

Digo: distributed golang

Hanxiao Lu (hl3424), Yufan Chen(yc3858), Sida Huang(sh4081), Wenqian Yan(wy2249)

February 2, 2021

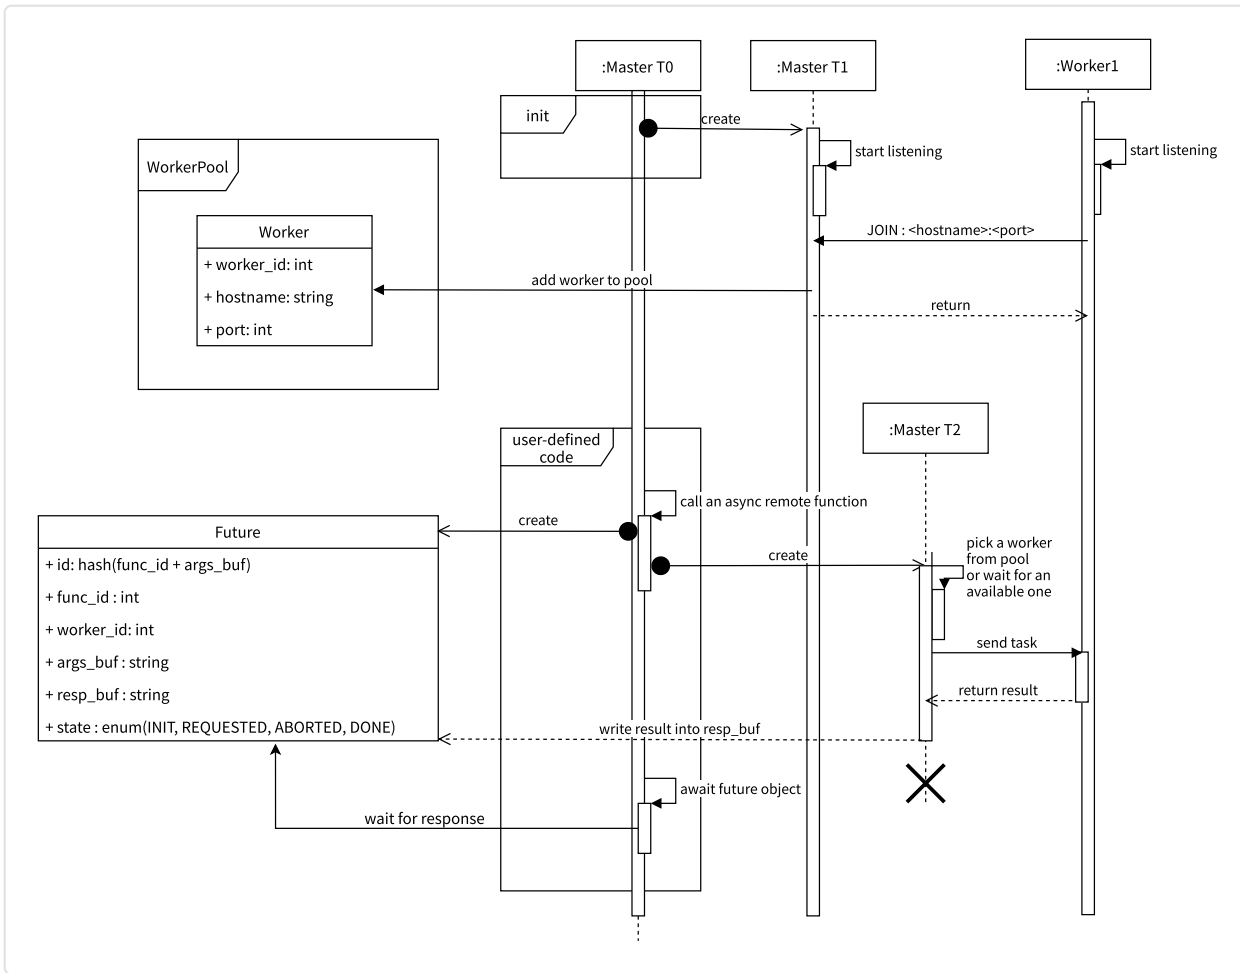
1 Language Overview

Digo is an imperative, statically typed, compiled programming language inspired by Golang, but with support for distributed routine. With Digo, we can easily create a master-worker model program that the master can distribute tasks to worker nodes.

To define a task, we use `async` as the keyword to create a new thread to execute a function. There are two types of async functions: `async func foo(...)` is a local task, and `async remote func foo()` is to define a task that is delivered to workers and execute remotely. We also designed `future` object to define a task.

To fit the master-worker distributed model, the compiled executable has two modes: master mode and worker mode, to specify when executing. An example to run with master mode: `./a.out -h 10.23.2.2 -p 80 -m master`. And to run code in the worker node, we need to specify its master's port and mode in args: `./a.out -h 10.23.2.2 -p 80 -m worker`.

The following plot exhibits the whole underlying process to make the complicated process clear:



As shown in the plot, the master keeps two major threads (T0 and T1): T0 is the main thread to execute the master function, and when it initiates, it creates T1 to create and maintain a socket sitting there and waiting for any connections from newly-added workers. When a worker (worker1) is added and starts running, it firstly listens to a port and sends `<hostname, port>` to the master. After the master(T1) receives this information, it records `<worker_id, hostname, port>` in WorkerPool. The worker can close the connection after it receives ACK from the master.

When executing the master function, everytime a future object is to be created by calling an async remote function, a new thread (T2 in the diagram) is created with it to:

- a. run scheduler to find a proper worker in the WorkerPool maintained by the master. If there is no available worker, this thread will stuck in a loop until it gets one.
- b. after it finds one worker to deliver the task to, it sends `<function_id + serialized parameters + future_object_id>` to the assigned worker.

Note that, by design, we only pass values instead of pointers to async remote functions and do not allow async tasks to access global variables.

When the worker receives task request from the master, it starts a new thread to look for the function specified by the `function_id` it received, and then execute the task. After it finishes the function, it sends a response containing `<result, future_object_id>` back to the master. After the thread (T2) in the master receives the response from the worker, it copies response to `resp_buf` in the corresponding future object. When the master executes `await` or `gather`, the underlying function enters a loop to check if the `resp_buf` in the future object is filled, otherwise it waits for its response in a blocking way.

An example of potential application is map-reduce model where tasks are computational intensive. Please refer to our word count sample program in the later part.

2 Language Details

2.1 Language Features

Type Scope	Static
Type Strength	Strongly Typed
Garbage Collection	Automatic
Type System	Statically typed
Evaluation	Strict evaluation
Type Inference	Yes

2.2 Data Types

Data Type	Operation	Examples	Passed by value/reference
string	<code>+, =, ==, >, <, >=, <=, !=</code>	<code>a := "helloworld"</code>	reference
int	<code>=, ==, +, -, *, /, >, <, >=, <=, %, +=, -=, !=, ++</code>	<code>a := 1</code>	value
float	<code>=, ==, +, -, *, /, >, <, >=, <=, %, +=, -=, !=,</code>	<code>a := 1.0</code>	value

	++		
bool	=, ==, !=, !, &&,	a := true	value
future	await, =	a := await gcd()	reference
slice	append(), len(), [begin:end]	a := []int{1,2,3} append(a, 4) b := a[0:2]	reference

2.3 Built-in Functions

Function Signature	Description
append(<slice>, <element>)	append a new element into a slice
len(<slice>)	get the length of a slice
gather(<future-slice>)	await all future objects in a slice. All future objects in the slice should return the same type of value.

2.4 Keywords

Go
<ol style="list-style-type: none"> 1 for, if, else, func, return, await, async, remote, 2 var, string, int, float, bool, continue, break

2.5 Control Flows

for loop

Go

```
1 for i := 0; i < 100; i++ {
2 }
3
4 for {}
5
6 for i < 100 {
7     i += 1
8     continue
9 }
```

If-else

Go

```
1 if a > b {
2 } else if a == b {
3 } else {
4 }
```

2.6 Functions

The normal function looks like Golang.

Go

```
1 func foo() {
2 }
```

The function named `master` will be treated as the entry function of the master:

Go

```
1 func master() {
2 }
```

Function with the `async` keyword will be treated as an async function to create a new thread to execute. When the master `await <async function>`, the master will check `res_buf` in future object and wait for the completion of this thread.

Go

```
1  async func task() int {
2  }
```

Function with the `async remote` keyword will be treated as a remotely-executed async function. When the master `await <async remote function>`, internally the program will pick a worker and let this worker execute this function, and wait for its response.

Go

```
1  async remote func task() int {
2  }
```

2.7 Futures

The return result of an asynchronous function is of `future` type. When a future object is awaited, the program will blockingly wait for the response of the worker that actually executes the function.

The layout of a future object looks like:

C

```
1  typedef enum {INIT, REQUESTED, ABORTED, DONE} state_t
2  typedef unsigned int id_t
3
4  struct future {
5      id_t id // hash(func_id + args_buf)
6      id_t func_id // id of the async function
7      // id of the worker the scheduler chose.
8      // If this is a local task, then worker_id = 0
9      id_t worker_id
10     string args_buf // a buffer to store the serialized arguments
11     string resp_buf // a buffer to store the serialized return result.
12     state_t state
13 }
```

2.8 Comments

Go

```
1 // inline comment
2 /* comment block */
```

3 Sample Program

Go

```
1 // word-count
2
3 // The workers execute this `async remote` function.
4 // Calling an `async remote` function returns immediately
5 // after the job is dispatched to a worker. It does not
6 // wait for the job to finish.
7 // An `async remote` function does not return with return values
8 // defined in the function prototype.
9 // Instead, it returns a future object, on which
10 // user can call await/gather to get the actual return values.
11 async remote func worker_count_word(words []string) ([]string, []int) {
12     // we count the words by first sorting them.
13     Sort(words)
14     resultWord := []string{}
15     resultCount := []int{}
16     if len(words) == 0 {
17         return resultWord, resultCount
18     }
19     word := words[0]
20     resultWord = append(resultWord, word)
21     resultCount = append(resultCount, 1)
22     for i := 1; i < len(words); i++ {
23         if words[i] == word {
24             resultCount[len(resultWord)-1]++
25         } else {
26             word = words[i]
27             resultWord = append(resultWord, words[i])
28             resultCount = append(resultCount, 1)
29         }
30     }
31     return resultWord, resultCount
32 }
```

```

33
34 // The master executes this `async` function in a new thread.
35 // Calling an `async` function returns immediately with
36 // a implicit future object.
37 async func count_word(workerWords [][]string) map[string]int {
38     futures := []future{}
39     for i := 0; i < len(workerWords); i++ {
40         // Calling the `async remote` worker_count_word function
41         // will automatically send the task to a worker.
42         // It does not block.
43         futures = append(futures, worker_count_word(workerWords[i]))
44     }
45     var words map[string]int
46     for i := 0 ; i < len(workerWords); i++ {
47         // Here, we are explicitly waiting for the remote task
48         // worker_count_word(workerWords[i])
49         // to finish.
50         resultWord, resultCount := await futures[i]
51         words[resultWord] += resultCount
52     }
53     return words
54 }
55
56 func split_file(file string, workerCount int) [][]string {
57     reader := FileReader(file)
58     workerWords := make([][]string, workerCount)
59     rotate := 0
60     for {
61         // Reads until the first occurrence of delim ' ' or '\n'
62         word, err := reader.ReadString([]string{" ", "\n"})
63         if err != nil && err != EOF {
64             panic("cannot read file")
65         }
66         // Dispatch words to workers evenly
67         worker := rotate % workerCount
68         workerWords[worker] = append(workerWords[worker], word)
69         rotate++
70         if err == EOF {
71             break
72         }
73     }
74     return workerWords
75 }
76
77 func word_count_entry() {

```



```

78     futures := []future{}
79     workerWords := split_file("book1.txt")
80     futures = append(futures, count_word(workerWords))
81     workerWords = split_file("book2.txt")
82     futures = append(futures, count_word(workerWords))
83     workerWords = split_file("book3.txt")
84     futures = append(futures, count_word(workerWords))
85
86     // Calling `async`/`async remote` functions does not block unless
87     // we explicitly wait for them to finish by calling await/gather.
88     // So word_count_entry() does not block until we `gather` here:
89     results := gather(futures)
90     // After gather returns, all tasks indicated by futures
91     // are finished.
92     print("Word Count result of book1.txt")
93     print(results[0])
94     print("Word Count result of book2.txt")
95     print(results[1])
96     print("Word Count result of book3.txt")
97     print(results[2])
98 }
99
100 func master() {
101     word_count_entry()
102 }

```

After having compiled this code, we will get an executable. In order to run this executable, we need to pass three required arguments:

Bash

- 1 -m [master|worker] specifies whether this is a master or a worker
- 2 -h <string> hostname of the master
- 3 -p <int> port of the master

4 Potential Application

Per the design of Digo, it can divide a complicated and computation-intensive task into separate parts and the master can assign each of them to available workers. Applications of our language include but do not restrict to the following examples in a distributed way:

- Image processing: The master wants to know the information about a given picture, it commands several workers to convolve with different kernels to do object recognition, object classification and etc.
- Language processing: The master wants to parse different ngrams on a given text, it commands different workers to do different ngram tasks. First worker to parse unigram, Second worker to parse bigram and etc.
- Matrix Multiplication: Matrix multiplication is often expensive. But with our language, master could command each worker to calculate each cell of the output. If task of matrix $A * B$ is given to master, master commands the first worker to calculate $A[0,:] * B[:,0]$, second worker to calculate $A[0,:] * B[:,1]$ and etc