

# GRAPLEr: A Distributed Collaborative Environment for Lake Ecosystem Modeling that Integrates Overlay Networks, High-throughput Computing, and Web Services

Kensworth Subratie  
University of Florida  
Gainesville, FL, USA  
kcratie@acis.ufl.edu

Saumitra Aditya  
University of Florida  
Gainesville, FL, USA  
saumitraaditya@acis.ufl.edu

Renato Figueiredo  
University of Florida  
Gainesville, FL, USA  
renato@acis.ufl.edu

Cayelan C. Carey  
Virginia Tech  
Blacksburg, VA, USA  
cayelan@vt.edu

Paul Hanson  
University of  
Wisconsin-Madison  
Madison, WI, USA  
pchanson@wisc.edu

## ABSTRACT

The GLEON Research And PRAGMA Lake Expedition – GRAPLE – is a collaborative effort between computer science and lake ecology researchers. It aims to improve our understanding and predictive capacity of the threats to the water quality of our freshwater resources, including climate change. This paper presents GRAPLEr, a distributed computing system used to address the modeling needs of GRAPLE researchers. GRAPLEr integrates and applies overlay virtual network, high-throughput computing, and Web service technologies in a novel way. First, its user-level IP-over-P2P (IPOP) overlay network allows compute and storage resources distributed across independently-administered institutions (including private and public clouds) to be aggregated into a common virtual network, despite the presence of firewalls and network address translators. Second, resources aggregated by the IPOP virtual network run unmodified high-throughput computing middleware (HTCondor) to enable large numbers of model simulations to be executed concurrently across the distributed computing resources. Third, a Web service interface allows end users to submit job requests to the system using client libraries that integrate with the R statistical computing environment. The paper presents the GRAPLEr architecture, describes its implementation and reports on its performance for batches of General Lake Model (GLM) simulations across three cloud infrastructures (University of Florida, CloudLab, and Microsoft Azure).

## Keywords

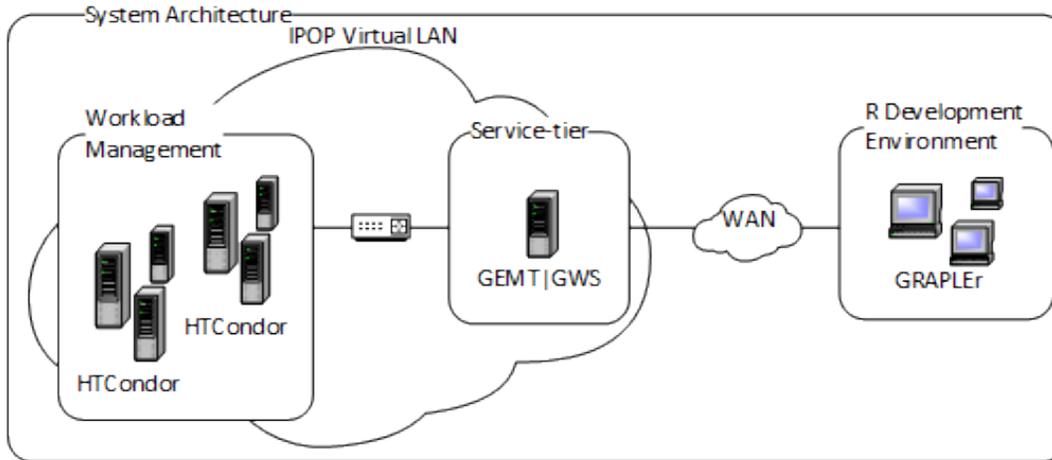
Climate Change, General Lake Model, Lake Modeling, HTCondor, Distributed Computing, IPOP, Overlay Networks

## 1. INTRODUCTION

The GLEON Research And PRAGMA Lake Expedition – GRAPLE – aims to improve our understanding and predictive capacity of water quality threats to our freshwater resources, including climate change. It is predicted that climate change will increase water temperatures in many freshwater ecosystems, potentially increasing toxic phytoplankton blooms [11, 1]. Consequently, understanding how altered climate will affect phytoplankton dynamics is paramount for ensuring the long-term sustainability of our freshwater resources. Underlying these consequences are complex physical-biological interactions, such as phytoplankton community structure and biomass responses to short-term weather patterns, multi-year climate cycles, and long-term climate trends [5]. New data from high-frequency sensor networks (e.g., GLEON) provide easily measured indicators of phytoplankton communities, such as in-situ pigment fluorescence, and show promise for improving predictions of ecosystem-scale wax and wane of phytoplankton blooms [18]. However, translating sensor data to an improved understanding of coupled climate-water quality dynamics requires additional data sources, model development, and synthesis, and it is this type of complex challenge that requires increasing computational capacity for lake modeling.

Searching through the complex response surface associated with multiple environmental starting conditions and phytoplankton traits (model parameters) requires executing and interpreting thousands of simulations, and thus substantial compute resources. Furthermore, the configuration, setup, management and execution of such large batches of simulations is time-consuming, both in terms of computing and human resources.

This puts the computational requirements well beyond the capabilities of any single desktop computer system, and to meet the demands imposed by these simulations it becomes necessary to tap into distributed computing resources. However, distributed computing resources and technologies are typically outside the realm of most freshwater science projects. Designing, assembling, and programming these systems is not trivial, and requires the level of skill typically available to



**Figure 1: System Architecture (GRAPLER).** Users interact with GRAPLER using R environments in their desktop (right). The client connects to a Web service tier (GWS) that exposes an endpoint to the public Internet. Job batches are prepared using GEMT and are scheduled to execute in distributed HTCondor resources across an IPOP virtual private network.

experienced system and software engineers. Consequently, this imposes a barrier to scientists outside information technology and computer science disciplines, and presents challenges to the acceptance of distributed computing as a solution to most lake ecosystem modelers.

GRAPLE is a collaboration between lake ecologists and computer scientists that aims to address this challenge. Through this inter-disciplinary collaboration, we have designed and implemented a distributed system platform that supports compute-intensive model simulations, aggregates resources seamlessly across an overlay network spanning collaborating institutions, and presents intuitive Web service-based interfaces that integrate with existing environments that lake ecologists are used to, such as R.

This paper describes GRAPLER, a cyberinfrastructure that is unique in how it seamlessly integrates a collection of distributed hardware resources through the IP-over-P2P [6, 8] overlay virtual network, supports existing models and the HTCondor distributed computing middleware [17], and exposes a user-friendly interface that integrates with R-based desktop environments through a Web service. As a multi-tiered distributed solution, GRAPLER incorporates several components into an application-specific solution. Some of these components are pre-existing solutions which are deployed and configured for our specific uses, while others are specifically developed to address unique needs.

The rest of this paper is organized as follows: Section 2 describes driving science use cases and motivates the need for the GRAPLER cyberinfrastructure. Section 3 describes the architecture, design, and implementation of GRAPLER. Section 4 describes a deployment of GRAPLER and summarizes results from an experiment that evaluates its capabilities and performance. Section 5 discusses related work, and Section 6 concludes the paper.

## 2. ARCHITECTURE AND DESIGN

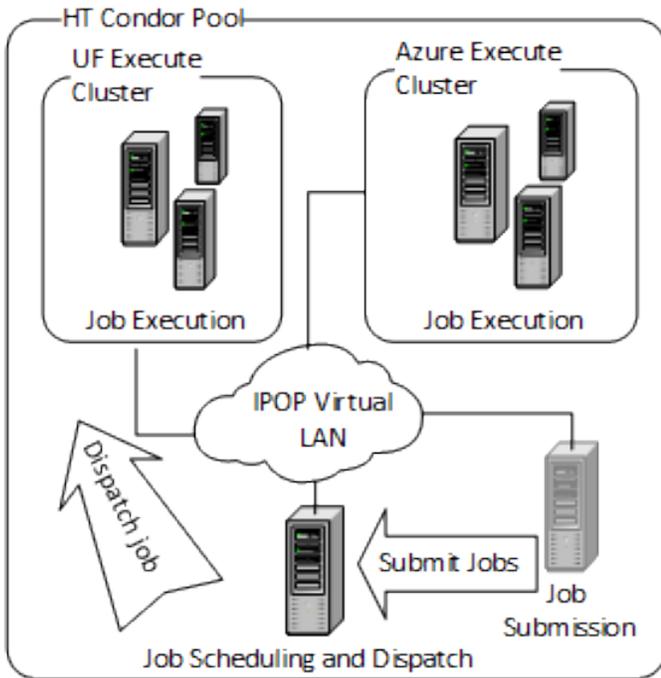
### 2.1 System Architecture (GRAPLER)

The system architecture of GRAPLER is illustrated in Figure 1. Starting from the user-facing components of GRAPLER, users interact with the system through client-side libraries that are called from an R development environment (e.g., R Studio) running on their personal computer. User requests are mapped by the R library to Application Programming Interface (API) calls that are then sent to the GRAPLER Web Service (GWS) tier. The GWS tier is responsible for interpreting the user requests, invoking the GRAPLER Experiment Management Tools (GEMT) to set up a directory structure for model inputs and outputs, and preparing and queuing jobs for submission to the HTCondor pool. The workload management tier is responsible for scheduling and dispatching model simulations across the compute resources, which are interconnected through the IPOP virtual network overlay. These are elaborated below.

### 2.2 Overlay Virtual Network (IPOP)

Rather than investing significant effort in development, porting, and testing new applications and distributed computing middleware, GRAPLER has focused on an approach in which computing environments are virtualized and can be deployed on-demand on cloud resources. While Virtual Machines (VMs) available in cloud infrastructures provide a basis to address the need for a user-provided software environment, another challenge remains: how to inter-connect VMs deployed across multiple institutions (including private and commercial cloud providers) such that HTCondor and the simulation models work seamlessly? The approach to address this problem is to apply virtualization at the network layer.

The IPOP [6] overlay virtual network allows GRAPLER to define and deploy its own virtual private network (VPN) that can span physical and virtual machines distributed across multiple collaborating institutions and commercial clouds. To accomplish this, IPOP captures and injects network traffic via a virtual network interface or “tap” device. The “tap”



**Figure 2: Workload Management (HTCondor).** GRAPLER supports unmodified HTCondor software and configuration to work across multiple sites (e.g., a private cloud at UF and a commercial cloud at MS Azure).

device is configured within an isolated virtual private address subnet space. IPOP then encrypts and tunnels virtual network packets through the public Internet. The “Tin-Can” [8] tunnels used by IPOP to carry network traffic use facilities from WebRTC (Web Real-Time Computing) to create end-to-end links that carry virtual IP traffic instead of audio or video.

To discover and notify peers that are connected to the GRAPLER “group VPN”, IPOP uses the eXtensible Messaging and Presence Protocol (XMPP). XMPP messages carry information used to create private tunnels (the fingerprint of an endpoint’s public key), as well as network endpoint information (IP address:port pairs that the device is reachable). For nodes behind network address translators (NATs), public-facing address:port endpoints can be discovered using the STUN (Session Traversal Utilities for NAT) protocol, and devices behind symmetric NATs can use TURN (Traversal Using Relays around NAT) to communicate through a relay in the public Internet. Put together, these techniques handle firewalls and NATs transparently to users and applications, and allow for simple configuration of VPN groups via an XMPP server.

### 2.3 Workload Management (HTCondor)

A key motivation for the use of virtualization technologies, including IPOP, is the ability to integrate existing, unmodified distributed computing middleware. In particular, GRAPLER integrates HTCondor [17], a specialized workload management system for compute-intensive jobs. Like other full-

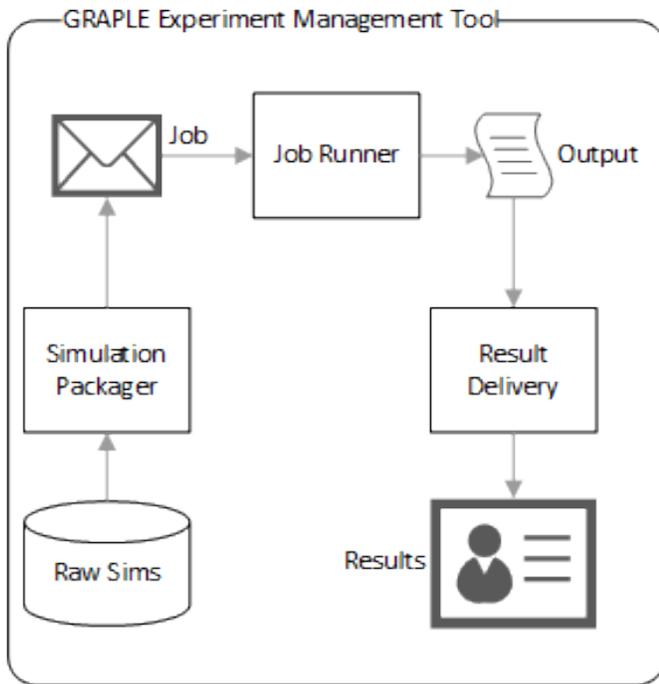
featured batch systems, HTCondor provides a job queuing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. Users submit their serial or parallel jobs to HTCondor, HTCondor places them into a queue, chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion. Figure 2 illustrates the structure of the HTCondor pool that is deployed for GRAPLER.

### 2.4 Experiment Management Tools (GEMT)

An HTCondor [17] resource pool running across distributed resources connected by IPOP provides a general-purpose capability where it is possible to run a variety of applications from different domains. Furthermore, application-tailored middleware can be layered upon this general-purpose environment to enhance the performance and/or streamline the configuration of simulations on behalf of users. GEMT (Figure 3) provides a suite of scripts for designing and automating the tasks associated with running General Lake Model (GLM) based experiments on a very large scale. Here, we use the term “Experiment” to refer to a collection of simulations that address a science use case question, such as determining the effects of climate change on water quality metrics. GEMT provides both the guidelines for the design and layout of individual simulations in the experiment. The primary responsibility of GEMT is to identify and target the task-level parallelism inherent in the experiment by generating proper packaging of executables, inputs, and outputs; furthermore, GEMT seeks to effectively exploit the distributed compute resources across the HTCondor pool by performing operations such as aggregation of multiple simulations into a single HTCondor job, compression of input and output files, and the extraction of selected features from output files.

For the simulations in an experiment, GEMT defines the naming convention used by the files and directories as well as their layout. The user may interact with GEMT in two possible ways: 1) directly, by using a desktop computer configured with the IPOP overlay software and HTCondor job submission software, or 2) indirectly, by issuing requests against the GRAPLER Web service. In the former case, once the user has followed the GEMT specification for creating their experiment, executing it and collecting the results becomes a simple matter of invoking two GEMT scripts. However, the user is left the responsibility of deploying and configuring both IPOP and HTCondor locally. Additionally, the user is now a trusted endpoint on the VPN which carries its own security implications. A breach of the user’s system is a potential vulnerable point to accessing the VPN. The latter case alleviates the user from both these concerns. This paper focuses on this latter approach, where GEMT scripts are invoked indirectly through the Web service.

There are three distinct functional modes for GEMT, which pertain to the different phases of the experiment’s lifetime. Starting with its invocation, on the submit node, GEMT selects a configurable number of simulations to be grouped as a single HTCondor job. The reason why multiple simulations may be grouped into a single HTCondor job is that, for short-running simulations, the costs of job scheduling and transfer of executables can be significant. By grouping



**Figure 3: GRAPLEr Experiment Management Tools (GEMT).** The GEMT Simulation Packager module takes a specification of the raw simulation inputs and groups them together into jobs; these are dispatched for execution through HTCondor, and their execution at the worker nodes is managed by the GEMT Job Runner module. The Result Delivery GEMT module collates results and presents to the user.

simulations into a single HTCondor job, redundant copies of the input can be eliminated to reduce the bandwidth transfer cost and only a single scheduling decision is needed to dispatch all the simulations in the job. The inputs and executables pertaining to a group of simulations are then compressed and submitted as a job to the HTCondor scheduler for execution. When this job becomes scheduled, GEMT is invoked in its second phase, this time on the HTCondor execute node. The execute-side GEMT script coordinates running each simulation within the job, and preparing the output so it can be returned to the originator. Finally, in its third phase, back on the submit node side, GEMT collates the results of all the jobs that were successful and presents them in a standard format to the end user.

GEMT implements user configurable optimizations to fine tune its operations for individual preferences. It can limit how many simulations are placed in a job, and it will compress these files for transfer. GEMT can also overlap the client side job creation with server side execution to minimize the wait time before results start being produced. These features can be set via a configuration file and combine to provide a simplified mechanism to execute large numbers of simulations.

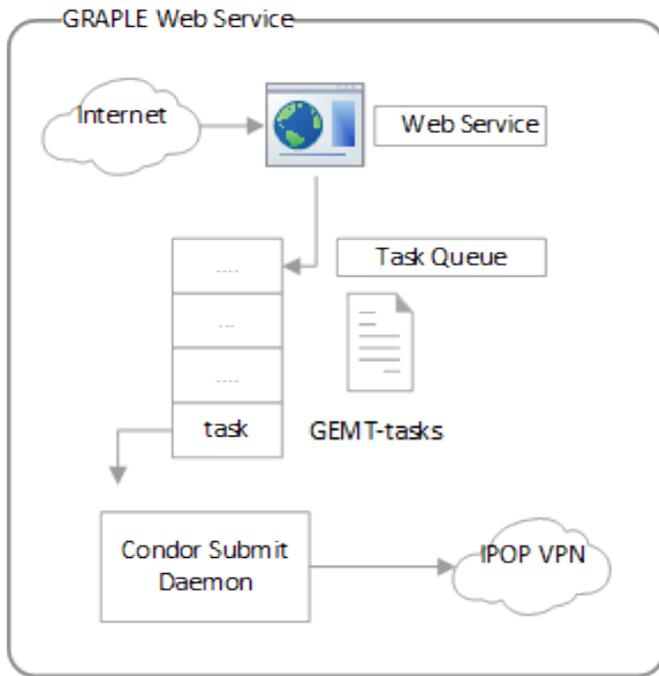
## 2.5 GRAPLEr Web Service (GWS)

The GWS module, as illustrated in Figure 4, is a publicly-addressable Web service available on the Internet, and serves as a gateway for requests to run experiments. GWS acts as a middleware service-tier which exposes an interface to R clients. Requests to run an experiment are made via this interface over the Internet using the HTTP protocol. The functionality provided by GWS is exposed to the R client’s user by means of publicly accessible endpoints, each of which is associated with a corresponding method that is invoked in the background. GWS utilizes the functionality of GEMT for simulation processing. GWS generates simulation input files as needed based on the user’s request (e.g., to vary air temperature according to a statistical distribution for a climate change scenario), configures and queues jobs, and consolidates and prepares results for download. GWS is co-located in the same host as the GEMT client. This host acts as the submit node to the HTCondor pool, where it monitors job submission and execution.

Representational State Transfer or REST, is an architectural style for networked hypermedia applications that is primarily used to build lightweight and scalable Web services. Such a Web service, referred to as RESTful, is stateless with a uniform interface and representation of entities, and communicates with clients via messages and address resources using URIs. GWS implements this paradigm and is designed to treat every job submission independently from any other. Note that there is per-experiment state that is managed by GWS, such as the status of each HTCondor job submitted by the GWS. The state of the experiment is maintained on disk, within the local filesystem, leaving the service itself stateless. GWS implements the public-facing interface using a combination of open-source middleware for Web service processing - Python Flask [7], and an asynchronous task queue - Python Celery [16]. The application is hosted using uWSGi (an application deployment solution) and supplemented by a Nginx reverse proxy server to offload the task of serving static files from the application. The employed technology stack facilitates rapid service development and robust deployment.

The GWS workflow begins when a request is received from an R client through the service interface, which is handled by Flask. The request to evaluate a series of simulations can be provided in one of several ways, as discussed in detail in the section covering the R Language Package. However, only data files are accepted as input - no user provided executable binaries or scripts are executed as part of the experiment. A single client-side request can potentially unfold into large numbers (e.g., thousands) of jobs, and GWS places these requests into a Celery task queue for asynchronous processing. Provisioning a task queue allows GWS to decouple the time-consuming processing of the input and task submission to HTCondor from the response to HTTP request.

A 40-character unique identifier (UID) is randomly-generated for each simulation request received by GWS; it is used as an identifier to reference the state of an experiment, and is thus used for any further interactions with the service for a given experiment. Using the UID returned by GRAPLEr, an R client can not only configure the job, but also monitor its status, download outputs, and abort the job. Once the input file has been uploaded to the service, GWS puts



**Figure 4: GRAPLER Web Service (GWS).** The GWS is responsible for taking Web service requests from users, interpreting them and creating tasks for remote execution using GEMT.

the task into the task queue and responds promptly with the UID. Therefore, the latency that the R developer experiences, from the moment the job is submitted to when the UID is received, is minimized. A GWS worker thread then dequeues GEMT tasks from the task queue, and processes the request according to the parameters defined by the user. Figure 4 shows the internal architecture and setup of GWS. A key feature of the GRAPLER service is to automatically create and configure an experiment by spawning a range of simulation scenarios by varying simulation inputs, based on the user’s request and application-specific knowledge. In particular, the service uses application-specific information to identify data in the input file (such as air temperature, or precipitation), and apply transformations to these data (i.e., adding or subtracting an offset to the base value provided by the user) to generate multiple simulation scenarios. GWS removes the burden from the user to generate, schedule, and collate the outputs of thousands of simulations within their own desktops, and allows them to quickly generate experiment scenarios from a high-level description that simply enumerates which input variables to consider, what function to apply to vary them, and how many simulations to create. The user also has the flexibility to retrieve and download only a selected subset of the results back to their desktops, thereby minimizing local storage requirements and data transfer times.

To illustrate this feature, consider API endpoint 9 in Figure 5. This endpoint exposes a method that enables the user to generate ‘N’ runs from a single baseline set of input files by drawing offsets to input values (e.g., air temperature)

#	Service Endpoint	Required Input	Response Description
1	service_status	N.A.	Online status message
2	GrapleRun	A tarball archive containing the simulation subfolders	Job UID
3	GrapleRunStatus	UID of previously submitted job	Job status message (completed/processing)
4	GrapleRunResults	UID of previously submitted job	URI to job results
5	GrapleRunMetOffset	A tarball archive containing base simulation file	Job UID
6	TriggerSimulation	UID*<Job definition string>	Submission status message (success/error)
7	download_file	UID*<File identifier>	URI to download results
8	return_dataframe	UID*<Define data-frame to be retrieved>	URI to download results
9	GrapleRunMetSample	A tarball archive containing simulation base file and job	Job UID
10	GrapleAbort	UID of previously submitted job	Status message (success/retry)

**Figure 5: GWS Application Programming Interface (API) Endpoints**

from a random distribution. With this API endpoint, the GRAPLER client can upload a single baseline set of input files, along with a short experiment description file. This file specifies which distribution (random, uniform, binomial, or Poisson) to choose samples from, the number of samples, the variable(s) to be modified, and the operation applied against a variable to each randomly-generated value (add, subtract, multiply, or divide). From this single input and description, GWS generates ‘N’ simulation input files, and calls GEMT Simulation Packager scripts to submit jobs to the HTCondor pool.

## 2.6 GRAPLER R Language Package

The user-facing component of GRAPLER is an R package that serves as a thin layer of software between the Web service and the R client development environment (IDE). It exposes an R language application programming interface which can be programmatically consumed by client programs wanting to utilize the GRAPLER functionality. GRAPLER is available on github and is installed on the client desktop, where it integrates into the R development environment. It acts as a proxy to translate user commands written in R into Web service calls. It also marshals data between the client and Web service as necessary. The following example illustrates a sequence of three R calls to submit an experiment to a GRAPLER service running on endpoint `graplerURL`, from a set of input files placed in sub-directories of a root directory folder on the client-side (`expDir`), check its status, and download results:

```
UID<-GrapleRunExperiment(graplerURL, expDir)
GrapleCheckExperimentCompletion(graplerURL, UID)
GrapleGetExperimentResults(graplerURL, UID)
```

The second example shows how a user can specify a parameter-sweeping simulation with 10,000 simulations which are derived from a baseline set of input files (stored in the `simDir` directory at the client) by modifying the `AirTemp` column time series in the GLM meteorological driver input data file `met_hourly.cvs`, in the range -10 to 30.

```

simDir=C:/Workspace/SimRoot/Sim0
driverFileName=met_hourly.csv
parameterName=AirTemp
startValue=-10
endValue=30
numberOfIncrements=10000
expUID<-GrapleRunExperimentSweep(graplerURL,
  simDir, driverFileName, parameterName,
  startValue, endValue, numberOfIncrements)
GrapleCheckExperimentCompletion(graplerURL, expUID)
GrapleGetExperimentResults(graplerURL, expUID)

```

To prevent the use of the Web service interface to execute arbitrary code, custom code – whether binary executables or R scripts – cannot be sent as part of the simulation requests; instead, users only provide input files and parameters for the GLM simulations. The scenarios that can be run are currently restricted to using GLM tools and our own scripts.

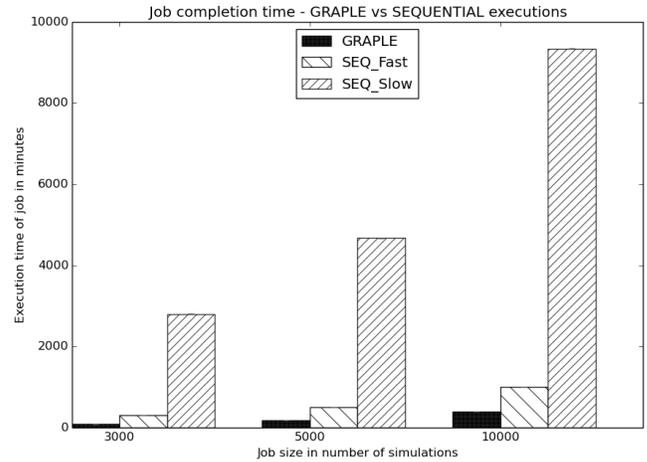
### 3. EVALUATION

In this section, we present a quantitative evaluation of a proof-of-concept deployment of GRAPLER. The goal of this evaluation is to demonstrate the functionality and capabilities of the framework by deploying a large number of simulations to an HTCondor pool. The HTCondor pool is distributed across multiple clouds and connected by the IPOP virtual network overlay. Rather than focusing solely on the reduction in execution times, we evaluate a setup that is representative of an actual deployment composed of execute nodes with varying capabilities.

A GLM simulation is specified by a set of input files, which describe model parameters and time-series data that drive inputs to the simulation, such as air temperature over time, derived from sensor data. The resulting output at the completion of a model run is a netCDF file containing time series of the simulated lake, with many lake variables, such as water temperatures at different lake depths. In our experiments, we use the 1-D GLM Aquatic Eco-Dynamics (AED) model. For a single example GLM-AED simulation of a moderately deep lake run for eight months at an hourly time step, the input folder size was approximately 3 MB, whereas the size of the resulting netCDF file after successful completion of the simulation was 90MB. The test experiment was designed to run reasonably quickly. However, we note that simulations run over decades and with output recorded more frequently may increase simulation time by 1 to 2 orders of magnitude.

We conducted simulation runs on different systems to obtain a range of simulation runtimes. With the baseline parameters, GLM-AED simulation times ranged from the best case of 6 seconds (on a CloudLab system with Intel Xeon CPU E5-2450 with 2.10GHz clock rate and 20MB cache) to 57 seconds (on a University of Florida system with virtualized Intel Xeon CPU X565 with 2.60GHz clock rate and 12MB cache). Note that individual 1-D GLM-AED simulations can be short-running; the GEMT feature of grouping multiple individual simulations into a single HTCondor job leads to increased efficiency.

Description of Experiment setup: The GRAPLER system



**Figure 6: Job runtimes for GRAPLER HTCondor pool, compared to sequential execution times on CloudLab and UF slots.**

deployed for this evaluation was distributed across three sites: University of Florida, NSF CloudLab, and Microsoft Azure. The GWS/GEMT service front-end, HTCondor submit node, and HTC-Central Manager were hosted on virtual machines running in Microsoft’s Azure cloud. We deployed three HTC-Execute nodes in total, with 16 cores each. Two nodes were hosted in virtual machines on a VMware ESX server at the University of Florida and one on a physical machine in the CloudLab Apt cluster at University of Utah. All the nodes in this experiment ran Ubuntu-14.04 and HTCondor version 8.2.8; nodes were connected by an IPOP GroupVPN virtual network, version 15.01. Each of the nodes was configured to have 16GB of memory allocated to them.

To conduct the evaluation, we carried out executions of three different experiments containing 3000, 5000 and 10000 simulations of an example lake with varying meteorological input data. Figure 6 summarizes the results from this evaluation. As a reference, we also present the estimated best-case sequential execution time on a single, local machine, taken the CloudLab and UF machines as a reference. For 10,000 simulations we achieved a speedup of 2.5 (with respect to sequential execution time of the fast workstation) and 23 (with respect to the sequential execution time at a UF virtual machine).

It is observed that the time taken to complete the job depended greatly on the way simulation tasks were allocated by the HTCondor scheduler. Note that the speedups are relatively modest compared to the best-case baseline, while significant compared to the worst-case baseline. The actual user-perceived speedup would be a function of which desktop environment a user would access the service from. Furthermore, because HTCondor is best-suited for simulations that are individually long-running, the raw user-perceived speedups of GRAPLER over local execution tend to increase as longer-running simulations are submitted through the service. We expect that, as demand for modeling tools by the lake ecology community increases, so will the complexity,

Job Size(# of simulations)	Service Response Time (seconds)	Input Processing Time (seconds)
3000	0.72	1790
5000	0.88	3039
10000	0.59	6854

**Figure 7: Input handling**

resolution and simulated epochs of climate change scenarios, further motivating users to move from a local processing workflow to remote execution through GRAPLER.

Submission of a job to the HTCondor pool involves processing of input (for sweep requests) and packaging of generated simulations into GEMT. In order to evaluate this step we carried out experiments to account for the time taken by GRAPLER to respond to a request to generate a given number of simulations and submit them for execution. The results are presented in Table 7. The metric `service_response` captures the time taken by GRAPLER to respond to a request with a UID, which is slightly more than the time required to upload the `base_input`. The metric `input_processing` captures the time taken to generate and compress all ‘N’ inputs for job submission.

Though not fully explored yet in the design of GRAPLER, another benefit of remote execution through a Web service interface is the leveraging of storage and data sharing capabilities of the collaborative infrastructure aggregated by distributed resources connected through the IPOP virtual network. For instance, the raw output size of the 10,000 simulation scenario described above is 900 GBytes. By keeping this data on the GRAPLER cloud and allowing users to share simulation outputs and download selected subsets of the raw data, the service can provide a powerful capability to its end users in enabling large-scale, exploratory scenarios, by both reducing computational time and relaxing local storage requirements at the client side.

#### 4. RELATED WORK

Several HTCondor-based high-throughput computing systems have been deployed in support of scientific applications. One representative example is the Open Science Grid (OSG [12]), which features a distributed set of HTCondor clusters. In contrast to OSG, which expects each site to run and manage its own HTCondor pool, GRAPLER allows sites to join a collaborative, distributed cluster by joining its virtual HTCondor pool via the IPOP virtual network overlay. This reduces the barrier to entry for participants to contribute nodes to the network – e.g., by simply deploying one or more VMs on a private or public cloud. Furthermore, GRAPLER exposes a domain-tailored Web service interface that lowers the barrier to entry for end users.

The NEWT [3] project also provides a RESTful-based Web service interface to High-Performance Computing (HPC) systems. NEWT is focused on providing access to a particular set of resources (NERSC), and does not address the need for a distributed set of (virtualized) computing resources to be interconnected by overlay virtual networks.

#### 5. CONCLUSION

GRAPLER, a distributed computing system which integrates and applies overlay virtual network, high-throughput com-

puting, and Web service technologies is a novel way to address the modeling needs of interdisciplinary GRAPLER researchers. The system’s contribution is its combination of power, flexibility, and simplicity for users who are not software engineering experts but who need to leverage extensive computational resources for scientific research. We have illustrated the system’s ability to identify and exploit parallelism inherent in GRAPLER experiments. Additionally, the system scales out, by simply adding additional worker nodes to the pool, to manage both increasingly complex experiments as well as larger number of concurrent users. GRAPLER is best suited for large numbers of long-running simulations as the distribution and scheduling overhead will increase the running time for such experiments. As lake models demand increased resolution and longer time scales to address climate change scenarios, GRAPLER provides a platform for the next generation of modeling tools and simulations to better assess and predict the impact to our planet’s water systems.

#### 6. ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation under Grants No. 1339737 and 1234983. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

#### 7. REFERENCES

- [1] J. D. Brookes and C. C. Carey. Resilience to blooms. *Science*, 334(6052):46–47, 2011.
- [2] S. R. Carpenter, E. H. Stanley, and M. J. Vander Zanden. State of the world’s freshwater ecosystems: physical, chemical, and biological changes. *Annual review of Environment and Resources*, 36:75–99, 2011.
- [3] S. Cholia, D. Skinner, and J. Boverhof. Newt: A restful service for building high performance computing web applications. In *Gateway Computing Environments Workshop (GCE), 2010*, pages 1–11. IEEE, 2010.
- [4] R. J. Figueiredo, P. O. Boykin, P. S. Juste, and D. Wolinsky. Integrating overlay and social networks for seamless p2p networking. In *Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises, 2008. WETICE’08. IEEE 17th*, pages 93–98. IEEE, 2008.
- [5] K. Flynn. Castles built on sand: dysfunctionality in plankton models and the inadequacy of dialogue between biologists and modellers. *Journal of Plankton Research*, 27(12):1205–1210, 2005.
- [6] A. Ganguly, A. Agrawal, P. O. Boykin, and R. Figueiredo. IP over P2P: Enabling self-configuring virtual IP networks for grid computing. In *International Parallel and Distributed Processing Symposium*, 2006.
- [7] M. Grinberg. *Flask Web Development: Developing Web Applications with Python*. O’Reilly Media, Inc., 1st edition, 2014.
- [8] P. S. Juste, K. Jeong, H. Eom, C. Baker, and R. Figueiredo. Tincan: User-defined p2p virtual network overlays for ad-hoc collaboration. *EAI*

*Endorsed Transactions on Collaborative Computing*,  
14(2), Oct 2014.

- [9] E. Litchman and C. A. Klausmeier. Trait-based community ecology of phytoplankton. *Annual Review of Ecology, Evolution, and Systematics*, pages 615–639, 2008.
- [10] H. W. Paerl, R. S. Fulton, P. H. Moisander, and J. Dyble. Harmful freshwater algal blooms, with an emphasis on cyanobacteria. *The Scientific World Journal*, 1:76–113, 2001.
- [11] H. W. Paerl and J. Huisman. Blooms like it hot. *Science*, 320(5872):5–7, 2008.
- [12] R. Pordes, D. Petravick, B. Kramer, D. Olson, M. Livny, A. Roy, P. Avery, K. Blackburn, T. Wenaus, F. Würthwein, et al. The open science grid. In *Journal of Physics: Conference Series*, volume 78, page 012057. IOP Publishing, 2007.
- [13] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015.
- [14] A. Rigosi, C. C. Carey, B. W. Ibelings, and J. D. Brookes. The interaction between climate warming and eutrophication to promote cyanobacteria is dependent on trophic state and varies among taxa. *Limnology and Oceanography*, 59(1):99–114, 2014.
- [15] V. H. Smith, G. D. Tilman, and J. C. Nekola. Eutrophication: impacts of excess nutrient inputs on freshwater, marine, and terrestrial ecosystems. *Environmental pollution*, 100(1):179–196, 1999.
- [16] A. Solem. *Celery: Distributed Task Queue*, 2013.
- [17] D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: The condor experience. *Concurrency-Practice and Experience*, 17(2-4):323–356, 2005.
- [18] K. Weathers, P. C. Hanson, P. Arzberger, J. Brentrup, J. Brookes, C. C. Carey, E. Gaiser, D. P. Hamilton, G. S. Hong, B. Ibelings, et al. The global lake ecological observatory network (gleon): the evolution of grassroots network science. 2013.