



Project Proposal Report

Team 15 – Ethan Grantz, Joshua Jeng, Garrett Mills, Abhigyan Saxena, QiTao Weng

Project Name

IntEarthNet using Swarm

Project Synopsis

Swarm will be a small domain-specific language for distributed computing, used to build IntEarthNet, an "internet health check" portal that shows global network speeds.

Project Description

Nearly all frameworks for distributed computing rely on external libraries for existing language to retrofit intra-machine communication using things like MPI. While these work, the integration can be clunky and requires the programmer to have some understanding of the inner working of the library they are using. Furthermore, most of these libraries have no logic for scheduling which nodes the code will run on by feature, instead relying on external schedulers to allocate the nodes beforehand. By contrast, Swarm is distributed by default. Code blocks and iteration is natively parallelized by default, and the logic for synchronizing results and sharing data structures is built into the language natively. Swarm also schedules its resources internally based on declarations of what resources are provided by which nodes. This not only removes the need for an external scheduler, but also requires the programmer to have less of an understanding of the underlying distribution logic and sync mechanisms. We want to use this DSL to build IntEarthNet, a web portal displaying internet access health around the world by using Swarm to run a suite of network tests on nodes around the world. For the final prototype for this class, Swarm should be able to lex, parse, and interpret basic code to execute an external script, aggregate the results, then store them in a shared database. IntEarthNet will read from that database and display the up-to-date metrics on a virtual globe.

Project Milestones

Sem 1:

1. Setup environment (tools, language, schema) (9/30)
2. DSL Feature List, Syntax, and Grammar (10/16)
3. Website UML Diagrams (10/16)
4. Complete Lexer (11/27)
5. Complete front-end UI (11/27)

Sem 2:

1. Complete Parser (1/29)
2. Complete Interpreter (3/30)
3. Complete ping report program (4/29)
4. Final Integration with global servers (5/6)

5. ML Proof of concept (5/6)

(See last page for Gantt Chart)

Project Budget

Development server

VPS	Monthly price	Months	Total
Google Cloud Compute	\$12	8	\$96

International servers according to most popular local hosting services (\$12/mo per server approx.) These are used in the final IntEarthNet report which will use VPS local to the region to ping and download pages from top sites in the region to measure latency and packet loss at any given moment.

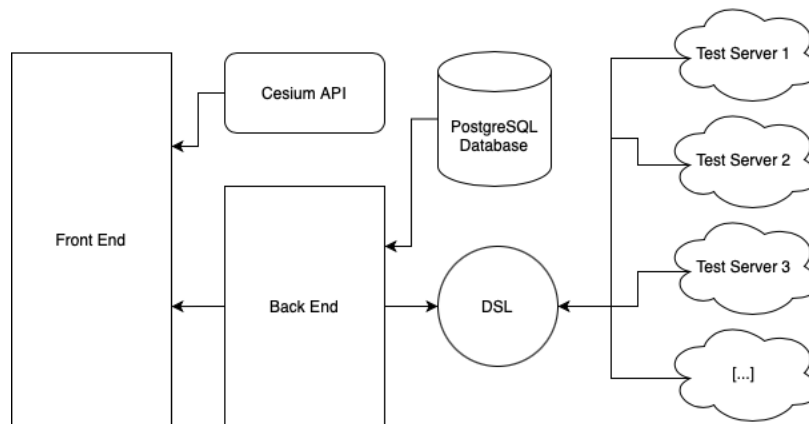
- At least one for each continent - except Antarctica
- Additionally subdivided by industrialization of different regions

Service	Approximate price
US AWS	\$5
Brazil HostGator	\$14
Germany Ionos	\$2
Australia OVHCloud	\$5
Ethiopia HahuCloud	\$8
Total	\$34/mo*

*Likely more as we add additional VPS for countries with vastly different economic development.

Total: \$130

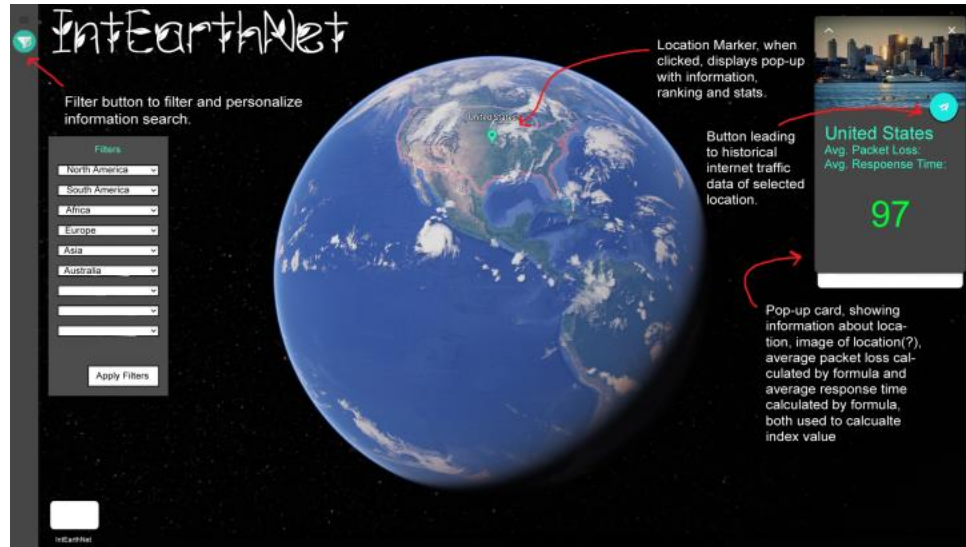
Preliminary Project Design



Website Structural UML Diagram

The IntEarthNet web application leverages an open-source API known as CesiumJS for many major aspects of its front-end. CesiumJS is effectively a JavaScript API that can display various 3D models and various data and labels related to that model. Each web server used to generate the internet weather report will be represented on the globe by a pin. Thus, to construct this, IntEarthNet’s user interface will use Cesium’s built-in model of Earth, as well as its pin builder and billboard. The remaining aspects of the UI will be designed from scratch in HTML/CSS/JS. In order to use the data gathered by the DSL, the web application has a PHP backend that queries a PostgreSQL database. The PostgreSQL database stores all information needed to display the data on the globe, including latitude and longitude coordinates,

location names, and the actual data gathered by the DSL. Finally, the website itself will be hosted on an Ubuntu-based virtual private server, on an Apache webserver.



Website mockup illustrating the pin filter

Users will be able to filter report data per diagnostic server by location. The pin filter feature will divide the pins into various sets, divided by continents. The user will be able to select which continents will show their data. In addition, basic statistical analyses of the selected data, such as averages or standard deviations, will be shown on the screen.

Besides the contemporary internet weather report, the web application will also be able to display historical information, allowing users to compare internet diagnostics over various points in time. As a result, the website will be able to display reports side-by-side across various points in time for easy visual comparison. This data will be stored in the PostgreSQL database, to which the DSL will add new data approximately every 5 minutes.

When performing its internet weather report, IntEarthNet will need to execute tasks across the various web servers we have across the globe. To perform these actions in parallel, the website will utilize a domain-specific language called Swarm, which we are writing ourselves from the ground up. Swarm's lexing and parsing will be built using Flex and Bison, respectively. These two tools have traditionally been the industry standard, so we can trust in their reliability.

The syntax is largely C-like but incorporates intentionally distinct keywords to denote the semantic differences between a synchronous language like C and an asynchronous, distributed language like Swarm. This syntax was chosen due to its high popularity – allowing for a low learning curve and easy adoption for use – and ease of parsing of its syntactic structure.

Most of the language constructs are as you would expect from a C-derived language – things like variable declarations and assignments, function calls, and block designations all track – but we have added our own constructs, like enumerations and resource blocks, which fit our particular use case. A cornerstone of the functionality of Swarm is the enumeration block. An example usage might look like:

```
enumerate mySet as myItem {  
    doSomething();  
}
```

The key difference between this enumeration block over a set and something like a for-loop iterator over a list in C++ is twofold:

1. The "list" type in Swarm is a set, which has no ordering.
2. Likewise, enumerations do not execute in a guaranteed order, and may execute in parallel.

So, the Swarm interpreter will take the items in the set and break up the processing within the block so it could potentially all happen in parallel among different workers – one “control stream” for each item in the set. This is perhaps analogous to all the iterations of a for-loop executing at once.

As a result, Swarm needs to introduce concepts to control two things: safe distributed access to shared data structures, and a way to request the execution of a block with access to specific computing resources.

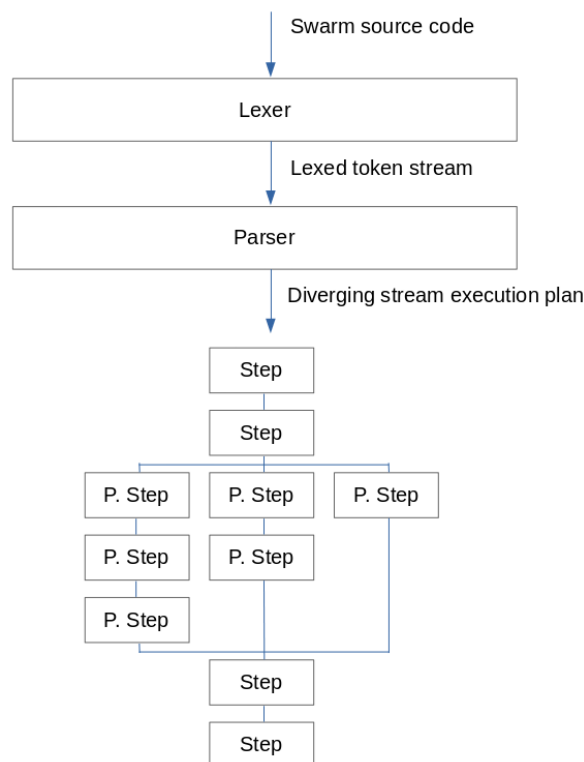
We will be using Redis to implement a distributed lock to ensure synchronization issues such as race conditions can be prevented natively through the Swarm itself. Swarm will allow users to create shared variables and data structures that will behave consistently when operations across multiple machines are executed concurrently. Swarm will be able to pass along variables and data structures during inter-process communications. Allowing the execution of the language to solve synchronization issues will make distributed computing much more easily accessible. Instead of booting up an instance of the same processed manually and iteratively on different computers and data centers and requiring the user to create a separate locking mechanism to prevent issues from arising when operating on the same data.

To solve the other case, Swarm introduces a concept of resource-blocks. These blocks request a computing resource that needs to be available to the executor for the code to execute correctly. For example, say we want a block of code to execute only if a particular file is available on the filesystem:

```
with file('/var/myfile.csv') as fileBuffer {  
    doSomething();  
}
```

In this case, the with-statement tells the interpreter that the code block within the braces can only be run on an executor that has the /var/myfile.csv file available. The interpreter finds an executor satisfying that criterion and the code within the with-block is run on that executor, which makes the file resource available in the fileBuffer variable.

This will be accomplished with a two-stage compiler-then-interpreter. The compiler will take the Swarm source code, lex it, then parse it into an AST. Finally, the compiler will create a diverging-streams execution plan. The structure of the plan reflects the ability of swarm to natively handle both linear and distributed-parallel code. For example:



Swarm (DSL) Structural Diagram

This execution plan is then passed to the interpreter, which steps through it piece-by-piece and sets up the environment and runs the code.

There are some design constraints we have to consider. The user interface for the IntEarthNet application will be built for a web-based platform, utilizing the CesiumJS framework. As part of the open-source Apache license, we will need to include an open-source licensing page on the aforementioned website. Besides the normal web languages (HTML, CSS, JavaScript, PostgreSQL, PHP, etc), the parallelized tasks will be done using our domain-specific language Swarm. While we do have expenses, namely the access to various worldwide servers required to check global internet speeds, the cost will be covered by the University of Kansas without any predefined budget constraint. Finally, as with all other EECS 581 projects, the project deadline will be in May 2022.

Ethical Issues

IntEarthNet is an application that pings various servers across the globe and then records data into an SQL table, displaying it onto a three-dimensional globe in the form of pins. Overall, we do not believe that IntEarthNet has any ethical issues per se. However, the data centers we are pinging are a major risk to the environment. Data centers are what can be equated to, factories of the information age, data centers make communication across the internet possible, but it comes at the cost of operating the sheer (and growing) amount of data centers using electricity. The energy consumption of digital technology is increasing rapidly over the year which is a major concern for the environment. We will be using Cesium.js API for displaying the globe on our website, we have not extensively researched cesium.js documentation as we will be using considerably basic elements from Cesium.js so we do not know what sort of ethical issues Cesium.js itself holds as an API.

Intellectual Property Issues

Much of the work we have planned in our project is original, so there are not many intellectual property issues concerning the originality of the code. However, we need to take care to credit and/or make original snippets of code used from examples or documentation for the libraries we are integrating. This is a broader issue in the software industry in general, but in our case that means crediting snippets of code used in the Swarm parser/interpreter that were taken from Flex, Bison, and class examples. Likewise, with IntEarthNet, we are building our UI based on the open-source Cesium project which means referencing their documentation quite heavily. As such, we need to take care to credit the original authors of the examples we reference, if any, in our final product. Finally, there's the question of who owns the code we write. This will be something we will hash out amongst ourselves as the project concludes. Different sections of Swarm and IntEarthNet were largely written by different authors so integrating the authors' IP is a concern.

Change Log

Google Earth turned out to be more difficult to use than anticipated, so we found an alternative in the form of CesiumJS. In addition, while CesiumJS does support using KML files, we were able to remove them from the design, as it is simpler to query data directly from the database and store it with JavaScript.

A feature to be able to filter the pins that are displayed was added.

Recognized the need to include an open-source license page to IntEarthNet.

IntEarthNet

University of Kansas
Team 15

Tue, 1/18/2022
Fri, 9/24/2021
1

Project Start:

Display Week:

TASK	ASSIGNED TO	PROGRESS	START	END	1/14	1/21	1/28	2/4	2/11	2/18	2/25	3/4	3/11	3/18	3/25	4/1	4/8	4/15	4/22	4/29	5/6	5/13
Semester 1																						
Setup environment (tools, language, schema)	All	50%	9/24	9/30																		
DSL Feature List, Syntax, and Grammar	Garrett, QITao, Joshua	20%	9/26	10/16																		
Website UML Diagrams	Ethan, Abhi	0%	9/30	10/16																		
Complete Lexer	Garrett, QITao, Joshua	0%	10/16	11/27																		
Complete front-end UI	Ethan, Abhi	0%	10/16	11/27																		
Semester 2																						
Complete Parser	Garrett, QITao, Joshua	0%	1/18	1/29																		
Complete Interpreter	Garrett, QITao, Joshua	0%	1/29	3/30																		
Complete ping report program	Ethan, Abhi	0%	1/18	4/29																		
Final Integration with global servers	Ethan, Abhi	0%	4/14	5/6																		
ML Proof of concept	All	0%	4/14	5/6																		
Winter Break																						