Network-Aware Resource Allocation in Distributed Clouds

Dissertation Research Summary

Atakan Aral

Thesis Advisor: Asst. Prof. Dr. Tolga Ovatman

Istanbul Technical University – Department of Computer Engineering

E-mail: aralat@itu.edu.tr

April 4, 2016

Short Bio

- **Research and Teaching Assistant** at Istanbul Technical University, Turkey
- **PhD Candidate in Computer Engineering** at Istanbul Technical University, Turkey

- 2011: **MSc in Computer Science and Engineering** at Politecnico di Milano, Italy
- 2009: **BSc in Computer Engineering** at Istanbul Technical University, Turkey

- Resource Management (Allocation, Placement, Discovery), Inter-Cloud, Distributed Software and Systems, Graph Theory, Discrete Optimization

Outline

Outline

Motivation

- In distributed computing environments, up to 85 percent of computing capacity remains idle mainly due to poor optimization of resources.
- Resource management in distributed cloud systems is a complex problem due to factors such as:
  - Multiple objectives (QoS, network latency, throughput, cost, . . .)
  - Constraints (SLAs, processing capacity, network bandwidth, energy consumption, . . .)
  - Multi-tenancy
  - Inter-Cloud, services scaling accross data centers or clouds
  - Heterogeneity and dynamicity
  - Geo-distributed access
  - Magnitude of data stored/processed in the cloud

Overview and Timeline

- We have identified three problems to tackle in increasing levels of distribution.

  A Single Data Center Allocation of VMs to hosts (PMs)

  Inter-Cloud Allocation of multi-VM clusters across several clouds (DCs)

  Edge Computing Placement of data replicas on STaaS providers/cloudlets

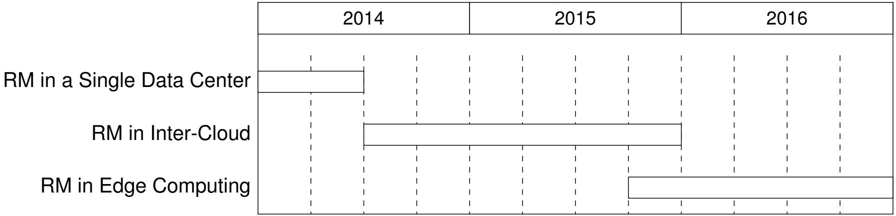

Outline

Single Data Center Scenario

Problem Definition

- Allocation of VM requests to hosts in order to **optimize utilization** and **minimize number of migrations.**

  Issues Heterogeneity, Host capacities, Unknown future requests

  Existing Solutions Linear Programming, Greedy algorithms, Heuristics

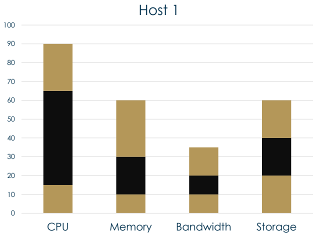

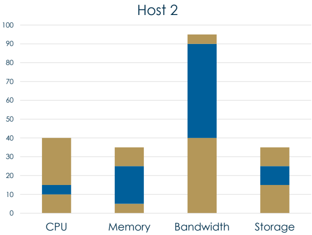

Single Data Center Scenario

Suggested Solution

- **Minimum span** evenness heuristic
  - Allocate each VM to the host whose evenness would become the greatest.
  - Evenness is measured as the minimum difference between the utilization of the most and the least utilized resources.
  - Baselines include skewness, standard deviation, etc.
- When a VM cannot be allocated, resort to MIP solution.
- Simulation results demonstrate effectiveness
  - Optimal placement in up to 10,8% of the cases, four times better than greedy
  - Postpone MIP algorithm up to 12,1% longer than greedy
  - MIP results in 34.5% less VM migrations than greedy

Inter-Cloud Scenario

Preliminary Information

Inter-Cloud / Federated Cloud

Mechanisms and policies for scaling hosted services across multiple, geographically distributed data centers and dynamically coordinating load distribution among these data centers.

- Allows seemingly infinite scalability,
- Provides better geographical coverage,
- Avoids vendor lock-in and eases hybridization,
- **Allows to scale VMs across multiple vendor clouds.**
  - fault tolerance, proximity to user base, vendor independence, cost benefits

Inter-Cloud Scenario

Problem Definition

- Allocation of VM clusters across Inter-Cloud in order to **improve QoS** and **reduce costs.** i.e. Virtual Network Embedding

  Issues: Heterogeneity, Resource capacities, Unknown future requests, Network latency and bandwidth

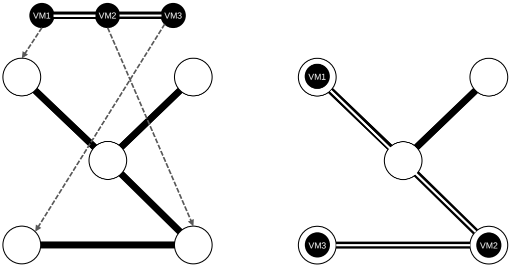

Inter-Cloud Scenario

Suggested Solution

- **Topology Based Mapping** algorithm
  - Map VM clusters to the subgraphs of the Inter-Cloud topology that is isomorphic to their topology.
  - Fall back to a heuristic in case of failure
- Simulation results demonstrate effectiveness
  - Decrease in deployment latency (by placing VMs close to the broker)
  - Decrease in communication latency (by placing connected VMs to the neighbour clouds)
  - Shorter execution time and increased throughput
  - Reduced resource costs (by balancing load and avoiding overload in any DC)

Inter-Cloud Scenario

Suggested Solution

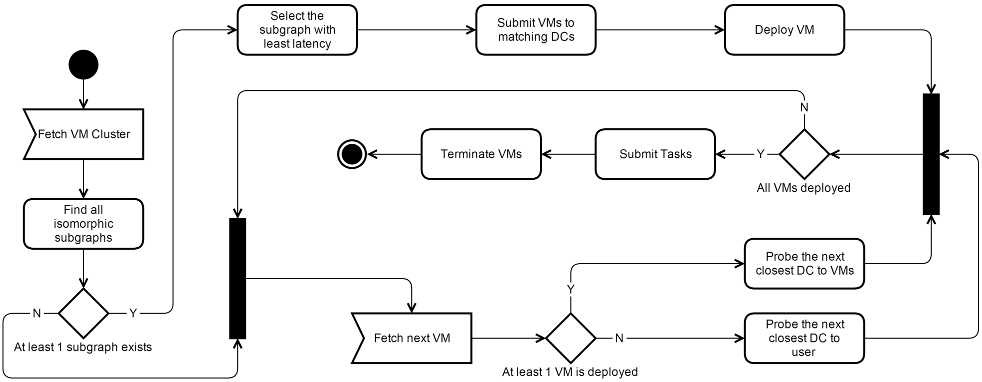

Inter-Cloud Scenario

Selected Results

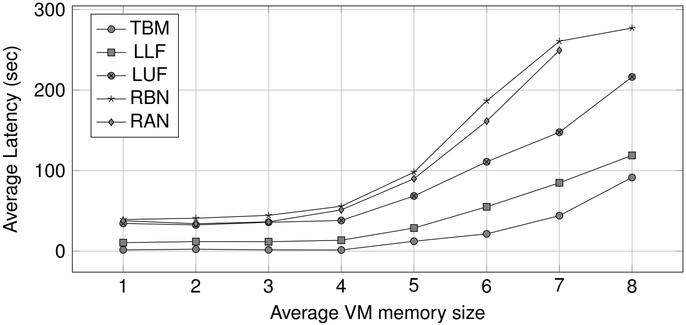

Figure: Evaluation results in terms of average latency with varying VM size.

Inter-Cloud Scenario

Selected Results

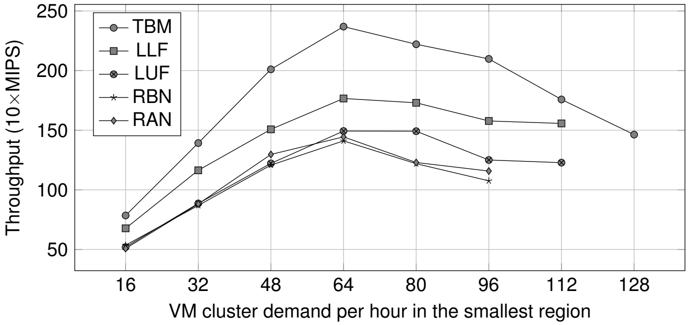

Figure: Evaluation results in terms of throughput with varying demand.

Edge Computing Scenario

Preliminary Information

Edge Computing

Pushing the frontier of computing applications, data, and services away from centralized nodes to the logical extremes of a network (e.g. mobile devices, sensors, nano data centers, routers, modems, ...)

- Provides low-latency access to computing resources for mobile code offloading.
- However, many applications still need to access data that is stored centrally
  - Due to the limited storage capacity of the edge entities, economic constraints, availability for offline analysis, simpler maintenance and concurrency control

Edge Computing Scenario

Problem Definition

- **Cost-effective** and **latency-aware** placement of data replicas
  - Issues: Cost-latency tradeoff, Internet topology, dynamic and mobile demand
  - Existing Solutions: Linear programming, centralized heuristics, caching
- An online, dynamic, distributed and light-weight replica placement algorithm is needed.

Edge Computing Scenario

Suggested Solution

- Distributed Facility Location Heuristic

Facility Location Problem

Finding a placement of facilities in order to serve the demands of geographically distributed customers with minimum cost (transportation and facility building)

- Demand for each replica i from each neighbour j: D_{ij}
- Average latency for each replica i from each neighbour j: L_{ij}
- Latency from each node k to each neighbour j: N_{jk}
- Cost of storing each replica i at each neighbour and current location j: C_{ij}
- User provided level of expansion: λ

Edge Computing Scenario

Suggested Solution

- Create a replica of object *i* at neighbour *j* iff:

$$L_{ij} D_{ij} \lambda > C_{ij}$$

- Remove the replica of the object *i* at *k* iff:

$$\sum_{\forall j} (L_{ij} D_{ij} \lambda) < C_{ik}$$

- Duplicate the replica of the object *i* from *k* to *l* iff:

$$L_{il} D_{il} \lambda > C_{il} \wedge \sum_{\forall j \neq l} (L_{ij} D_{ij} \lambda) > C_{ik}$$

Edge Computing Scenario

Suggested Solution

- Migrate the replica of the object *i* from *k* to *l* iff:

$$\sum_{\forall j}(L_{ij}D_{ij}\lambda) - \left(\sum_{\forall j \neq l}\left((L_{ij}+N_{kl})D_{ij}\lambda\right) + (L_{il}-N_{kl})D_{il}\lambda\right) > C_{il} - C_{ik}$$

- A special case where $\exists! j[D_{ij} > 0]$:

$$N_{kl}D_{il}\lambda > C_{il} - C_{ik}$$

Edge Computing Scenario

An Example

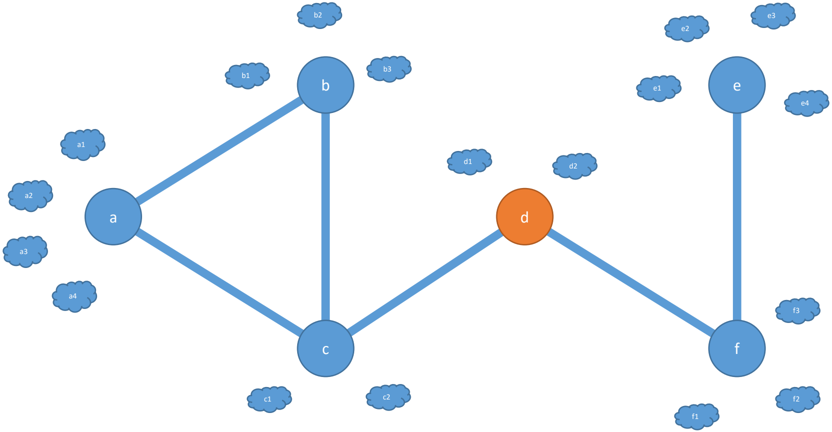

Edge Computing Scenario

An Example

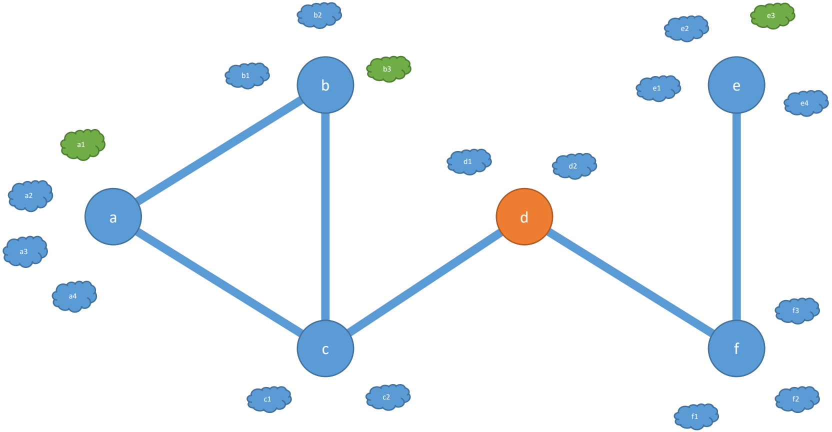

Edge Computing Scenario

An Example

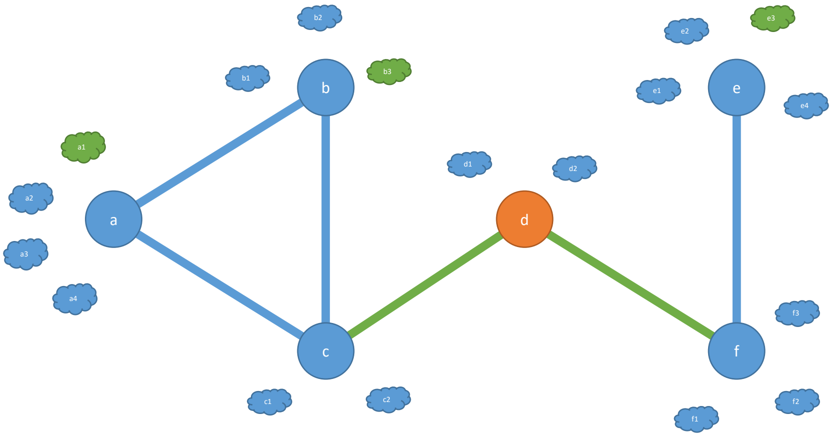

Edge Computing Scenario

An Example

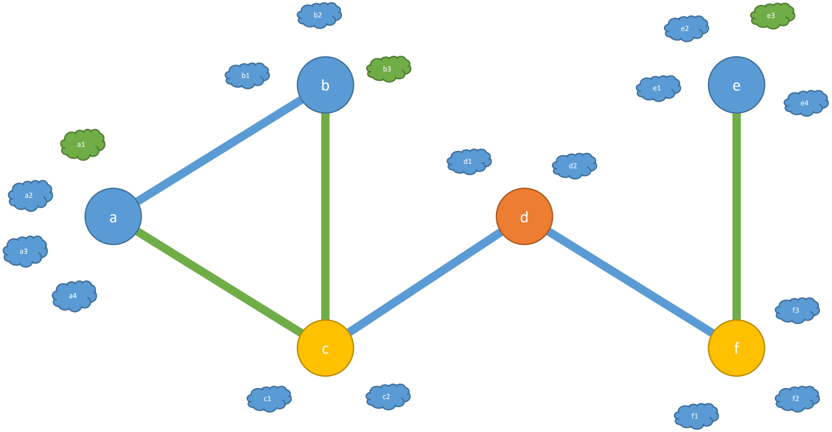

Edge Computing Scenario

An Example

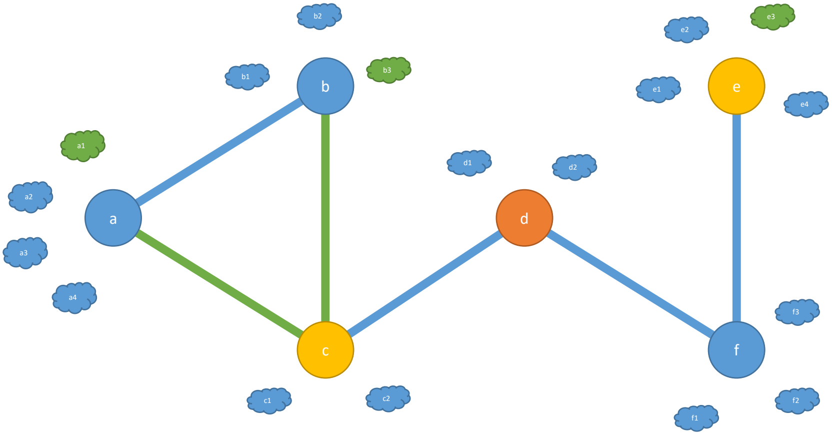

Edge Computing Scenario

An Example

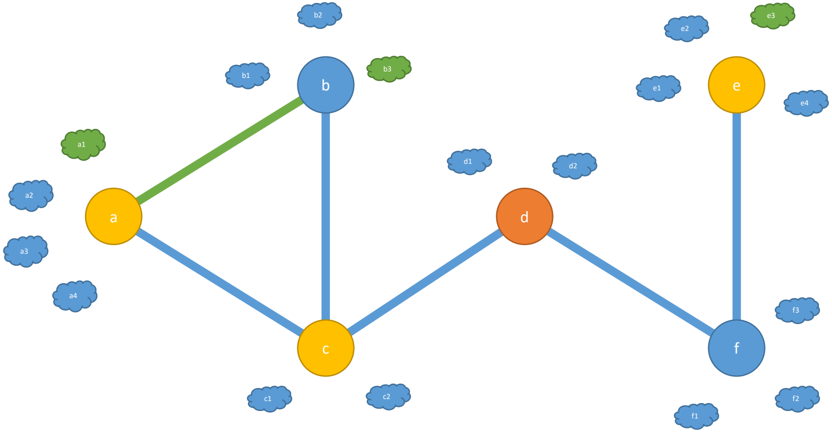

Edge Computing Scenario

An Example

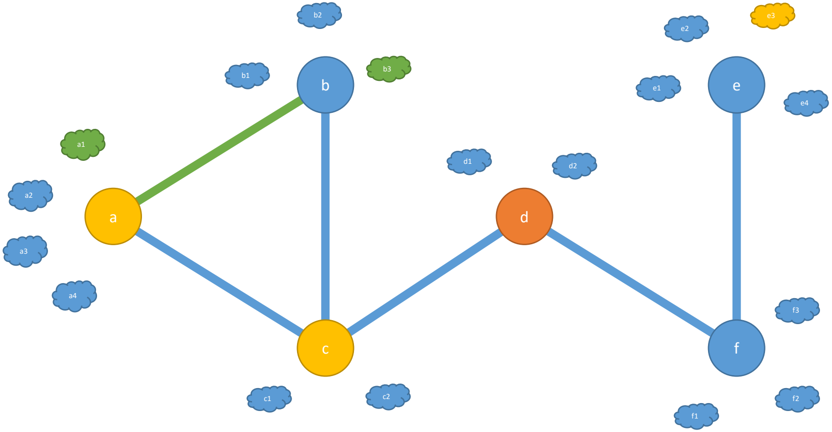

Edge Computing Scenario

An Example

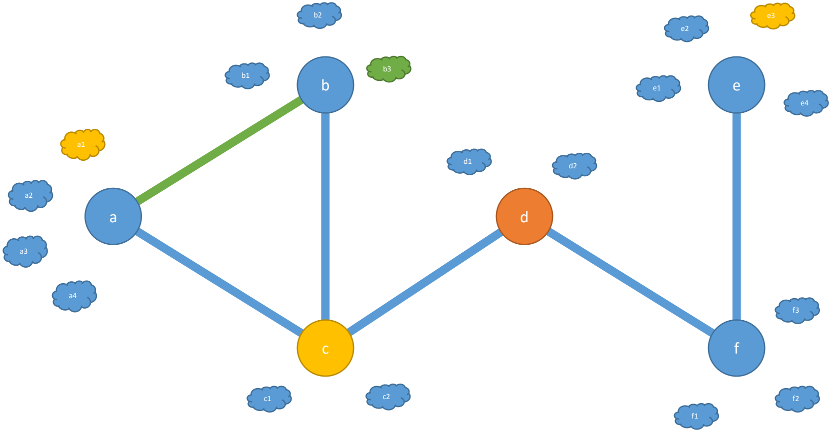

Edge Computing Scenario

An Example

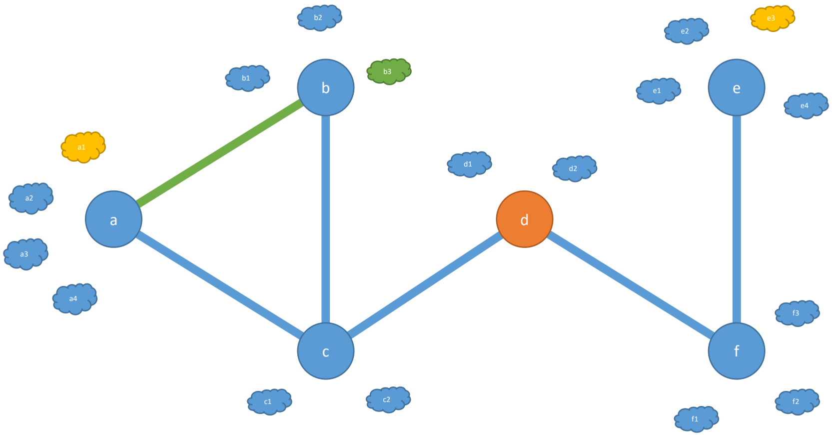

Outline

Conclusion

Summary

A Single Data Center Allocation of VMs to hosts (PMs)

- Published in CLOSER 2014

Inter-Cloud Allocation of multi-VM clusters across several clouds (DCs)

- Published in IEEE CLOUD 2015, manuscript in revision for the Journal of Systems and Software

Edge Computing Placement of data replicas on STaaS provides/cloudlets

- Ongoing research, manuscript in preperation

Conclusion

Future Work

- Demonstration of geographical and temporal locality
- Operation effectiveness and precedence
- Local maxima
- Bandwidth and network overhead consideration
- Replica location discovery
- Messaging protocol between replica locations
- Consistency control

Thank you!