with provisioning.
- Tailored DHT-based discovery.
- Best-effort semantics.



Conclusion

- **Angler** supports resource allocation from *dark pools*.
 - Manages the complexity of finding nearby providers.
 - Provisions Kubernetes namespaces on provider infrastructure.
- Protects privacy of providers and requestors.
- Multiparty Computation (MPC) keeps confidential information secret.
- Secure resource allocation has low overhead.



angler-dex.com



Angler: Dark Pool Resource Allocation

James Choncholas, Ketan Bhardwaj, Vlad Kolesnikov, Ada Gavrilovska

