#### University of Nebraska - Lincoln DigitalCommons@University of Nebraska - Lincoln

Theses, Dissertations, & Student Research in Computer Electronics & Engineering

Electrical & Computer Engineering, Department of

Summer 8-2-2013

# Adaptive Co-Scheduler for Highly Dynamic Resources

Kartik Vedalaveni University of Nebraska-Lincoln, kartik@cse.unl.edu

Follow this and additional works at: http://digitalcommons.unl.edu/ceendiss Part of the <u>Computational Engineering Commons</u>, and the <u>Other Computer Engineering</u> <u>Commons</u>

Vedalaveni, Kartik, "Adaptive Co-Scheduler for Highly Dynamic Resources" (2013). *Theses, Dissertations, & Student Research in Computer Electronics & Engineering*. Paper 24. http://digitalcommons.unl.edu/ceendiss/24

This Article is brought to you for free and open access by the Electrical & Computer Engineering, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Theses, Dissertations, & Student Research in Computer Electronics & Engineering by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

#### ADAPTIVE CO-SCHEDULER FOR HIGHLY DYNAMIC RESOURCES

by

Kartik Vedalaveni

#### A THESIS

Presented to the Faculty of

The Graduate College at the University Of Nebraska

In Partial Fulfilment of Requirements

For the Degree of Master Of Science

Major: Computer Science

Under the Supervision of Professor David Swanson

Lincoln, Nebraska

August, 2013

#### ADAPTIVE CO-SCHEDULER FOR HIGHLY DYNAMIC RESOURCES

Kartik Vedalaveni, M.S

University Of Nebraska, 2013

Adviser: David Swanson

There are many kinds of scientific applications that run on high throughput computational (HTC) grids. HTC may utilize clusters opportunistically, only running on a given cluster when it is otherwise idle. These widely dispersed HTC clusters are heterogeneous in terms of capability and availability, but can provide significant computing power in aggregate. The scientific algorithms run on them also vary greatly. Some scientific algorithms might use high rates of disk I/O and some might need large amounts of RAM. Most schedulers only consider cpu availability, but unchecked demand on these associated resources that aren't managed by resource managers may give rise to several issues on the cluster.

On the grid there could be different schedulers on different sites and we cannot rely upon features of one kind of scheduler to characterize the nature and features of every job. Most state of the art schedulers do not take into account resources like RAM, Disk I/O or Network I/O. This is as true for the local schedulers as much it is for the grid. Often there is a need to extend these schedulers to solve situations arising from new and/or complex use cases either by writing a plugin for existing schedulers or by CoScheduling. A key issue is when resources like RAM, Disk I/O or Network I/O are used in an unchecked manner and performance degrades as a result of it. Further scheduling jobs that claim the degraded resources could overwhelm the resource to an extent that the resource will finally stop responding or theWe solve system will crash.

We schedule based on minimum turnaround time of the sites which will help increase

the throughput of the overall workload for the Co-Scheduler, which also is a good loadbalancer in itself. The Co-Scheduler is driven by the fact that turnaround time increases when concurrent jobs accessing these resources reach a threshold value which in-turn causes degradation and this is the basis for this work.

With an increase in the number of entities concurrently using the resource, there is a need to monitor and schedule concurrent and unmanaged access to any given resource to prevent degradation. These issues that we encounter in real life at the Holland Computing Center are the basis and motivation for tackling this problem and for developing an adaptive approach for scheduling. This co-scheduler must be aware of multi-resource degradation, balance load across multiple sites and run clusters at high efficiency and share resources fluidly. An initial implementation tested at HCC will be evaluated and presented.

#### ACKNOWLEDGMENTS

First and foremost, I would like to thank my advisor Dr. David Swanson. Throughout my research and writing he's been a constant source of support for me. His guidance and teachings have helped not only to shape my thesis but also in all walks of my life. I thank David for identifying this interesting research problem and bringing it to me. His hard work has been a constant source of inspiration. I would like to thank Derek Weitzel for working with me to solve the issues with grid submissions and for helping me with his technical expertise.

I'd like to take this opportunity to thank the entire team at Holland Computing Center for providing me with the required resources, support and continuous advise even after I broke their systems. It has been a wonderful experience working with these folks.

# Contents

Co	onten	ts		v		
Li	List of Figures vi					
1	Intr	oductio	n	1		
2 Background			d	4		
	2.1	High-	Throughput Computing	• 4		
	2.2	HTCo	ndor	• 5		
	2.3	Grid G	Computing	. 6		
	2.4	Globu	s	· 7		
	2.5	Open	Science Grid	• 9		
3	Rela	elated Work				
	3.1	Comp	arison of existing mechanisms	. 11		
		3.1.1	HTCondor Flocking	. 11		
		3.1.2	HTCondor Job Router	. 13		
		3.1.3	OSG Match Maker	. 15		
4	Ada	ptive C	Co-Scheduler For Highly Dynamic Resources	16		
	4.1	Introd	uction	. 16		

	4.2	Degradation detection and Capacity based scheduling	18				
	4.3	Multi-site load distribution based scheduling	18				
	4.4	Implementation of Capacity based scheduling	21				
	4.5	Implementation of Multi-site load distribution based scheduling	22				
	4.6	Programming APIs	24				
		4.6.1 HTCondor Log Reader and User API	24				
		4.6.2 Synchronization of Co-Scheduler Code	27				
5	Eval	luation	30				
	5.1	Introduction	30				
	5.2	Evaluation of Co-Scheduler on Tusker cluster with Lustre	32				
	5.3	Evaluation of Co-Scheduler on independent NFS server	34				
	5.4	Scalability Analysis of the Co-Scheduler	38				
6	Con	clusion	39				
	6.1	Future work	40				
Bi	Bibliography 42						

# List of Figures

3.1	JobRouter: Transformation of Jobs	13
3.2	Configuration to implement Job Routing	14
4.1	Degradation occurrence: Co-Scheduling	17
4.2	Multi-site Job distribution function	19
4.3	Multi-site scheduling algorithm overview, classification of sites into slower	
	and faster	24
4.4	Mutex on Number Of Jobs, sumOfK variable	27
4.5	Synchronization of threads after reading turnaround time	28
5.1	The program that generates I/O load for a given period of time	32
5.2	Histogram showing turnaround time distribution of 375 Jobs, when run on	
	Tusker cluster via bulk <i>Condor_G</i> submission	33
5.3	Histogram showing turnaround time distribution of 375 Jobs, when run on	
	Tusker via Co-Scheduler	33
5.4	Goodput graph showing the turnaround time distribution of all the jobs	
	submitted on Tusker via $Condor_G$ bulk submission	35
5.5	Goodput graph showing the turnaround time distribution of 375 jobs submit-	
	ted on Tusker with the Co-Scheduler	35

5.6	Histogram showing turnaround time distribution of 300 Jobs, when run via	
	<i>Condor_G</i> bulk submission on a custom small scale independent NFS server	
	on Tusker	36
5.7	Histogram showing turnaround time distribution of 300 Jobs, when run using	
	a Co-Scheduler on a custom small scale independent NFS server on Tusker	36
5.8	Goodput graph showing the turnaround time distribution of 300 Jobs, when	
	run via <i>Condor_G</i> bulk submission on a custom small scale independent NFS	
	server on Tusker	37
5.9	Goodput graph showing the turnaround time distribution of 300 jobs submit-	
	ted, when run using a Co-Scheduler on a custom small scale independent NFS	
	server on Tusker	37

## Chapter 1

## Introduction

Ian Foster states in his paper The Anatomy Of the Grid [5] that *Grid computing is about controlled sharing of resources with resource owners enforcing policies on the owned resources.* The resources come in the form of hardware and software that allows us to submit jobs, run the jobs and monitor the jobs on the grid. Universities usually have multiple clusters across their campus and these are usually owned by different departments but stand united under the banner of the university. Campus grids are mini grids where jobs are spanned across multiple clusters based on the need of the user and available resources. There is a provision for jobs to overflow to the national grid infrastructure[17].

Modern schedulers used in clusters provide numerable features for policy making, resource management and scheduling. The problem of cluster performance degradation that occurs when one of the resources is overwhelmed is a problem that hasn't been addressed. This problem occurs when many jobs are scheduled on a single system. Some of the schedulers like maui [7] are sophisticated enough to take into account contention of other resources like RAM but ultimately convert the 2D vector values of CPU and RAM into a single scalar value. In practice, this hard-codes the value or presents these resources in a fixed ratio which makes us question the effectiveness of such scheduling

mechanisms.

At Holland Computing Center of the University of Nebraska-Lincoln, we're tackling this issue of cluster degradation caused by over exploitation of one or more resources. The proposed solution adaptively schedules such highly dynamic resources on the grid and across multiple sites by adaptively scaling with respect to the performance of a given cluster. We schedule based on minimum turnaround time of the sites which will help increase the throughput of the overall workload for the Co-Scheduler, which also is a good load-balancer in itself.

Existing schedulers depend on the availability of resources and frequent polling of it via resource manager to determine the slots for scheduling. It should be noted that state of the art schedulers like maui/torque, slurm, HTCondor take into account only CPU as a resource and resources like RAM, Disk I/O or Network I/O are either ignored or their resource equivalent is converted into a scalar value which isn't an effective way of tackling the multiple resource scheduling problem. This might result in scheduling excessive jobs on a single machine. In some schedulers, for example, it will be pre-assigned on every node that 1 CPU-CORE will have 2 GB of memory. This kind of pre-assignment seems like it is addressing the resource degradation of RAM in addition to CPU but suffers from the above mentioned problems. We take a turnaround time approach to measure degradation. We define a resource to be degraded if and when we submit jobs that result in increasing the turnaround time by 25%. It must be noted that there isn't explicit measurement of status of individual resources like RAM, Disk I/O, Network I/O, since there are no guarantees the resource managers keep track of these resource allocations. The Co-Scheduler is driven by the fact that turnaround time increases when concurrent jobs accessing these resources reach a threshold value which in-turn causes degradation and this is the basis for this work.

One aspect to Co-Scheduler is degradation detection and management by adapting

to the throttling of a resource. Another aspect is to efficiently distribute jobs across multiple sites which increases the throughput of the workload. The Co-Scheduler adapts to degradation by backing off submission rates and waiting for a period of time (till the detection of a target capacity) for the next incremental submission. It efficiently distributes the jobs and balances load across multiple sites, thus submitting more jobs to a site with less turnaround time and also ensuring to submit lesser jobs to the sites with larger turnaround time. This efficiently balances load across multiple sites on a grid and improves throughput because the loads may vary dynamically. The capacity of jobs the cluster can efficiently run without degradation needs to be detected and updated again and again after a given period of time.

Finally, since the grid environment is a heterogeneous environment, the scheduler implementation must use APIs that are available on all the systems across the grid. To implement such a Co-Scheduler we limit ourselves to HTCondor based clusters and utilize libcondorapi that needs to be present on the submission site and we can still submit jobs to sites without HTCondor scheduler via *Condor\_G* interface. The result is a threaded Co-Scheduler that submits to multiple sites concurrently, is aware of multiple-resource degradation, and balances load efficiently across multiple sites of the grid.

Please note that all the references to the Co-Scheduler in this thesis refers to the adaptive Co-Scheduler designed by us.

## Chapter 2

# Background

## 2.1 High-Throughput Computing

High Throughput Computing (HTC) is defined as a computing environment that delivers large amounts of computational power over a long period of time. The important factor being over a long period of time which differentiates HTC from HPC which focuses on getting a large amount of work done in a small amount of time. The workloads that run on HTCondor systems don't have an objective of how fast the job can be completed but how many times can the job be run in the next few months. In another definition of HTC, European Grid Infrastructure defines HTC as a computing paradigm that focuses on the efficient execution of a large number of loosely coupled tasks [15]. One such distributed computing software providing High-Throughput Computing services is HTCondor developed by the team at University of Wisconsin-Madison[15].

### 2.2 HTCondor

HTCondor is a distributed system developed by HTCondor team at the University of Wisconsin-Madison. It provides an HTC environment to sites that foster research computing and enables sites to share computing resources on otherwise idle computers at a given site. HTCondor includes a batch queuing system, scheduling policy, priority scheme, and resource classifications for a pool of computers, which is mainly used for compute-intensive jobs. HTCondor runs on both UNIX and Windows based workstations that are all connected by a network, although there are other batch schedulers out there for dedicated machines. The power of HTCondor comes from the fact that the amount of compute power represented by the sum total of all the non-dedicated desktop workstations sitting on people's desks is sometimes far greater than the compute power of a dedicated central resource. There are many unique tools and capabilities in HTCondor which make utilizing resources from non-dedicated systems effective. These capabilities include process checkpoint and migration, remote system calls and ClassAds. HTCondor also includes a powerful resource manager with an efficient match-making mechanism that is implemented via ClassAds, which makes HTCondor understandable when compared with other compute schedulers[15]. With the numerous features of HTCondor, it is used for grid computing on the scale of a large number of massive computers that are loosely coupled on the open science grid. We architected a bioinformatics program iSG, indel sequence generator, which simulates the evolutionary events of highly diverged DNA and protein sequences. iSG is a high memory footprint program. We architected it so as to make it run on the grid and use less amount of peak memory [14]. We retained the order of execution of individual jobs by making extensive use of HTCondor's DAGMAN. DAGMAN allows one to manage a progressing of ordered execution of jobs, the jobs are in the form of a directed acyclic graph.

## 2.3 Grid Computing

As the cost of computers and the network connecting them is decreasing, there is movement towards a paradigm shift in computing with the clustering of geographically distributed resources. The goal of this is to provide a service oriented infrastructure. Organizations like the Grid community and Global Grid Forum are constantly investing effort in making Grid a platform with standard protocols [5], to provide seamless and secure discovery and access to infrastructure and interactions among the resources and services.

With the advent of parallel programming and distributed systems it became obvious that loosely coupled computers can be used for computing purposes and this network of workstations gave rise to the notion of distributed computing. The Globus project began in 1996 and Argonne National Laboratory was responsible for a process and middleware communication system called Nexus which provides remote service requests across heterogeneous machines. The goal of Globus was to build a global Nexus that would provide support for resource discovery, data access, authentication and authorization.

At this point "Grid" replaced the use of "meta-computer" and *researchers from a collaboration of universities termed Grid as integrated resources with integration of many computational visualization and information resources into a coherent infrastructure.* 

The following are goals and visions of Grid Computing:

#### Seamless Aggregation of Resources and Services Aggregation involves three aspects,

- 1. Aggregation of geographically distributed resources
- 2. Aggregation of capacity
- 3. Aggregation of capability

The key factors to enable such aggregation includes protocols and mechanisms to secure discovery, access to and aggregation of resources for the realization of virtual organization and applications that can exploit such an environment.

- **Ubiquitous Service-Oriented Architecture** This is the ability of the grid environment to do secure and scalable resource and data discovery along with scheduling and management based on a wide variety of application domains and styles of computing.
- **Autonomic Behaviors** The dynamism, heterogeneity and complexity of the grid has made many researchers rethink their systems. This new trend aims at system configuration and maintenance of the grid with the least human effort and has led to numerous projects like autonomic grids, cognitive grids and semantic grids.

A key computational grid in the USA is the Open Science Grid, described further in section 2.5. At the heart of the grid computing tools provided by the open science grid lies the globus toolkit, that provides features and interfaces that enable grid computing.

## 2.4 Globus

The Globus project[4] was intended to accelerate the then meta-computing that was built on distributed and parallel software technologies. The term meta-computing was used before the term grid was actually coined which refers to a networked virtual supercomputer, constructed dynamically from geographically distributed resources linked by high speed networks [4]. Scheduling and managing a large group of heterogeneous resources is a challenging and daunting task on the grid. The Globus toolkit provides a framework for managing and scheduling of grid resources across heterogeneous environments. The Globus toolkit consists of a set of modules, each module provides an interface for the provision of an implementation of low level functionalities of the toolkit:

- Resource location and allocation
- Communications
- Unified resource information service
- Authentication interface
- Process creation
- Data Access

There are multiple Virtual Organizations across Open Science Grid. One such VO is Nebraska's HCC VO at Holland Computing Center. A VO provides researchers with computing resources and enables sharing of resources across other VO's. It additionally signs the user's certificate so that a user is identified with a given VO.

We delegate our identification on the grid. *voms-proxy-init* is used to generate a voms proxy at the submit host, we can specify voms as our virtual organization, hcc:/hcc in this case, and hours as the number of hours the proxy would be active/valid. It must be ensured that the proxy period is approximately greater than the length of period of runtime of jobs for successful running of all the jobs. A proxy provides a secure way to access grid resources, by creating a proxy and delegating our identity to it for a short period of time. It enables us to sandbox the theft of the proxy credentials and limit the damage to the short lived proxy and keep our original set of credentials safe.

HTCondor-G extends HTCondor to the grid environment. Jobs are sent across the grid universe using Globus software. The Globus toolkit provides support for building grid systems but submitting, managing and executing jobs have the same capabilities in

both the HTCondor and the HTCondor-G world. Globus provides fault tolerant features to HTCondor-G jobs. GRAM is the Grid Resource Allocation and Management protocol, which supports remote submission of computational request.

gt2 is an initial GRAM protocol which is used in Globus Toolkit version 1 and 2. gt2 is also referred to as the pre-web services GRAM or GRAM2. Similarly, gt5 is the latest GRAM protocol, which is an extension of GRAM2 and is intended to be more scalable and robust, referred to as GRAM5.

#### 2.5 Open Science Grid

The Open Science Grid(OSG) [11], provides service and support for resource providers and scientific institutions using a distributed fabric of high throughout computational services. OSG was created to facilitate data analysis from the Large Hadron Collider[11]. OSG provides resources and directions to Virtual Organizations(VO's) for the purposes of LHC experiments and HTC in general.OSG doesn't own resources but provides software and services to users and enables opportunistic usage and sharing of resources among resource providers. The main goal of OSG is to advance science through open distributed computing. The OSG provides a multi-disciplinary partnership to federate local, regional, community and national cyber-infrastructures to meet the needs of research and academic communities at all scales.

Building an OSG site requires analysis of the requirements and careful planning. The major components of an OSG site include a Storage Element and a Compute Element. Storage elements (SE) manage physical systems, disk caches and hierarchical mass storage systems. A SE is an interface for grid jobs to the Storage Resource Management protocol (SRM) and the Globus GridFTP protocol and others. A storage element requires an underlying storage system like hadoop, xrootd and a GridFTP server and an SRM

interface.

A Compute Element(CE) allows grid users to run jobs on your site. It provides a large number of services when run on the gatekeeper. The basic components include the GRAM and GridFTP on the same CE host to successfully enable file transfer mechanisms of HTCondor-G.

## Chapter 3

# **Related Work**

### 3.1 Comparison of existing mechanisms

There are situations where we can have multiple HTCondor pools exist and some of the pools have many idle slots that are available for utilization. To efficiently utilize the resources across pools condor provides mechanisms like HTCondor flocking and HT-Condor job router. These mechanisms provide similar functionality to the Co-Scheduler designed here. In the following sections, we compare and contrast these mechanisms with the design and functionality of Co-Scheduler in grid environments.

#### 3.1.1 HTCondor Flocking

Flocking refers to a mechanism where a job that cannot run on its own HTCondor pool due to a lack of resources runs in another HTCondor pool where resources are available. HTCondor flocking enables load sharing between pools of computers. As pointed out by Campus Grids Thesis[17] flocking helps balance large workflows across different pools because of the scavenging and greedy nature of the HTCondor scheduler.

To accomplish flocking HTCondor uses multiple components. condor\_schedd adver-

tises that it has idle jobs to the remote condor\_collector. During the next phase of negotiation, if it's found that there are computers available then they are allotted to the jobs in the matchmaking phase and the jobs are then run on the remote pools. It so appears, and the local job queue's maintained as if the jobs are running locally.

Although HTCondor flocking has workflow balancing features across multiple HTCondor pools, it isn't aware of the slower and faster sites/pools. The adaptive Co-Scheduler keeps tabs on turnaround time and is aware of which site is faster or slower. If we look at the idea that scavenging idle resources increases throughput, it does. However, if we look at the case where jobs are greater than the available slots across all the pools, HTCondor flocking stops working here as deeper understanding of sites would be required here to push more jobs at faster sites and importantly push less jobs at slower sites. Flocking's role is to scavenge idle computing slots across multiple pools.

Another aspect where HTCondor flocking isn't designed to perform well is when jobs are contending for the same resource (CPU, RAM, Disk I/O & Network I/O). Even though there will be idle slots on remote pools, HTCondor flocking might not be able to use those slots as they'd be degraded because of the contention. In this case again Co-Scheduler comes in handy. Based on the turnaround time, Co-Scheduler relaxes the load on the resource by submitting fewer jobs this helps reduce degradation and doesn't put excessive job load on the degraded resource. It diverts the jobs elsewhere to perhaps another site, which in-turn increases the throughput.

To conclude, we can say that HTCondor flocking provides features of balancing large workflows and doesn't include features that would detect degradation in the cluster. It also doesn't keep track of information of site performance, which might be exploited to increase the overall throughput of the system.

#### 3.1.2 HTCondor Job Router

The HTCondor manual defines the functions of job router to be the following:

The HTCondor Job Router is an add-on to the condor\_schedd that transforms jobs from one type into another according to a configurable policy[15]. This process of transforming the jobs is called job routing.

HTCondor Job Router can transform vanilla universe jobs to grid universe jobs and as it submits to multiple sites, the rate at which it starts submitting equals the rate at which the sites execute them. This provides a platform to balance large workflows across multiple grid sites and replenish the jobs at faster sites once they get done. Job router sends more jobs to a site if the jobs submitted are not idle and stops submitting jobs if the submitted jobs sit idle on the remote cluster. Job router is not aware about which site is faster of slower.



Figure 3.1: JobRouter: Transformation of Jobs

```
# Now we define each of the routes to send jobs on
1
  JOB_ROUTER_ENTRIES = \setminus
2
      [ GridResource = "gt5 ff-grid.unl.edu/jobmanager-pbs"; \
3
        name = "Firefly"; \
4
     ]
5
      [ GridResource = "gt5 tusker-gw1.unl.edu/jobmanager-pbs"; \
6
        name = "Tusker"; \
7
     ]
       8
      [ GridResource = "gt5 pf-grid.unl.edu/jobmanager-condor"; \
9
        name = "Prairiefire"; \
10
     ] \
11
```

Figure 3.2: Configuration to implement Job Routing

A job is transformed to the grid universe by making a copy of the original job ClassAd and modifying some attributes of the job. This copy is called the routed copy and this routed copy shows up in the job queue with a new job id[15].

HTCondor job router utilizes a routing table which contains the listing of sites the job must be submitted to and the name of the potential grid resources. Routing is processed via a HTCondor config file which is defined by the new ClassAds.

HTCondor Job Router appears to be a step up from HTCondor Flocking in terms of scavenging resources and sending the extra jobs to another HTCondor pool. HTCondor Job Router also maintains the submission rate of the jobs on submit hosts equal to that of remote clusters. But the job router does not keep track of how fast each HTCondor cluster is. Thus HTCondor job router does not optimize the displacement of jobs from a slower cluster. Since it maintains the rate of submission of jobs, we can be sure that if a job is completed at a faster site, it's immediately replaced by the next one. The same is true for a job at a slower site. This might result in increasing the degradation if there exists some. For example, if the job router kept track of slower sites, it could send less jobs to slower sites and thereby increase the overall throughput of the given workflow. A preliminary examination shows that HTCondor job router does seem to have degradation detection features, since it does not submit to a pool that already has idle jobs. However, close examination reveals that even though it submits to a pool with non-idle jobs, it isn't aware of the fact that these jobs in the queue would have undergone degradation due to contention on some resource and thus isn't a viable mechanism for degradation detection.

#### 3.1.3 OSG Match Maker

OSG Match maker is a tool that was created to be distributed among small to medium sized VOs to be used as a powerful submit interface to OSG. It is designed to retrieve VO specific site information and also to verify and maintain the sites by regularly submitting verification jobs to make sure a site can continue receiving jobs. OSG Match maker is another tool that provides information about the grid sites to the HTCondor scheduler. HTCondor scheduler does all the matchmaking with the ClassAds attributes. OSG Match Maker is no longer supported on OSG.

# Chapter 4

# Adaptive Co-Scheduler For Highly Dynamic Resources

## 4.1 Introduction

Due to the heterogeneous nature of the grid it is challenging to schedule the dynamic resources on the grid. We propose a solution that primarily solves the problem of degradation due to contention among jobs for any given resource like RAM, I/O and the network. The contention among jobs for the same given resource gives rise to degradation. For example, contention on large scale job submissions can crash the I/O resource. Efficient and high throughput distribution of jobs across multiple clusters is another challenging problem that we have addressed in this particular Co-Scheduler Solution.



Figure 4.1: Degradation occurrence: Co-Scheduling

In figure 4.1, with C1 called job propagation factor and k the initial number of jobs submitted, we can see that on each iteration k = C1 \* k number of jobs are added. However, if we have a degraded system, we submit k = k/C1 number of jobs by backing off on the amount of jobs submitted.

### 4.2 Degradation detection and Capacity based scheduling

In the first part of the solution we target the basic problem of degradation. To give an example, suppose we have a file server that can serve 100 MBps of data which is available for concurrent access and a user has submitted 200 jobs each consuming 10 MBps of data I/O. If 200 CPU slots are available, a typical state of the art Scheduler schedules these 200 jobs without actually looking at the I/O load. This results in a degraded system over time if more jobs are Scheduled to utilize this file server. This file server might even crash. We would need a degradation handling mechanism that would adaptively scale with the load and exponentially back off during a high contention period. Thus we need an intervention in the form of a Co-Scheduler that intelligently handles degradation. We start the capacity algorithm by submitting jobs to the cluster and by measuring the turnaround time of each iteration. If we find that the current turnaround time of an iteration is 25% greater than the previous iteration we term it as a degradation. We now try to find the best capacity of jobs between the current and the previous iteration by exponentially backing off between these two iterations. We keep finding the target capacity dynamically for each iteration as the contention may change and target capacity keeps varying. Once the target capacity is found then it is retained for job submissions for a given amount of time limited to a maximum of ten iterations of jobs, thus solving the problem of performance degradation.

#### 4.3 Multi-site load distribution based scheduling

Since we're already tracking site performance, the other problem that we tackled was that of efficient distribution of work load to multiple sites of the grid, referred in section 4.2. Some of these sites could be error prone, some of them could be faster and some

```
C1 -> Job propagation constant
multiSite -> list storing turnaround times of multiple sites.
low -> min(multiSite)
high -> max(multiSite)
avg -> average(multiSite)
```

$$f(C1) = \begin{cases} C1 * 4 & multiSite_i = low \\ C1/2 & multiSite_i = high \\ C1 * 2 & multiSite < avg \end{cases}$$

Figure 4.2: Multi-site Job distribution function

can be slower. We need to take an approach that improves the overall throughput of the workload, thus we keep track of average AgTime which equals (Queue wait time + Run-time) for each batch of jobs submitted to the cluster. We take the average AgTime of the batch of jobs that is submitted across multiple clusters. We group AgTimes of these batches of jobs based on the value of the AgTime to be greater than, less than, or equal to the average AgTime of all batches of jobs. This gives us a classification of sites that run faster and also that are more available. We exploit this information in our scheduler. To implement multi-site job submission and load balancing, we start submitting one job to each cluster and measure the AgTime once the job returns. Now we have the information of AgTime of all the clusters. Continuing to submit jobs concurrently in a similar manner would result in the completion of the workload in a non-efficient way. Out of the list of all the AgTimes of different sites we take the sites that have lower than average AgTime of all batches of jobs and increase the job propagation factor for these jobs. This in-turn increases the number of jobs submitted to that site and helps us efficiently distribute the workload to faster clusters. Submitting to a faster site would enable us to submit more jobs to the particular site. As the AgTime is lower we would finish jobs quickly and make room for more jobs if we submit less to slower sites. Slower sites can still help

us improve the throughput to a greater extent as long as we submit less jobs to them and prevent performance degradation at those sites. The error handling mechanism is gracefully handled by HTCondor and by this approach we have an efficient system with increased throughput and with proper distribution of load across multiple clusters.

The co-scheduler requires the presence of a HTCondor installation and libcondorapi. It is written in the C++ language and has extensively made use of pthreads for synchronization and multithreading. The co-scheduler has two components to it: the first one is a capacity detection algorithm which is found by measuring degradation and the other is a multi-site workload distribution algorithm. The following paragraphs detail the engineering aspects of the design. The program begins by taking the number of jobs, the site's information and submit script information as its input. The number of jobs are the ones submitted across multiple sites. The next input file contains information on the list of sites that can be used for load balancing in the GridResource format of the HTCondor ClassAds API, e.g. tusker-gw1.unl.edu/jobmanager-pbs. It is assumed that all the sites listed in the sites input file are working and do not have any misconfiguration issues. The third and the final option to the scheduler is the submit description file. All the job ClassAd information and requirements can be written in this section and all of these will be applicable on the grid when a particular job is scheduled. There are many different kinds of universe in condor including Vanilla, java, grid. vanilla enables job submissions to the local condor pool, java enables submissions of java programs and grid enables direct submission to the grid. The two main important attributes required in the submit description file are *universe* which needs to be "grid" all the time as we are submitting to grid sites via HTCondor-G, and the other most important thing is the *grid-proxy* which is elaborated in section 2.4. condor\_wait [15] can be used to wait for a certain job or number of jobs to complete, which watches the log file and sees if the completion entry of the job is made in the log file that is generated by the *log* command in the HTCondor

submit description file. There are two options we can specify to condor\_wait. One is -num, number-of-jobs that waits till the number-of-jobs are completed and the other option is to specify wait, seconds that waits for specified amount of time. If nothing is specified condor\_wait waits indefinitely till the job(s) are complete.

## 4.4 Implementation of Capacity based scheduling

This algorithm detects degradation and once we detect degradation we find the target capacity:

```
Algorithm 1 Algorithm for determining target capacity by detecting degradationc2 \leftarrow 1.25 {c2: Degradation factor}while true doif T2 < c2^*T1 thenjobSubmission(k) {Submit k jobs}T1 = (T1 + T2)/2end ifif T2 > c2^*T1 thendegradation_high \leftarrow kdegradation_low \leftarrow k/2optimalCapacity(degradation_high, degradation_low)end ifend while
```

#### Algorithm 2 Algorithm for determining target capacity by detecting degradation

```
mid \leftarrow (high + low)/2

k \leftarrow mid

jobSubmission(k) {Submit k jobs}

if T<sub>2</sub> < c<sub>2</sub> * T<sub>1</sub> then

T1 \leftarrow (T1 + T2)/2

optimalCapacity(mid,high);

end if

if T<sub>2</sub> > c<sub>2</sub> * T<sub>1</sub> then

optimalCapacity(low,mid);

end if

return mid
```

Algorithm 1 sets up the stage for degradation detection. T1 is the turnaround time for previous iteration of job submission and T2 is the turnaround time for current iteration of job submission. C1 is the job propagation factor that is used to vary the amount of jobs submitted and C2 is the degradation factor that is arbitrarily set to 25%. Co-Scheduler submits more jobs if T2 doesn't exceed T1 by the value of C2 \* T1. That is, if the turnaround time hasn't increased by 25% or else the algorithm calls the target capacity algorithm that detects the target capacity.

Algorithm 2 is passed two parameters high and low. "high" represents the upper bound where degradation has taken place and "low" the lower bound where degradation hasn't taken place. We have another variable "mid" which is the average of high and low. The algorithm submits "mid" amount of jobs to a site and if the degradation is detected we find capacity in the lower range of values from low to mid. If the degradation is not detected we find the capacity in the upper range of mid to high. The algorithm works like a modified binary search and takes O(log n) iterations to find the target set of jobs that a site can handle.

# 4.5 Implementation of Multi-site load distribution based scheduling

In figure 4.3, multi-site load distribution based scheduling uses the agtime (queue wait time + run-time) at each site for sending future jobs. Its a threaded algorithm. The following is the extracted function of per thread execution:

In Algorithm 3, each thread is responsible for submitting jobs to each site. Each thread stores the information of average AgTime of the batch of jobs submitted to the given site. Each thread finds the average AgTime from other sites and averages this information

Algorithm 3 Algorithm for distribution of workflow load across multiple sites on the grid

```
c1 \leftarrow 2 {c1: Job propagation factor} {multiSite is a list storing AgTimes of multiple
sites.}
low = min(multiSite)
high = max(multiSite)
average = \sum_i multiSite_i / Sizeof(multiSite_i)
if AgTime_Thread == low then
c1 \leftarrow c1 * 4
end if
if AgTime_Thread == high then
c1 \leftarrow c1/2 > 1?(c1/2:1)
end if
if AgTime_Thread < average then
c1 \leftarrow c1 * 2
end if
```

to find the average distribution of AgTime across the sites. Once we find this average we compare it with the site average to classify different sites into faster and slower sites based on the performance of the sites.



Figure 4.3: Multi-site scheduling algorithm overview, classification of sites into slower and faster

## 4.6 **Programming APIs**

#### 4.6.1 HTCondor Log Reader and User API

HTCondor provides Job Log Reader API [15] that polls logs for job events by giving us API access to the events and outcomes. The following is the constructor for initializing a ReadUserLog object.

Constructor: ReadUserLog reader(fp,false,false);

ULogEventOutcome (defined in condor\_event.h):

Status events for job detection:

ULOG\_OK: Event is valid ULOG\_NO\_EVENT: No event occurred (like EOF) ULOG\_RD\_ERROR: Error reading log file ULOG\_MISSED\_EVENT: Missed event ULOG\_UNK\_ERROR: Unknown Error

All the job log entries are named as events and these events could range from being ULOG\_OK where the event has taken place and is valid to ULOG\_UNK\_ERROR where an error has taken place.

The following pseudo-code is extracted from the logReader module that first detects all the valid events and then based on the data-structure of the event object detects if the event kind is ULOG\_EXECUTE, meaning the job has begun executing, via eventNumber data member. Finally, we detect the ULOG\_JOB\_TERMINATED event where the job has successfully terminated. We also cast the more general event object into JobTerminatedEvent to access data members of the JobTerminatedEvent. The comments in the listed pseudo-code provides more details on this extracted code.

Job Submission Pseudo-Code:

```
void logReader(string hostFile, args *data, int nSites) {
FILE *fp;
ReadUserLog reader(fp,false,false);
ULogEvent *event = NULL;
while(reader.readEvent(event)==ULOG_OK) {
if((*event).eventNumber==ULOG_EXECUTE ) {
```

```
//Cast into Execute Event
7
                    ExecuteEvent *exec
8
                    = static_cast < ExecuteEvent * > (event);
9
                    //condor_wait -num K, where K is the
10
                    //amount of jobs completed till the wait.
11
                         char tmp[100];
12
                         sprintf(tmp,"condor_wait -num
13
                         \%d \%s",count,hostFile.c_str());
14
                }
15
                if((*event).eventNumber==ULOG_JOB_TERMINATED)
                                                                     {
16
                //Cast into Job Terminated Event
17
                    JobTerminatedEvent *term
18
                    = static_cast<JobTerminatedEvent*>(event);
19
20
                    if(term->normal)
                                          {
21
                         //on Normal termination,
22
                         //works only for local jobs,
23
                         //find the CPU time of local jobs
24
                    }
25
                }
26
           }
27
  }
28
```

#### 4.6.2 Synchronization of Co-Scheduler Code

There are critical sections in the code where synchronization becomes absolutely necessary. One shared variable is number-of-jobs. The number of jobs executed across the sites should remain fixed and the value must match the value that has been given as input. In a threaded system where each thread is executing jobs on a different cluster it becomes necessary to define N, the number of jobs executed or executing as a critical section. Here we define a pthread mutex and lock it for all write accesses to N. We serialize the access to N and make conditional checks during job submission, so as not to allow job submissions when N is greater than the input value. By serializing the access across multiple threads the total jobs executed remains equal to the input number of jobs. The following code block demonstrates the use of mutexes for serialization of N among the threads.

```
1 pthread_mutex_t mymutex = PTHREAD_MUTEX_INITIALIZER;
2 pthread_mutex_lock (&mymutex);
3 sumOfK+=k;
4 pthread_mutex_unlock (&mymutex);
```

Figure 4.4: Mutex on Number Of Jobs, sumOfK variable

Another section of code where synchronization becomes important is while distributing the jobs to multiple sites. When measuring which cluster is faster, we need information of turnaround time from all the sites before proceeding. This means all the threads need to be executing the same line of code before proceeding with the further program. Thus, the absolute need for synchronization. To handle this problem we make use of a conditional pthread variable. The following code demonstrates the use of conditional wait and conditional signal variable that clears the block on all the threads waiting based on the given condition.

```
pthread_mutex_t syncMutex = PTHREAD_MUTEX_INITIALIZER;
1
  pthread_cond_t synchronize_cv = PTHREAD_COND_INITIALIZER;
2
3
  pthread_mutex_lock(&syncMutex);
4
  multiSite.push_back(stats[thread_id].T1);
5
  tid_multiSite.insert(std::pair<int,int>
6
   (data->tid, stats[thread_id].T1)
7
  );
8
9
  if ( multiSite.size() < numOfSites )</pre>
10
   {
11
    pthread_cond_wait(&synchronize_cv, &syncMutex);
12
  }
13
  else
14
  {
15
  pthread_cond_broadcast(&synchronize_cv);
16
  }
17
  pthread_mutex_unlock(&syncMutex);
18
```

Figure 4.5: Synchronization of threads after reading turnaround time

We need to have the turnaround time of the first job, which enables us to measure how fast each cluster is. To do so, we need to wait for each job to complete on the first submission thereafter the jobs are submitted asynchronously and the average turnaround time is updated in the list. In this pseudocode the threads beginning first add the turnaround time to the list and conditionally wait if the size of the list is less than the number of sites. The last thread comes in and the same condition is voided and executes a conditional broadcast to unblock all the waiting threads and the scheduling proceeds further to asynchronously schedule further jobs. In this example of synchronization code there is a possibility that all of the threads wait continuously if the availability of the sites is very low. The above code needs availability on clusters to prevent threads waiting for long time. The Co-Scheduler will wait till the timeout period and if the jobs do not execute we remove the jobs on that site and proceed further with Co-Scheduling threeby preventing a deadlock.

# Chapter 5

## **Evaluation**

## 5.1 Introduction

We evaluate the degradation detection algorithm by the distribution of turnaround time. The turnaround time of all the jobs when run using co-scheduler and the turnaround time of all the jobs when run directly on the cluster using the HTCondor-G grid interface were both investigated.

If a NFS server has the capacity to serve 50 clients without degradation it implies the turnaround time of all the jobs on the 50 clients are within the limits that would not be a degraded turnaround time. If the same NFS server is forced to serve 200 clients then we would see a degradation and would see an increase in the turnaround time of the jobs that are now contending for the I/O. This would result in degraded turnaround time and the same jobs would have to take lot more time to complete which is the usual scenario that occurs at Holland Computing Center when large amount of I/O bound jobs are submitted to the cluster and there is no check on the usage and the capacity of the I/O resource.

In this evaluation, we've submitted 375 jobs to the cluster "Tusker" using bulk

HTCondor-G submission and then submitting it through the co-scheduler so that we can keep track of degradation, detect it and prevent it. The submit scripts for the co-scheduler are generated dynamically and the grid universes are populated from the input sites file. The machine used for jobs submission was a HTCondor developmental virtual machine, that had to be configured for the grid submissions before its use for the experiment. 375 Jobs were run on the lustre filesystem and the results for this run are in figure 5.2 is of "Tusker". We were able to see degradation when we submitted jobs in bulk using the HTCondor-G interface but we couldn't be sure if the degradation was the result of just our I/O based test job. Thus, separate individual NFS servers were setup on both the clusters sandhills and tusker and tests were run again. The results of the tests are presented in figure 5.6 and figure 5.7.

We're making use of histograms of binned turnaround time and good put graphs to study the effect of degradation. A histogram that projects degradation has many jobs that have higher turnaround times compared to a non-degraded graph. A histogram that doesn't reflect degradation will have most of its jobs within the permissible limits of turnaround time and won't have large variance in terms of the turnaround time. We also make use of goodput graphs that plot an area graph w.r.t the sorted list of turnaround time. This provides a good way to visualize the turnaround time. Ideally, the area under the curve must be minimum in case of a non degraded graph but in a graph that projects degradation the area under the curve is larger than its non-degraded counterpart.

# 5.2 Evaluation of Co-Scheduler on Tusker cluster with Lustre

Figure 5.2 shows how there exists a large variance in turnaround time by having a larger turnaround time than that present in figure 5.3. The former was the result of bulk submission whereas the later is the result of a co-scheduler run. The associated goodput graphs in figure 5.4 and figure 5.5 show their respective areas with degraded run taking up larger area. When we look at the turnaround time distribution on independent NFS servers we can infer many things, that the degradation is significant in the figure 5.6 and the co-scheduler resolves degradation to a larger extent as noted in figure 5.7.

```
int m,n,i,j;
2
3
   n = atoi(argv[1]);
4
   n = n * 100;
5
   m = n+1;
6
   ofstream ofile(argv[2],ios::app);
7
8
   for(j=0;j<m;j++)</pre>
                            {
9
        for(i=0;i<m/100;i++)</pre>
                                      {
10
   ofile<<i<<j;
11
   }
12
   ofile.flush();
13
   }
14
15
   ofile.close();
16
```

1

Figure 5.1: The program that generates I/O load for a given period of time



Figure 5.2: Histogram showing turnaround time distribution of 375 Jobs, when run on Tusker cluster via bulk  $Condor_{-}G$  submission



Figure 5.3: Histogram showing turnaround time distribution of 375 Jobs, when run on Tusker via Co-Scheduler

# 5.3 Evaluation of Co-Scheduler on independent NFS server

The test I/O job used in this experiment continuously does disk write operations for a given period of time, the usual time taken by the program being 300-500 seconds when run on different resources. It generated approximately 3 GB of data file in that period. Designing our own test I/O program, facilitated the change of the path to the NFS servers where we'd be doing our benchmark to test the servers. The test I/O job takes two parameters, one is the size N and the other is the path where the I/O must take place. There are two clusters under consideration, Sandhills and Tusker. 375 jobs are run on each of the clusters and then a separate set of 375 jobs are scheduled using Co-Scheduler. We measure the extent of degradation on each cluster and quantitatively determine how degradation is handled by the Co-Scheduler.

In figure 5.2 considerable jobs take more than 2500 seconds to complete. When the resource is not degraded the test job takes approximately 300-500 seconds of time. Tusker is currently running Lustre filesystem. If the capacity of the filesystem permitted to serve 375 jobs (or the number of jobs running concurrently) then most jobs would be completed within 500 seconds of time but that isn't the case with Tusker. We can conclude the resultant graph is the cause of degradation of the disk I/O resource, but there needs to be clarification about the source of the degradation. It could either be caused by the test I/O jobs that were submitted or can be caused by a set of jobs that are already running on the cluster. To clarify this issue and get a clear picture of what caused the degradation the result from the figure 5.6 sheds more information.



Figure 5.4: Goodput graph showing the turnaround time distribution of all the jobs submitted on Tusker via  $Condor_G$  bulk submission



Figure 5.5: Goodput graph showing the turnaround time distribution of 375 jobs submitted on Tusker with the Co-Scheduler



Figure 5.6: Histogram showing turnaround time distribution of 300 Jobs, when run via *Condor\_G* bulk submission on a custom small scale independent NFS server on Tusker



Figure 5.7: Histogram showing turnaround time distribution of 300 Jobs, when run using a Co-Scheduler on a custom small scale independent NFS server on Tusker



Figure 5.8: Goodput graph showing the turnaround time distribution of 300 Jobs, when run via *Condor\_G* bulk submission on a custom small scale independent NFS server on Tusker



Figure 5.9: Goodput graph showing the turnaround time distribution of 300 jobs submitted, when run using a Co-Scheduler on a custom small scale independent NFS server on Tusker

In the figures 5.8 and 5.9 we can see most of the jobs take about 500 seconds of time to complete. We've prevented excessive degradation here and the I/O resources serve the concurrent jobs without a degraded turnaround time. This reduces the total time taken for the execution of this batch of jobs making way for the use of cluster for other jobs that might need computing.

A Multi-Site load distribution algorithm distributes the job based on agtime (queue wait time + run-time) and load on the cluster. The following was the distribution of jobs when run on Co-Scheduler with 750 Jobs. The cluster with lower AgTime will get more amount of jobs to run. For Tusker the throughput was 47.2 CPU time hours/elapsed time hours and for Sandhills it was 3.12 CPU time hours/elapsed time hours. Which means there was greater availability on Tusker and more jobs were submitted to tusker than Sandhills.

## 5.4 Scalability Analysis of the Co-Scheduler

The Co-Scheduler is architected in a way so that each thread is responsible for dispatching jobs on a given site. The Co-Scheduler can potentially spawn thousands of threads and serve them efficiently but there are only hundreds of OSG sites to run on. It can certainly scale to all the sites of the OSG.

Co-Scheduler can take a large number of jobs as input, of the order of hundreds of thousands of jobs and dispatch the jobs efficiently with the load-balancing algorithm and ensure these jobs run without incurring degradation.

Co-Scheduler is a robust program written in C++ that schedules based on turnaround time to detect degradation and schedules based on (queue wait time + run-time) for the sake of load balancing.

## Chapter 6

## Conclusion

Co-Scheduling on the Grid is challenging with the presence of dynamically varying resources like Network, RAM, compute resources, all of which are available opportunistically. The current Co-Scheduler detects degradation and taxes degraded resources less in terms of scheduling jobs to the degraded resources. The capacity of the cluster is a varying quantity that determines the amount of concurrent jobs that may run without degradation. Co-Scheduler successfully finds the capacity of the cluster and maintains the capacity up to 10 successive iterations of the scheduler before recalculating the capacity.

Scheduling based on agtime enables us to submit more jobs to the clusters having greater availability. Co-Scheduler improves throughput of the job set which handles the case of degradation by virtue of a multi-site job distribution algorithm. Because Co-Scheduler directly effects the load on the resource, Co-Scheduler ensures the average overall turnaround time per job is lower.

Resource management is completely abstracted. We do not need to consider the allocation and management of disk I/O or RAM or Network I/O. The Co-Scheduler is driven by the fact that whenever a resource is overloaded it results in degradation and subsequently does not respond optimally or crashes. A simple and efficient way

of tracking degradation is implemented which monitors the load of resources on the cluster. As much as this is a positive aspect of the Co-Scheduler, the downside is that there isn't a way to tell which of the resource is the cause of the degradation. Because the Co-Scheduler is used with large workloads is usually known based on the workload the resource that is under stress.

While Co-Scheduling offers multiple enhancements for the users of large workloads there are some downsides of it. Co-Scheduler submits the jobs stepwise, in the sense that it submits a smaller set of jobs and waits for its outcome to decide upon the quantity of jobs on the next iteration. This nature of the Co-Scheduler can accumulate a lot of queue wait time if the availability of the compute resources are lower. In the multi-site load distribution algorithm, if one of the sites has very low availability the Co-Scheduler would accumulate the queue wait time of that particular site before it can complete its execution. To solve this issue we schedule based on AgTime (queue wait time + run-time) so that clusters with availability get more number of jobs than the ones with lower availability. When a batch of jobs is submitted to multiple grid sites, the sites not responding timeout after a defined period of time and clean up the jobs that were submitted so that the Co-Scheduler may further schedule these jobs on faster and more available clusters.

#### 6.1 Future work

Degradation factor, C<sub>2</sub> plays an important role in the detection of degradation and in determining the capacity of the cluster. An important future work is to derive C<sub>2</sub> specifically for a given resource. It would be of importance for scheduling if we had the ability to classify the nature of degradation either by taking input from the user or detect dynamically from a degraded resource at run-time.

To protect the resources at HCC its necessary to enforce Co-Scheduler like scheduling

mechanism. To enforce Co-Scheduler like scheduling on protected resources, further work needs to be done to integrate Co-Scheduler code in the local resource manager. Another important aspect would be to include the features of Co-Scheduler in Job Router. Further study of Co-Scheduler's effect on throughput would provide us directions to tweak multi-site load balancing algorithm.

# Bibliography

- [1] Douglas Thain David Cieslak Nitesh Chawla. Condor log analyzer. http://condorlog.cse.nd.edu, 2009.
- [2] Monika Choudhary and Sateesh Kumar Peddoju. Turnaround time based job scheduling algorithm in dynamic grid computing environment. In *Proceedings of the CUBE International Information Technology Conference*, CUBE '12, pages 490–493, New York, NY, USA, 2012. ACM.
- [3] D.H.J. Epema, M. Livny, R. van Dantzig, X. Evers, and J. Pruyne. A worldwide flock of Condors: Load sharing among workstation clusters. *Future Generation Computer Systems*, 12(1):53–65, 1996.
- [4] I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *International Journal of High Performance Computing Applications*, 11(2):115, 1997. 2.4
- [5] Ian Foster, Carl Kesselman, and Steven Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. *Int. J. High Perform. Comput. Appl.*, 15(3):200–222, August 2001. 1, 2.3
- [6] J. Frey, T. Tannenbaum, M. Livny, I. Foster, and S. Tuecke. Condor-G: A computation management agent for multi-institutional grids. *Cluster Computing*, 5(3):237–246, 2002.

- [7] Cluster Resources Inc. Cluster resources :: Products TORQUE Resource Manager:, Jaunary 2011. http://www.clusterresources.com/pages/products/ torque-resource-manager.php. 1
- [8] Miron Livny and Rajesh Raman. High-throughput resource management. In Ian Foster and Carl Kesselman, editors, *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
- [9] LLNL. Slurm reference.
- [10] Holland Computing Center. Holland computing center, January 2011. http://hcc. unl.edu/.
- [11] R. Pordes, D. Petravick, B. Kramer, D. Olson, M. Livny, A. Roy, P. Avery, K. Blackburn,
   T. Wenaus, et al. The Open Science Grid. In *Journal of Physics: Conference Series*,
   volume 78, page 012057. IOP Publishing, 2007. 2.5
- [12] I. Sfiligoi. glideinWMS-a generic pilot-based workload management system. In *Journal of Physics: Conference Series*, volume 119, page 062044. IOP Publishing, 2008.
- [13] J.V. Sumanth, D.R. Swanson, and Hong Jiang. Adaptive load balancing for longrange md simulations in a distributed environment. In *Parallel Processing*, 2006. ICPP 2006. International Conference on, pages 135–146, 2006.
- [14] Kartik Vedalaveni Cory Stroupe Ashu Guru David Swanson. Architecting indelsequence-generator for htcondor., September 2011. 2.2
- [15] HTCondor Team. Condor manual, section 5.6. 2.1, 2.2, 3.1.2, 3.1.2, 4.3, 4.6.1
- [16] D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: The Condor experience. *Concurrency and Computation: Practice and Experience*, 17(2-4):323– 356, 2005.

- [17] Derek Weitzel. Campus grids: A framework to facilitate resource sharing. Master's thesis, University Of Nebraska-Lincoln, 2012. 1, 3.1.1
- [18] M. Zvada, D. Benjamin, and I. Sfiligoi. CDF GlideinWMS usage in Grid computing of high energy physics. In *Journal of Physics: Conference Series*, volume 219, page 062031. IOP Publishing, 2010.