

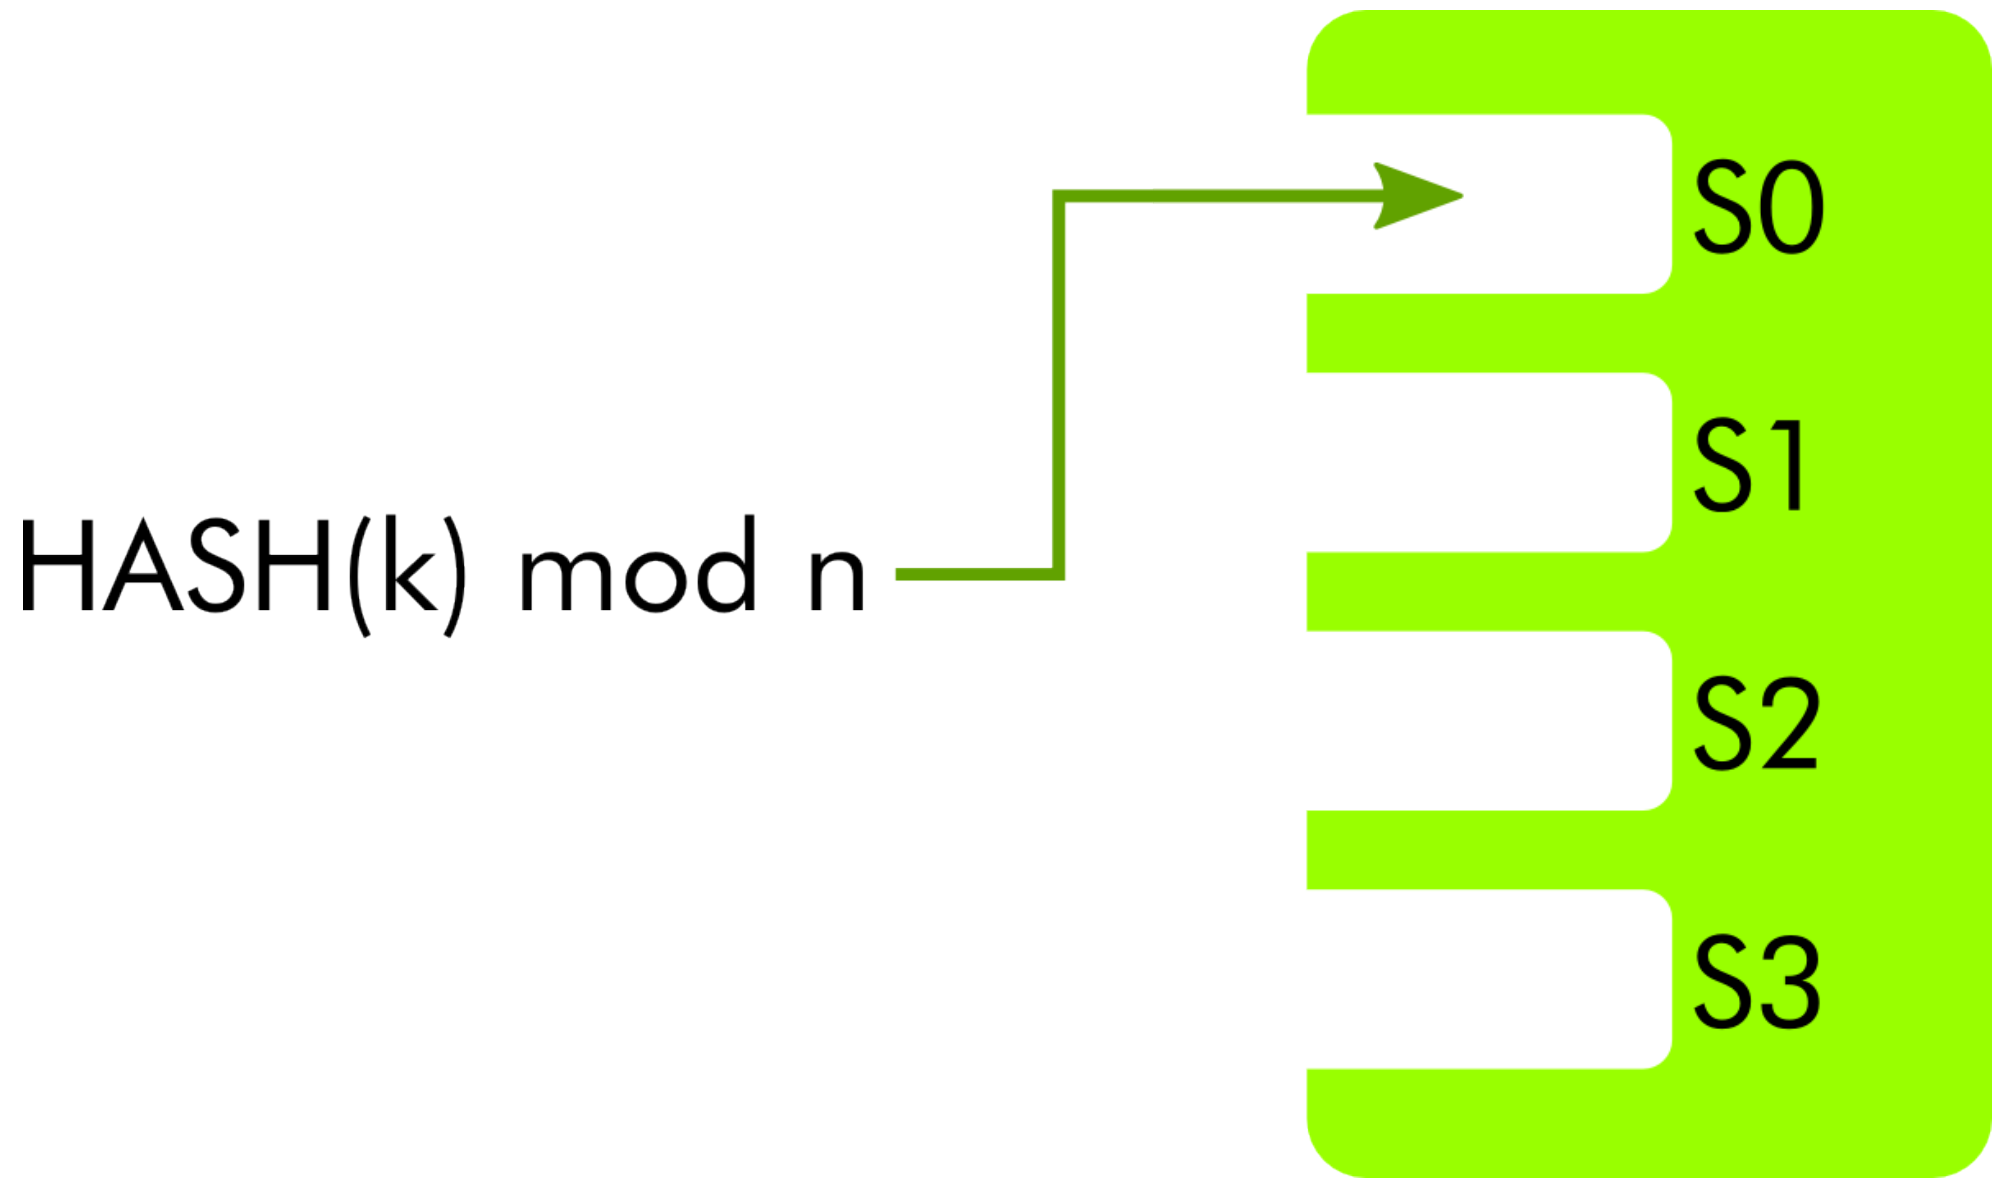
Riak and Consistent Hashing

John Leach
@johnleach
john@brightbox.com

Scaling vertically

Scaling horizontally

Hash table



$\text{HASH}(k) \bmod n$



S0

S1

S2

S3

S4

$\text{HASH}(k) \bmod n$



S0


S1


S2


S3


S4

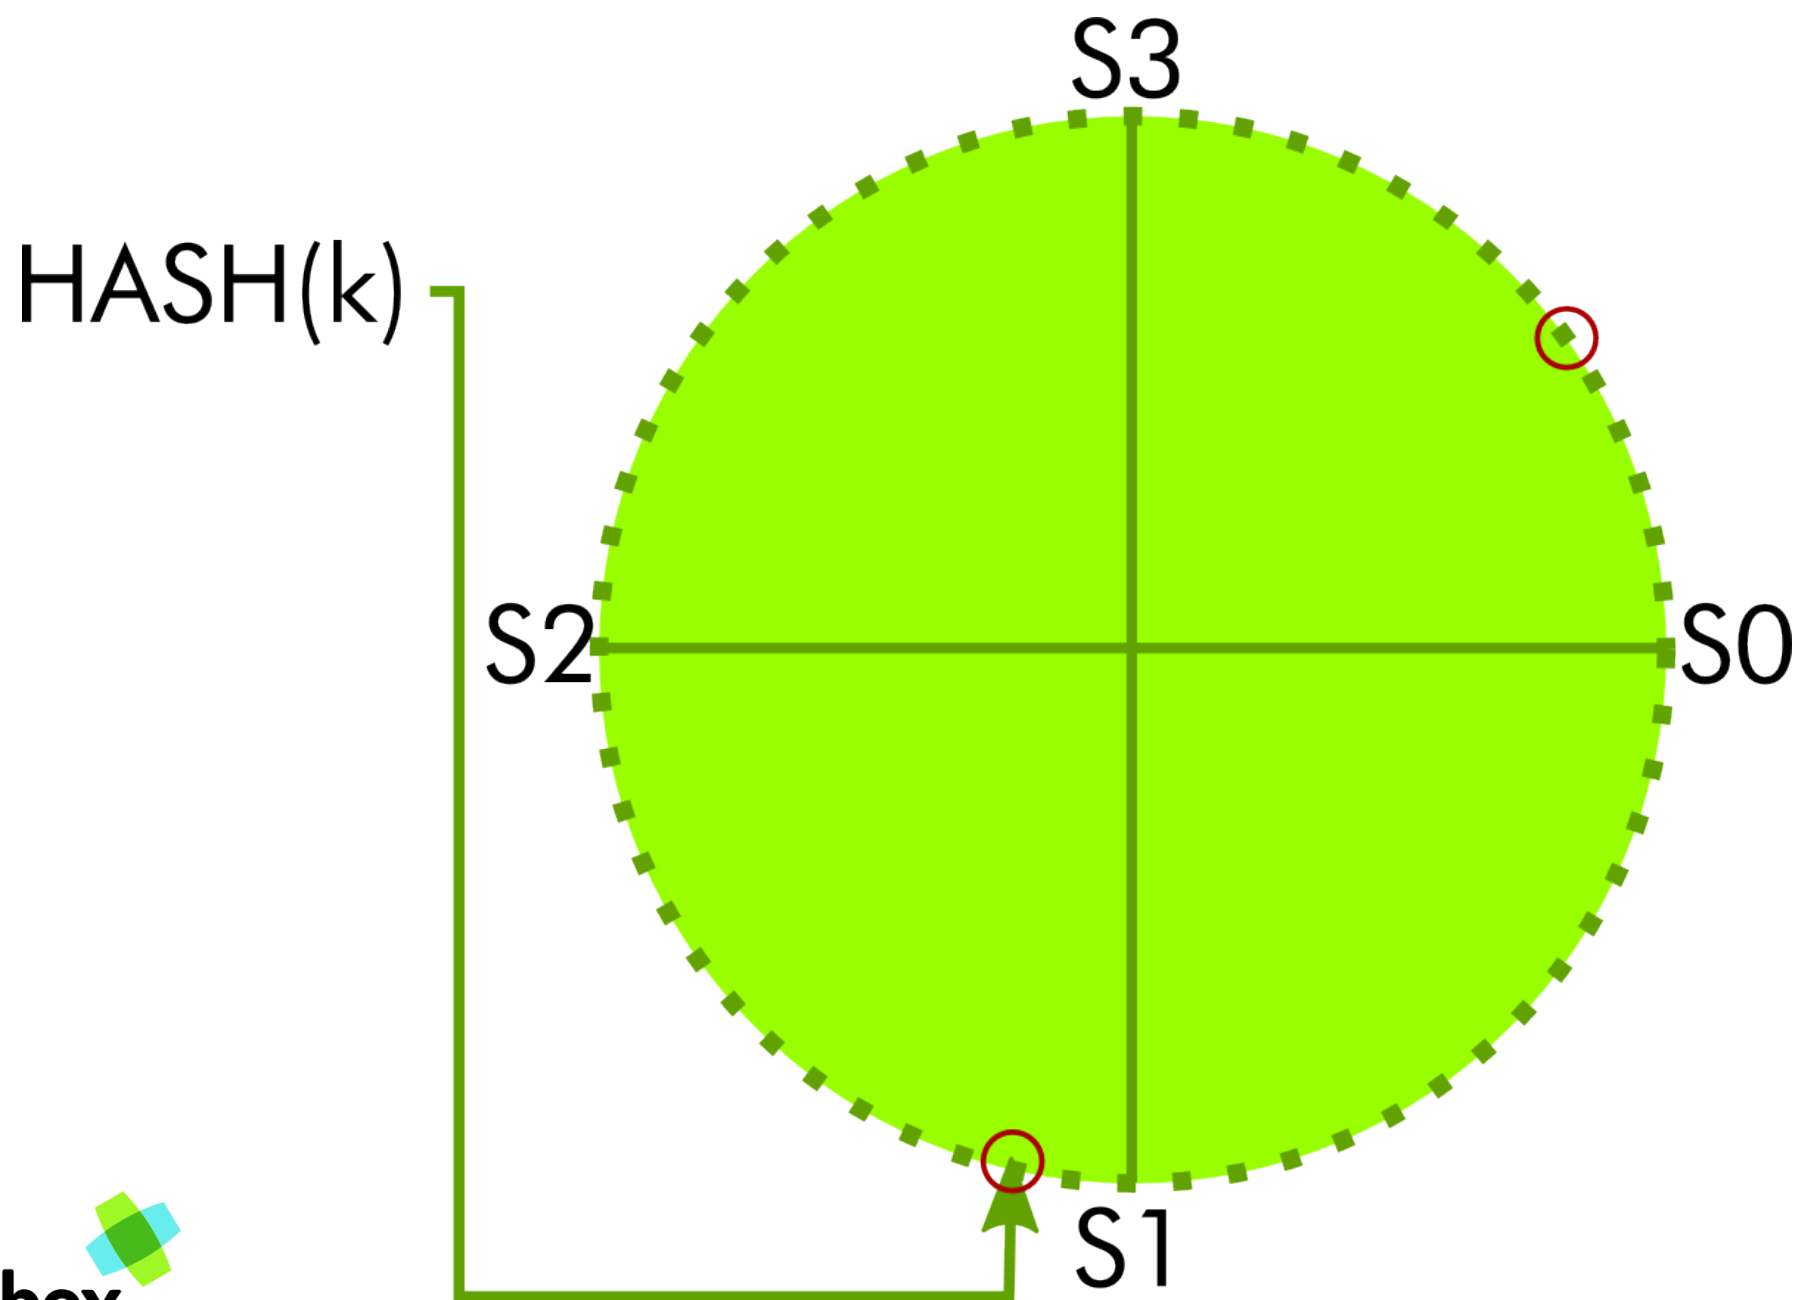
Consistent hash ring

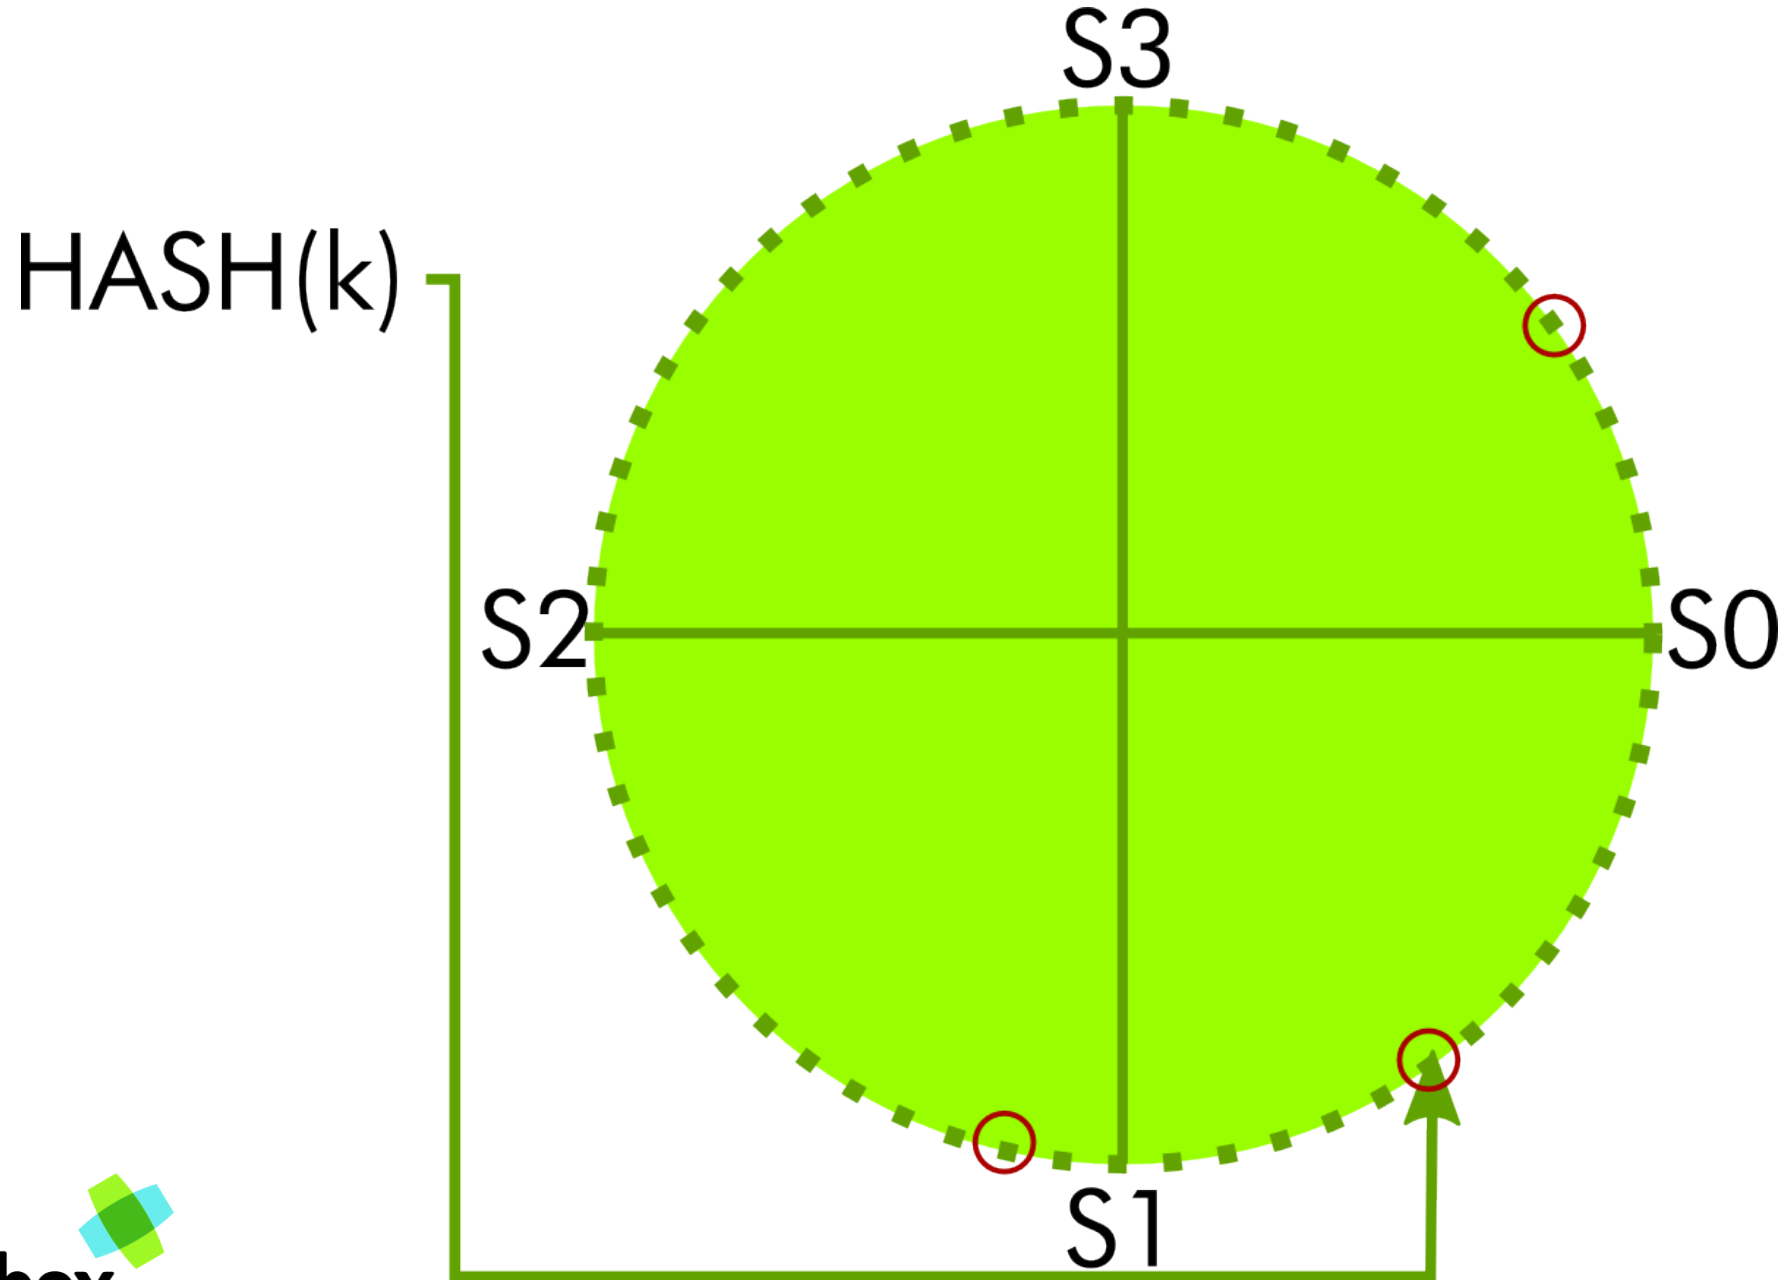
























Vnodes






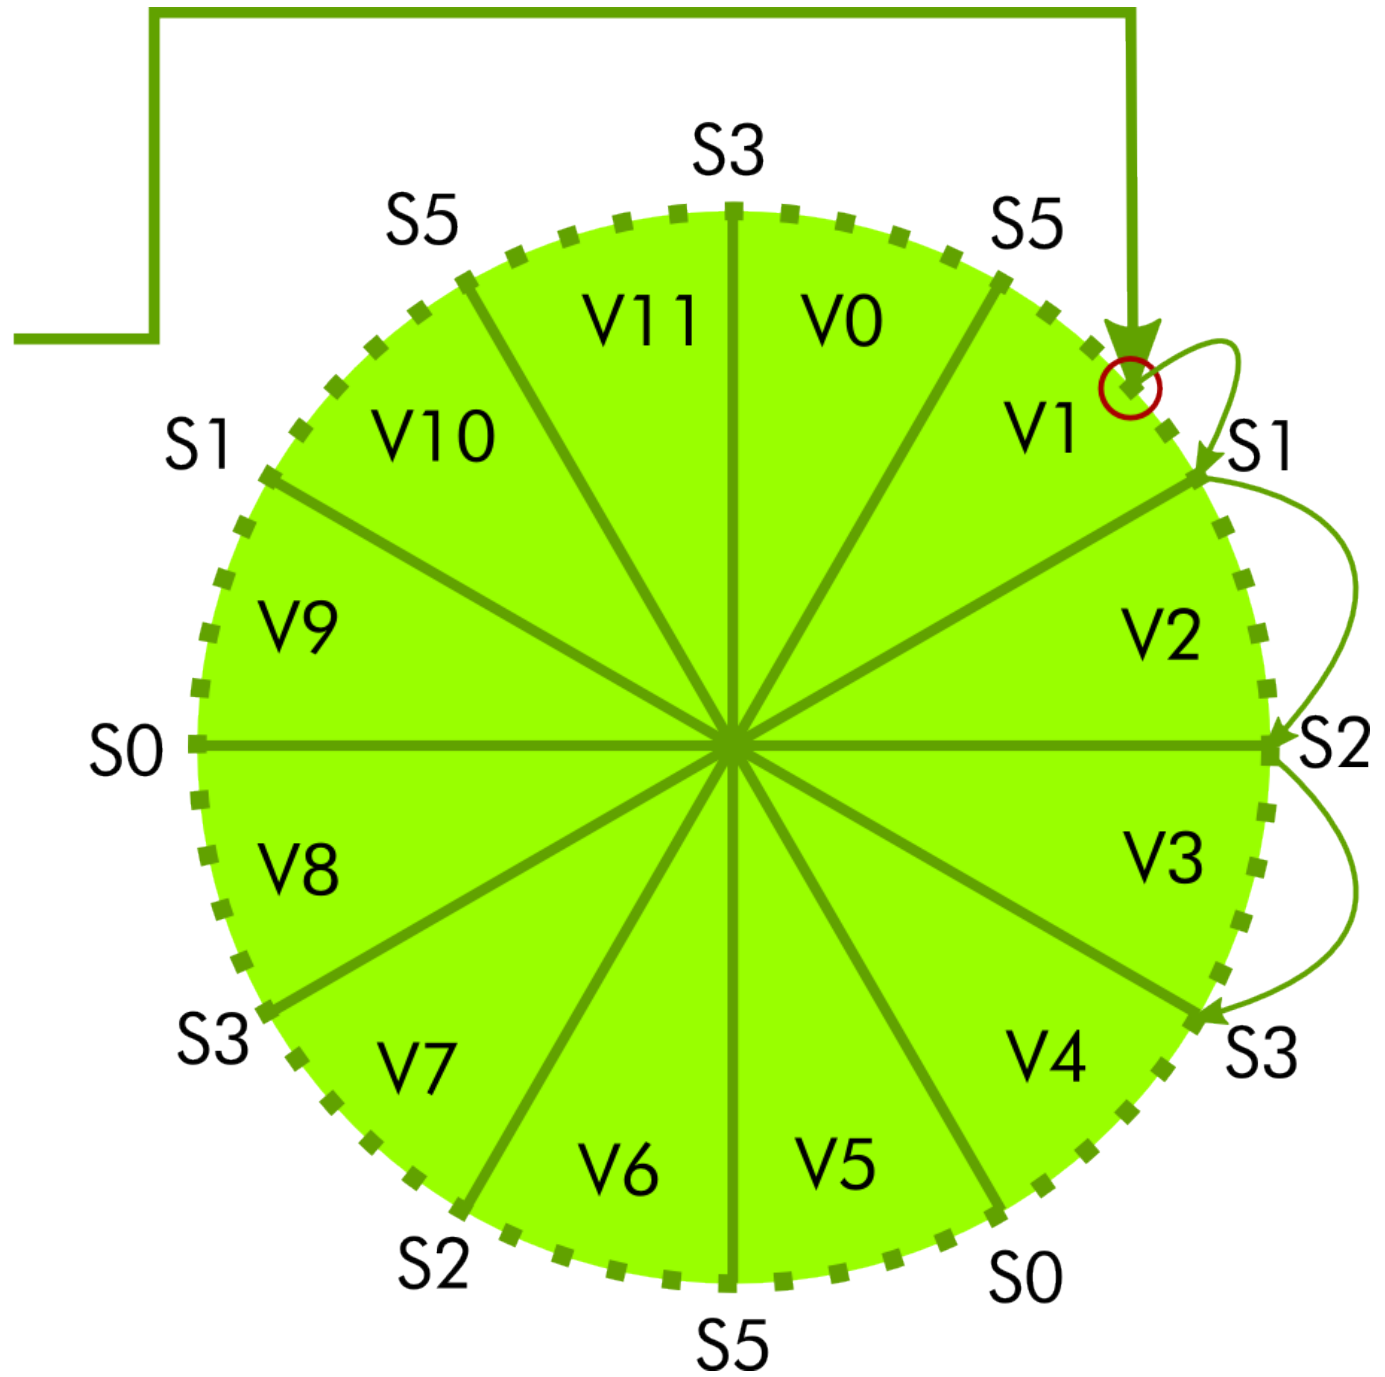


Replication

$\text{HASH}(k)$



$\text{HASH}(k)$



Consistent hashing

- Still have to migrate data
- Minimized the number of migrations
- Can't make intelligent placing decisions like a Hadoop NameNode
- Assumes uniform distribution of keys
- Doesn't solve uneven value size problem