



LECTURE 5: MAKING DHTS DO MAGIC TRICKS!

Ken Birman
Spring, 2022



TODAY'S AGENDA: TWO PARTS

- Understanding how to put “anything at all” into a DHT for scalability and high performance: DHTs can hold data *in memory*.
- Limitations and factors the DHT developer must keep in mind.

CENTRAL QUESTIONS WE WILL FOCUS ON

How spread out do I want my data to be?

How do deal with very large, complex, data structures (with pointers!)

What to do about the small but non-zero risk of key collisions




REMINDER: DHT BENEFITS (AND WHY)

The DHT idea can be traced back to work by people at Google, and to papers like the Jim Gray paper on scalability.

We take some service and structure it into shards: sub-services with the identical API, but handling disjoint subsets of the data.


We need some way to know where to place each data item. We use a key here: the key is a kind of unique name for the data item, and by turning it into an integer modulo the number of shards, we find the target shard.

DHT PICTURE



A DHT assumes that the data-center network is very fast. 10us delays for a gRPC are typical – whereas the client application has 100ms to send a reply back to the human end-user (customer)

SIDE REMARK



In fact it isn't horrendously costly that the items are scattered around

- Those 10us - 100 us delays are very small and you might even be able to fetch all your data in parallel, by issuing concurrent **put/get** requests (this requires an *asynchronous put/get*, or multiple threads).
- A file system would have the same overheads, and more system calls: to do a put or get would require *multiple* file system operations, if each object is in a file by itself. So in fact **put/get** is relatively cheaper.


On the other hand, it still is faster to do talk to a server on the same machine

COLLOCATION

Sometimes we actually can run programs on the same machines that host the DHT, and may want to avoid those network delays entirely.

This can lead us to think about ways to ensure that certain (key,value) pairs end up *collocated* where our program will run.

LOCKING LIMITATION



We are not given any way to do locking or 2-phase commit. In fact Jim Gray showed us that locking across shards would be ineffective.

A **get** or **put** is an atomic action on a single shard. For fault-tolerance, the shard update can use state machine replication (atomic multicast or Paxos).

With this approach we get “unlimited” scaling, and we can even keep all the data in memory (as long as the number of shards is big enough).

DHTS WORK BEST FOR DATA THAT DOESN'T CHANGE AFTER IT IS INITIALLY STORED

Once a web page has been uploaded, we probably won't update it again. A web page that won't change is an example of *immutable data*.

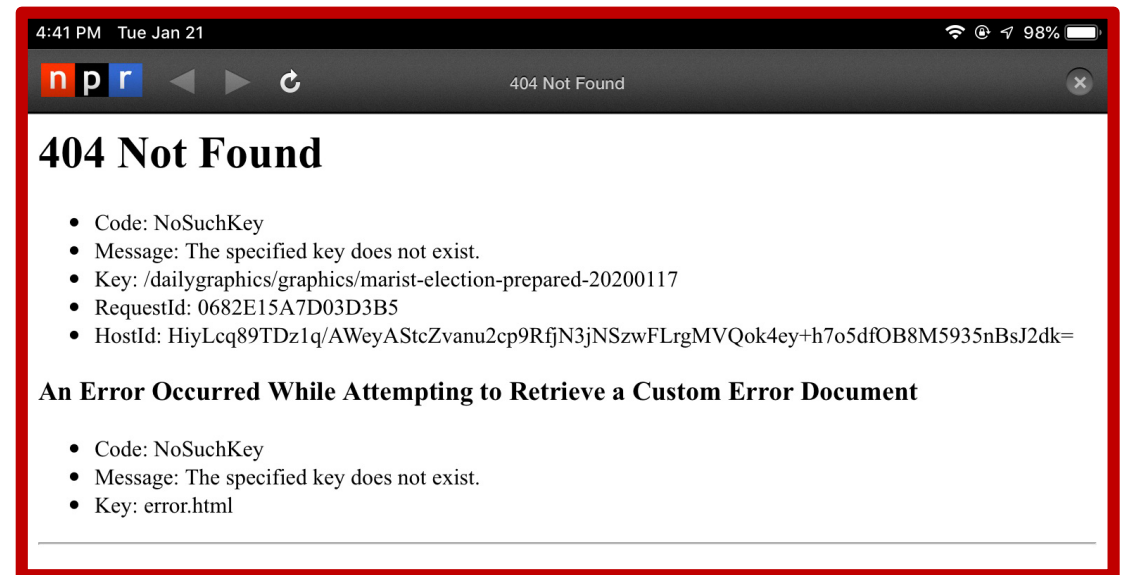
A DHT is ideal for this kind of data. Locking isn't useful for a big read-only data set even if we didn't have sharding!

This is one reason DHTs are universally popular for caching.

EXAMPLE: AN NPR NEWS ARTICLE

This error message from a popular news site, NPR, is clearly caused by not finding data in a DHT!

They probably stored their articles in the DHT, but somehow got an error when trying to fetch this article to build my web page. It could be an example of CAP: When a DHT resizes elastically, sometimes it makes errors for a few seconds afterwards...



BUT AN NPR NEWS STORY IS TOO EASY.

For the first few years, big search companies focused on just downloading snapshots of web pages and offering quick ways to find things.

Over time, however, there was an appreciation that the social network is the bigger opportunity. And these evolve rapidly over time.

So we saw a growing need to store data that does change.



HOW TO STORE “ANYTHING” AT LARGE SCALE

These data sets are huge – MUCH too big for any single computer.

Yet not only do we want to hold the data for access, we want superfast access: we want the data to be in memory, not on disk!

A DHT can solve this for us. The network I/O cost is a factor, but is still much faster than disk I/O. And modern datacenter networks, with the fastest software, can push network delays down to the 1-2us range.



HOW TO STORE “ANYTHING” AT LARGE SCALE

Puzzle: A DHT officially just holds (key, value) data:

- The key is an integer. Some permit various sizes: 64 bits, 128, 256.
- The value is generally either another integer, or a byte array.
- Some DHTs are specialized for (integer, byte[]), and some for (integer,integer). These often are used “together” for flexibility.

So how can we come up with a key for “anything at all”? And how can we use this value model to store “anything at all”?

Solution: We use *serialization*. This converts an object to a byte array.

COMING UP WITH SUITABLE DHT KEYS

You need a unique name for the objects you are storing.

For example: Ken's dog was named Biscuit. But "Biscuit" is not a unique name. The DHT could have some other object with that name too.



On the other hand, `"/users/Ken Birman/pet/Biscuit"` is a unique key, and we can hash it with SHA64 or SHA128 into a unique integer.

BASICALLY... TWO CASES TO CONSIDER

[1] Objects at the “root” of some kind of structure need a name that a client application will know a-priori. We use a file pathname as a key.

- It will be hashed automatically, turning it into a numerical key.
- Collisions are very rare, but not “impossible”. Must keep this in mind!

[2] Objects that are really internal to a structure, like “rows in a table” or “photos in a list of pets,” need some sort of generated unique key. Here we use a *key generating service* to make life easier.



... KEY GENERATING SERVICES.

They provide a genuinely unique key

- Microsoft and AWS both have “registry” services.
- The resulting keys will have no obvious meaning to a human user, but is unique.
- These are for objects we will fetch “algorithmically”, not for external programs to use directly.

EXAMPLE: MICROSOFT REGISTRY

Registry Editor

File Edit View Favorites Help

Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\BGFX

Name	Type	Data
(Default)	REG_SZ	(value not set)
AnimationTotal	REG_DWORD	0x000031eb (12779)
BPP	REG_DWORD	0x00000020 (32)
CompressBitmaps	REG_DWORD	0x00000000 (0)
FadeCpu	REG_DWORD	0x00000076 (118)
FadeFrames	REG_DWORD	0x00000007 (7)
FadeHigh	REG_DWORD	0x00000025 (37)
FadeLow	REG_DWORD	0x00000023 (35)
FadeMemory	REG_DWORD	0x001cda80 (1890944)
FadeOverlap	REG_DWORD	0x00043874 (276596)
FadeProlog	REG_DWORD	0x000030ec (12524)
FadeTotal	REG_DWORD	0x000031eb (12779)
Height	REG_DWORD	0x00000300 (768)
LogoSize	REG_DWORD	0x0004514a (282954)
ProgressFrames	REG_DWORD	0x00000000 (0)
ProgressHigh	REG_DWORD	0x00000000 (0)
ProgressLow	REG_DWORD	0x00000000 (0)
ProgressManual	REG_DWORD	0x00000000 (0)
ProgressMemory	REG_DWORD	0x0007be80 (507520)
ProgressProlog	REG_DWORD	0x00000000 (0)
ProgressTotal	REG_DWORD	0x00000000 (0)
ResidentSize	REG_DWORD	0x0001dc90 (122000)
Resume	REG_DWORD	0x00000001 (1)
Width	REG_DWORD	0x00000400 (1024)

CAN YOU PICK A KEY THAT WILL MAP TO A SPECIFIC SHARD?

In fact, you probably can! It may take a few tries.

You would need to know the DHT hashing scheme and be free to vary the keys (like by appending a number or something). Then you could just retry until the key hashes to the shard you want!

Some DHT developers actually use ideas like this. Some DHTs even let you just specify where to **put** your objects (but those need a form of lookup table, to be able to find the objects later for **get**).

NAME SPACES AND KEYS

“What is in a name? That which we call a rose
By any other name would smell as sweet...”

Juliet, in Shakespeare's Romeo and Juliet

A ***name space*** is some sort of user-oriented, semantically sensible, place to store names of objects. We could actually have one object in many name spaces, if the same object makes sense in different situations.

The namespace is used as a “service” to map from a name that makes sense to the user, to a unique internal key that makes sense in a DHT.

A cloud file system always has a namespace server as one component. We think of the storage servers as a separate, distinct, component.

KEN'S PETS

So we could, for example, have a kind of table listing all the pets Ken has had, with information about them

Pet Name	Period	Species	Photo List	Health Status
Nerd	1961-1962	Gerbil	<i>Empty</i>	<i>Deceased</i>
Susie	1970-1986	Keeshund	IMG-17171, ...	<i>Deceased</i>
Biscuit	2003-2013	Golden Retriever	IMG-22187, ...	Deceased

This table would be a “name space” if the photo list is a list of keys

MANY THINGS CAN BE GIVEN UNIQUE KEYS

We could have a unique key for each row in a table.

We could have a unique key for each photo in a photo album. The album itself could be “named” but also have a key: its value would be a list of the keys for photos in the album.

Cloud systems use this approach very broadly!

COULD KEYS “COLLIDE”?

Yes, if the keys don't have a large enough range of values, or the random number generator isn't very effective.

Most cloud systems favor fairly large keys, like 64 bits. And some key generators use a variety of tricks to make sure that they won't give out the same key twice. A random number generator wouldn't necessarily work.

Collisions would cause problems because two different objects would end up sharing what should be a unique name – a serious inconsistency.



WHAT PROBABLY HAPPENED IN THE NPR SITE?

It probably wasn't a key collision.

In fact, I get news alerts for certain kinds of news stories, like confirmed first-encounters with space aliens. So... NPR posts a first-encounter story. And I receive an immediate alert!

The story was saved into a DHT, but maybe the DHT replication scheme is a bit slow, or it was resizing just at that moment for elasticity reasons. Until it “settles”, the key is correct but the story just can't be found (yet).

DO YOU REMEMBER “BASE”

Mentioned in lecture 3. BASE means a ‘**B**asically **A**vailable **S**ervice with **E**ventual Consistency’.

A DHT often uses the BASE guarantee. Updates aren’t instantly visible or even guaranteed to “stick” if a crash happens right after the update. But problems are rare.

This is a frustrating but common cloud property – many services do this.

HOW DID THIS RELATE BACK TO CAP?

CAP was all about how we often face choices:

- Give a super fast response on a web page, or an instant notification
- But this means not waiting for the elastic reconfiguration to finish, or for the replicated update to fully propagate to all the replicas.
- We might operate in an inconsistent way, briefly!

CAP says: you can't have all three from {C,A,P} at the same time.

Relax consistency to get better availability and respond immediately even if a service you would have liked to check with is temporarily not responsive.




SUPPOSE WE HAD A LIST OF PHOTOS

The list itself is a data structure of objects that are linked by pointers.

The objects could be something like a photo description, and the photo (or other kinds of photo properties: “meta-data”).

HOW TO STORE “ANYTHING” AT LARGE SCALE

```
class myPet {  
  int uid;  
  animal species;  
  hashset<string,photo> photo_collection;  
  ....  
}
```



```
myPet biscuit {  
  uid := 5731,  
  species := animal::dog,  
  photo_collection := { ["on a rug", ●],  
                       ["in the woods", ●] }
```




These objects all have type names the application understands, data that matches.

But the DHT wouldn't know these object types! Amazon created the DHT long before you wrote the code defining this class.

NEXT ISSUE: CLOUD DATA FORMS A HUGE, COMPLICATED GRAPH



Ken belongs to Entrepreneurs' Org



The entrepreneurs shared a viral (completely exaggerated, basically fake) story about a complete cure for cancer.

SERIALIZATION CONCEPT

With a single object, a serializer just emits a self-describing data structure as a byte array, listing the field-types and their values.

If an object has sub-fields, it uses a recursive descent to serialize the inner object too. The self-describing byte array has a way to represent: “this is an object too, and the next 184 bytes describe it”.

The serializer output is often larger than the original object

ALSO, RECURSION ONLY HELPS FOR WHAT IS LOGICALLY A “SINGLE NODE”

The serializer knows how to do recursive serialization but only for what is logically a single object (and its subobjects).

```
myPet biscuit {  
  uid := 5731,  
  species := animal::dog,  
  photo_collection := { ["on a rug"],  
                      ["in the woods", ]}  
}
```



Key = “Ken Birman/pets/Photos of Biscuit”

Value = [0xFF 0xA6 0x1B 0x00 0x99 0x11 0x03 0xFF 0xFF ...]

A deserializer is used to recreate the data structure. This is also sometimes called marshalling and demarshalling.

REASONS SERIALIZATION CAN BE COSTLY

We generally like the byte array to be in a “hardwareindependent format”, meaning that both Intel and ARM (and other devices) can reconstruct the object into their local data representation.

- There are several opinions about the best byte-order for integers
- Floating point formats are hardware-specific
- Some systems null-terminate strings; others just view a string as bytes
- Memory “alignment” rules differ from machine to machine.
- Compilers can make additional optimization choices



WHAT ABOUT A MORE COMPLEX STRUCTURE?


When we talked about Facebook we focused on caching images. But what if we also wanted to use a DHT to hold a social network graph

- A social network has lots of structure: entities *related* by edges.
- Such a graph is *much* too large to hold on one computer.

The edges are like pointers. We can use unique keys for this role, but now we may run into the issue of collocation.

A SIMPLE CASE

Let's think about a graph with just three nodes



SHOULD WE COLLOCATE THE ROOT, A AND B?

If we can put them on one shard, and our access program can run on that same node, we gain speed.

But we also load more data into that single shard, and could cause a less smooth distribution of data.

Also, for a really big graph, the shard would run out of space.



("Ken Birman/pets/Biscuit",

```
myPet biscuit {  
  uid := 5731,  
  species := animal::dog,  
  photo_collection := { ["on a rug", 1],  
                      ["in the woods", 1] }  
}
```



OUR BASIC CHOICE

Pick some sort of rule for chunking the graph into “subgraphs” that will reside on a single shard. Think of the node id’s as keys.

For pointers within the same shard, collocate the child objects with their parent. Here some form of local key will be fine.

For pointers that lead to nodes in other chunks, hence other shards, we use true keys that tell us which shard has that chunk.

MODIFIED BEHAVIOR?

We would need to check: *is this node-id for a local node, or one in some other DHT shard?*

- If local, return a pointer to it.
- If remote, fetch it via a network RPC. Allocate memory (temporarily) to hold the copy. If the application modifies the object, write it back using another RPC. Free the memory when the access is finished.

WHAT WILL THIS COST?

Now every node access might require a DHT **get** operation.

We can do a bit better: if we run our tree logic on the same computers that run our DHT, then **local** get operations will be free. Only remote ones will be costly.

In Homework 2 we explore the value of recoding a tree search so that it runs “inside” a DHT instead of “outside” the DHT.

DHTS CAN BE “TOO GENERAL” FOR SOME USES

At Facebook, the early work on social networking used one shared DHT.

But the social networking graph was huge, very complicated, and very heavily used. Facebook decided that it was just not an efficient solution.

They redesigned it, and later in the course we will learn about their solution: Facebook TAO, a specialist for social network graph storage.

A WORRY ABOUT GARBAGE COLLECTION!

When we use a distributed solution, such as a DHT, we often need to make copies of objects. For example, server A uses `get` to fetch a copy of some node that resides on server B.

But these copies then occupy memory, and we need to free that memory when finished with the copy. Otherwise, memory will leak and the server will eventually crash.

Fortunately, modern languages automate this step (in C++, use shared pointer template for the same behavior).

WHAT ABOUT STORAGE IN THE DHT ITSELF?

A DHT doesn't necessarily know how to garbage collect the (key,value) objects stored into it!

It does keep track: each DHT object has a user-id (who created it), and the user's account gets charged for the space consumed.

- Benefits? Total control, plus keys only need to be unique for the user or the application, not across the whole cloud.
- Problem: Many people are careless about freeing up the space!

DHT CLEANUP SUGGESTIONS

Always give an expiration time for every (key,value) tuple. If you want to keep an object longer, use a longer expiration time, but not infinite.

In your code, try to explicitly delete any temporary content you load into a DHT. Put it there, run for a while, but then delete it.

Commercial products offer fancier and powerful features, such as “delete these objects when such-and-such a server shuts down.” Use them!

THE CLOUD CAN BE A VERY EXPENSIVE PLACE

When used properly, a cloud is often cheaper than owning hardware.

This is because you are sharing costs such as buying it and managing it with other users, and your share is potentially much lower.

But carelessness in storage management can leave all sorts of junk that might never be deleted, and you *will* be charged for it!

OTHER THINGS TO THINK ABOUT

Some DHT products allow you to control the hash function they will use. But this is not a standard thing – many do not.

Every DHT allows you to pick your own keys. And most DHTs tell you what hash function they will use, so you can pick keys “intelligently” if you wish.

Hot spots can be an issue. Even uniformly random inserts might be clumpy. And you can't know which keys will be queried: some may be very popular

ONE DHT? OR ONE PER APPLICATION?

Amazon and Azure both urge you to use their DHT products: AWS Dynamo and Azure Cosmos.

But this means that many applications and even many users would potentially share the same storage infrastructure!

Is this a bad thing?

ISSUES TO THINK ABOUT



The DHT server will be more efficient in its use of space if shared by many users, so it will be ecologically greener and hence cheaper to use.

It will also stay busy all the time. If we plan to own a server and power it up, keeping it busy makes a lot of sense. But a shared server could become a hot spot because of some other user who pounds on some specific DHT item and overloads that shard.

- Your performance would suffer, and yet you have no way to know why!

SHARING CAN CREATE SECURITY ISSUES

Another form of leakage arises if data from one application or one user becomes visible to some other application or user, without permissions.

A DHT with distinct key spaces shouldn't leak information, but there could be software bugs or even performance behaviors that actually do reveal sensitive content, unintentionally.

We will discuss this in a future lecture. It is not a huge risk, but it is worth being aware of it.

SUMMARY

Almost anything can be stored into a DHT. The cloud does this. But it isn't free: you need to be clever to encode your application into a DHT.

Think about keys, object sizes, access patterns, costs. Be wary of “leakage” (neglecting to delete temporary data) or you will get a BIG monthly bill!