Technical Disclosure Commons

June 09, 2017

Shape matching scheduler

Xiao Zhang

Richard Wu

Follow this and additional works at: http://www.tdcommons.org/dpubs_series

Recommended Citation

Zhang, Xiao and Wu, Richard, "Shape matching scheduler", Technical Disclosure Commons, (June 09, 2017)
http://www.tdcommons.org/dpubs_series/549

Shape matching scheduler

ABSTRACT

Scheduling jobs with different CPU and memory requirements often leads to resources being stranded or underutilized across at least one dimension e.g., CPU or memory. This leads to sub-optimal utilization of data center or cloud infrastructure and wastage of valuable resources. Solutions to address the resource stranding problem that involve pooling resources across different machines are complex and impose extra hardware requirements. This disclosure describes a shape matching optimization technique that is a flexible software-based solution to preferentially allocate jobs to machines by matching the shape of the job to the available shape of the machine. The shape matching scheduler improves resource utilization.

KEYWORDS

- virtual machines

- resource stranding

- job scheduling

- job allocation

BACKGROUND

To schedule a computing job, a machine should satisfy resource requirements for the job across multiple dimensions, e.g., CPU, memory, etc. If a machine runs out of CPU resources while executing a set of jobs, despite having abundant memory resources available, additional jobs cannot be scheduled on that machine. The abundant but unused memory here is referred to as a stranded resource. Resource stranding impacts data center or cloud utilization. For example, shortage of memory due to machines with low memory to CPU ratio can lead to CPU stranding, which is expensive. Although resource stranding can be alleviated by building an

aggregated resource pool such that a job can use resources across machines, this solution is both complex and hardware dependent.

DESCRIPTION

This disclosure address the resource stranding problem through matching and optimization of resources. Shape property of a job or a machine is defined as a ratio of resources, e.g., memory to CPU resources. In case of more than two resources, shape can be defined as a vector instead of a ratio. Shape matching optimization preferentially matches a job's shape to the machine's shape such that both CPU and memory are allocated and consumed at roughly the same rate. By implementing shape-matching optimization, resource stranding on a particular dimension (CPU or memory) is effectively reduced and resource utilization is improved. Otherwise, when memory is in low supply, expensive CPU resources are stranded and wasted. Alternatively, memory is sub-optimally used (stranded) if CPU resources are unavailable.

In a cluster, there are both machines and jobs with heterogeneous shapes. More specifically, machines have different memory to CPU ratios. This can lead to the resource stranding problem. Suppose machine A has 128 GiB memory and 16 CPUs, and jobs with memory to CPU ratio of 10 (for every CPU, 10 GiB of memory is required) are to be scheduled on machine A. At most 12.8 CPUs can be optimally utilized for such jobs on machine A, thereby consuming 128 GiB and 12.8 CPUs. After scheduling such jobs, that consume all of the available 128GiB memory, machine A still has 3.2 (16 - 12.8) CPUs available. These resources that end up going unused is a considerable loss. However, if the memory to CPU ratio of the jobs is 8 (128/16) instead, both memory and CPU would be fully used. Shape matching strives to match jobs to machines with similar shape.

Virtual machines (e.g., in a cloud datacenter) can be configured with varying memory to CPU ratios. Some VMs are CPU-intensive while other VMs are memory-intensive. Shape-matching optimization helps match virtual machines to physical machines of similar memory to CPU ratios, thereby improving and optimizing cloud utilization. Such optimization is completely transparent to virtual machines.
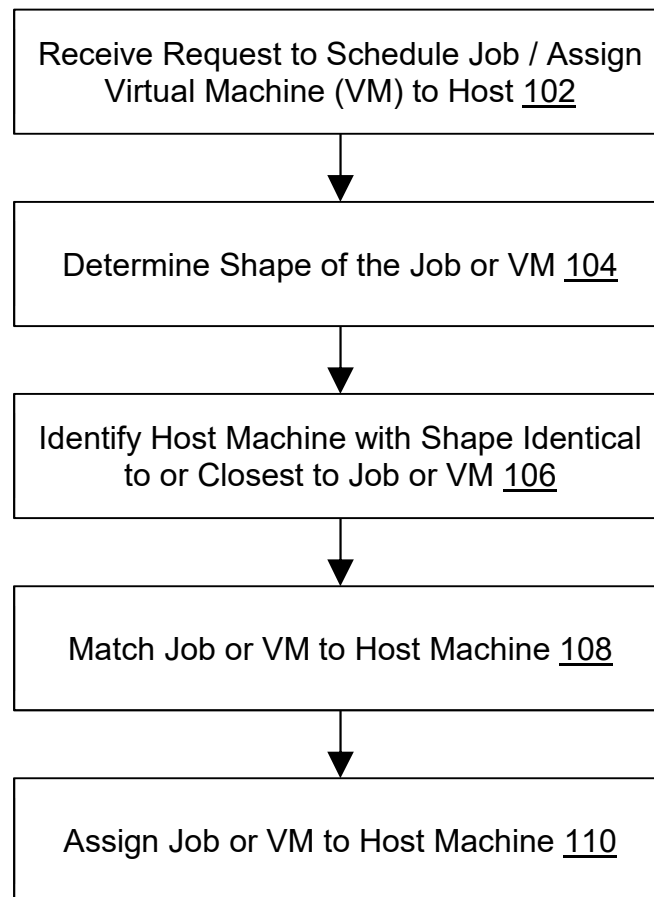
```
┌─────────────────────────────────────────┐
│  Receive Request to Schedule Job / Assign │
│    Virtual Machine (VM) to Host 102       │
└─────────────────────────────────────────┘
                    │
                    ▼
┌─────────────────────────────────────────┐
│   Determine Shape of the Job or VM 104    │
└─────────────────────────────────────────┘
                    │
                    ▼
┌─────────────────────────────────────────┐
│  Identify Host Machine with Shape Identical │
│    to or Closest to Job or VM 106         │
└─────────────────────────────────────────┘
                    │
                    ▼
┌─────────────────────────────────────────┐
│    Match Job or VM to Host Machine 108    │
└─────────────────────────────────────────┘
                    │
                    ▼
┌─────────────────────────────────────────┐
│    Assign Job or VM to Host Machine 110   │
└─────────────────────────────────────────┘
```

Fig. 1

Fig. 1 displays an example process to schedule a job or assign a virtual machine ("VM") to a physical host machine. The shape matching optimization scheduler receives a request to schedule a job or assign a virtual machine to a physical host machine (102). Assuming all machines can accommodate the job or VM request, the scheduler then determines the shape of

the job or the VM (104). After determining the shape of the job, the scheduler searches available

host machines to identify the one with shape identical to or closest to the job or the VM (106)

and matches the job or the VM to the host machine with the closest shape (108). The scheduler

then assigns the job or the VM to the host machine (110). The shape matching scheduler

optimizes the host machine resources by matching the host machine shapes with the shapes of

the jobs or shapes of the VMs assigned to them. A cloud data center that implements a shape

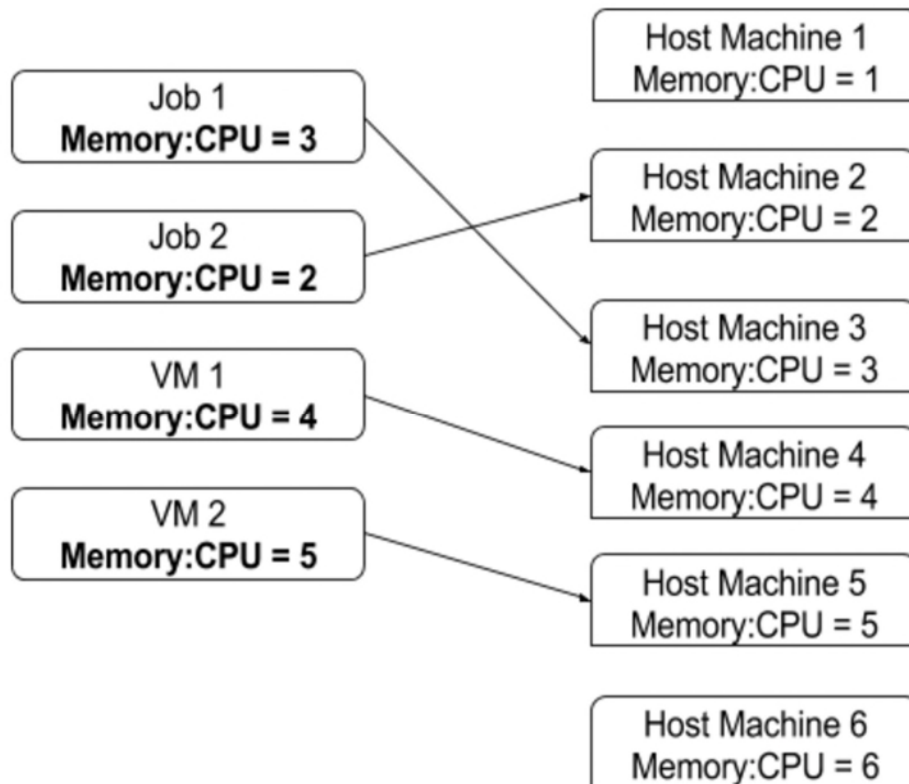matching scheduler can improve overall utilization of the cloud.

Fig. 2

Fig. 2 illustrates a sample matching of 2 jobs and 2 virtual machines, each with a

different memory to CPU ratio with a host machine of the same shape (same memory to CPU

ratio) from host machines in a cluster. Fig. 2 illustrates that through use of the described solution, no resource is stranded within host machines 2 to 5 and that the CPU as well as memory resources of the host machines can be fully utilized in optimal match. While Fig. 2 shows an exact match between the memory and CPU ratios of the jobs/ VMs and the host machines, the shape-matching scheduler can perform inexact matches to available host machines that bear a shape close to the job to be scheduled. By implementing a shape-matching scheduler, resource stranding is reduced, thereby improving utilization of computing resources.

CONCLUSION

This disclosure describes techniques to preferentially schedule jobs on machines where the resource shape of a job matches the available resource shape of the machine. Implementing these techniques reduces resource stranding and improves utilization of machines in a cluster or cloud.