wuala

P2P Online Storage
http://wua.la

CERN, June 2008

Dominik Grolimund

large, reliable, and secure
distributed online storage

harness idle resources of
participating computers

old dream of computer science

> “The design of a world-wide, fully transparent distributed file system for simultaneous use by millions of mobile and frequently disconnected users is left as an exercise for the reader.”

A. Tanenbaum, Distributed Operating System, 1995 lots of research projects

OceanStore (UC Berkeley)
Past (Microsoft Research)
CFS (MIT)

we were inspired by them

wanted to make it work

first step: closed alpha

upload any file in any size

access from anywhere
share with friends and groups
publish to the world

free and simple application
Win, Mac, Linux

start from the web,
no installation required

start with 1 GB provided by us

if you want more,
you can trade or buy storage

online storage

with the “power of P2P”

fast downloads

no file size limit

no traffic limit

privacy

all files are encrypted on your computer

your password never leaves your computer

so no one, not even we, can see your files

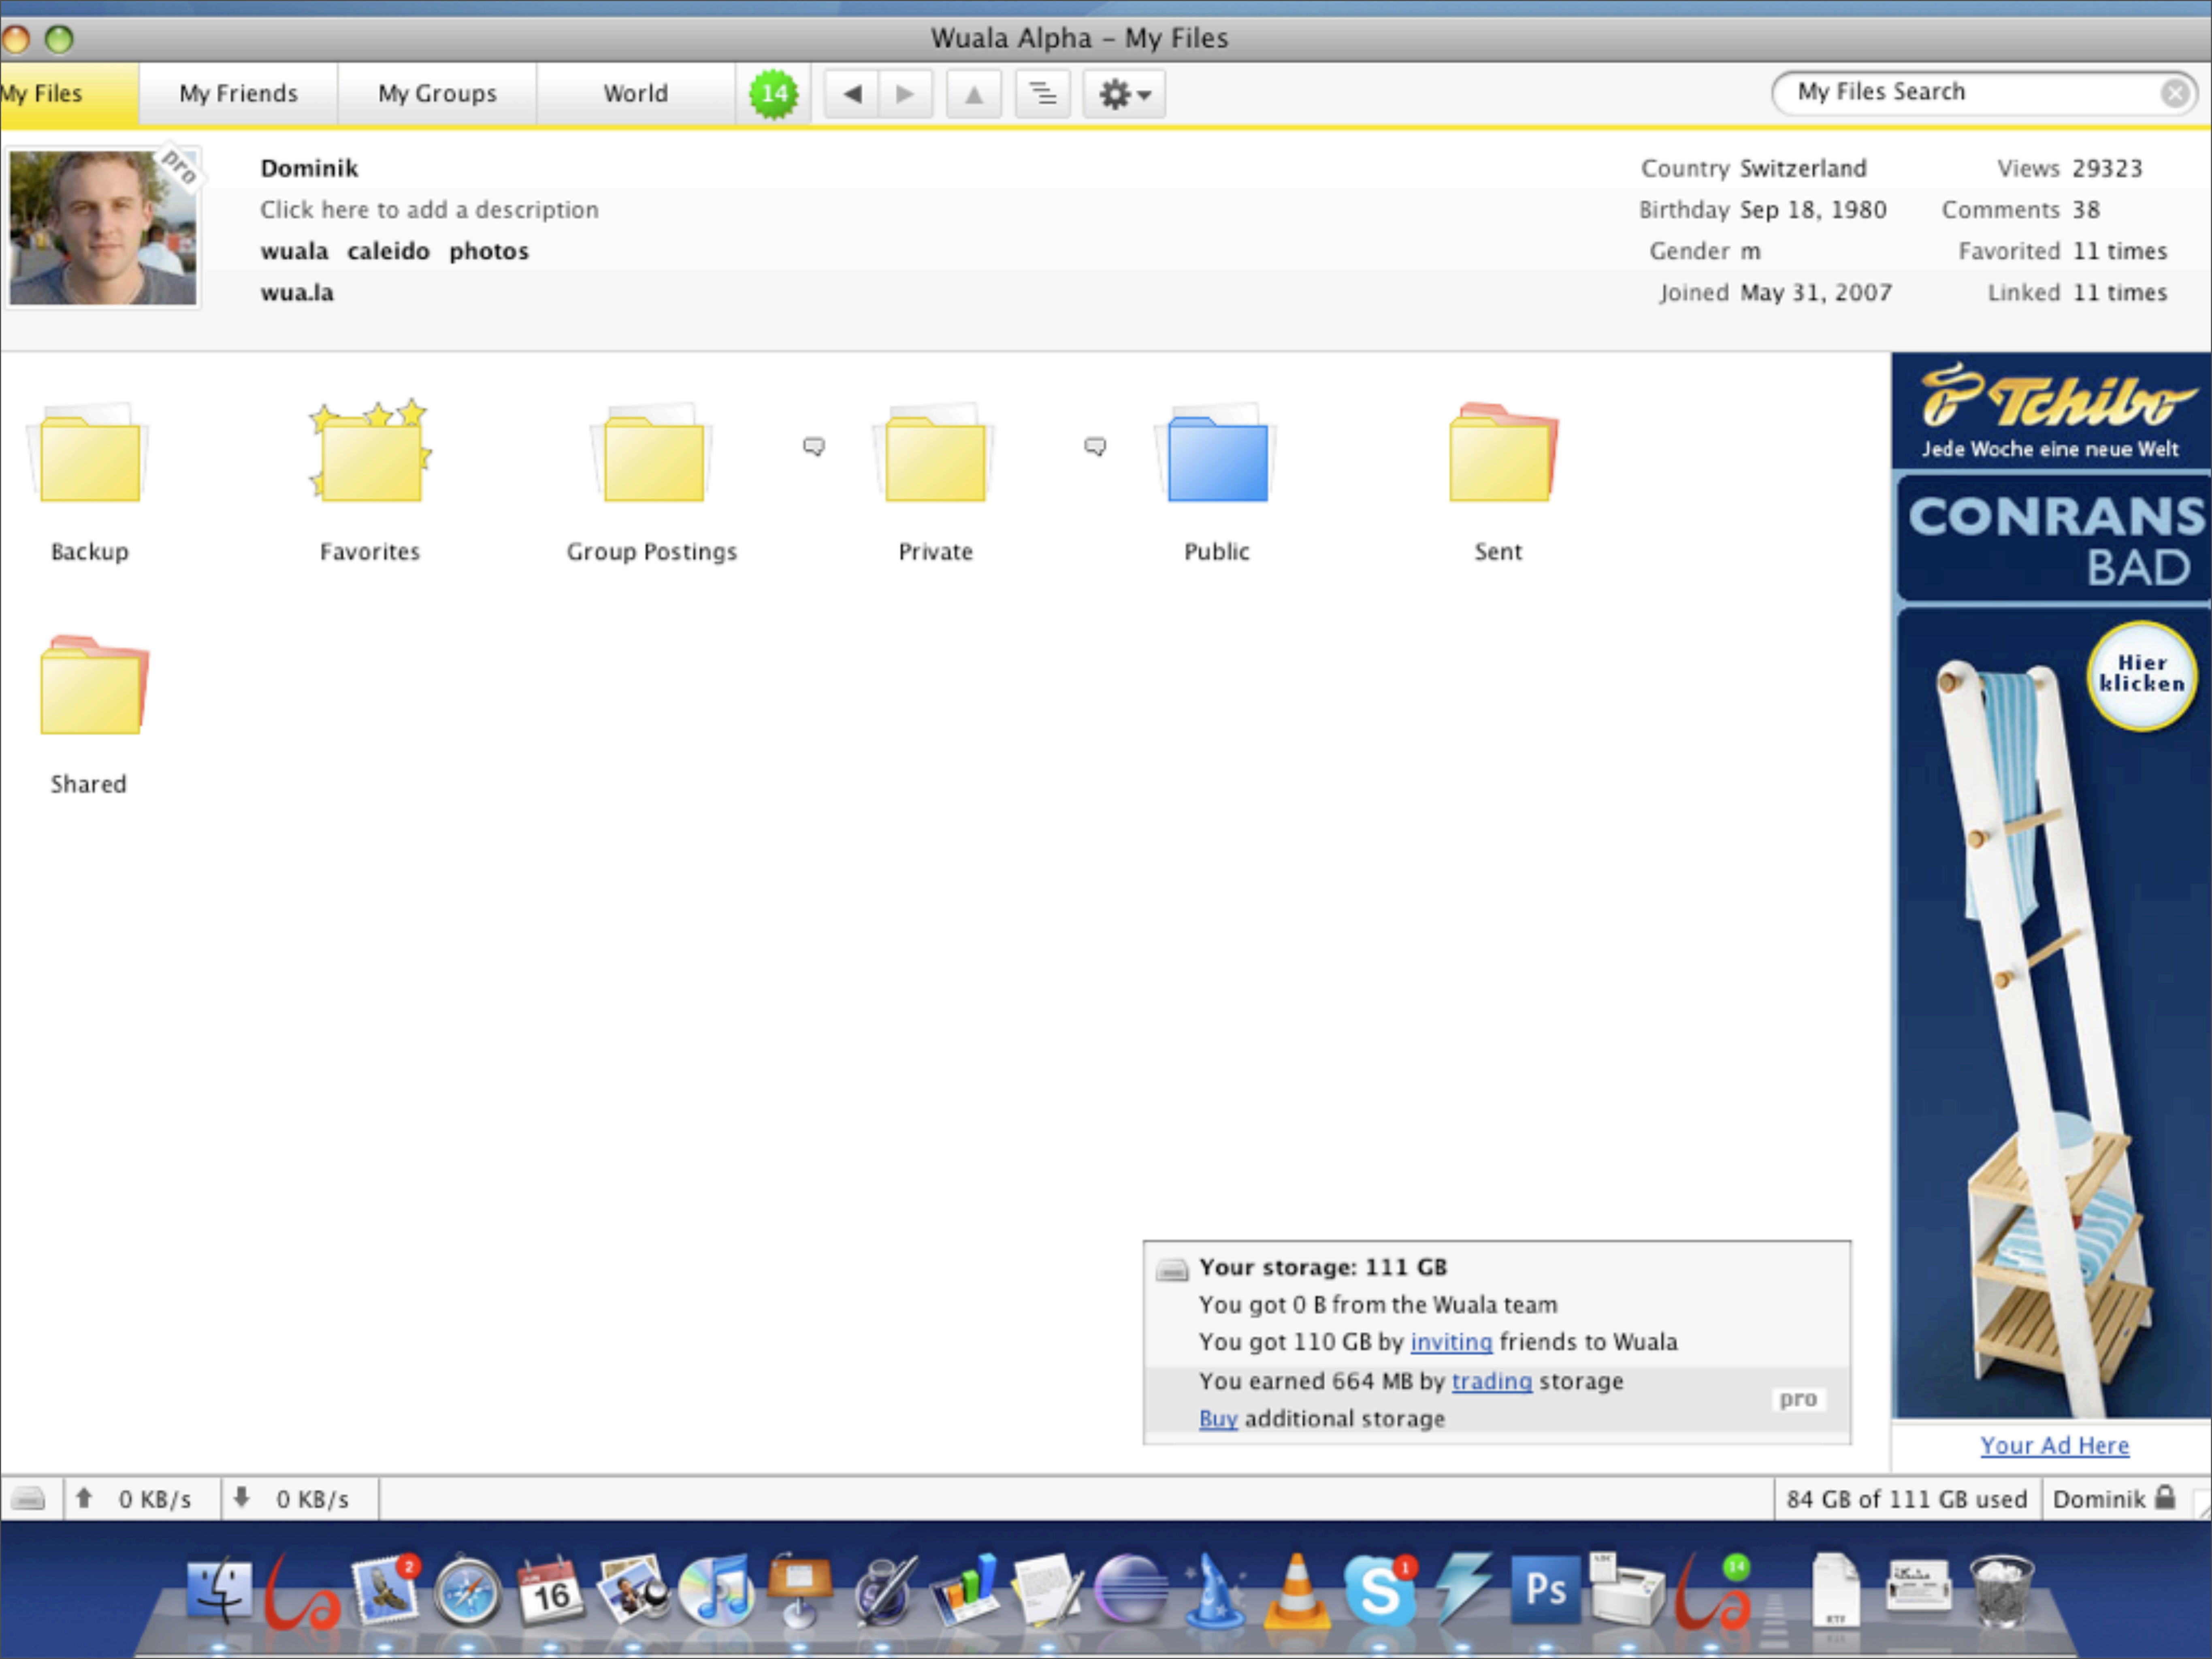

Dominik
Click here to add a description
wuala caleido photos
wua.la

Country	Switzerland	Views	29323
Birthday	Sep 18, 1980	Comments	38
Gender	m	Favorited	11 times
Joined	May 31, 2007	Linked	11 times

Your storage: 111 GB
You got 0 B from the Wuala team
You got 110 GB by inviting friends to Wuala
You earned 664 MB by trading storage pro
Buy additional storage

0 KB/s 0 KB/s 84 GB of 111 GB used Dominik

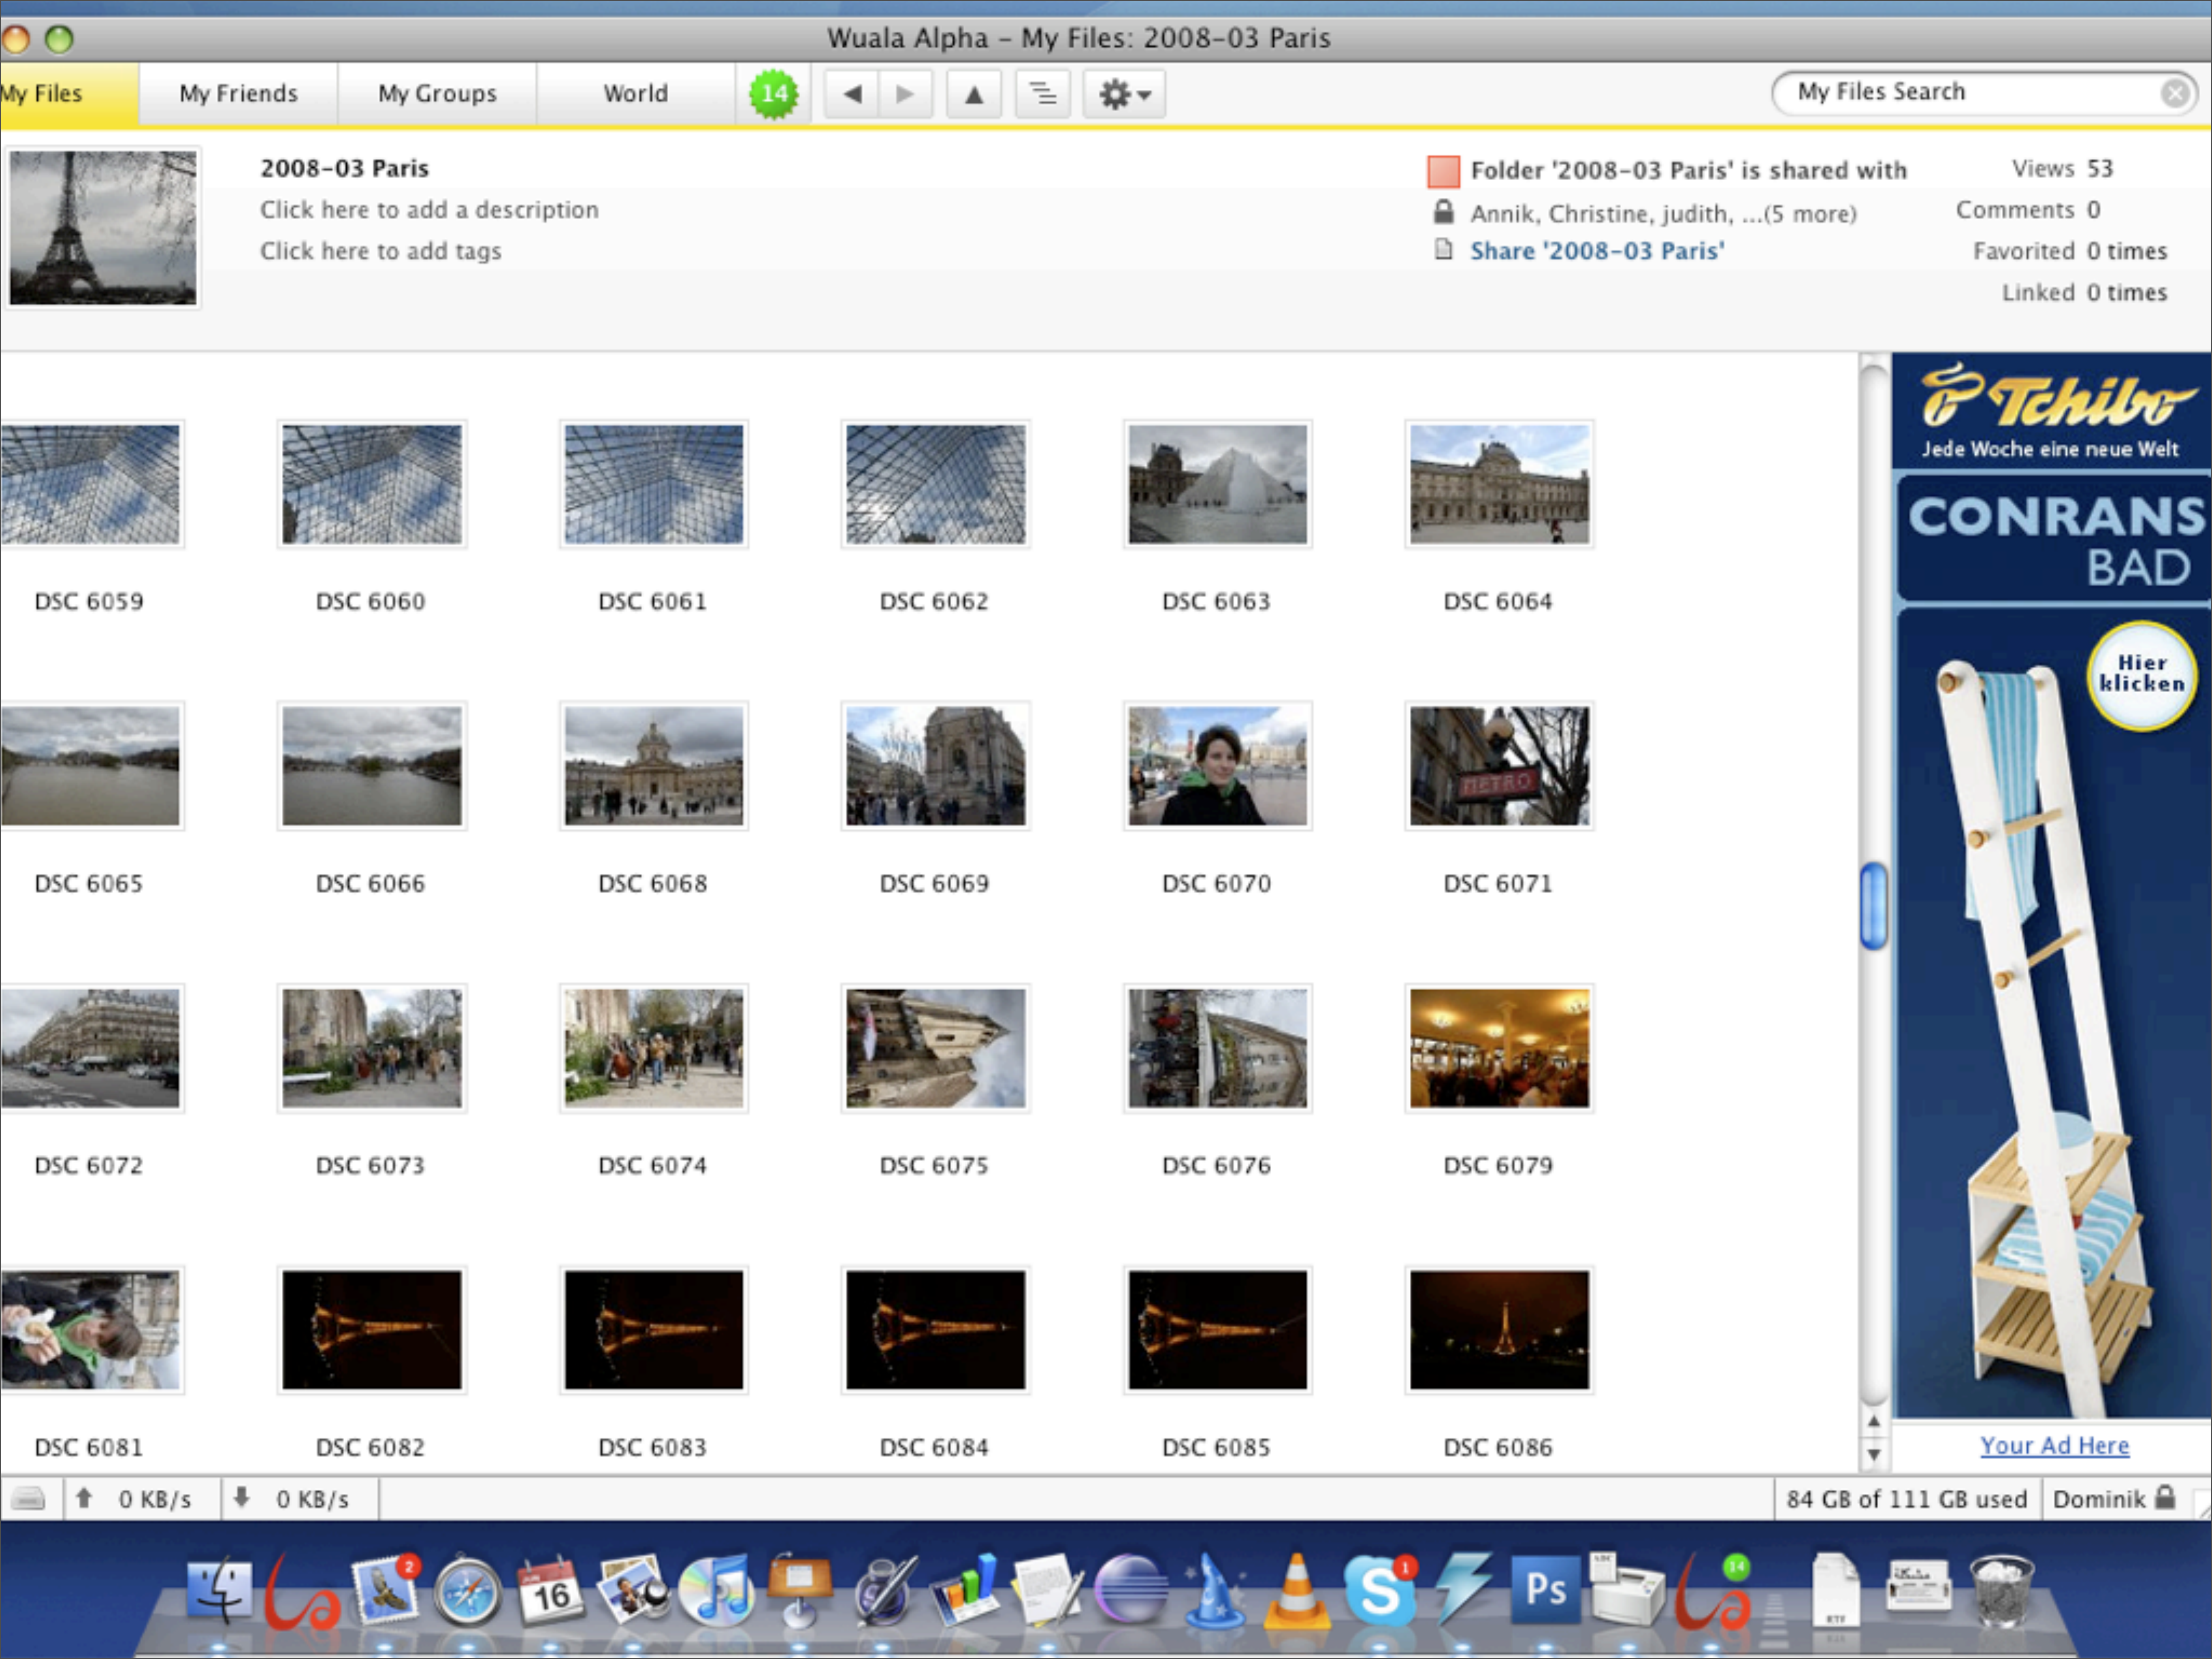

Wuala Alpha – My Files: 2008-03 Paris

2008-03 Paris

Click here to add a description

Click here to add tags

Folder '2008-03 Paris' is shared with

Annik, Christine, judith, ...(5 more)

Share '2008-03 Paris'

Views	53
Comments	0
Favorited	0 times
Linked	0 times

DSC 6059 DSC 6060 DSC 6061 DSC 6062 DSC 6063 DSC 6064

DSC 6065 DSC 6066 DSC 6068 DSC 6069 DSC 6070 DSC 6071

DSC 6072 DSC 6073 DSC 6074 DSC 6075 DSC 6076 DSC 6079

DSC 6081 DSC 6082 DSC 6083 DSC 6084 DSC 6085 DSC 6086

0 KB/s 0 KB/s

84 GB of 111 GB used

Dominik

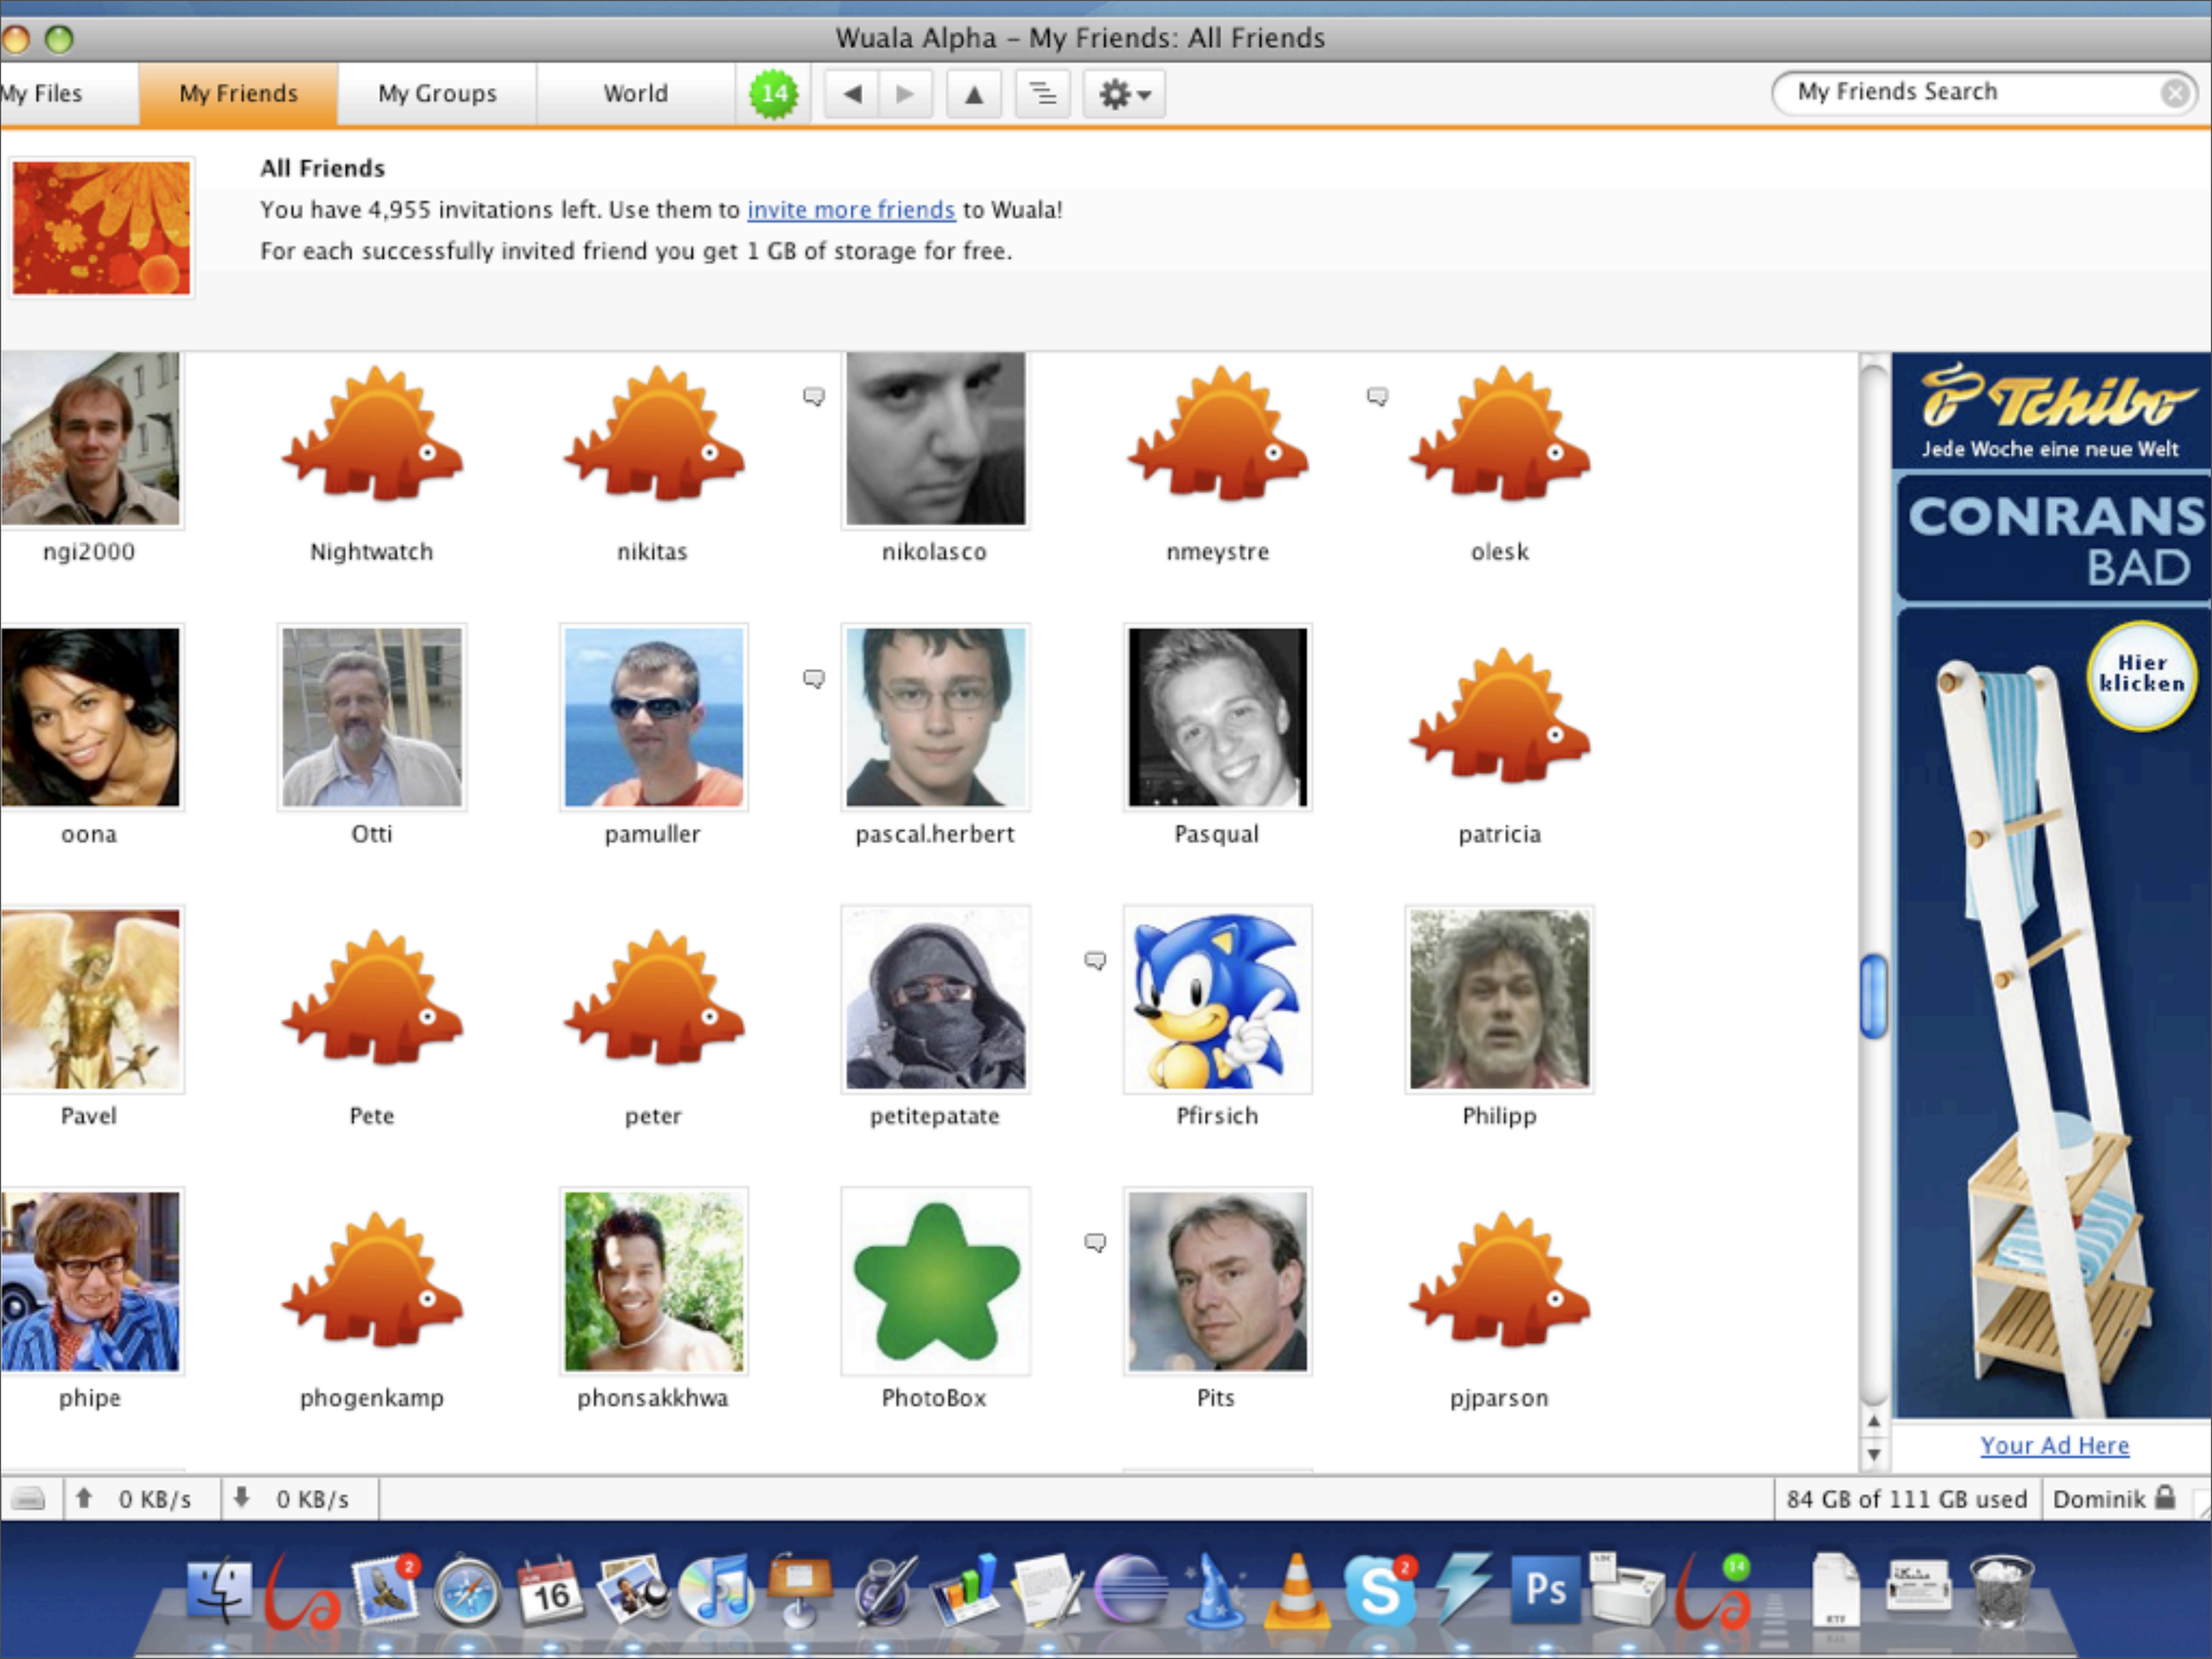

All Friends

You have 4,955 invitations left. Use them to invite more friends to Wuala!

For each successfully invited friend you get 1 GB of storage for free.

ngi2000 Nightwatch nikitas nikolasco nmeystre olesk

oona Otti pamuller pascal.herbert Pasqual patricia

Pavel Pete peter petitepatate Pfirsich Philipp

phipe phogenkamp phonsakkhwa PhotoBox Pits pjparson

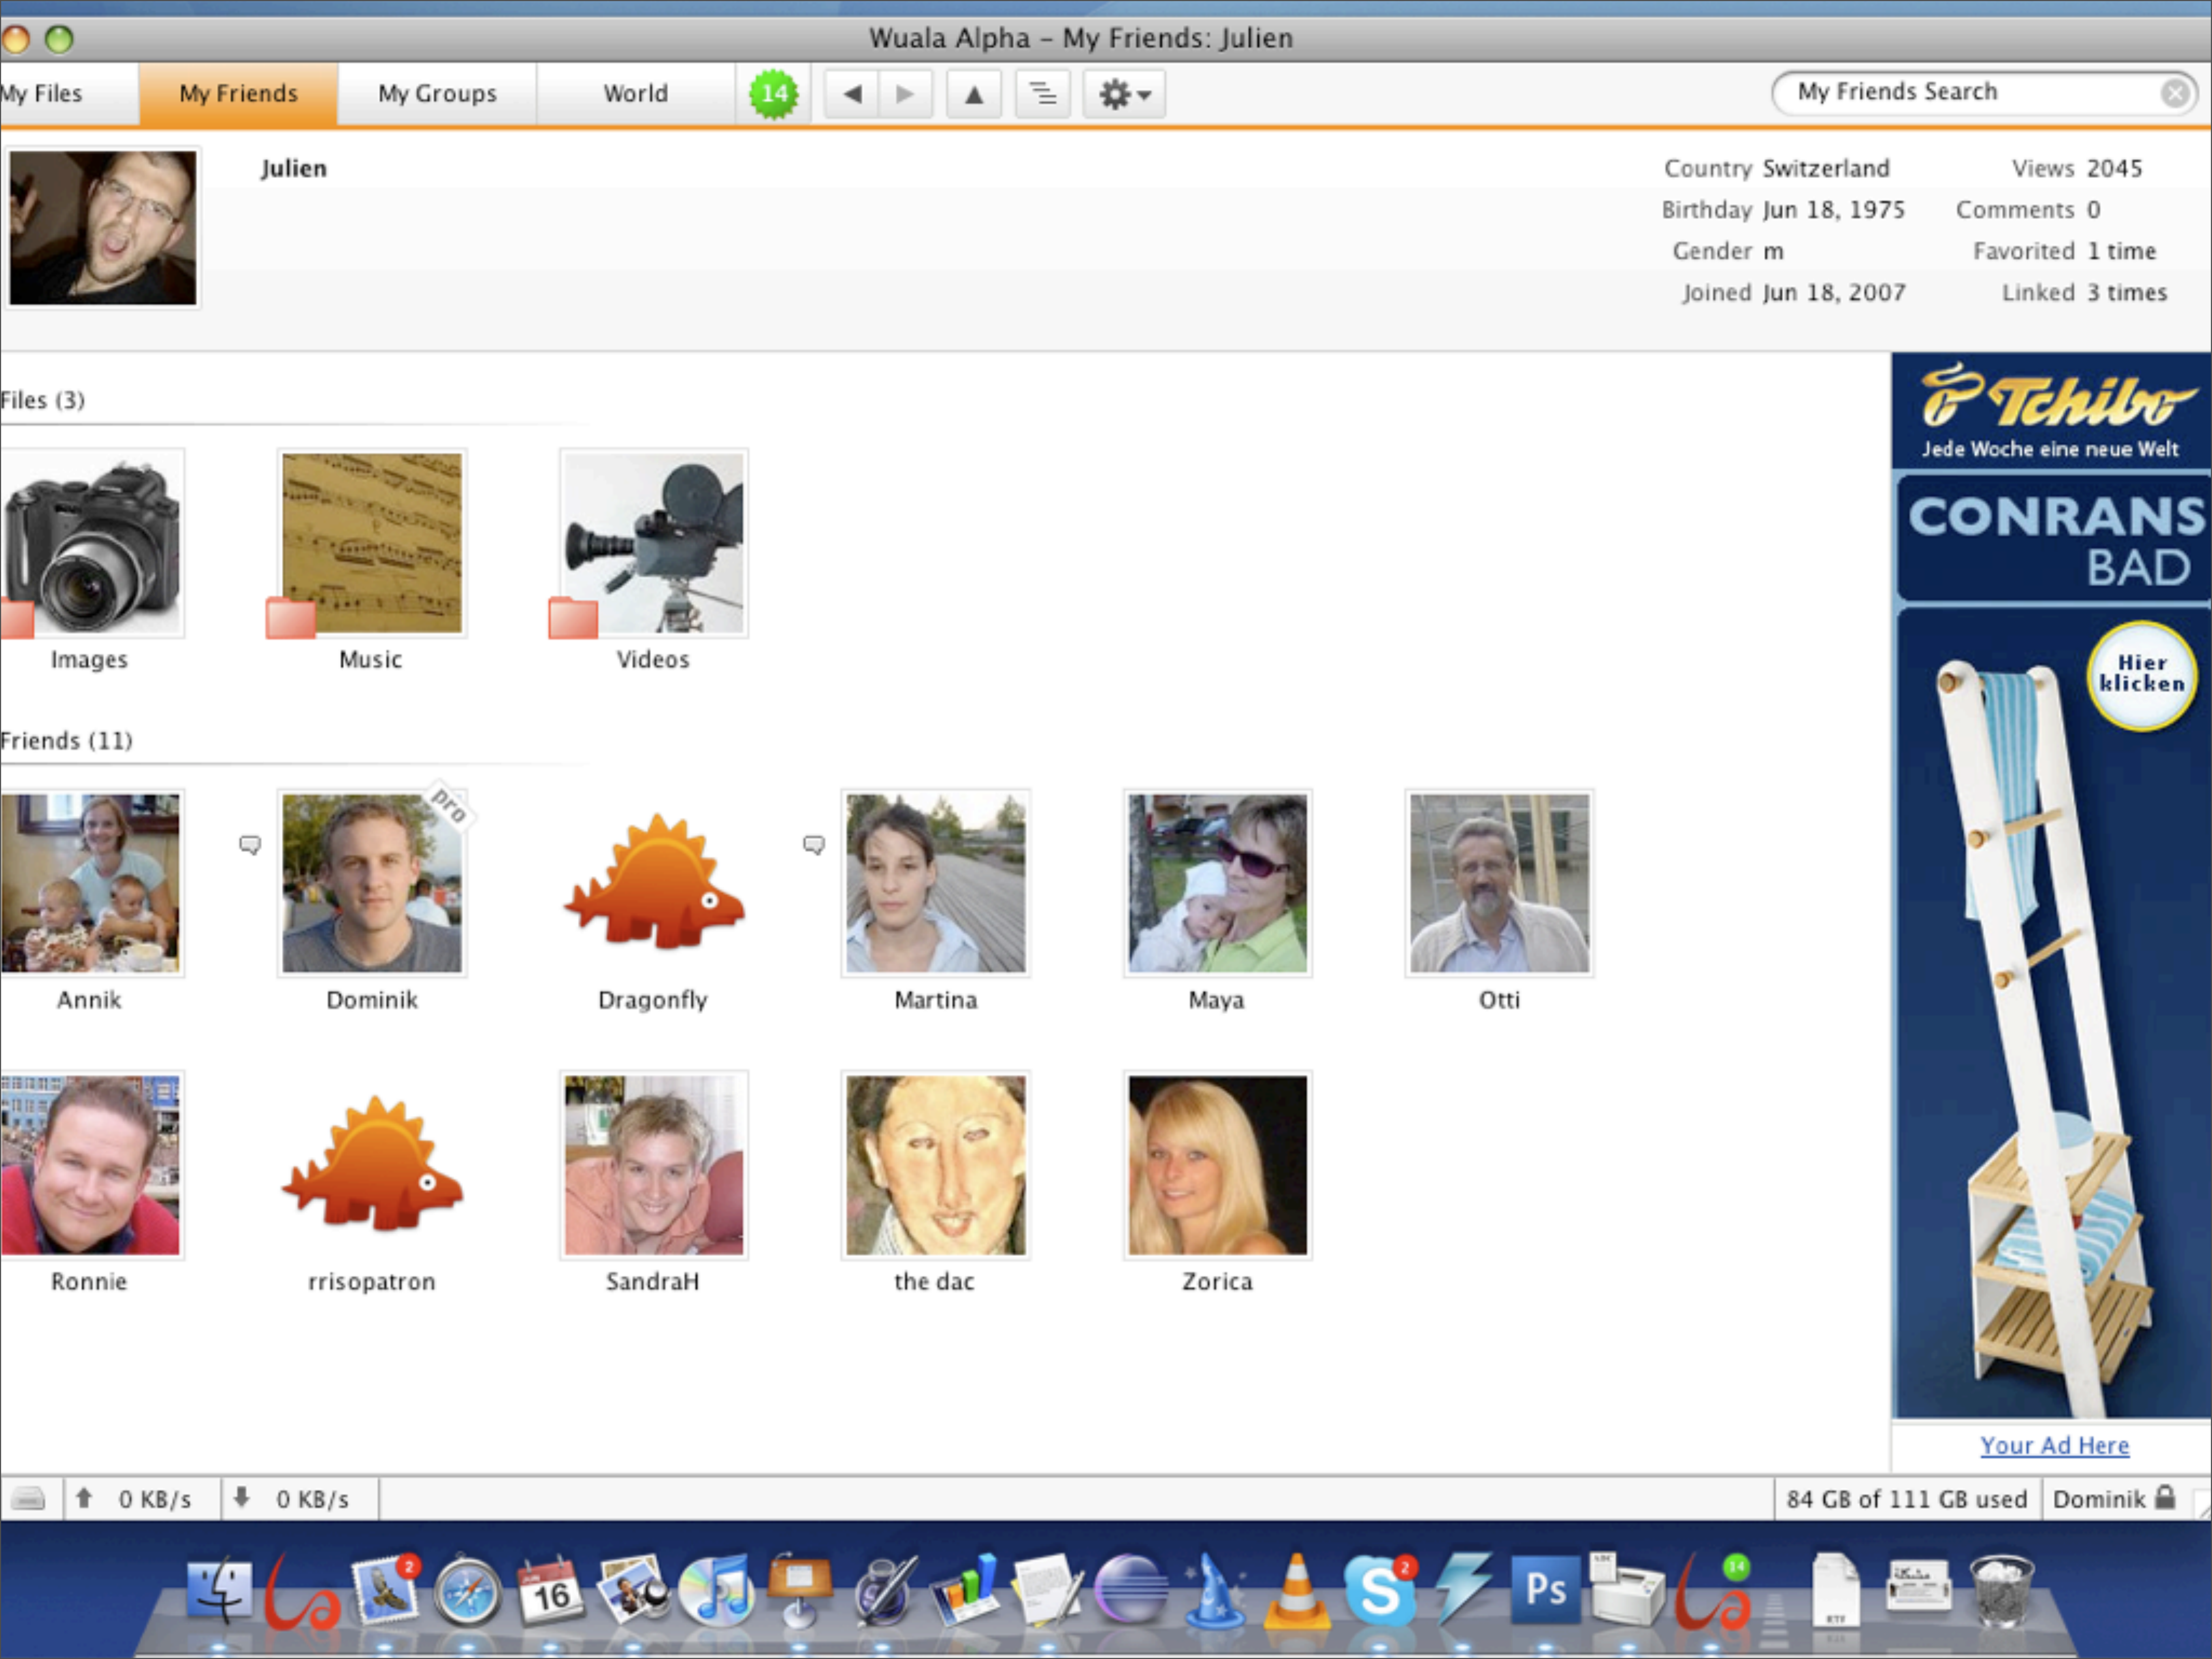

Julien

Country	Switzerland	Views	2045
Birthday	Jun 18, 1975	Comments	0
Gender	m	Favorited	1 time
Joined	Jun 18, 2007	Linked	3 times

Files (3)

Images

Music

Videos

Friends (11)

Annik

Dominik

Dragonfly

Martina

Maya

Otti

Ronnie

rrisopatron

SandraH

the dac

Zorica

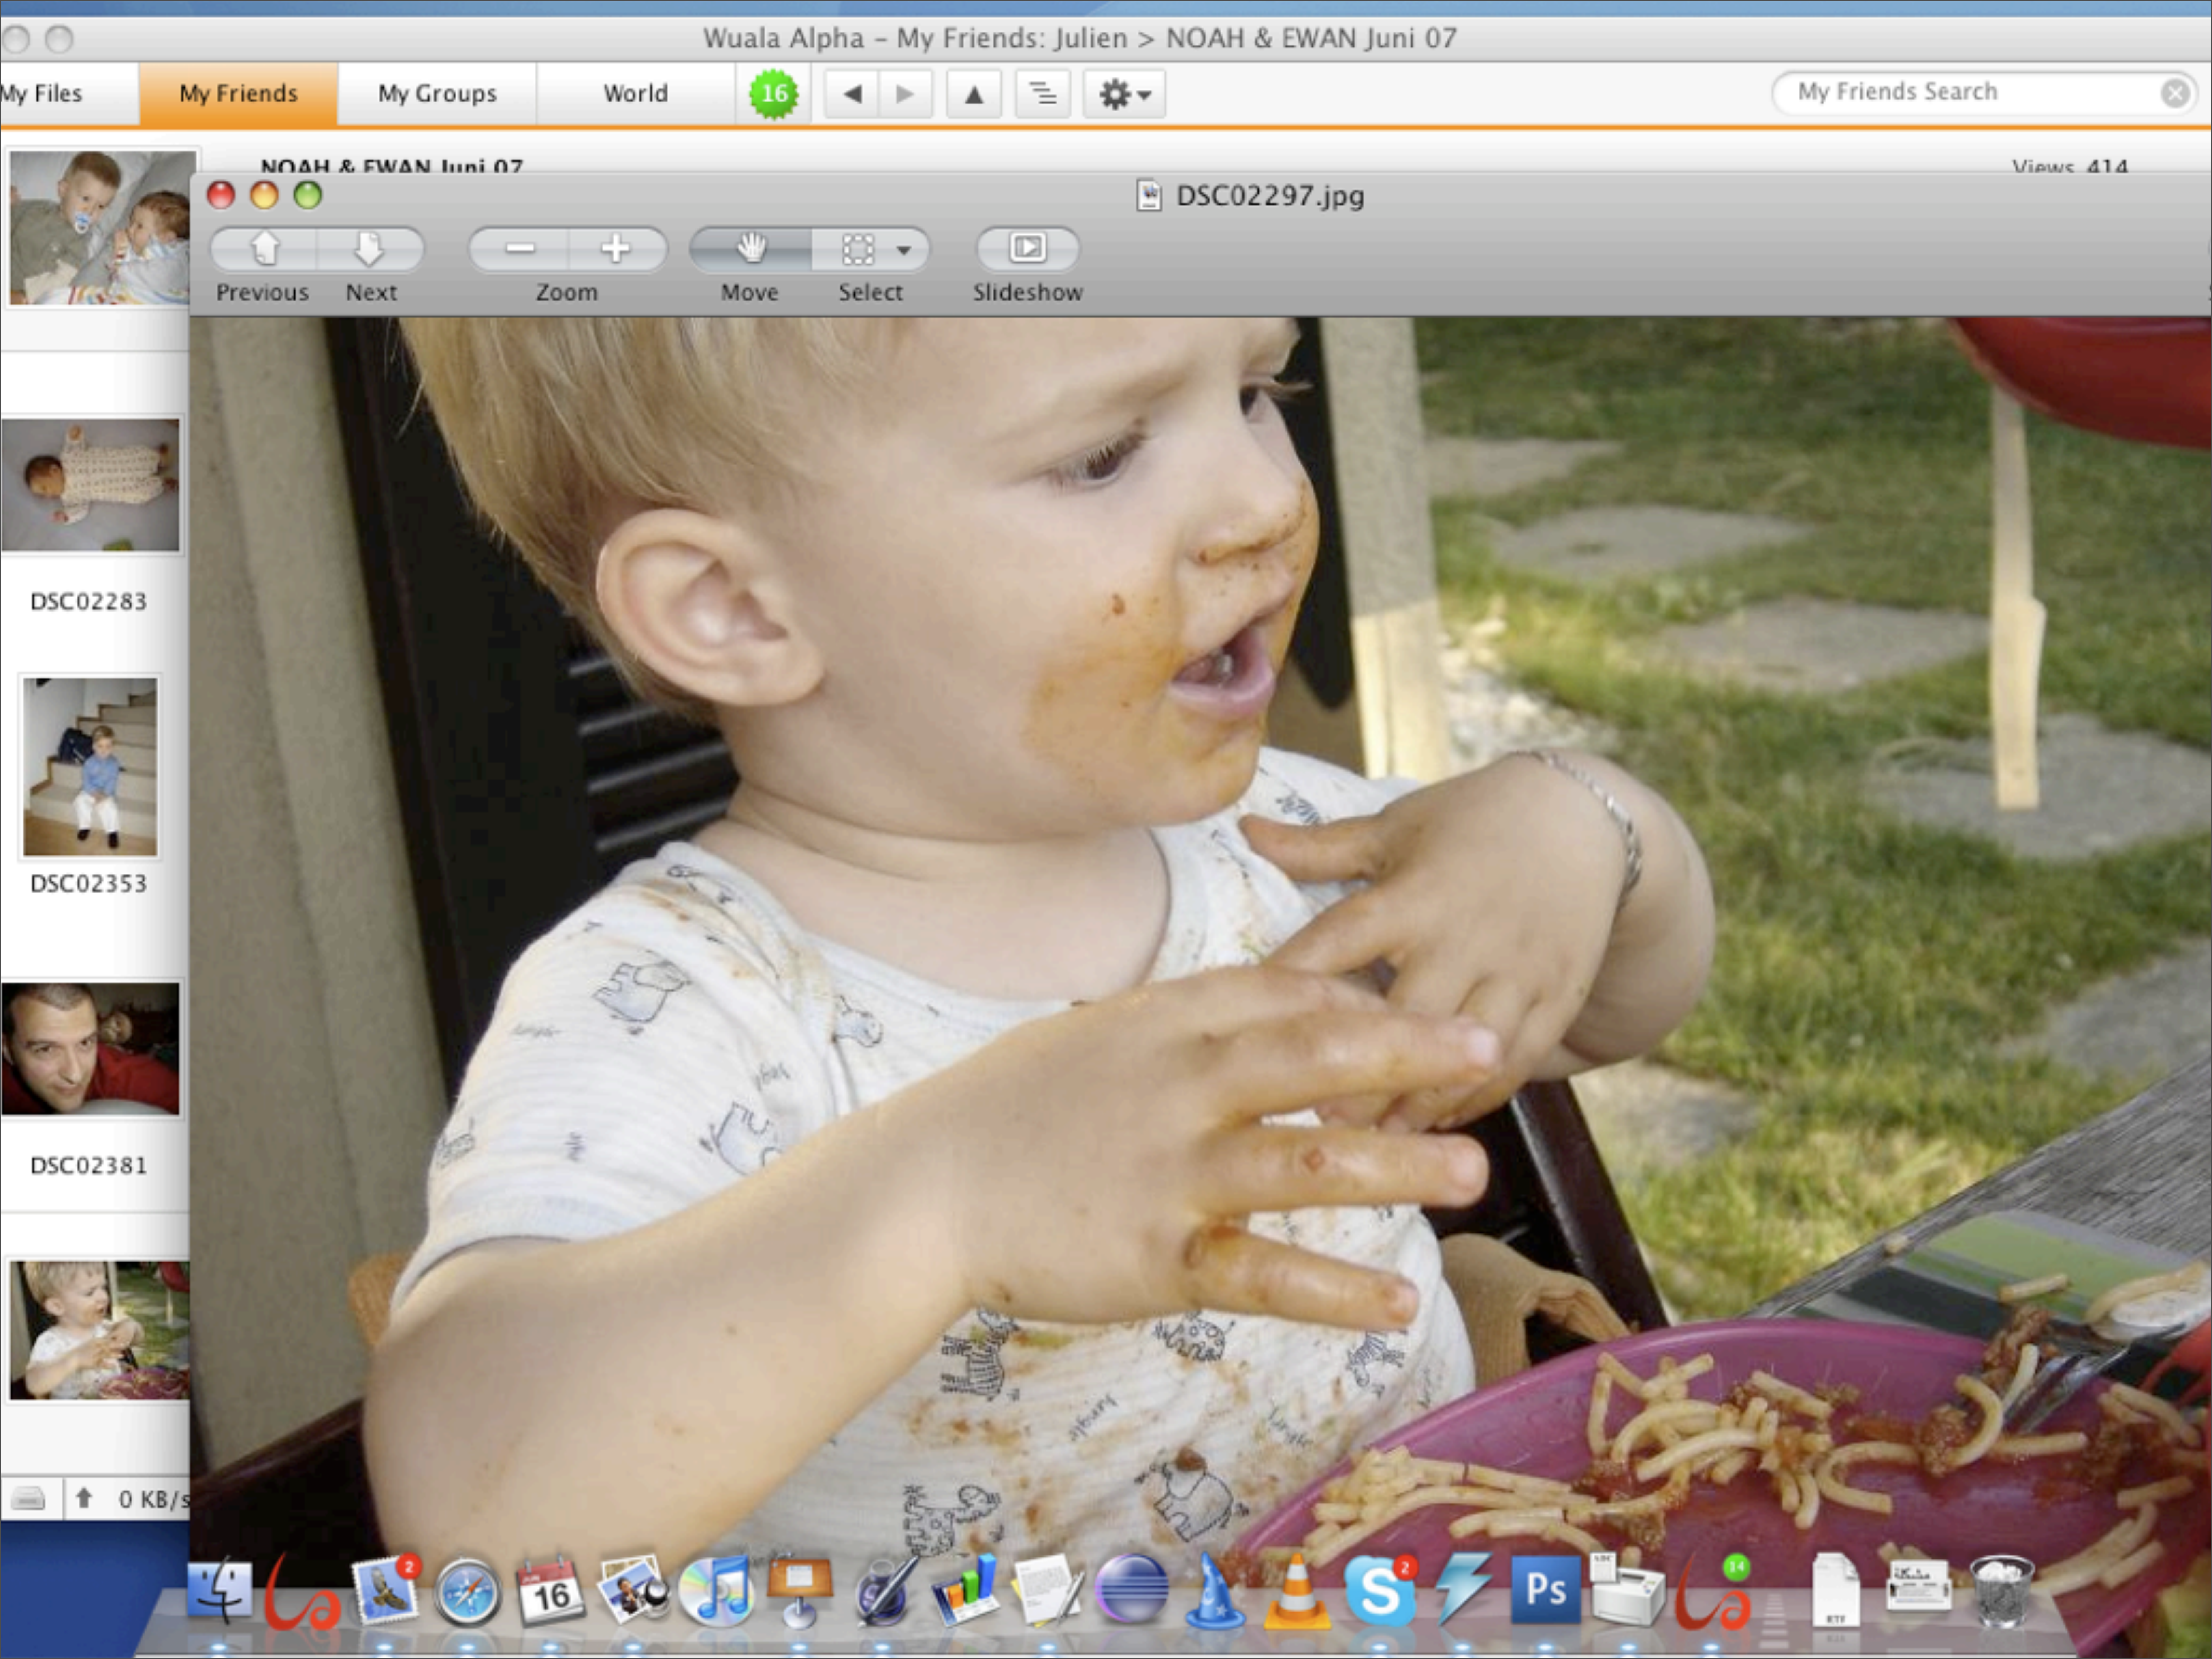
Wuala Alpha – My Friends: Julien > NOAH & EWAN Juni 07
My Files
My Friends
My Groups
World
16
My Friends Search
NOAH & EWAN Juni 07
Views 414
DSC02297.jpg
Previous
Next
Zoom
Move
Select
Slideshow
DSC02283
DSC02353
DSC02381
0 KB/s

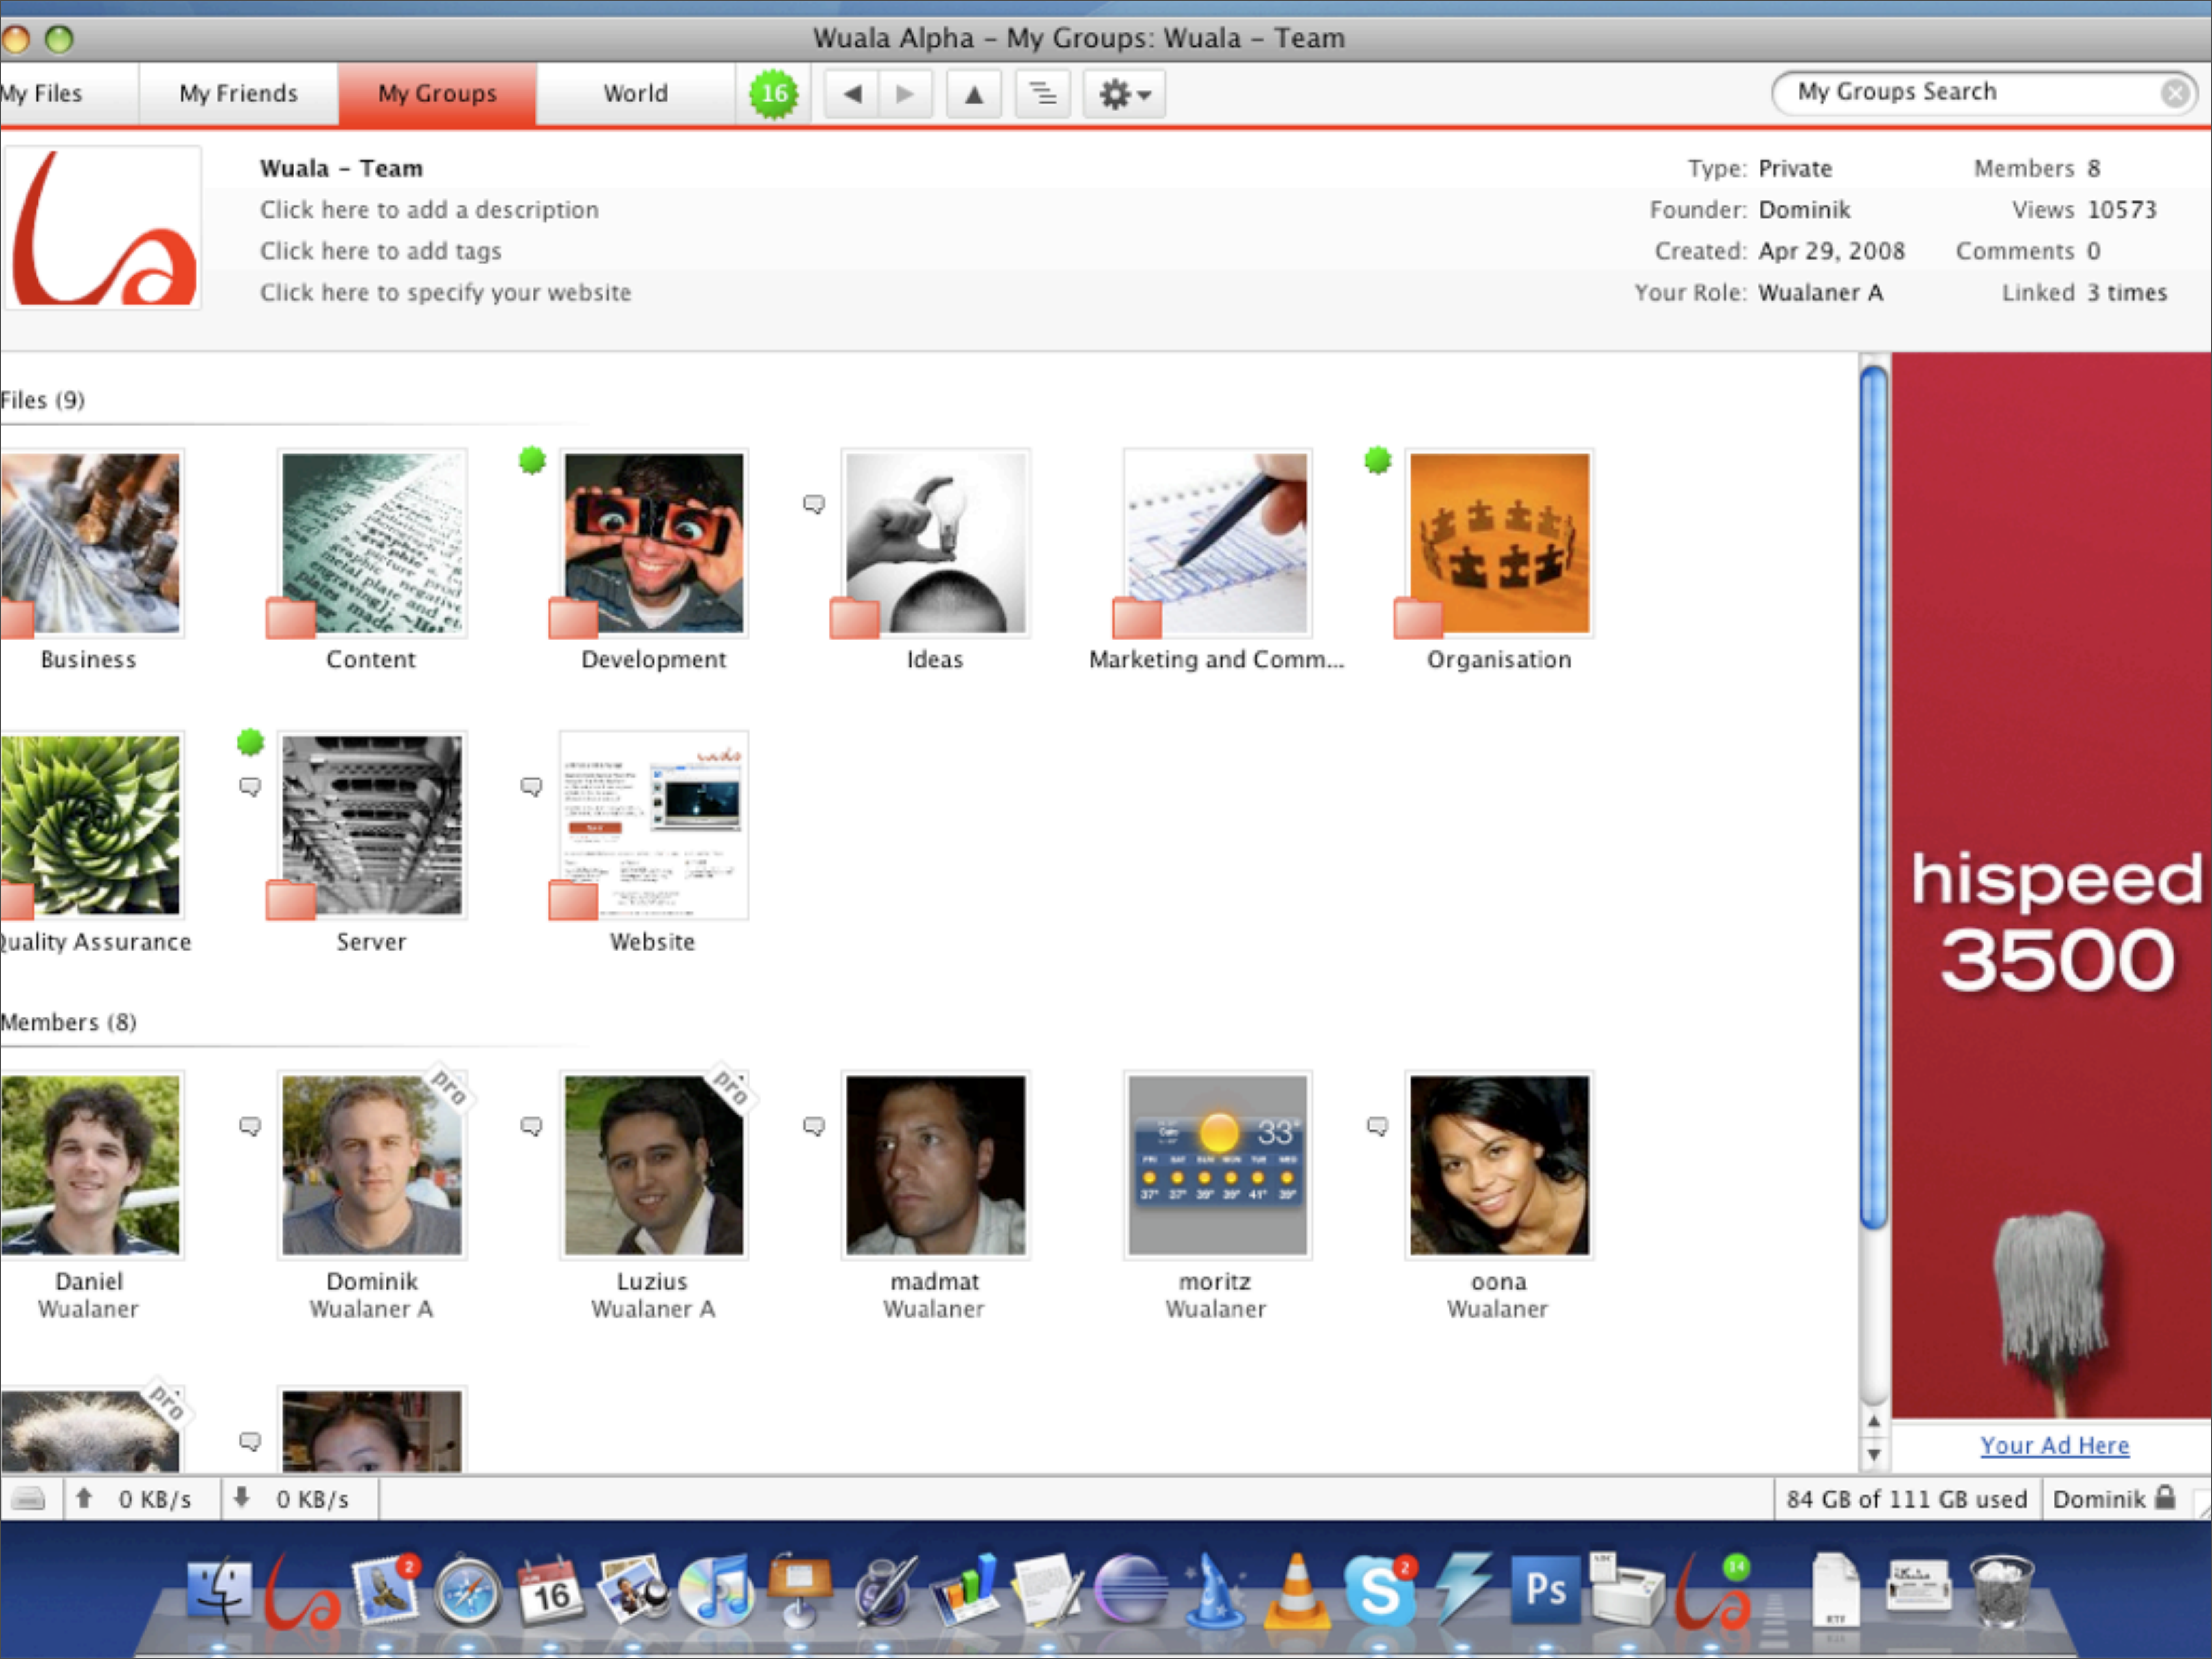

Wuala Alpha – My Groups: Wuala – Team

Wuala – Team

Click here to add a description

Click here to add tags

Click here to specify your website

Type:	Private	Members	8
Founder:	Dominik	Views	10573
Created:	Apr 29, 2008	Comments	0
Your Role:	Wualaner A	Linked	3 times

Files (9)

Members (8)

Daniel	Wualaner
Dominik	Wualaner A
Luzius	Wualaner A
madmat	Wualaner
moritz	Wualaner
oona	Wualaner

0 KB/s

0 KB/s

84 GB of 111 GB used

Dominik

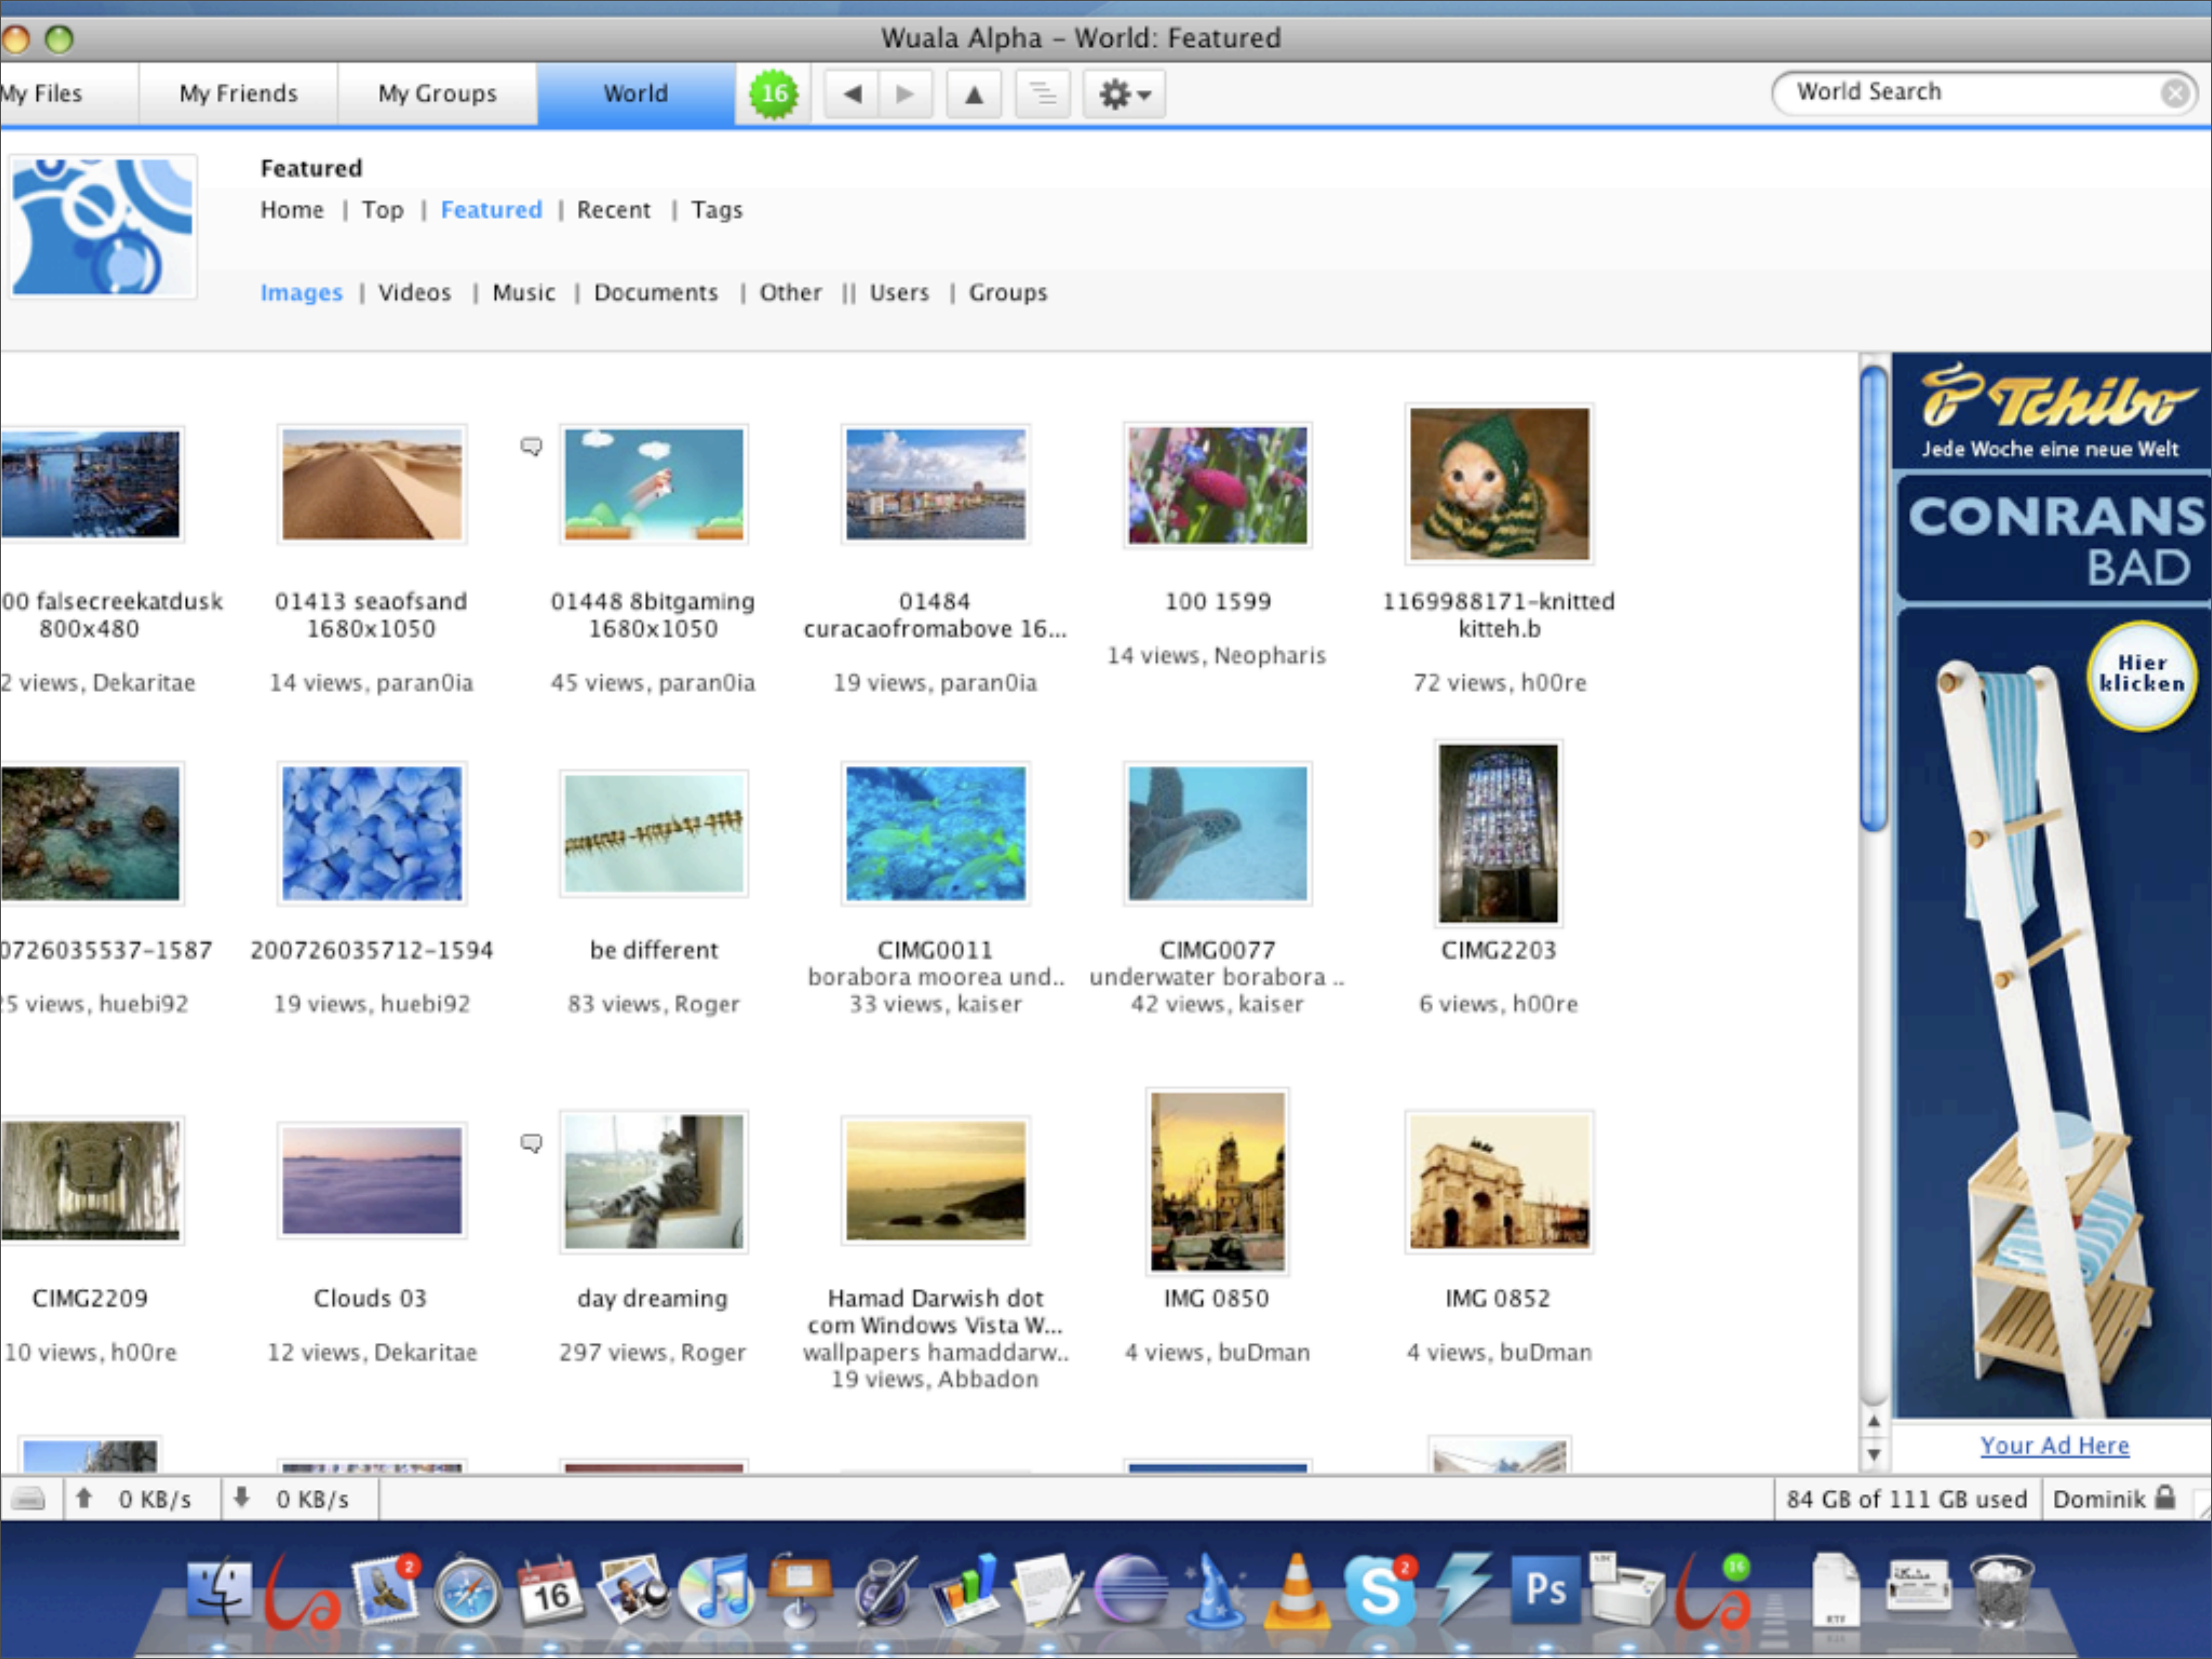
Wuala Alpha – World: Featured
My Files
My Friends
My Groups
World
16
World Search
Home | Top | Featured | Recent | Tags
Images | Videos | Music | Documents | Other || Users | Groups
0 KB/s
0 KB/s
Dominik
16

Featured

00 falsecreekatdusk 800x480
2 views, Dekaritae

01413 seaofsand 1680x1050
14 views, paran0ia

01448 8bitgaming 1680x1050
45 views, paran0ia

01484 curacaofromabove 16...
19 views, paran0ia

100 1599
14 views, Neopharis

1169988171-knitted kitteh.b
72 views, h00re

0726035537-1587
5 views, huebi92

200726035712-1594
19 views, huebi92

be different
83 views, Roger

CIMG0011
borabora moorea und..
33 views, kaiser

CIMG0077
underwater borabora ..
42 views, kaiser

CIMG2203
6 views, h00re

CIMG2209
10 views, h00re

Clouds 03
12 views, Dekaritae

day dreaming
297 views, Roger

Hamad Darwish dot com Windows Vista W...
wallpapers hamaddarw..
19 views, Abbadon

IMG 0850
4 views, buDman

IMG 0852
4 views, buDman

84 GB of 111 GB used

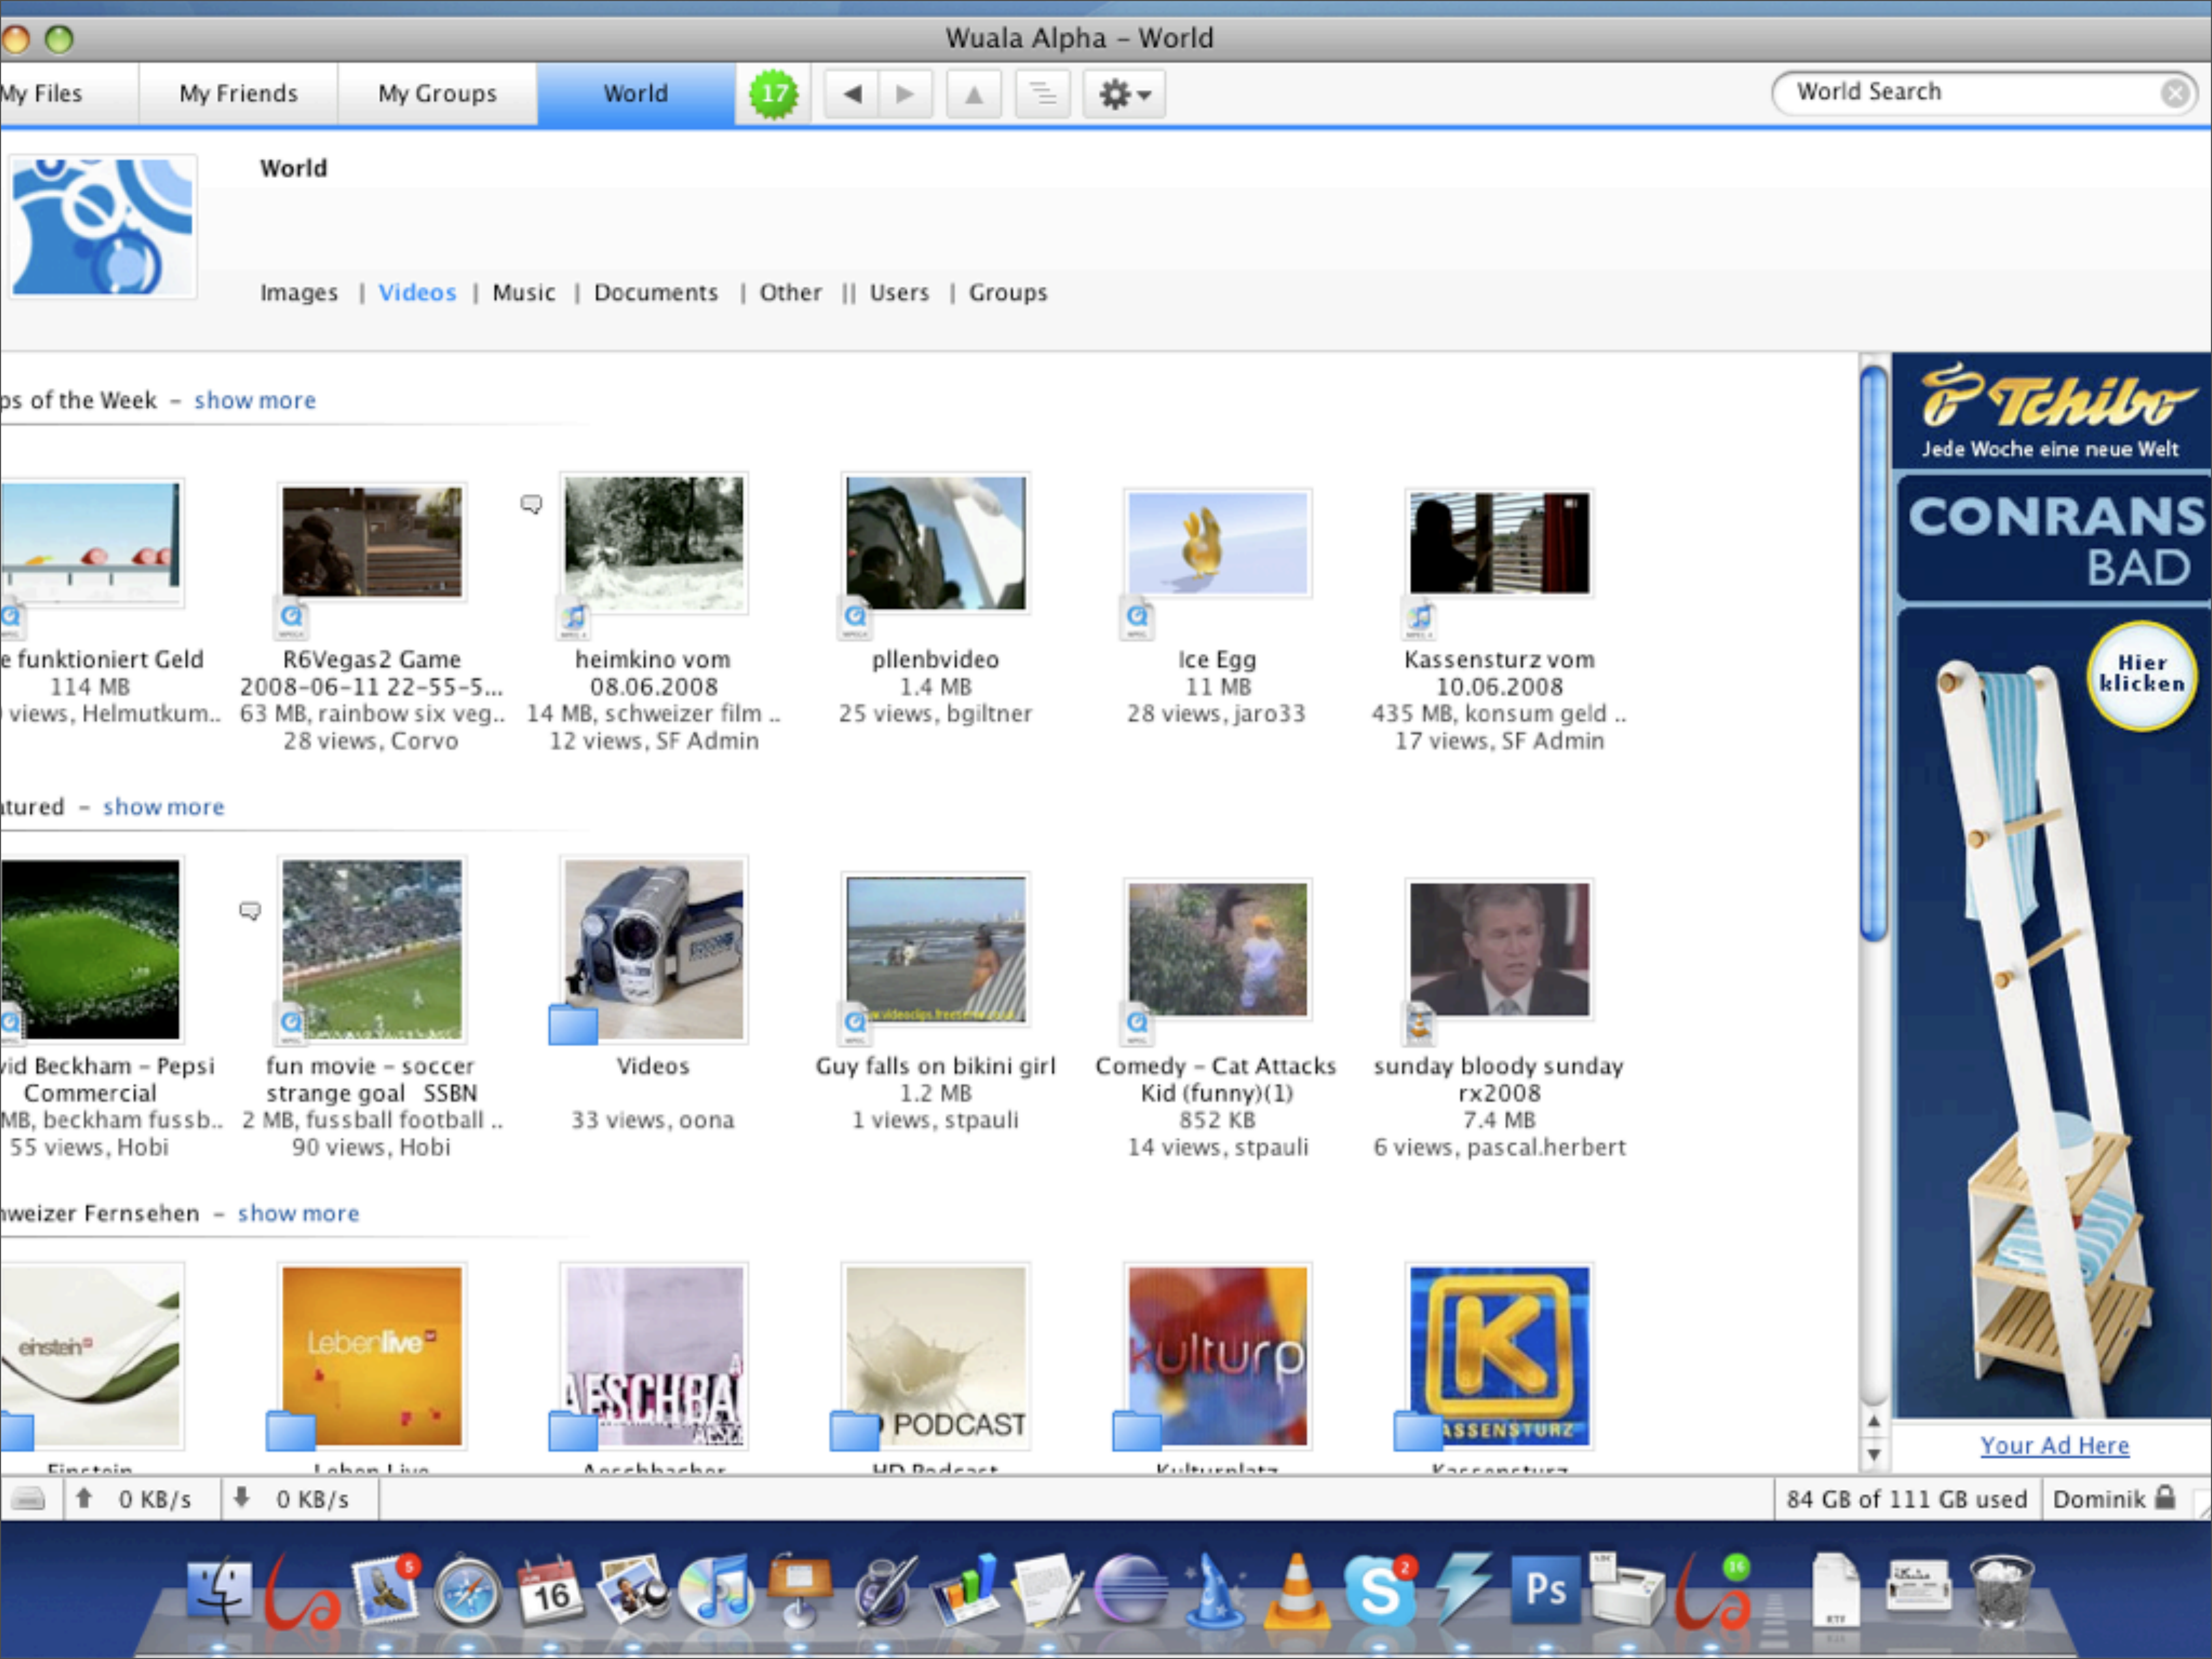

ps of the Week – show more

e funktioniert Geld
114 MB
 views, Helmutkum..

R6Vegas2 Game
2008-06-11 22-55-5...
63 MB, rainbow six veg..
28 views, Corvo

heimkino vom
08.06.2008
14 MB, schweizer film ..
12 views, SF Admin

pllenbvideo
1.4 MB
25 views, bgiltner

Ice Egg
11 MB
28 views, jaro33

Kassensturz vom
10.06.2008
435 MB, konsum geld ..
17 views, SF Admin

tured – show more

vid Beckham – Pepsi
Commercial
MB, beckham fussb..
55 views, Hobi

fun movie – soccer
strange goal SSBN
2 MB, fussball football ..
90 views, Hobi

Videos

33 views, oona

Guy falls on bikini girl
1.2 MB
1 views, stpauli

Comedy – Cat Attacks
Kid (funny)(1)
852 KB
14 views, stpauli

sunday bloody sunday
rx2008
7.4 MB
6 views, pascal.herbert

weizer Fernsehen – show more

Einstein

Leben Live

Aeschbacher

HD Podcast

Kulturplatz

Kassensturz

0 KB/s
0 KB/s
84 GB of 111 GB used
Dominik

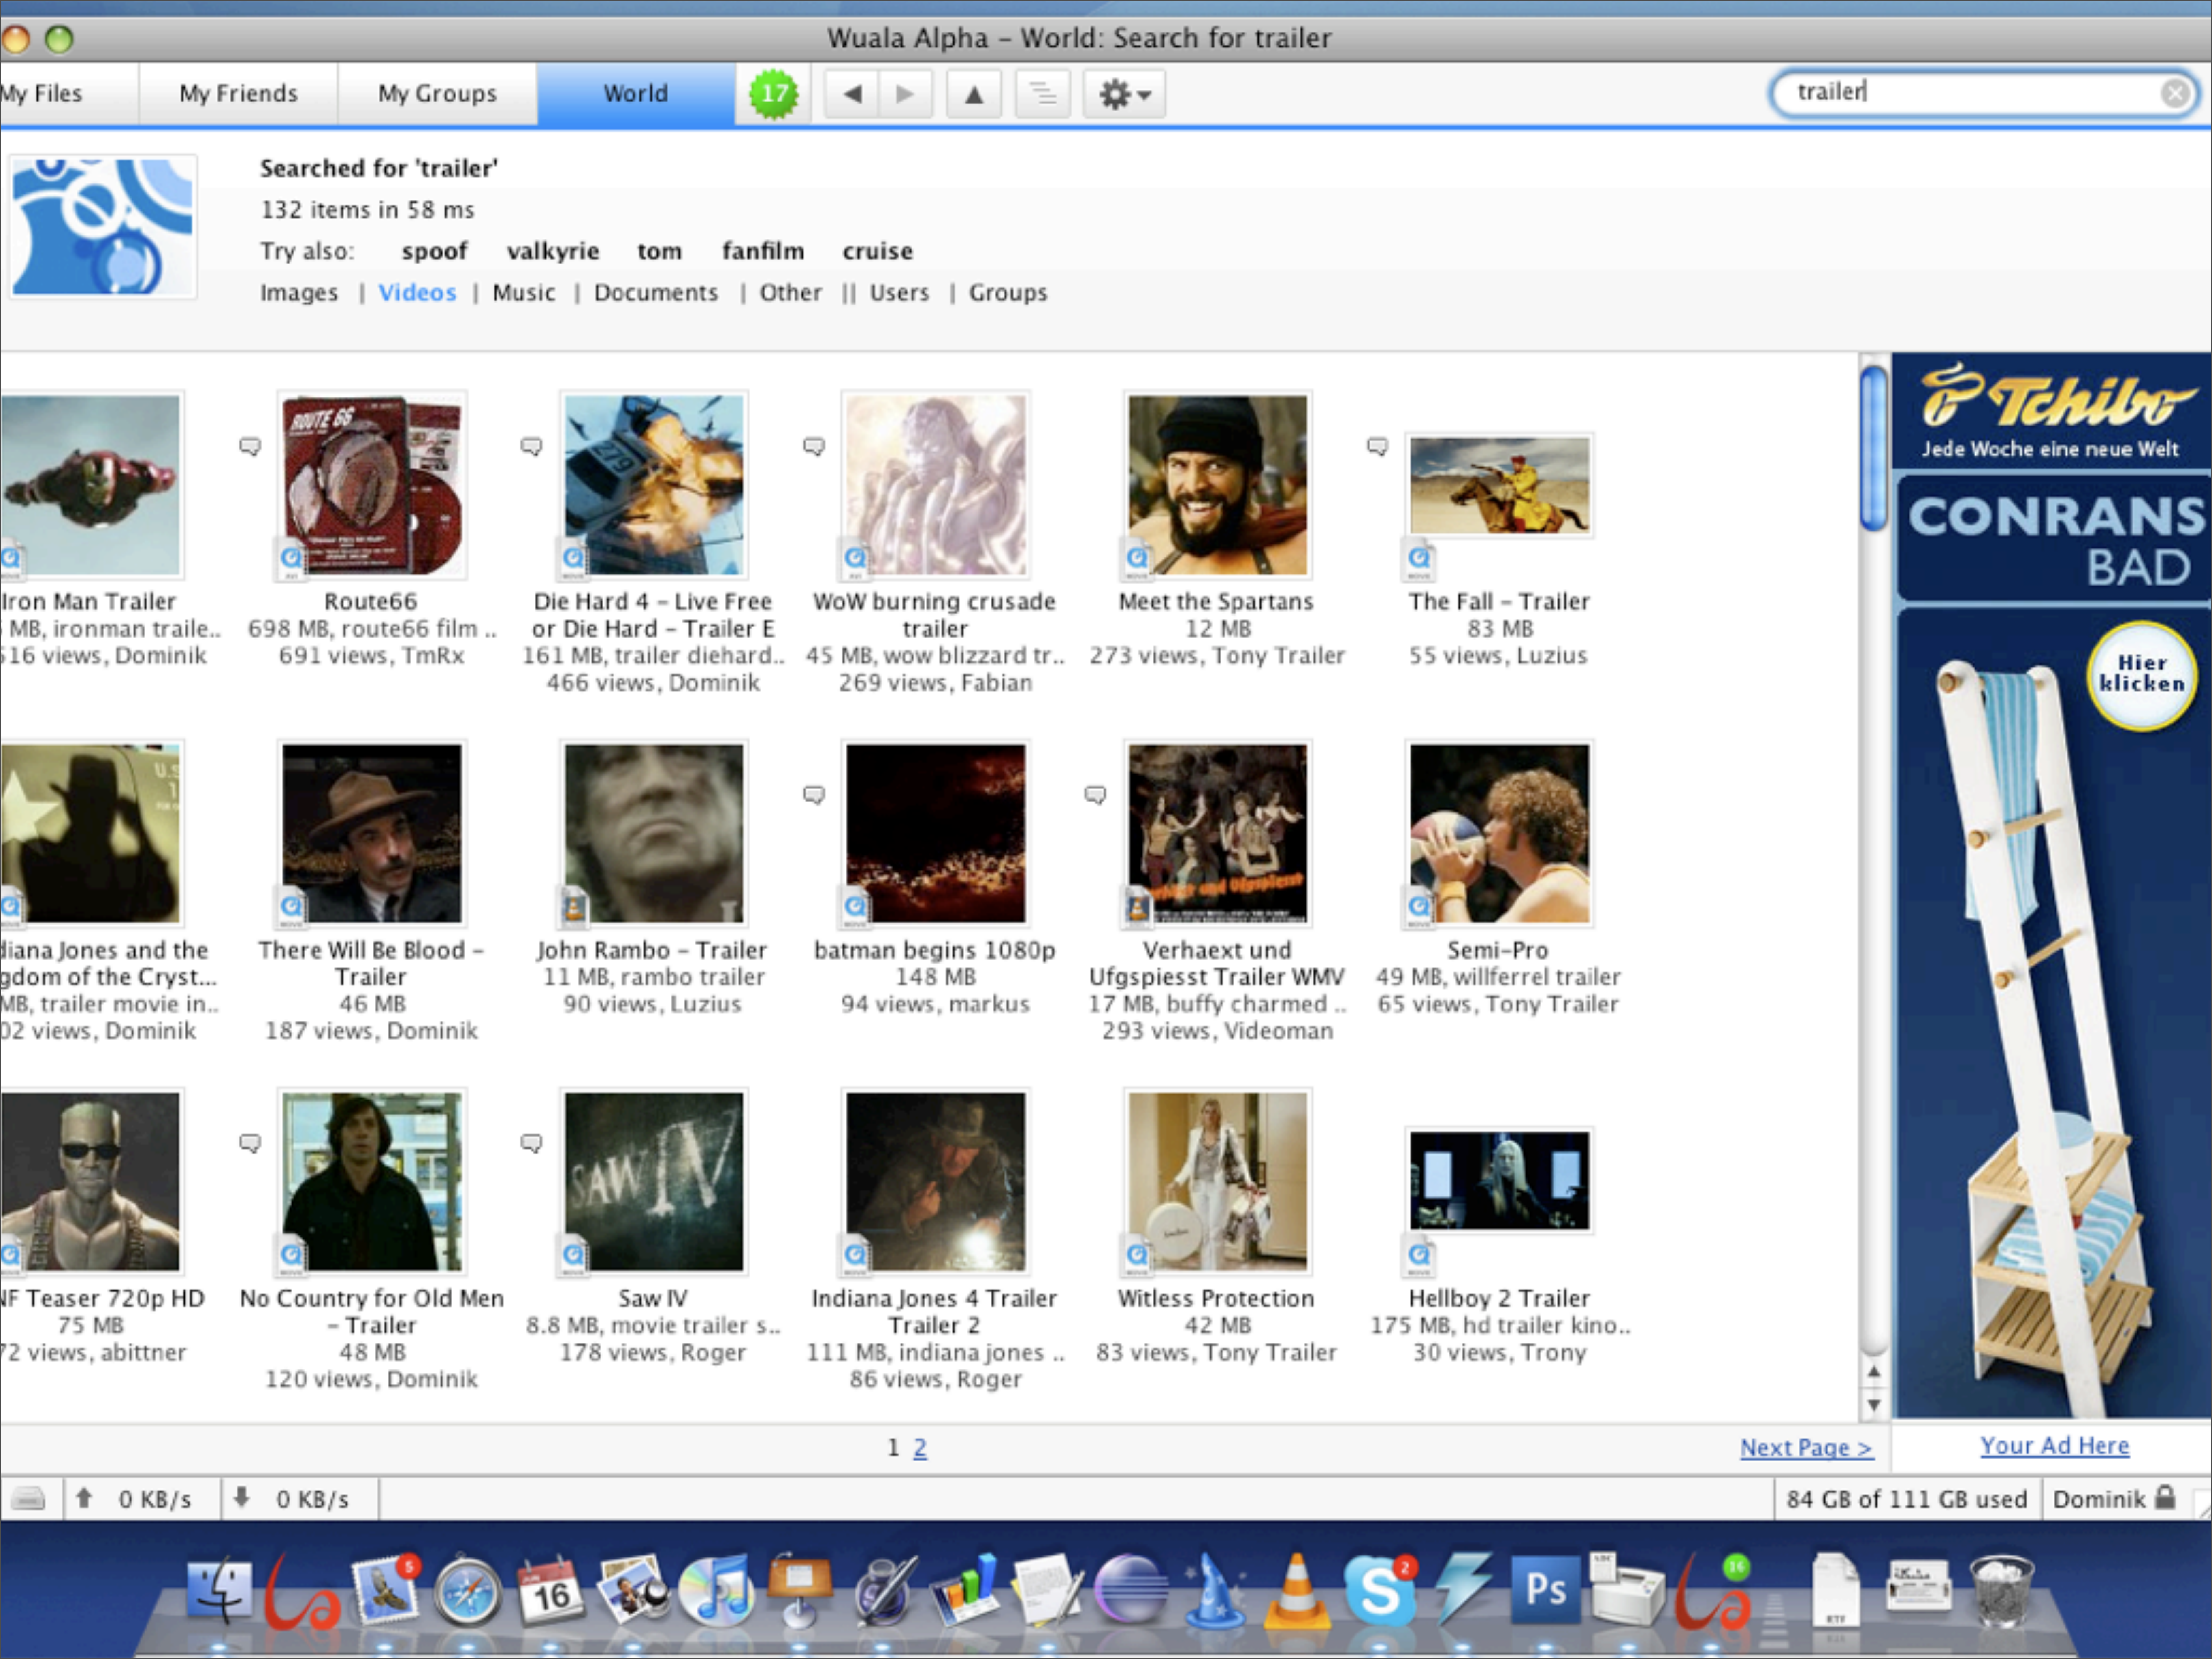
Wuala Alpha – World: Search for trailer
My Files
My Friends
My Groups
World
17
trailer
Jede Woche eine neue Welt
CONRANS
BAD
Hier klicken

Searched for 'trailer'

132 items in 58 ms

Try also: **spoof** **valkyrie** **tom** **fanfilm** **cruise**

Images | Videos | Music | Documents | Other || Users | Groups

Iron Man Trailer
MB, ironman traile..
16 views, Dominik

Route66
698 MB, route66 film ..
691 views, TmRx

Die Hard 4 – Live Free or Die Hard – Trailer E
161 MB, trailer diehard..
466 views, Dominik

WoW burning crusade trailer
45 MB, wow blizzard tr..
269 views, Fabian

Meet the Spartans
12 MB
273 views, Tony Trailer

The Fall – Trailer
83 MB
55 views, Luzius

diana Jones and the gdom of the Cryst...
MB, trailer movie in..
02 views, Dominik

There Will Be Blood – Trailer
46 MB
187 views, Dominik

John Rambo – Trailer
11 MB, rambo trailer
90 views, Luzius

batman begins 1080p
148 MB
94 views, markus

Verhaext und Ufgspiesst Trailer WMV
17 MB, buffy charmed ..
293 views, Videoman

Semi-Pro
49 MB, willferrel trailer
65 views, Tony Trailer

NF Teaser 720p HD
75 MB
72 views, abittner

No Country for Old Men – Trailer
48 MB
120 views, Dominik

Saw IV
8.8 MB, movie trailer s..
178 views, Roger

Indiana Jones 4 Trailer Trailer 2
111 MB, indiana jones ..
86 views, Roger

Witless Protection
42 MB
83 views, Tony Trailer

Hellboy 2 Trailer
175 MB, hd trailer kino..
30 views, Trony

1 2

Next Page >

Your Ad Here

0 KB/s 0 KB/s

84 GB of 111 GB used

Dominik

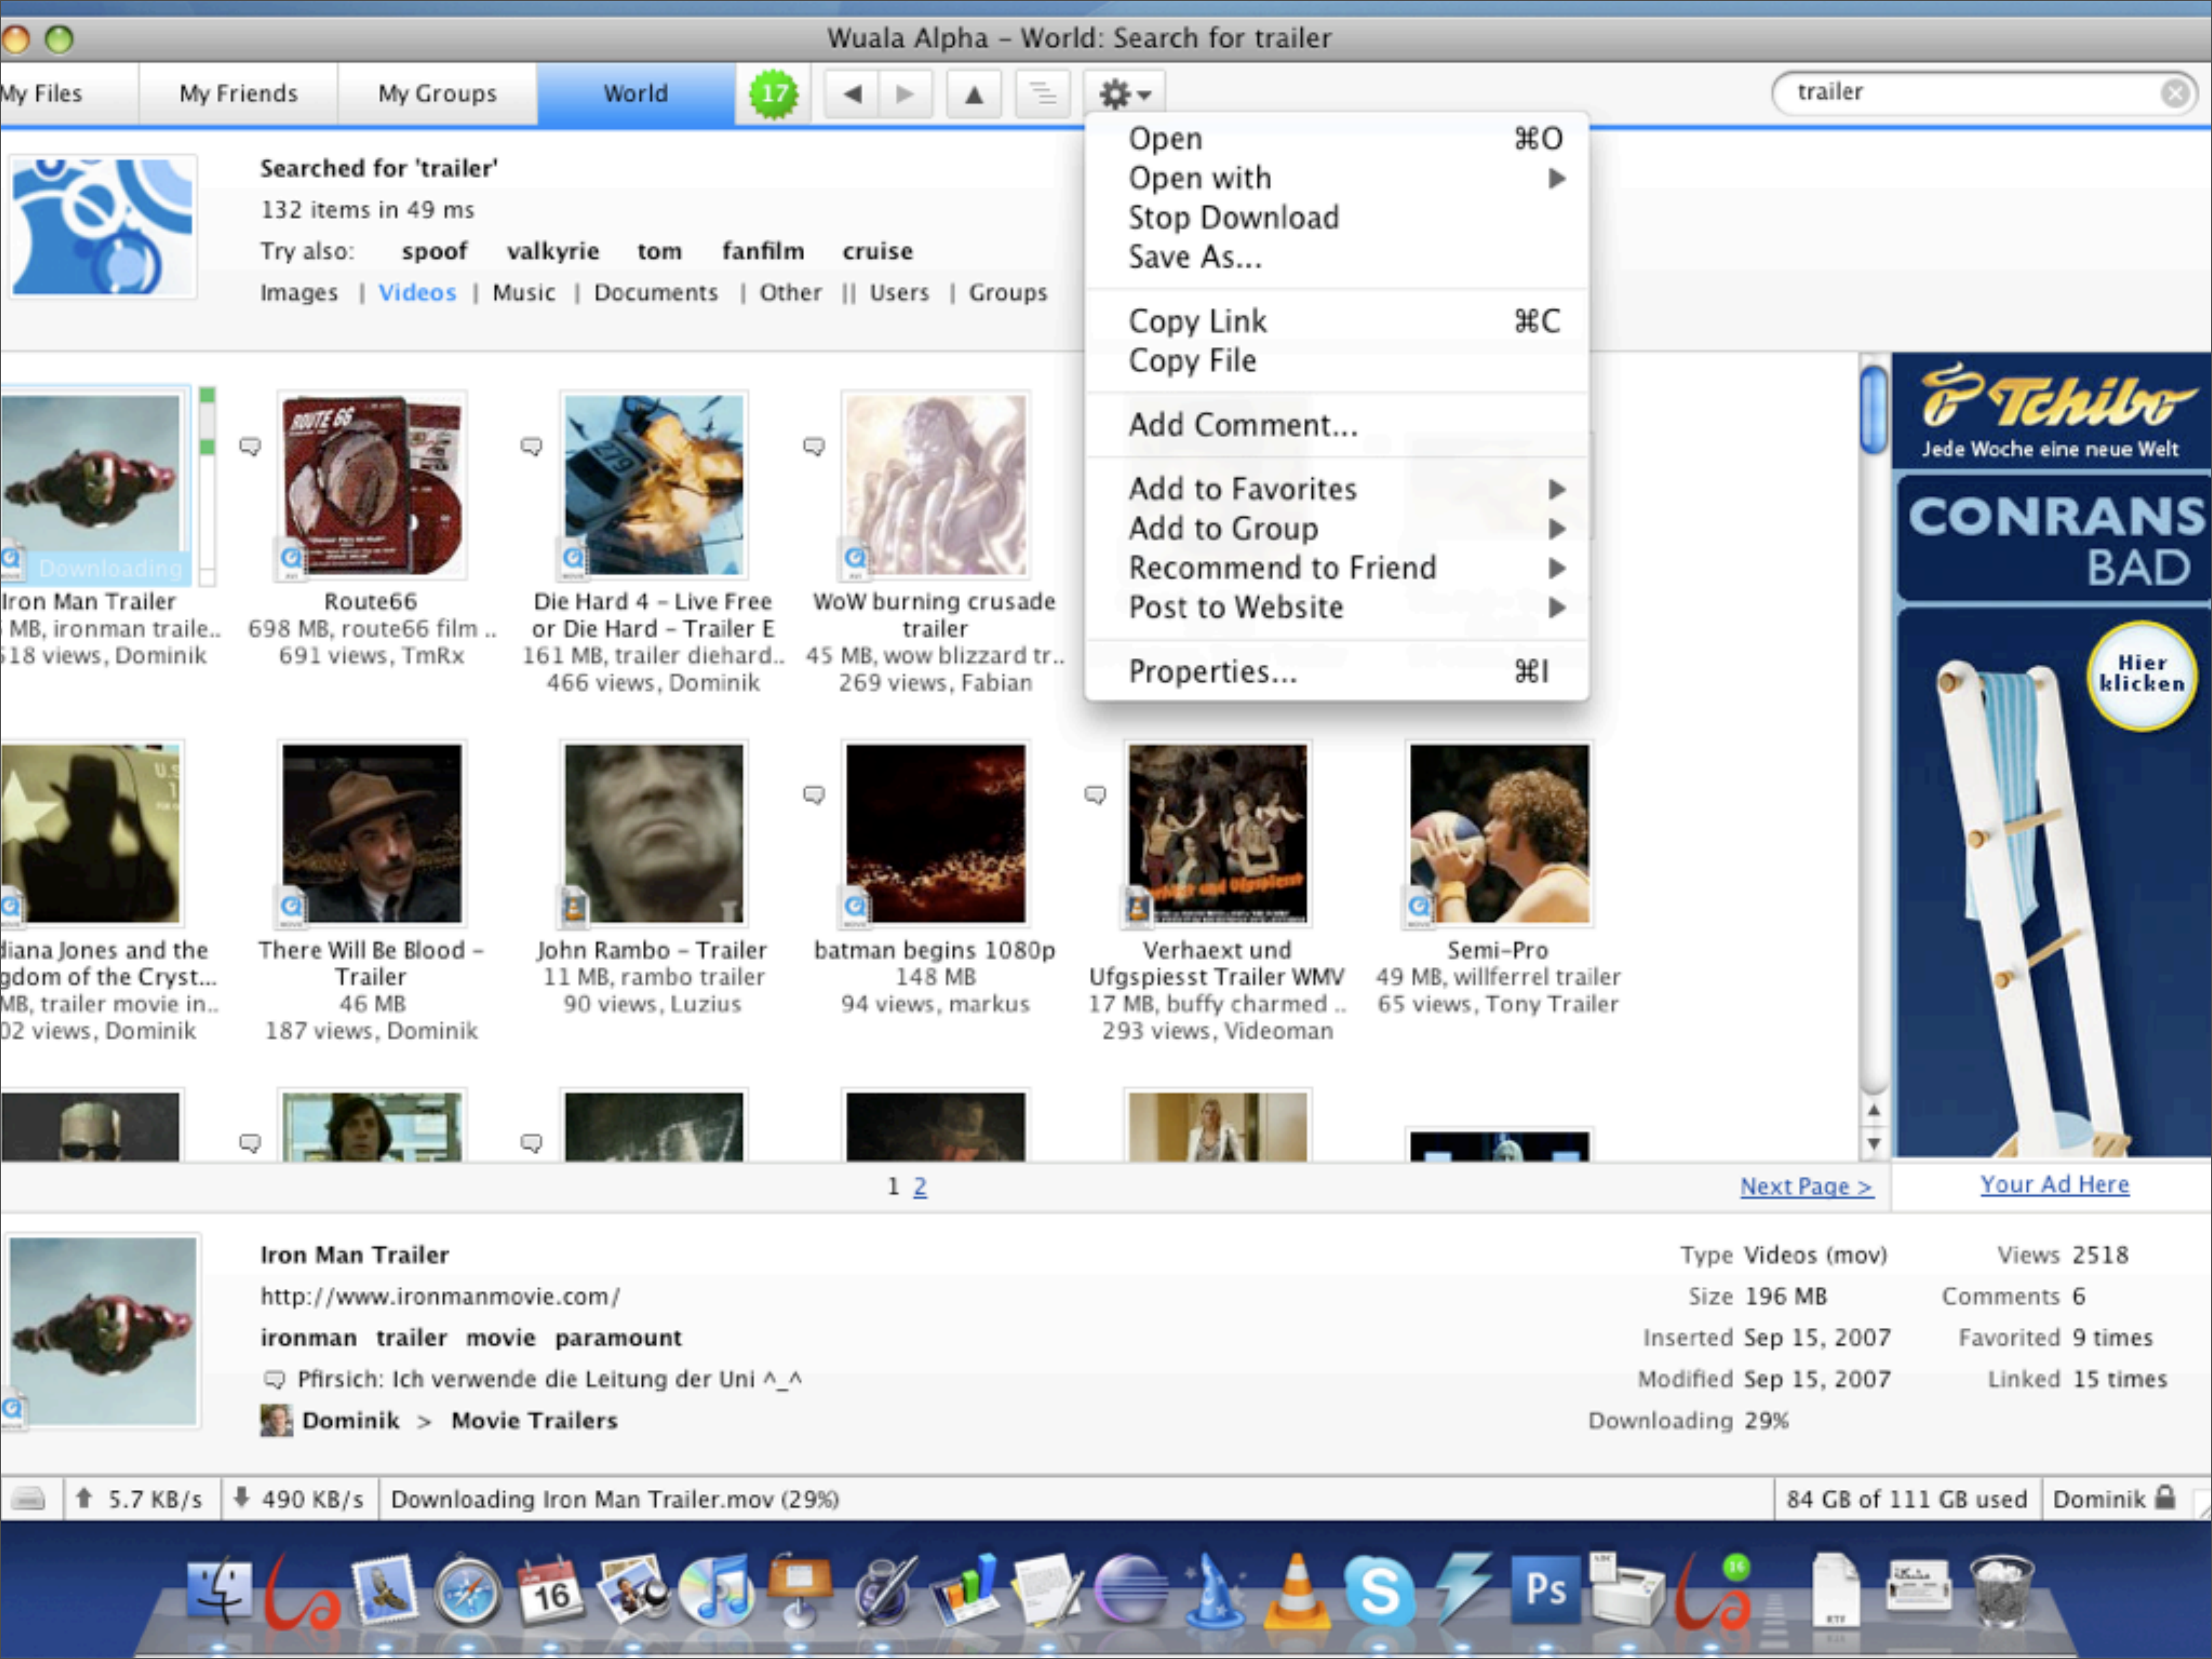

Wuala Alpha – World: Search for trailer

Searched for 'trailer'

132 items in 49 ms

Try also: spoof valkyrie tom fanfilm cruise

Images | Videos | Music | Documents | Other || Users | Groups

Iron Man Trailer
MB, ironman traile..
18 views, Dominik

Route66
698 MB, route66 film ..
691 views, TmRx

Die Hard 4 – Live Free or Die Hard – Trailer E
161 MB, trailer diehard..
466 views, Dominik

WoW burning crusade trailer
45 MB, wow blizzard tr..
269 views, Fabian

diana Jones and the
gdom of the Cryst...
MB, trailer movie in..
02 views, Dominik

There Will Be Blood – Trailer
46 MB
187 views, Dominik

John Rambo – Trailer
11 MB, rambo trailer
90 views, Luzius

batman begins 1080p
148 MB
94 views, markus

Verhaext und Ufgspiesst Trailer WMV
17 MB, buffy charmed ..
293 views, Videoman

Semi-Pro
49 MB, willferrel trailer
65 views, Tony Trailer

1 2 Next Page >

Iron Man Trailer

http://www.ironmanmovie.com/

ironman trailer movie paramount

Pfirsich: Ich verwende die Leitung der Uni ^_^

Dominik > Movie Trailers

Type	Videos (mov)	Views	2518
Size	196 MB	Comments	6
Inserted	Sep 15, 2007	Favorited	9 times
Modified	Sep 15, 2007	Linked	15 times
Downloading	29%		

↑ 5.7 KB/s ↓ 490 KB/s Downloading Iron Man Trailer.mov (29%) 84 GB of 111 GB used Dominik

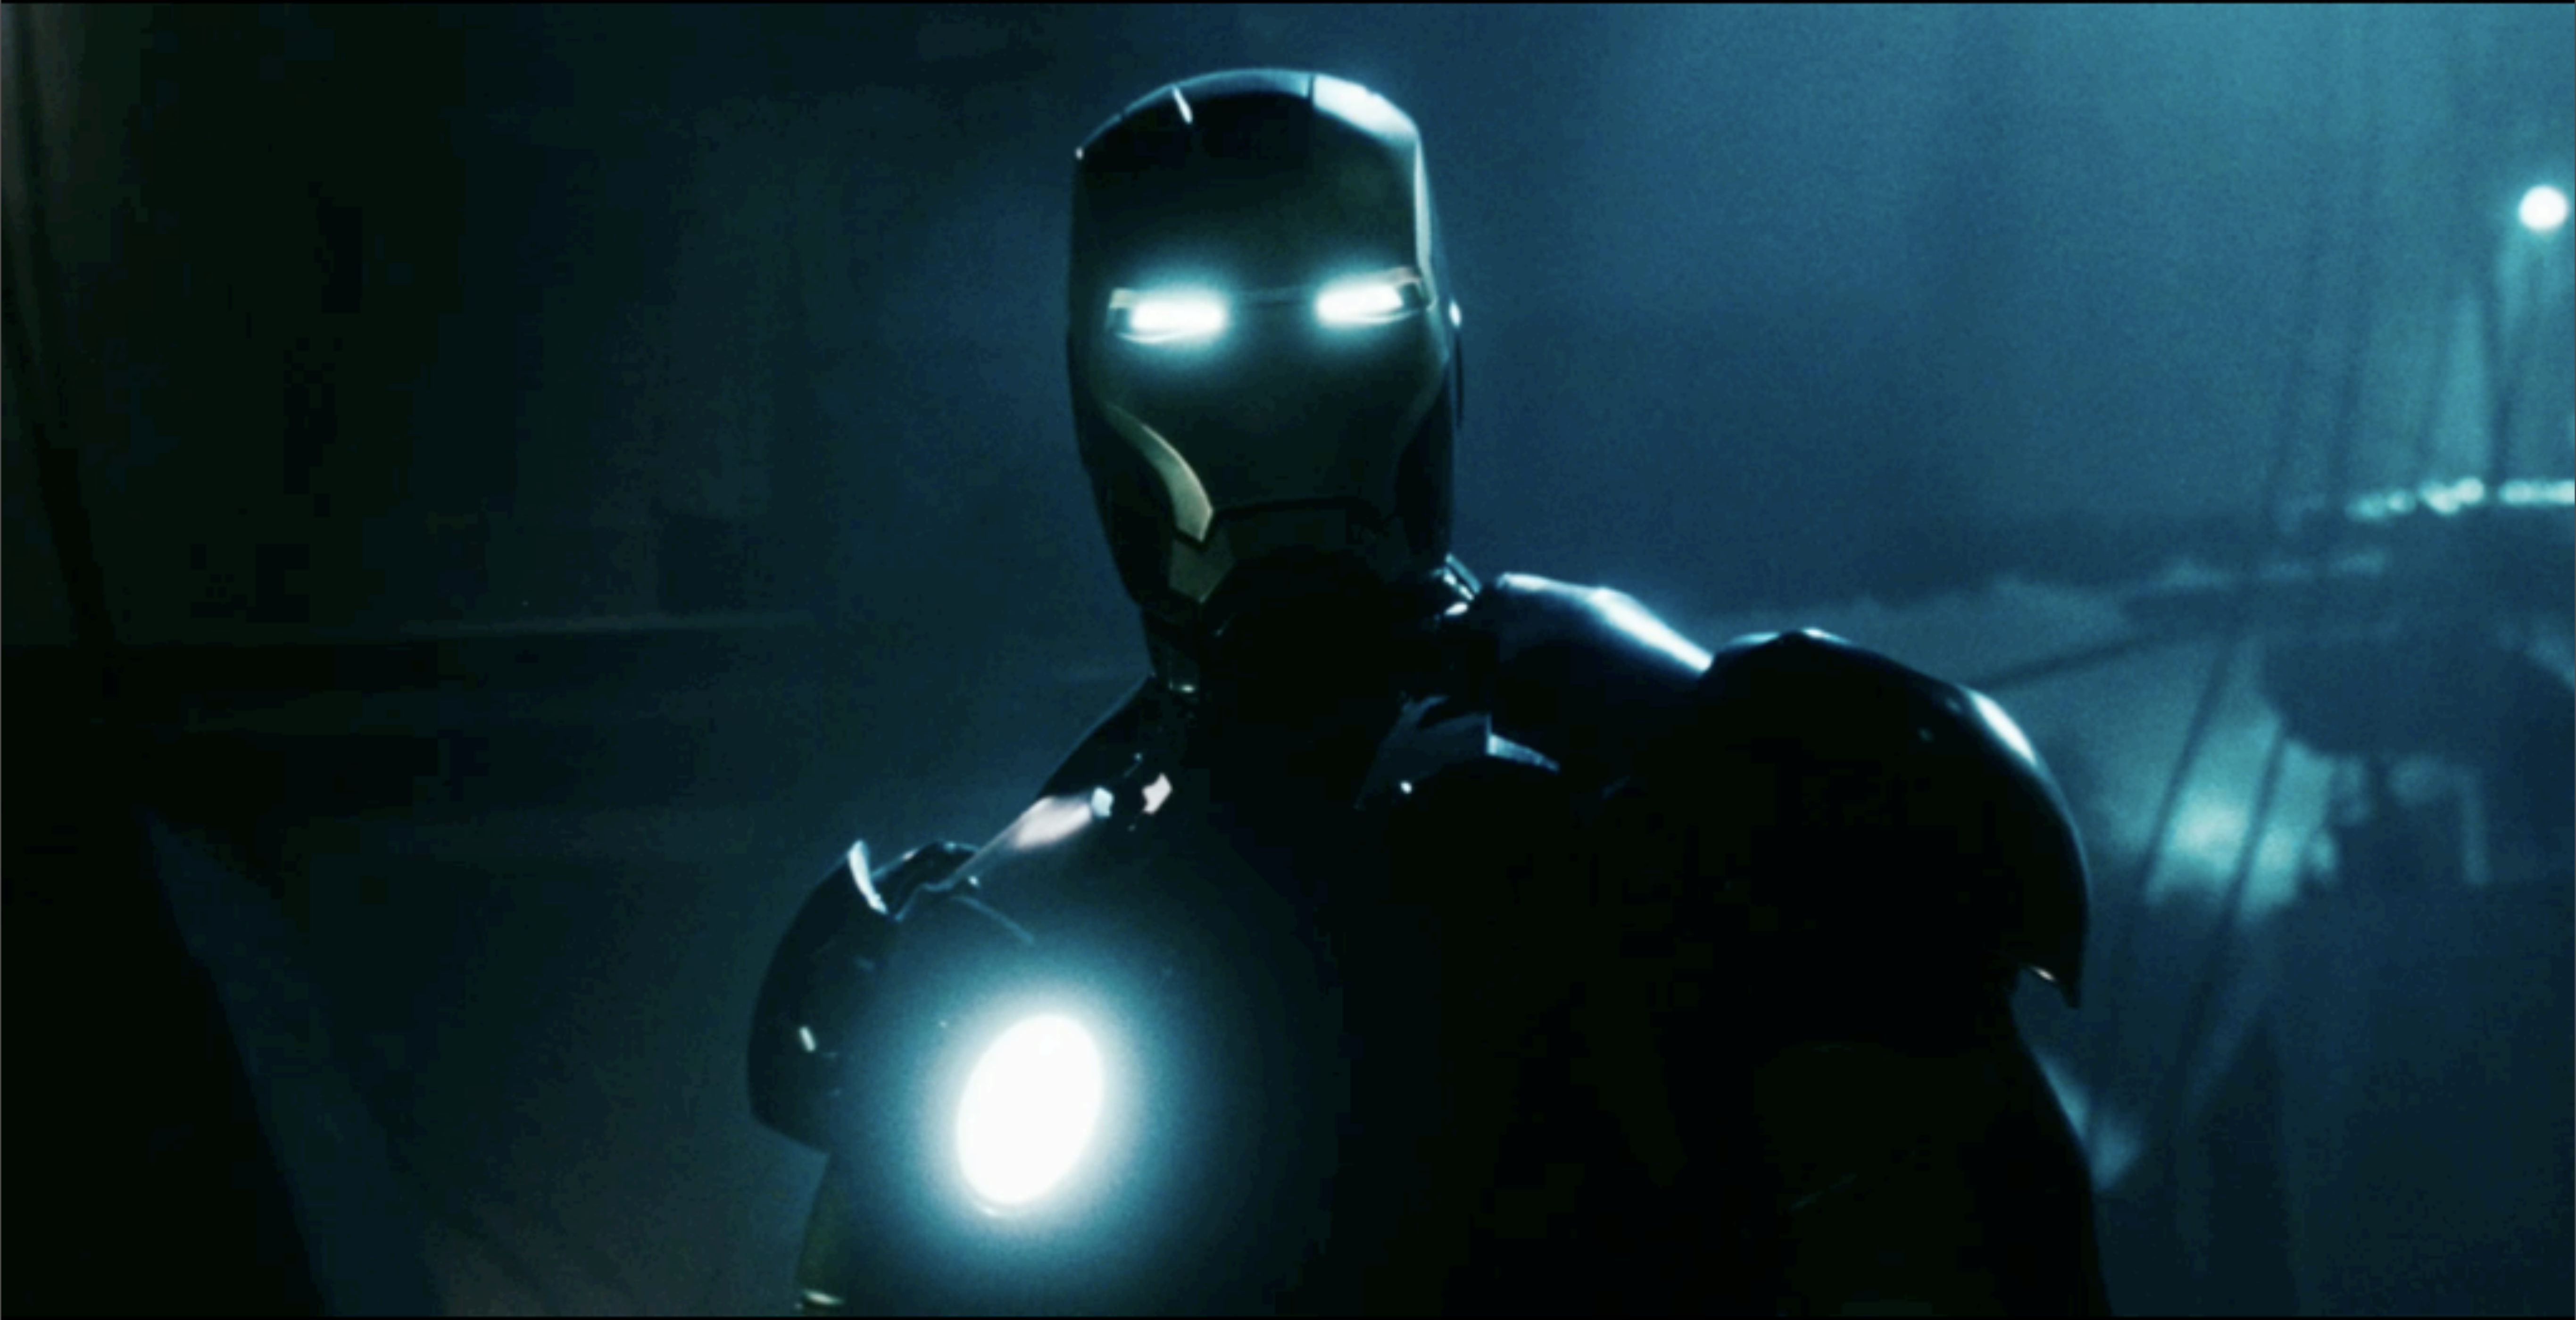

how does it work?

data stored in the p2p network

users’s computer can be offline

how to ensure availability
(persistent storage)?

two approaches

1. make sure the data is always in the network

move the data when a computer goes offline

bad idea for lots of data and high churn rate

2. introduce redundancy

redundany = replication?

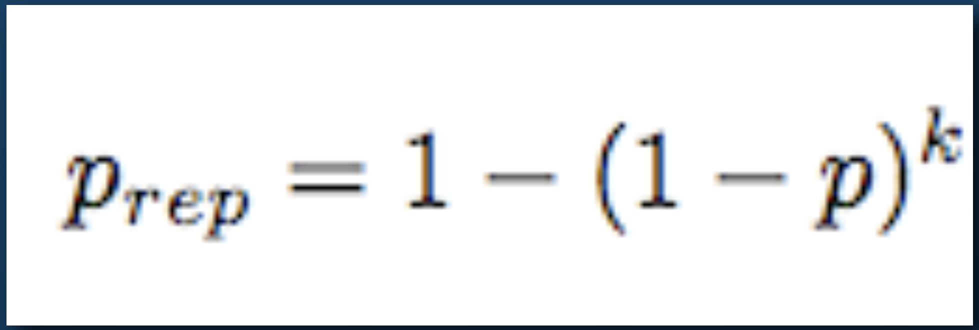

$$p_{rep} = 1 - (1 - p)^k$$

p = node availability

k = redundancy factor

p_{rep} = file availability

redundany = replication?

$$p_{rep} = 1 - (1 - p)^k$$

example

$p = 0.25$

$k = 5$

$p_{rep} = 0.763$ → not enough

redundany = replication?

$$p_{rep} = 1 - (1 - p)^k$$

example

$p = 0.25$

$k = 24$ ⟶ unrealistic

$p_{rep} = 0.999$

erasure codes

encode m fragments into n

need **any** m out of n to reconstruct

reed-solomon (optimal codes)

RAID storage systems

(vs. low-density-parity-check need (1+e) * m,

where e is a fixed, small constant)

availability

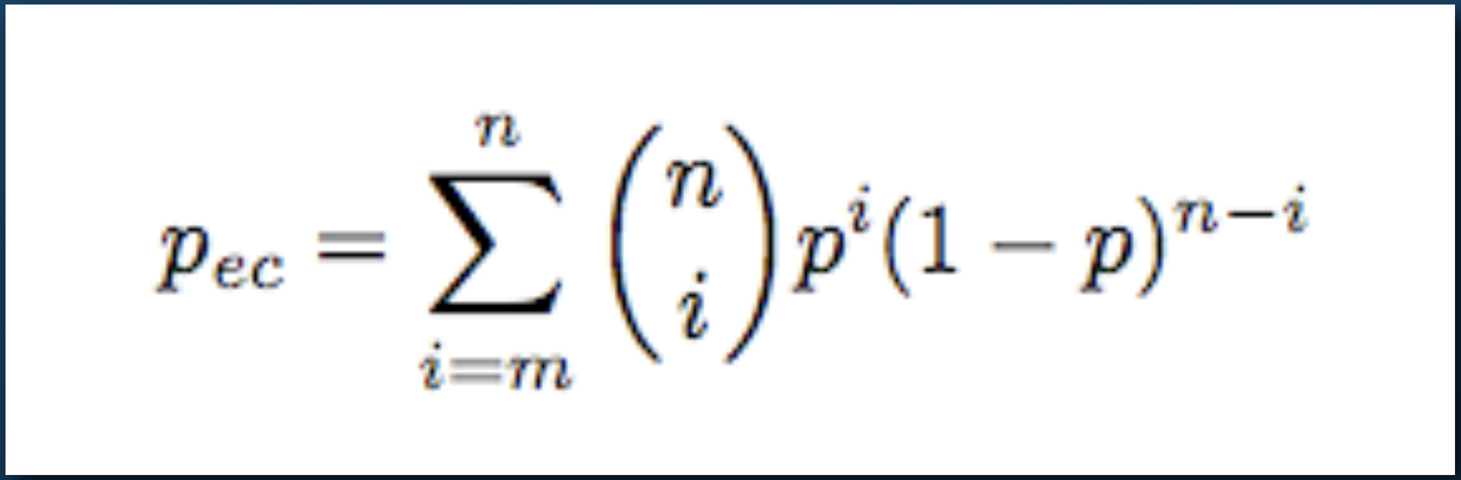

$$p_{ec} = \sum_{i=m}^{n} \binom{n}{i} p^i (1-p)^{n-i}$$

$p = 0.25$

$m = 100$, $n = 517$, $k = n/m = 5.17$

$p_{ec} = 0.999$

$k = n/m = 5.17$ vs. $k = 24$ using replication

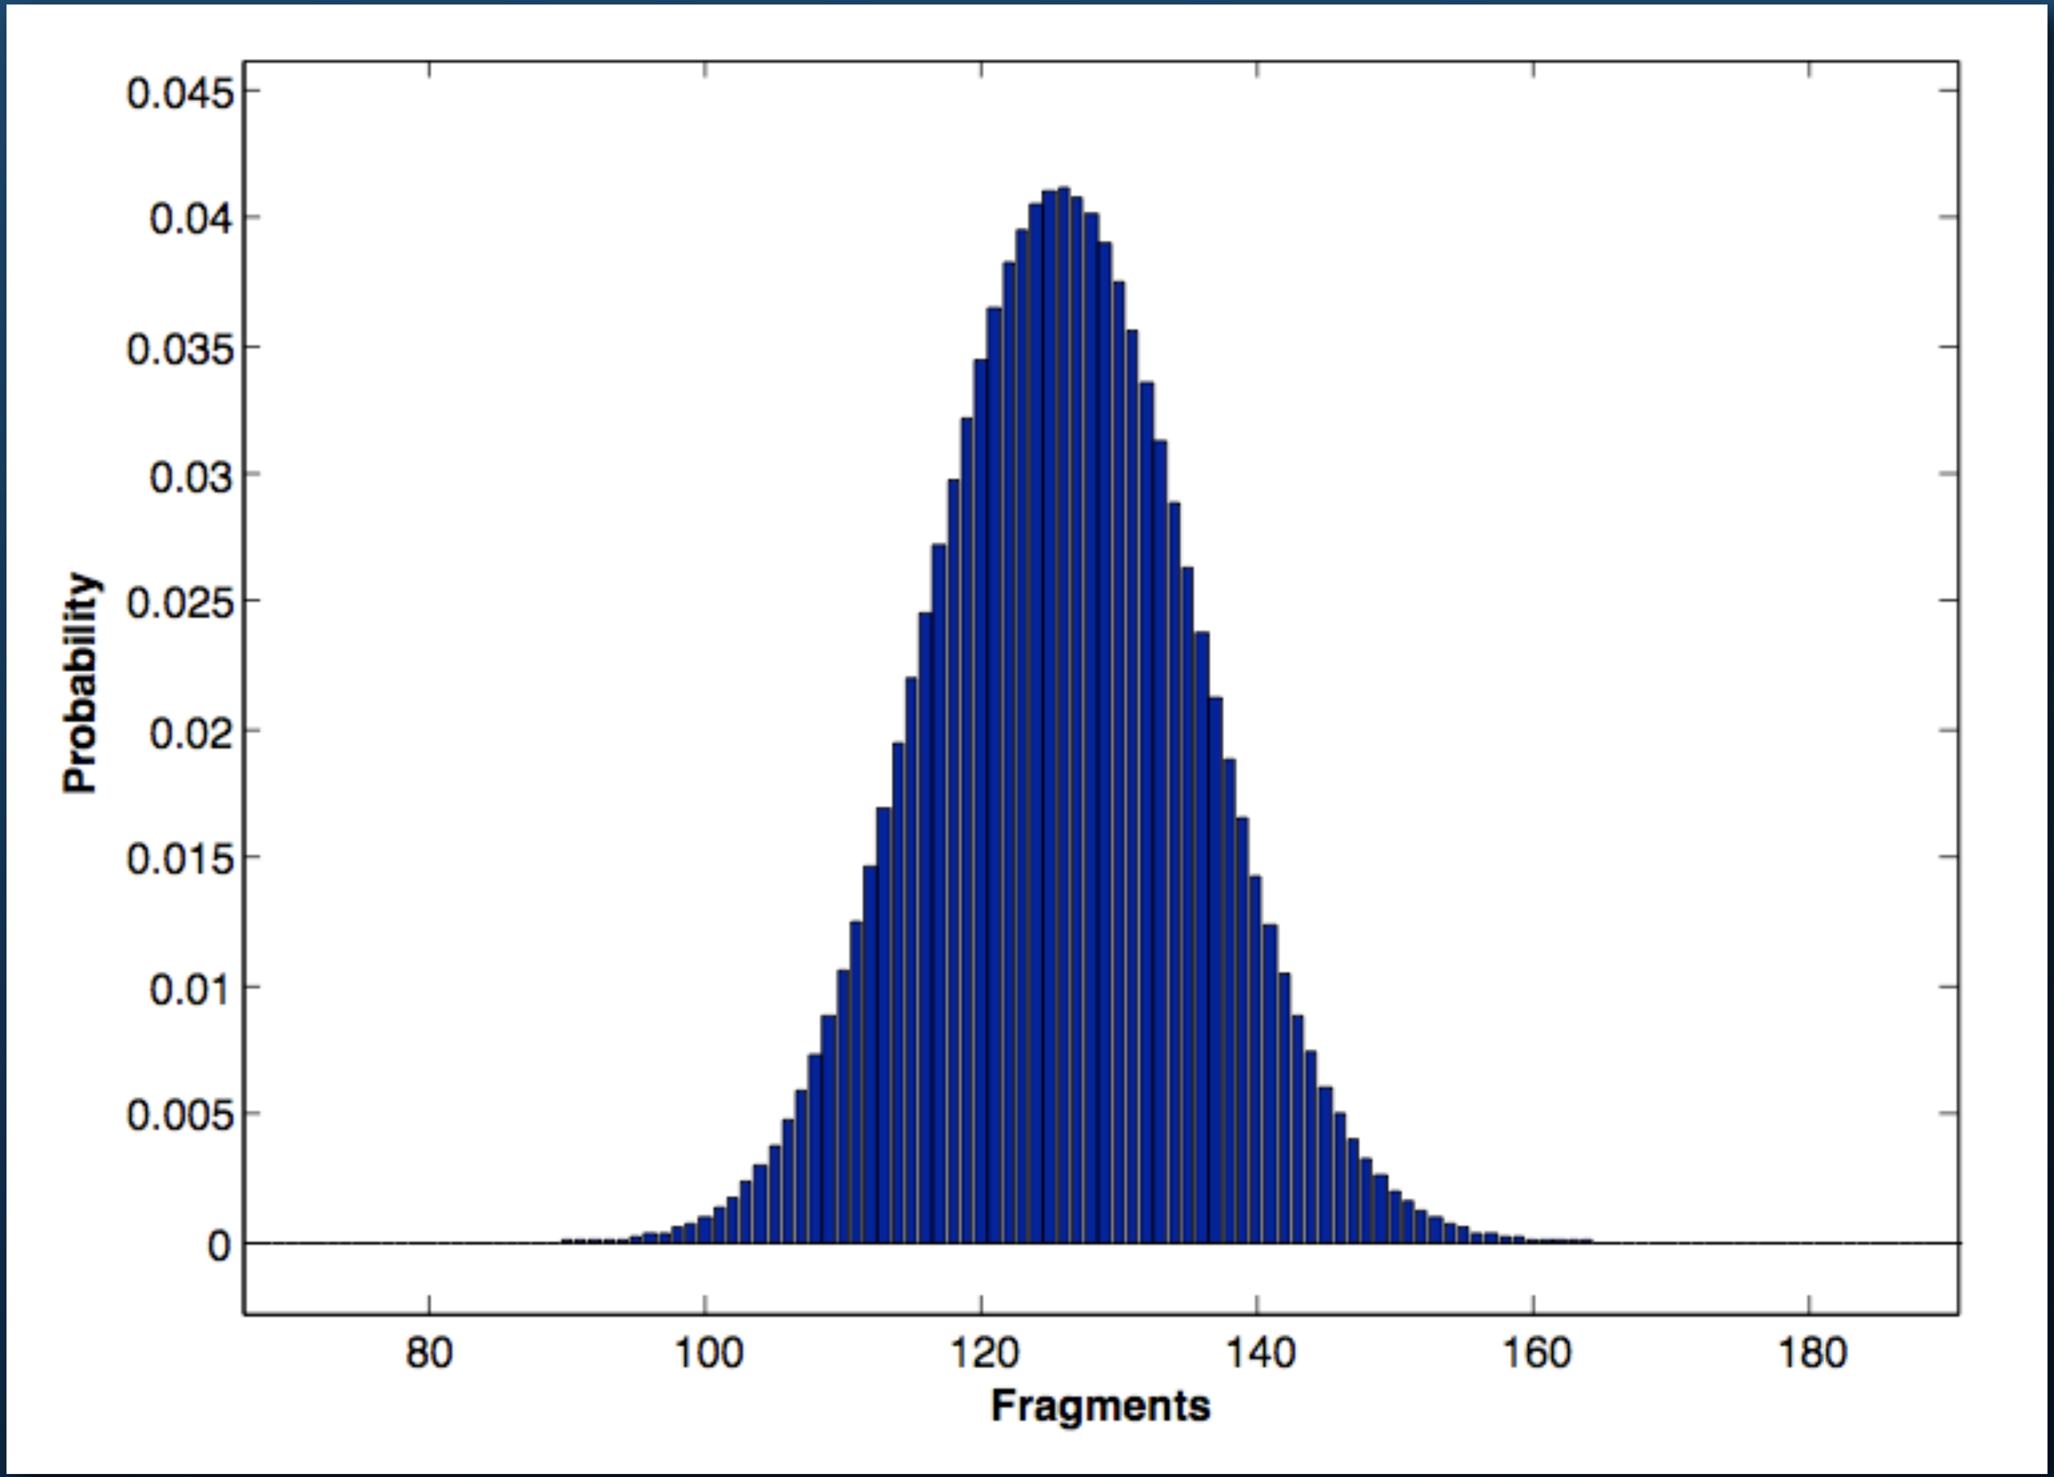
0.045
0.04
0.035
0.03
0.025
0.02
0.015
0.01
0.005
0
Probability
80
100
120
140
160
180
Fragments

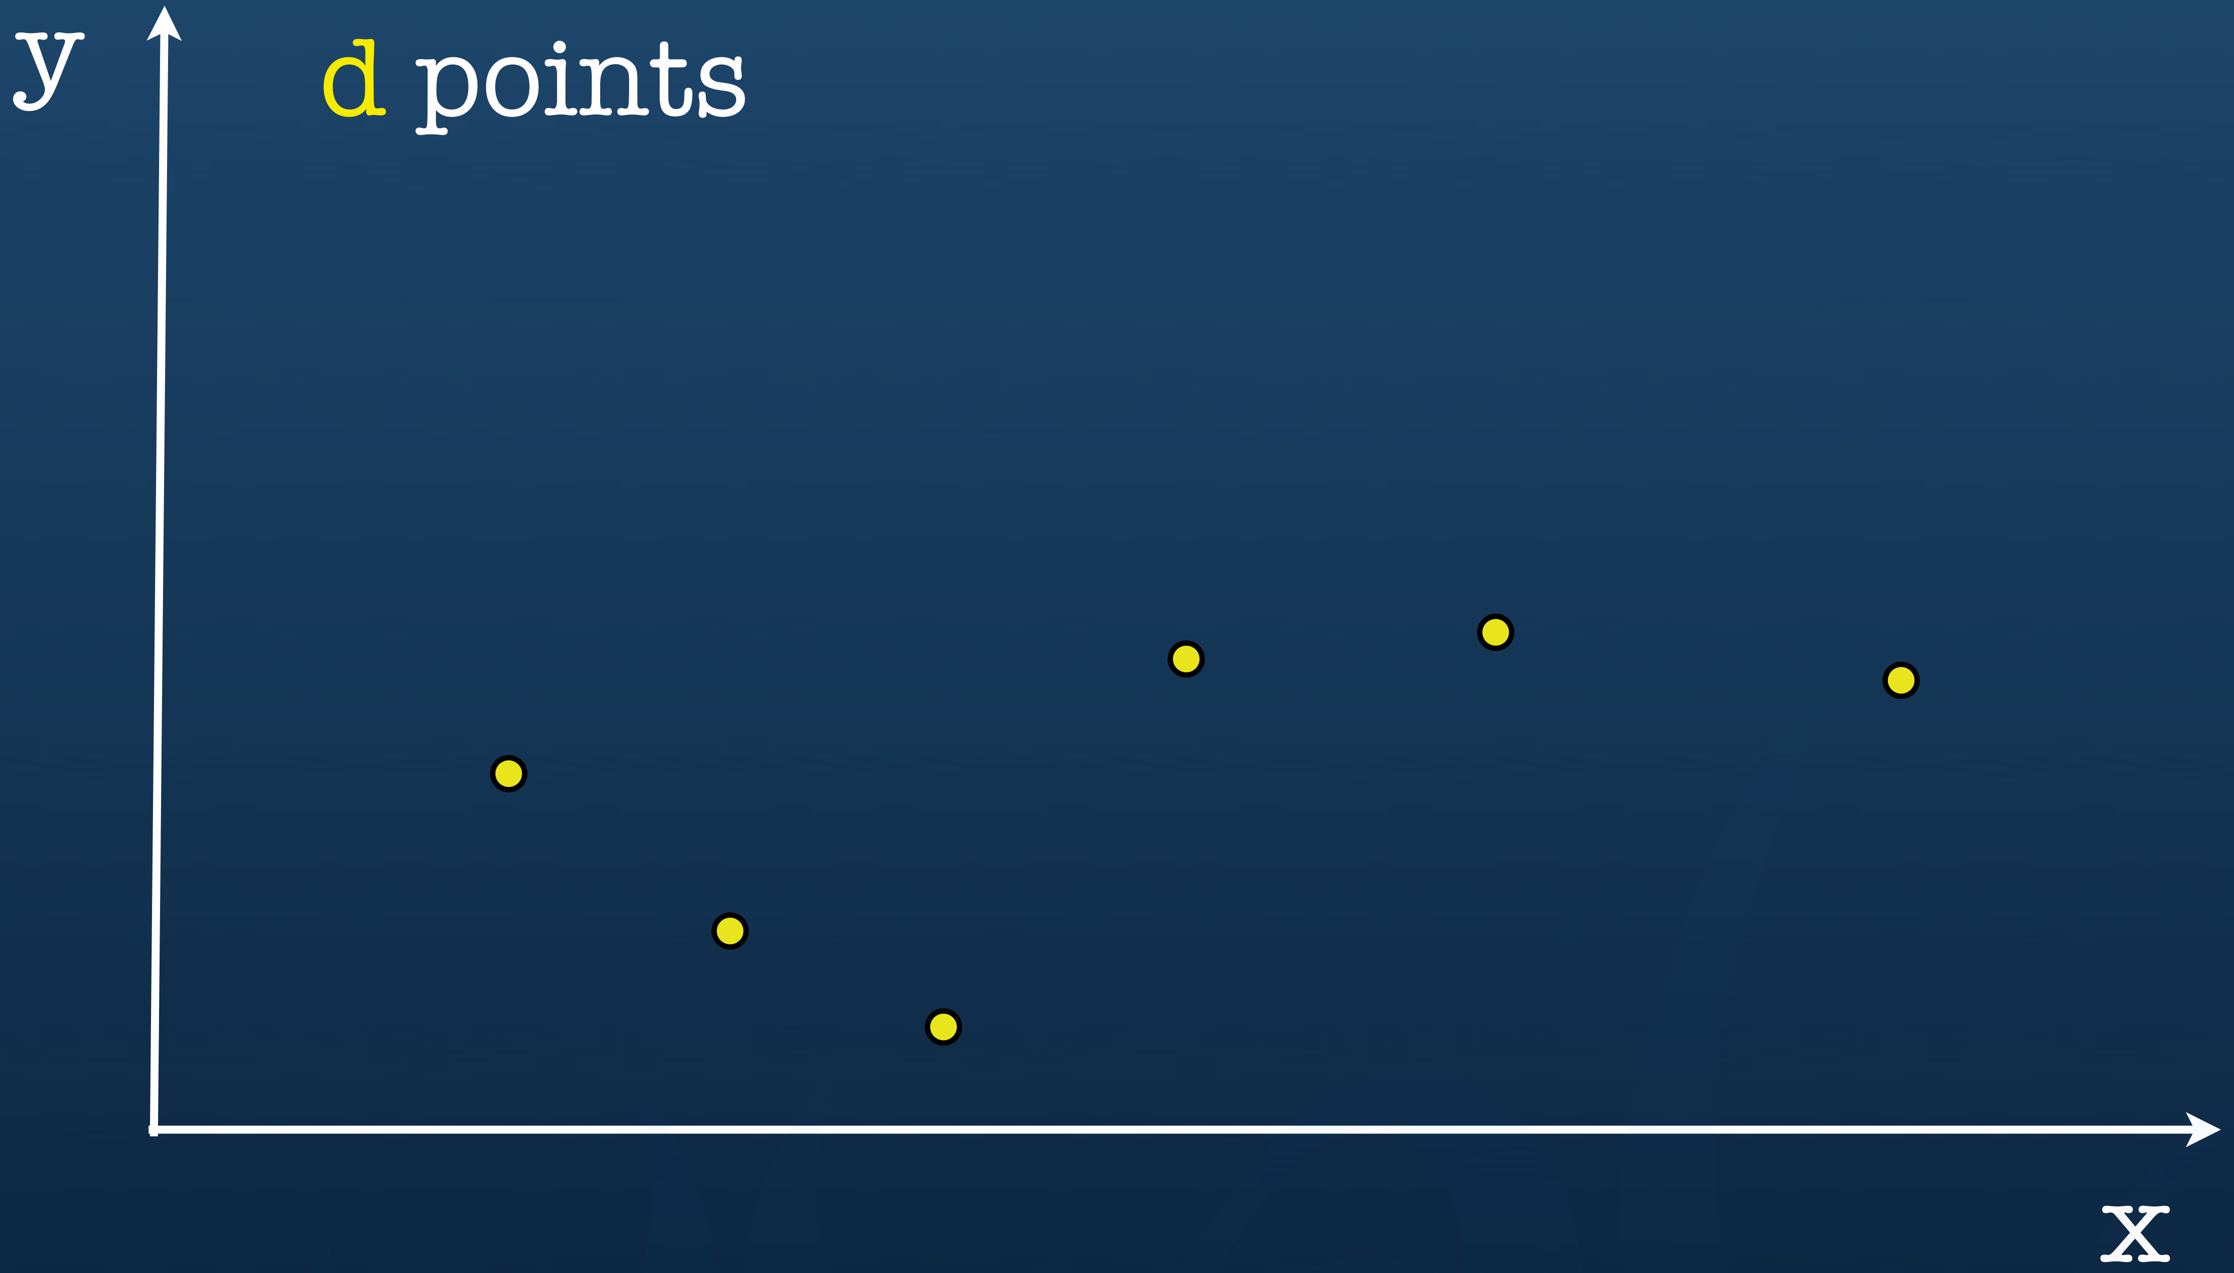
y
d points
x

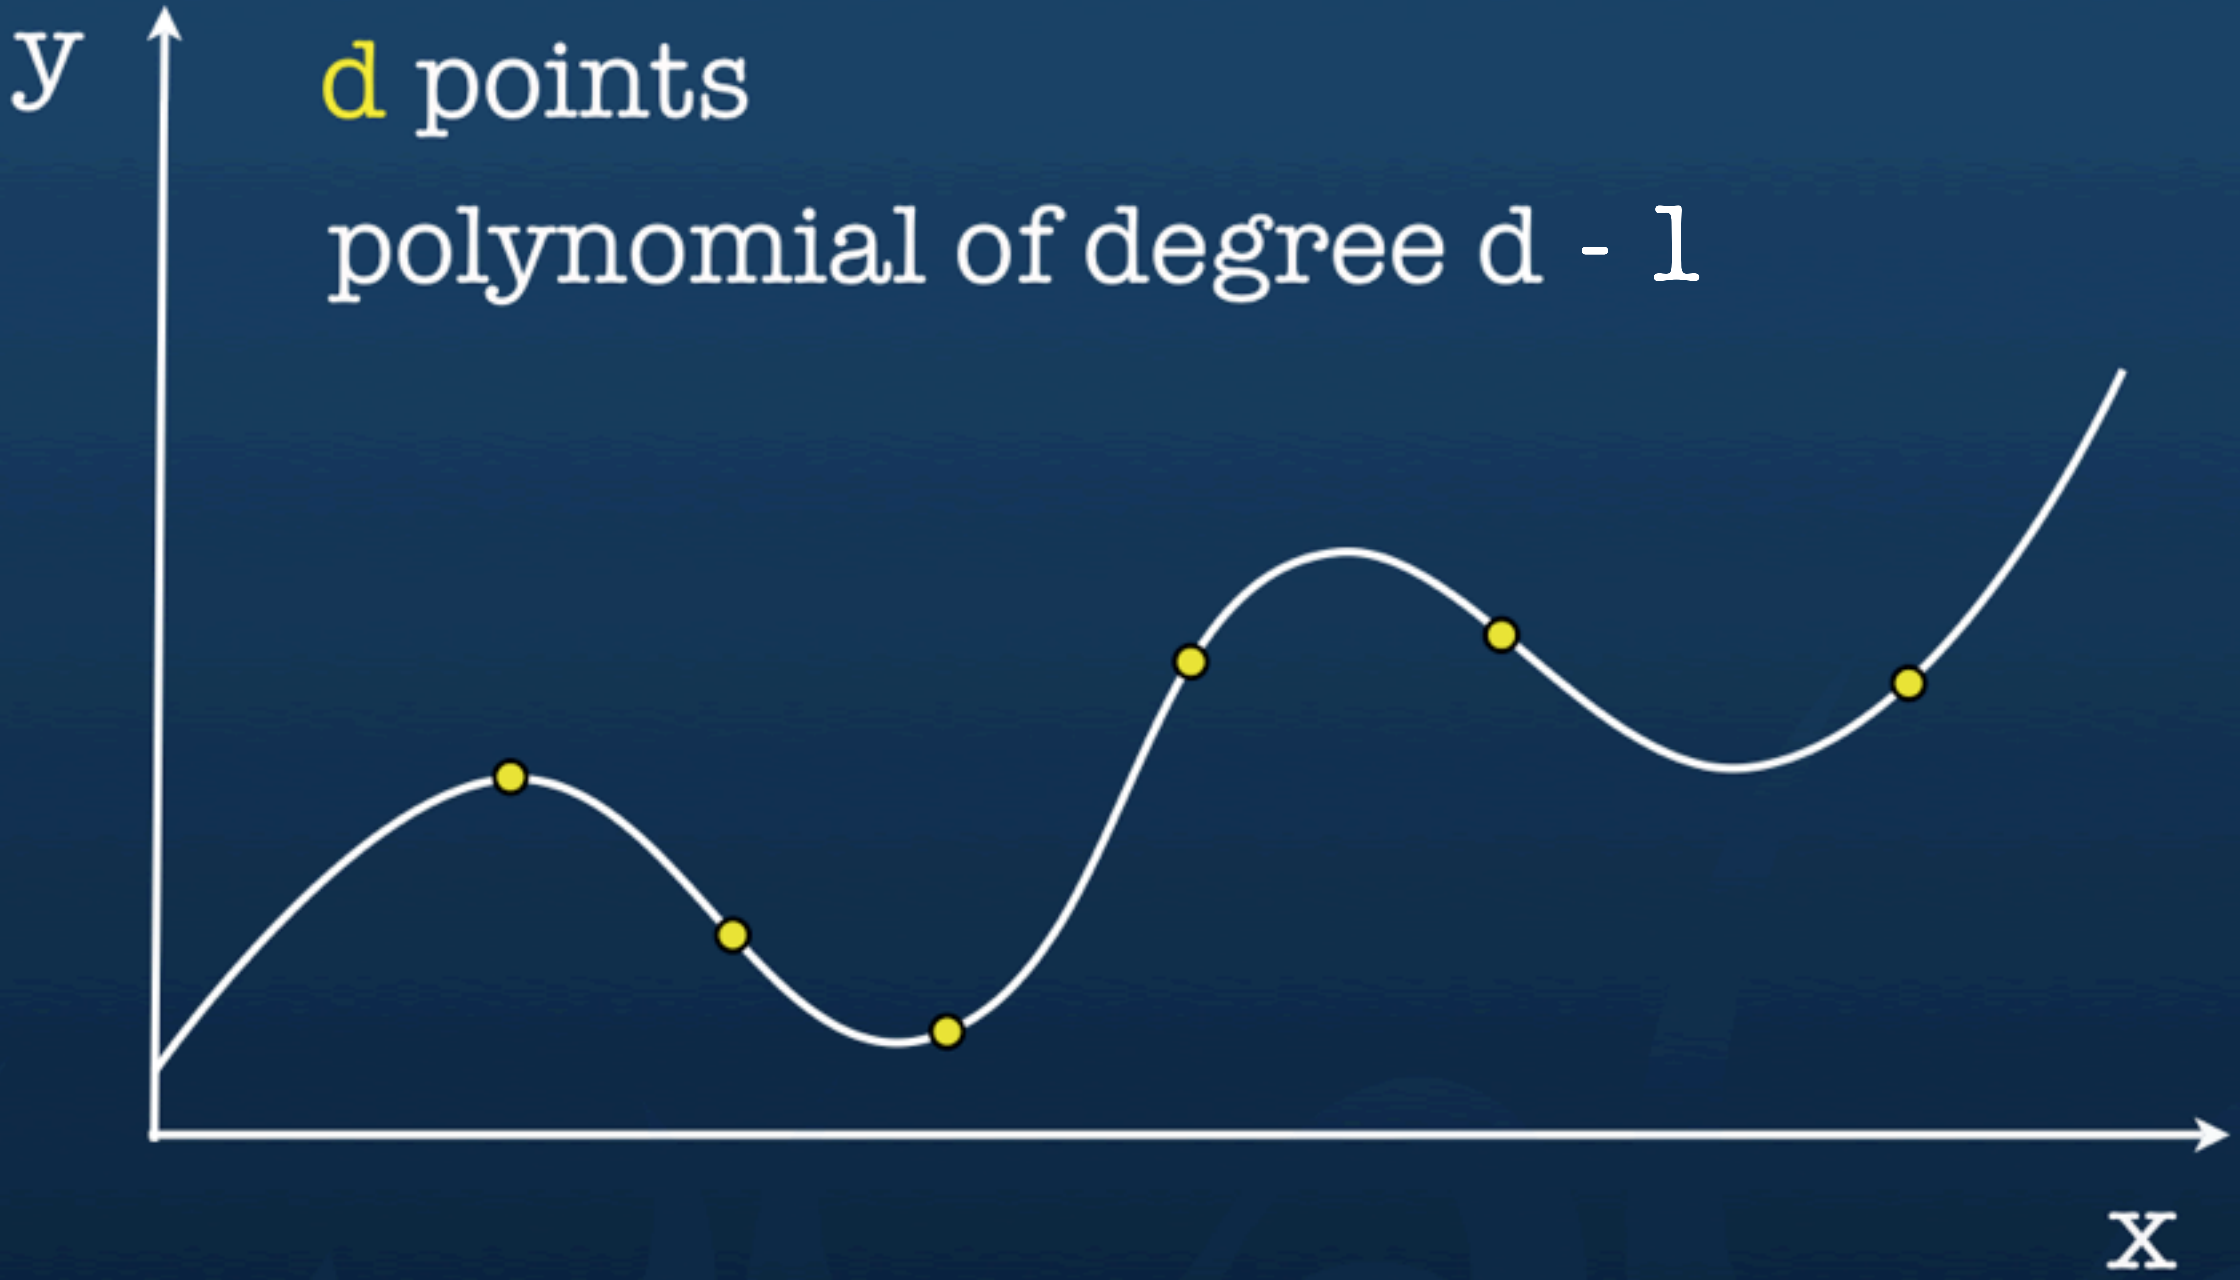

d points

polynomial of degree d - 1

d points

polynomial of degree d - 1

any d points

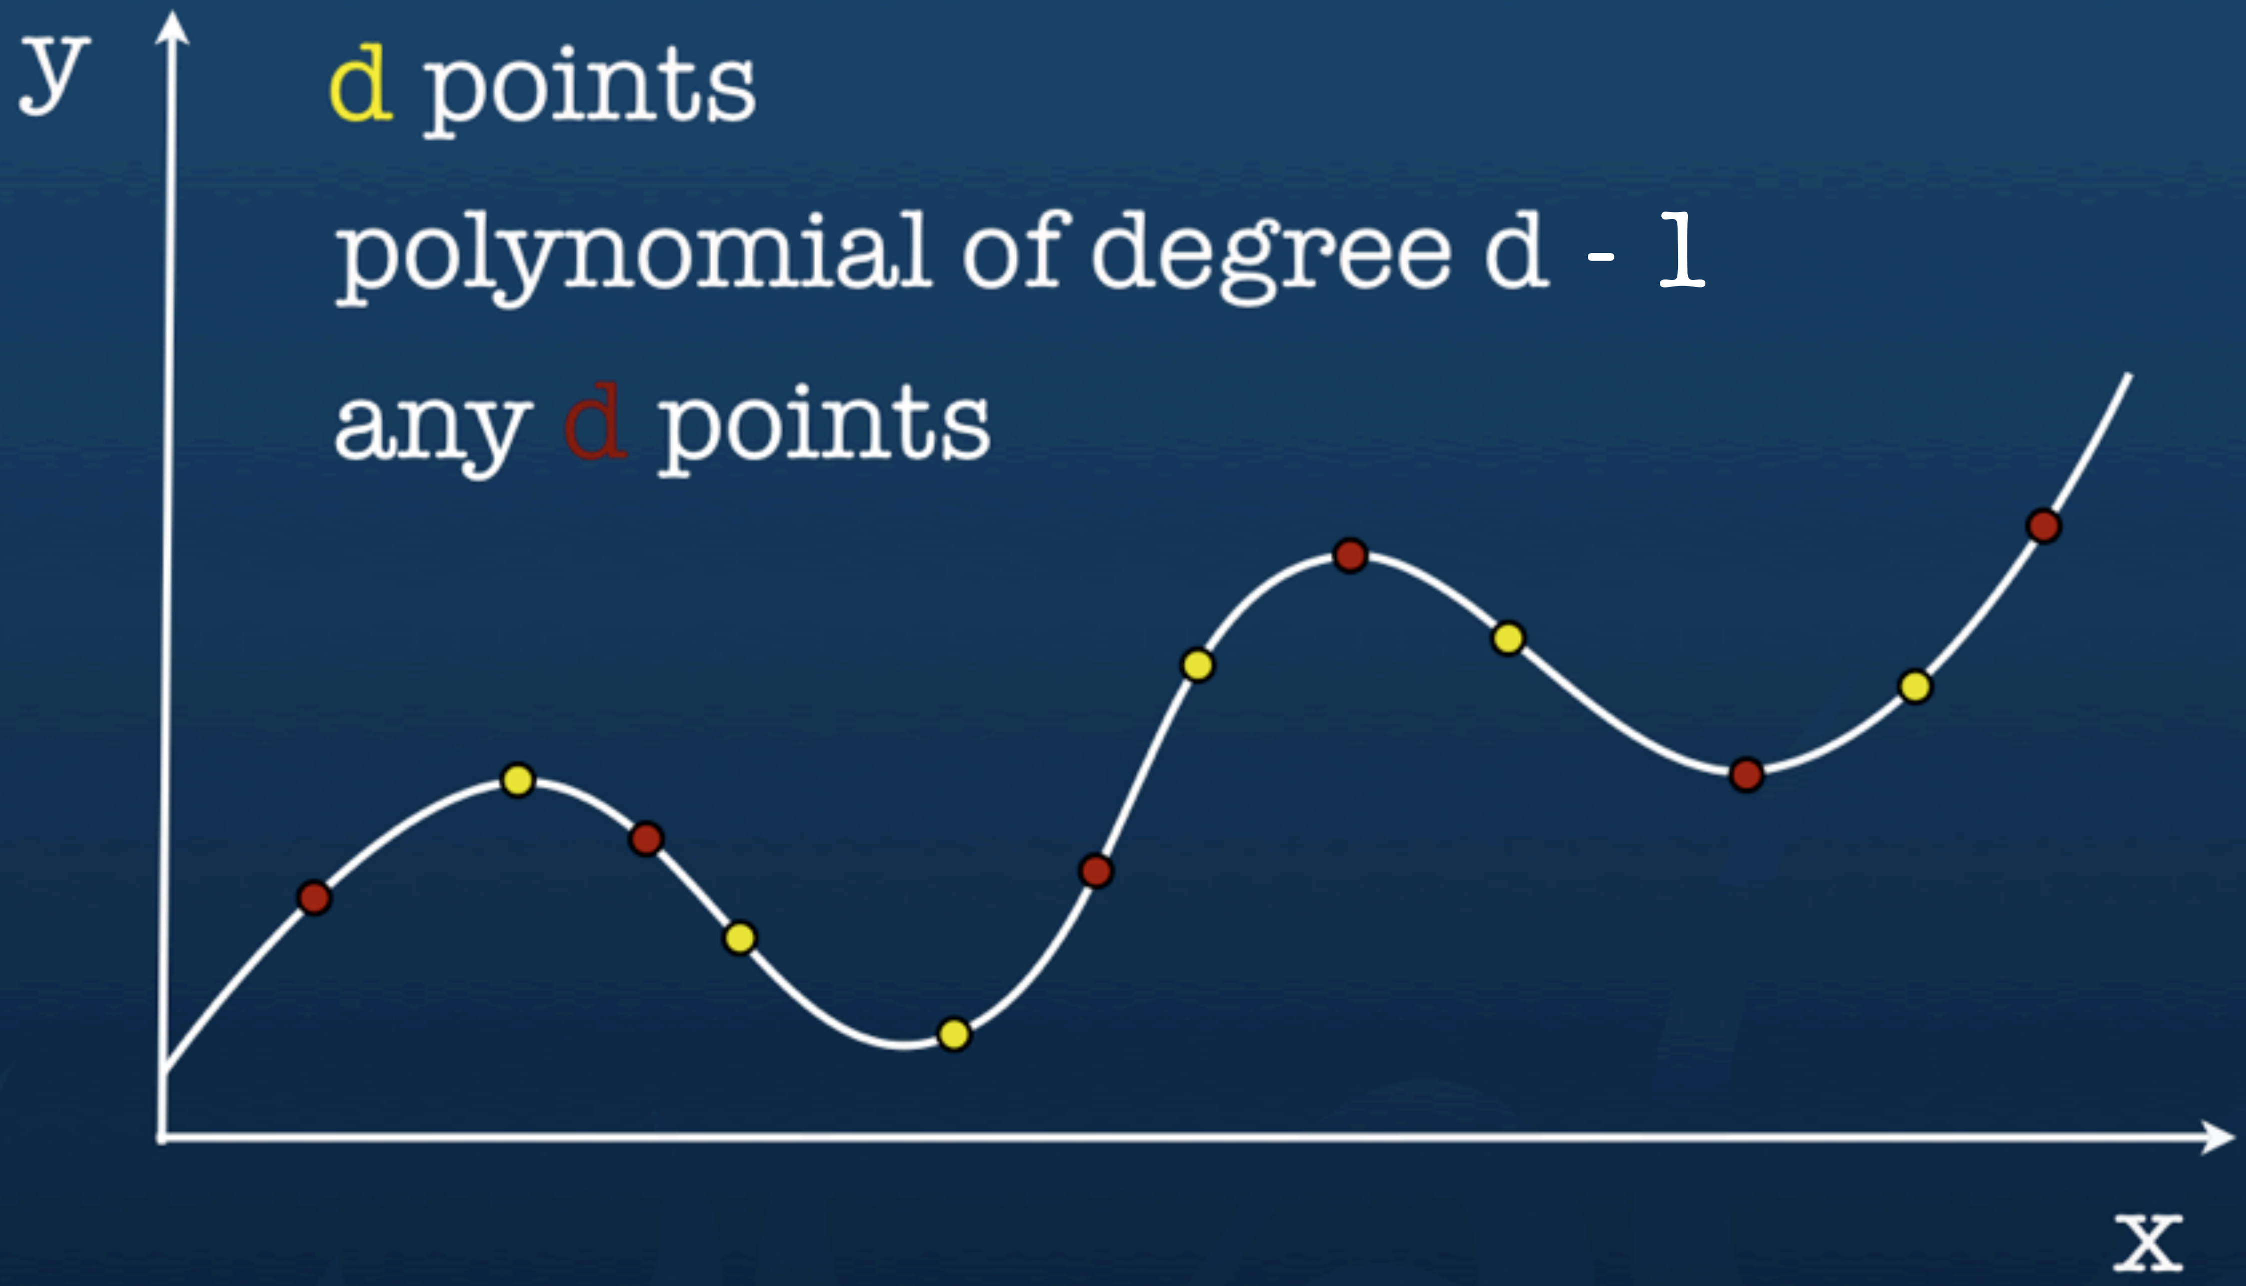

alice stores a file

roadtrip.mpg

alice drags roadtrip.mpg into wuala

1. encrypted on alice's computer (128 bit AES)

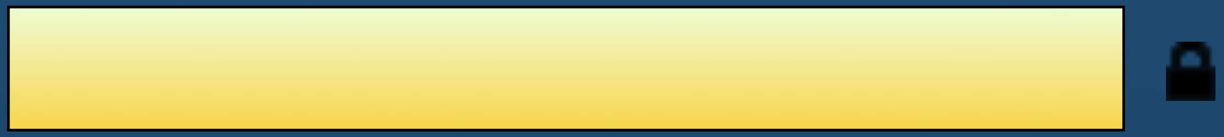

1. encrypted on alice’s computer (128 bit AES)

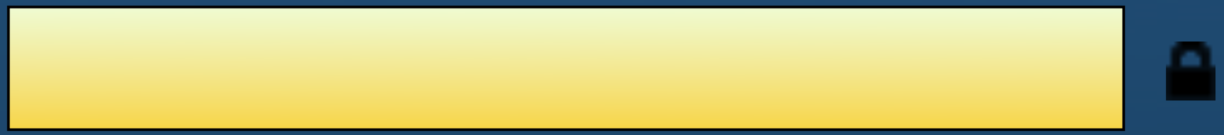

2. encoded into redundant fragments

1. encrypted on alice's computer (128 bit AES)

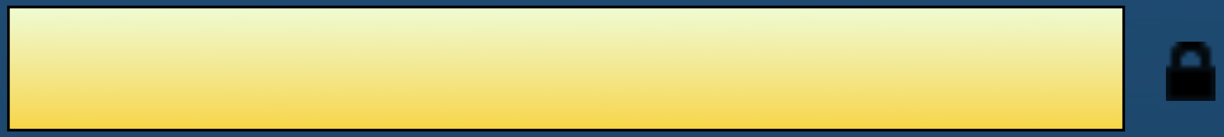

2. encoded into redundant fragments

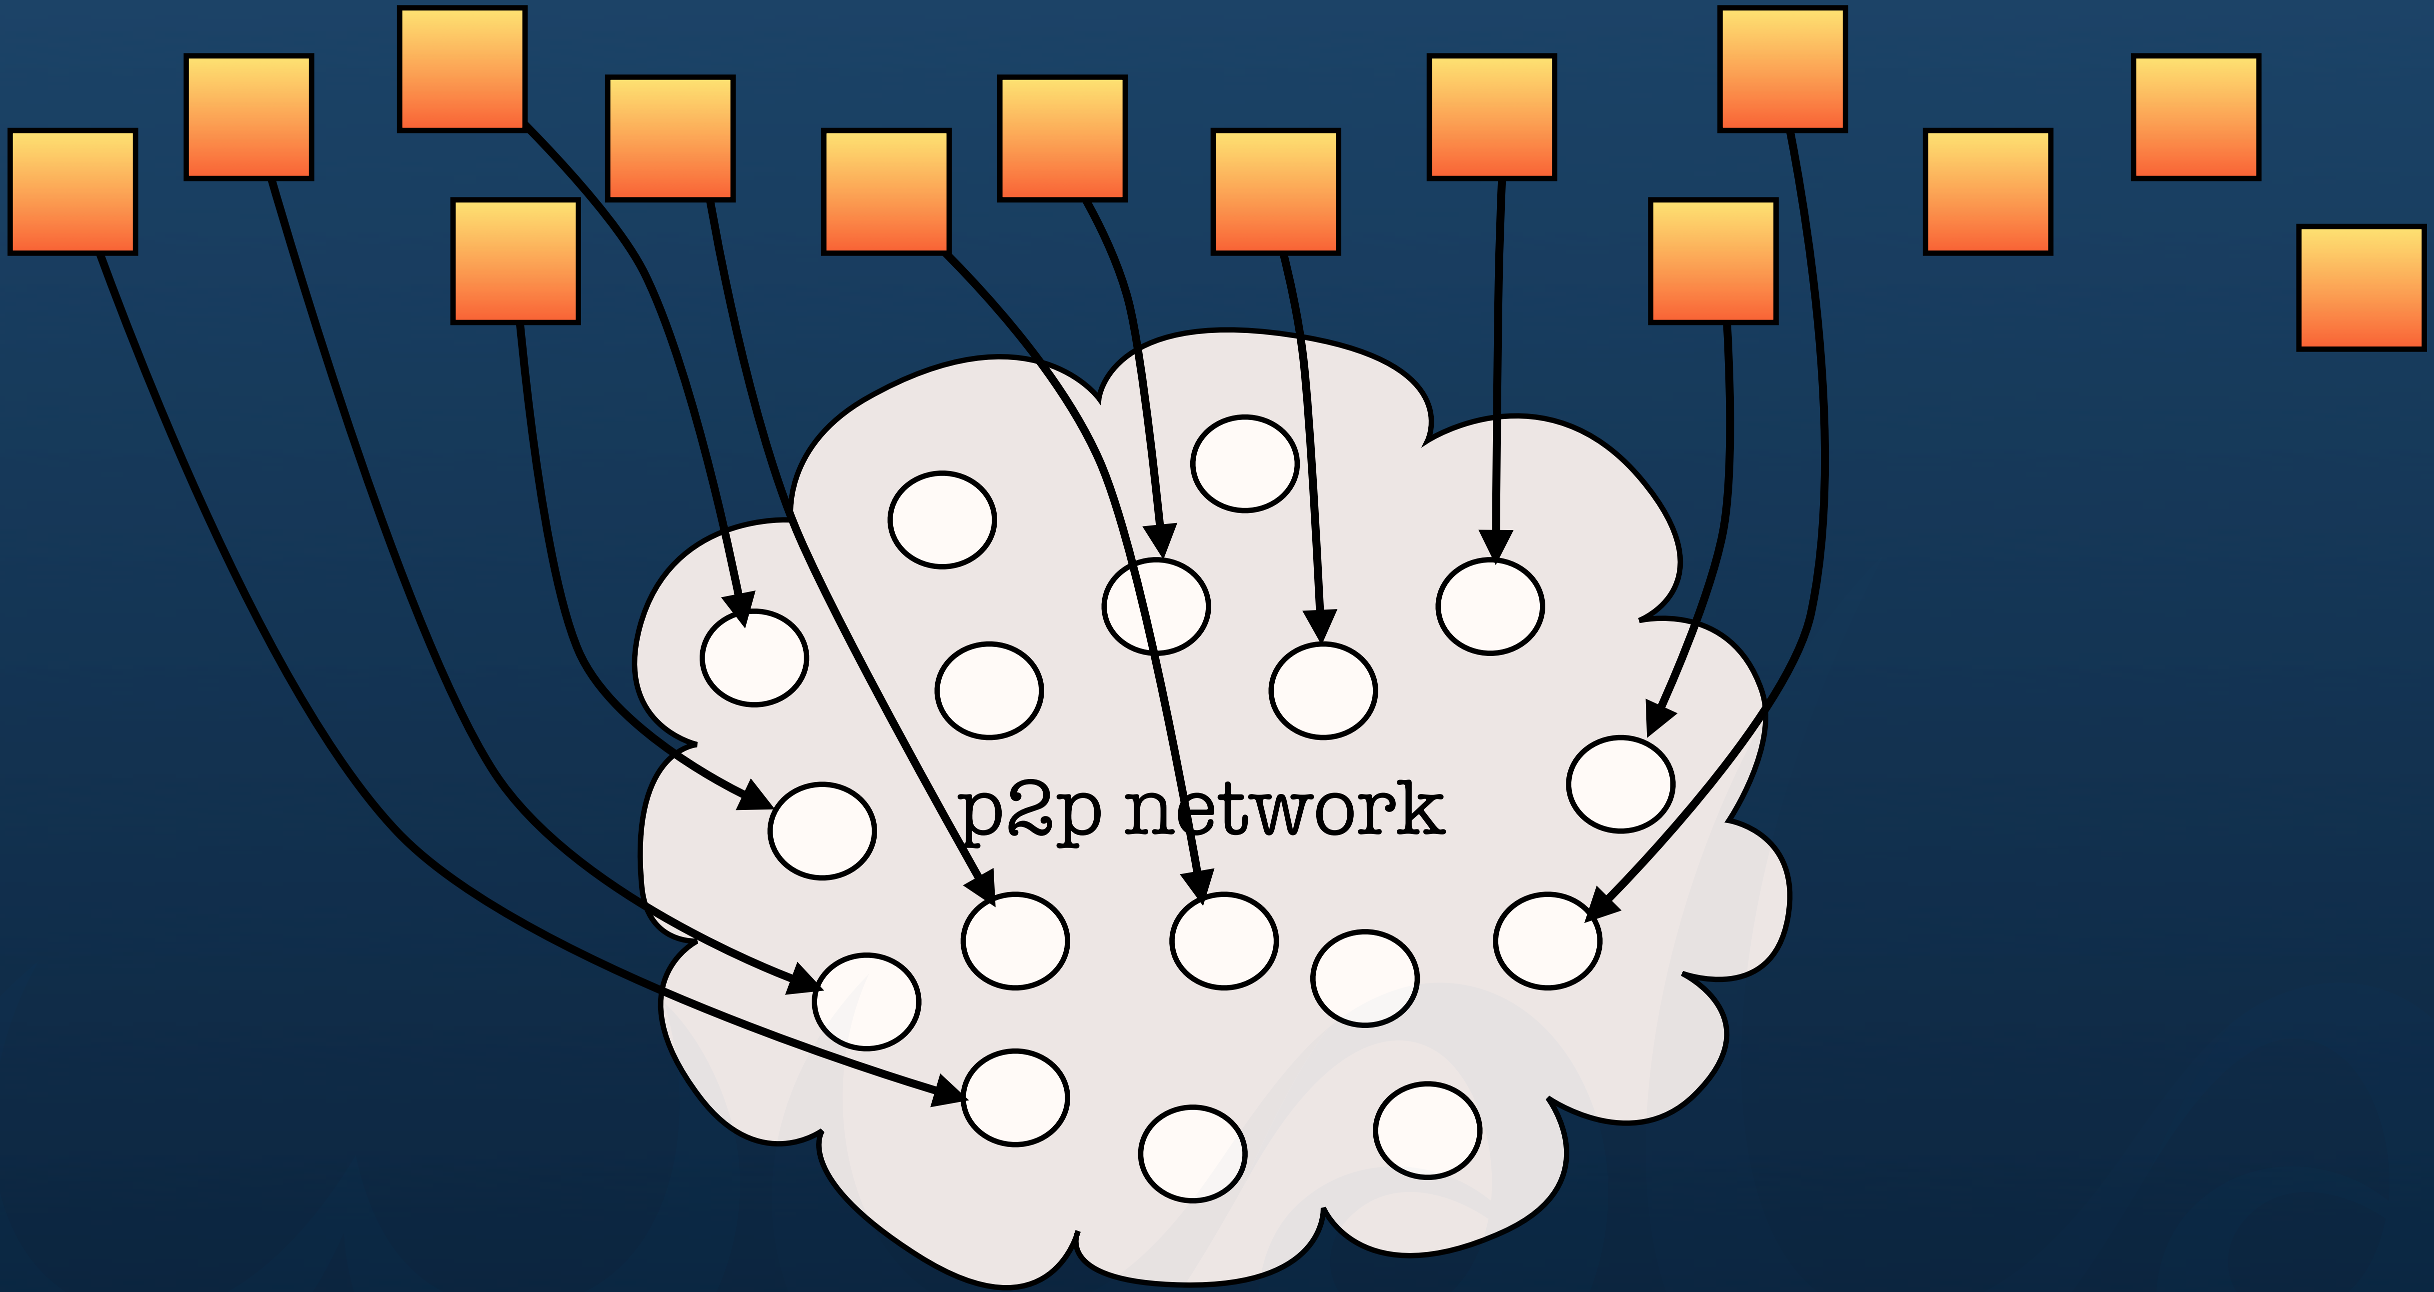

3. uploaded into the p2p network

1. encrypted on alice’s computer (128 bit AES)

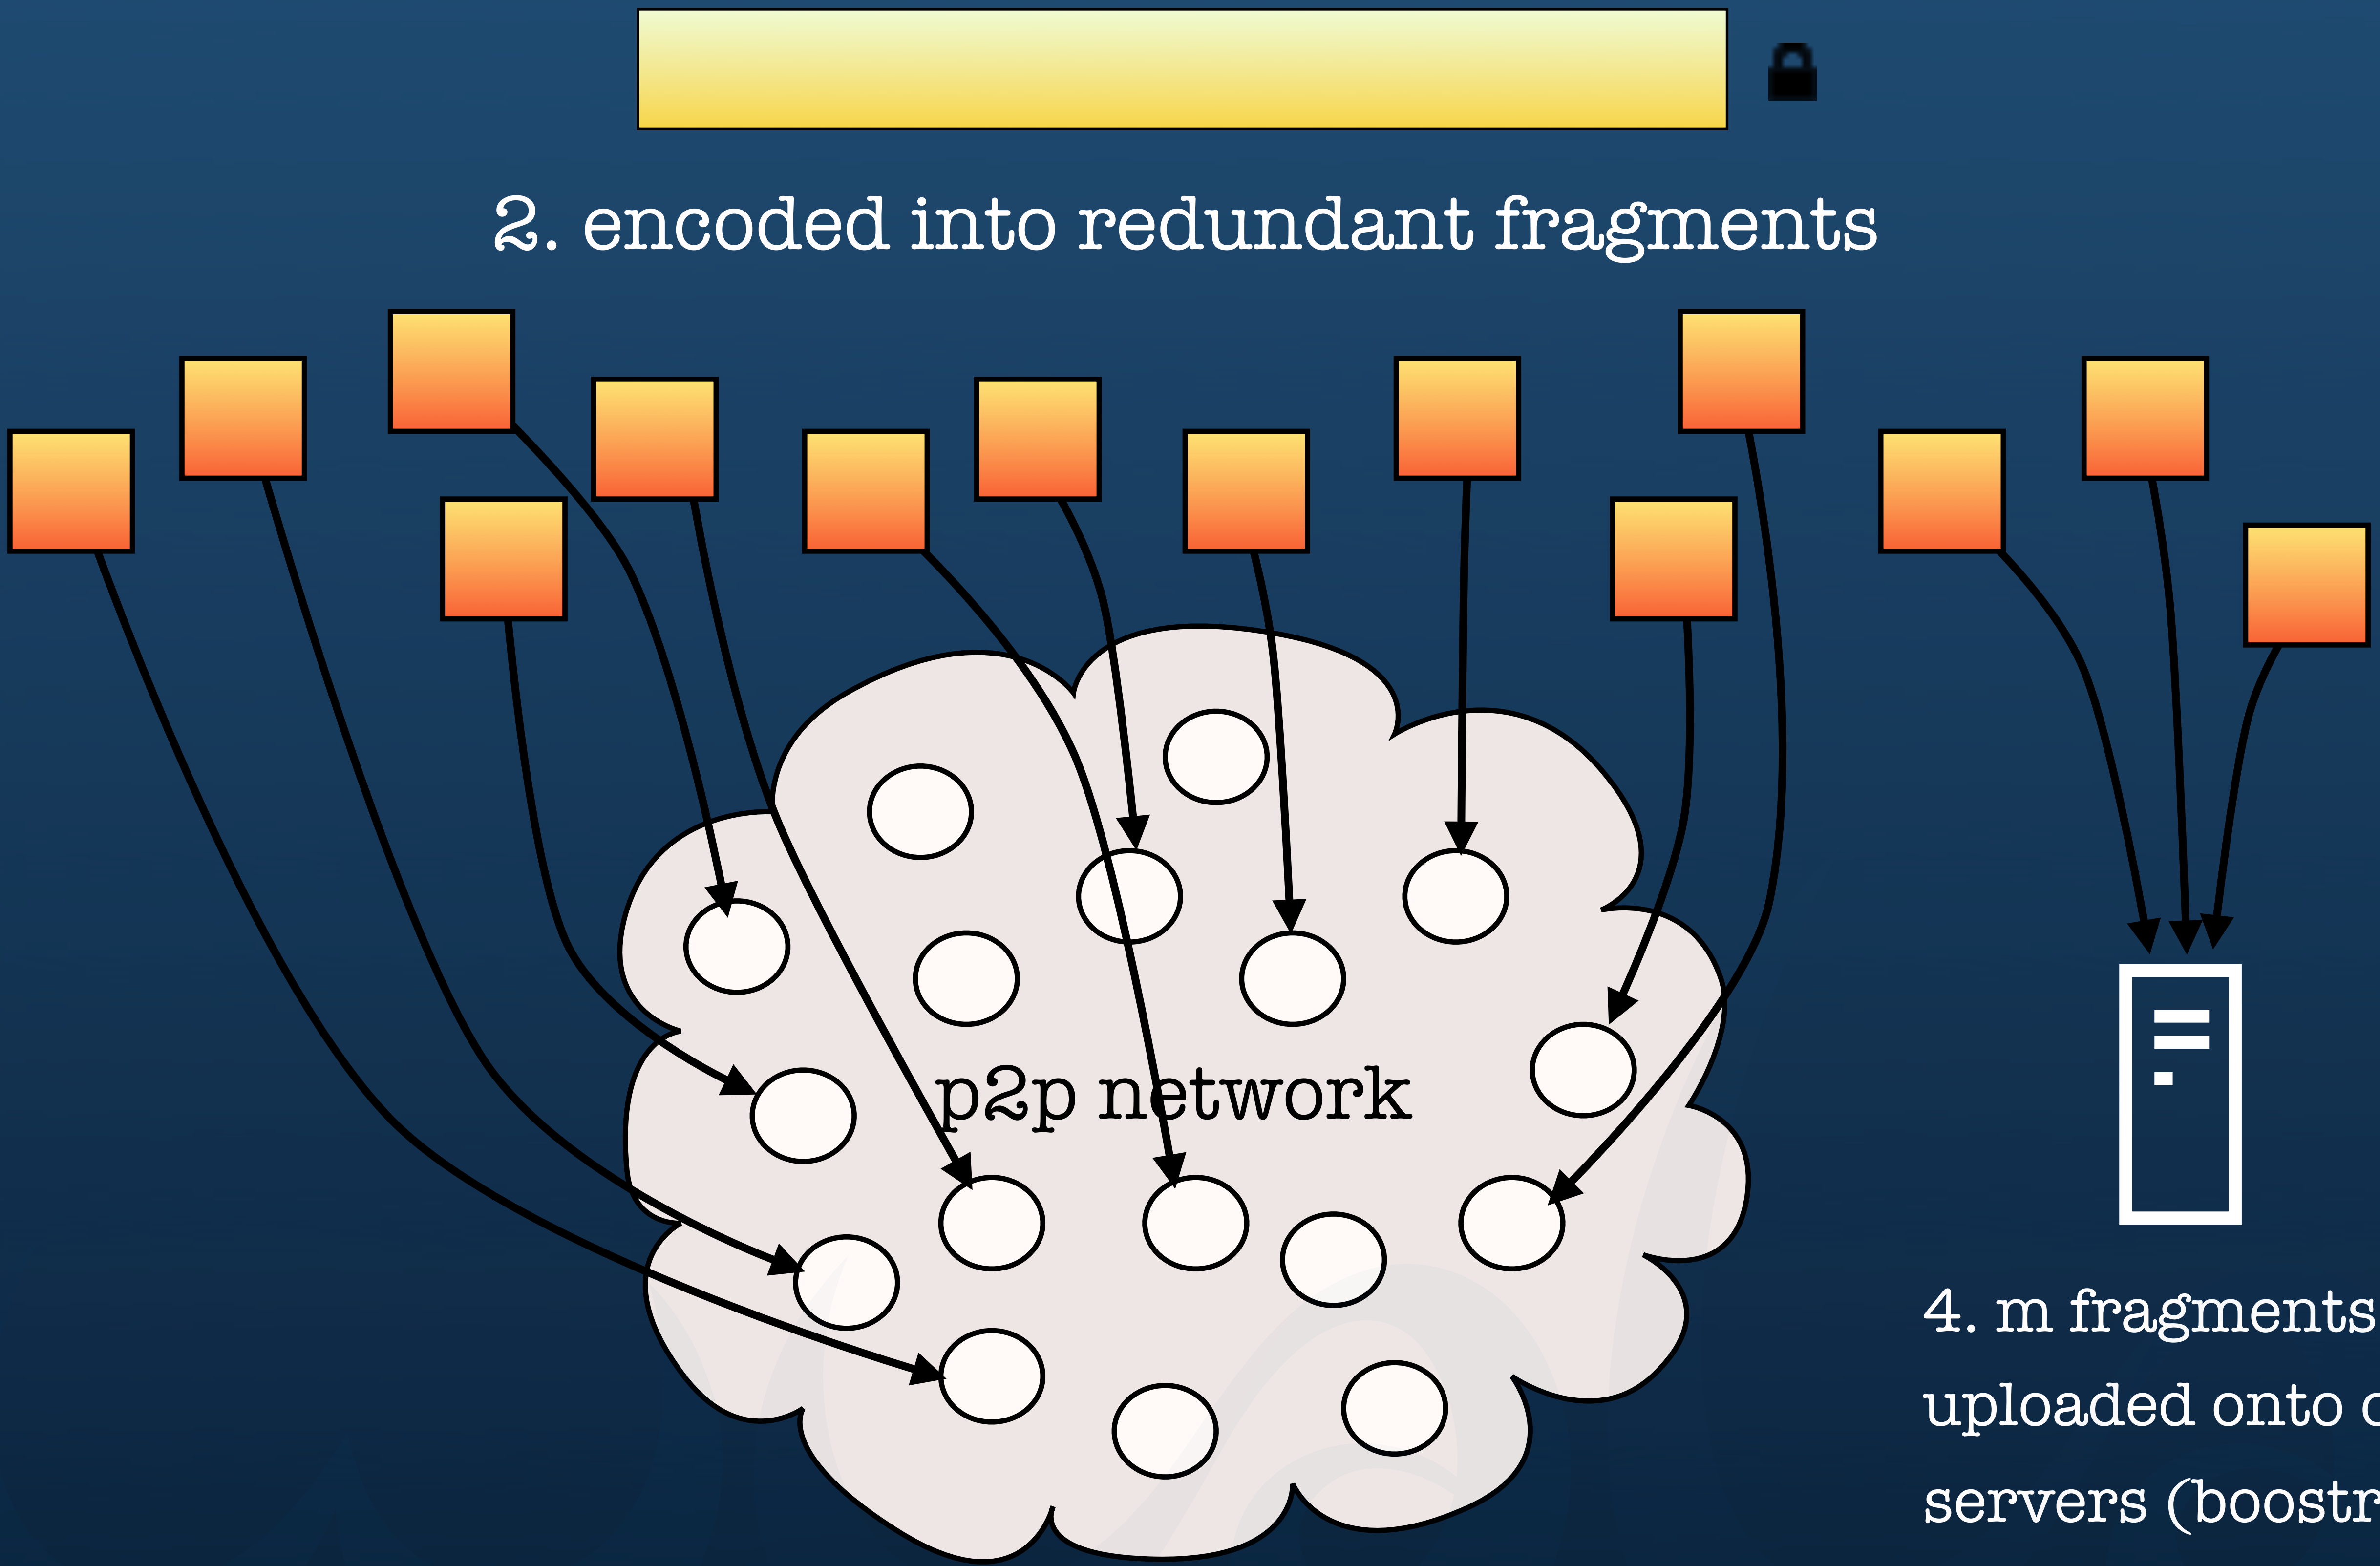

2. encoded into redundant fragments

3. uploaded into the p2p network

4. m fragments uploaded onto our servers (boostrap, backup)

alice shares the file with bob

alice and bob have friendship key

alice encrypts file key and exchanges it with bob

bob wants to download the file

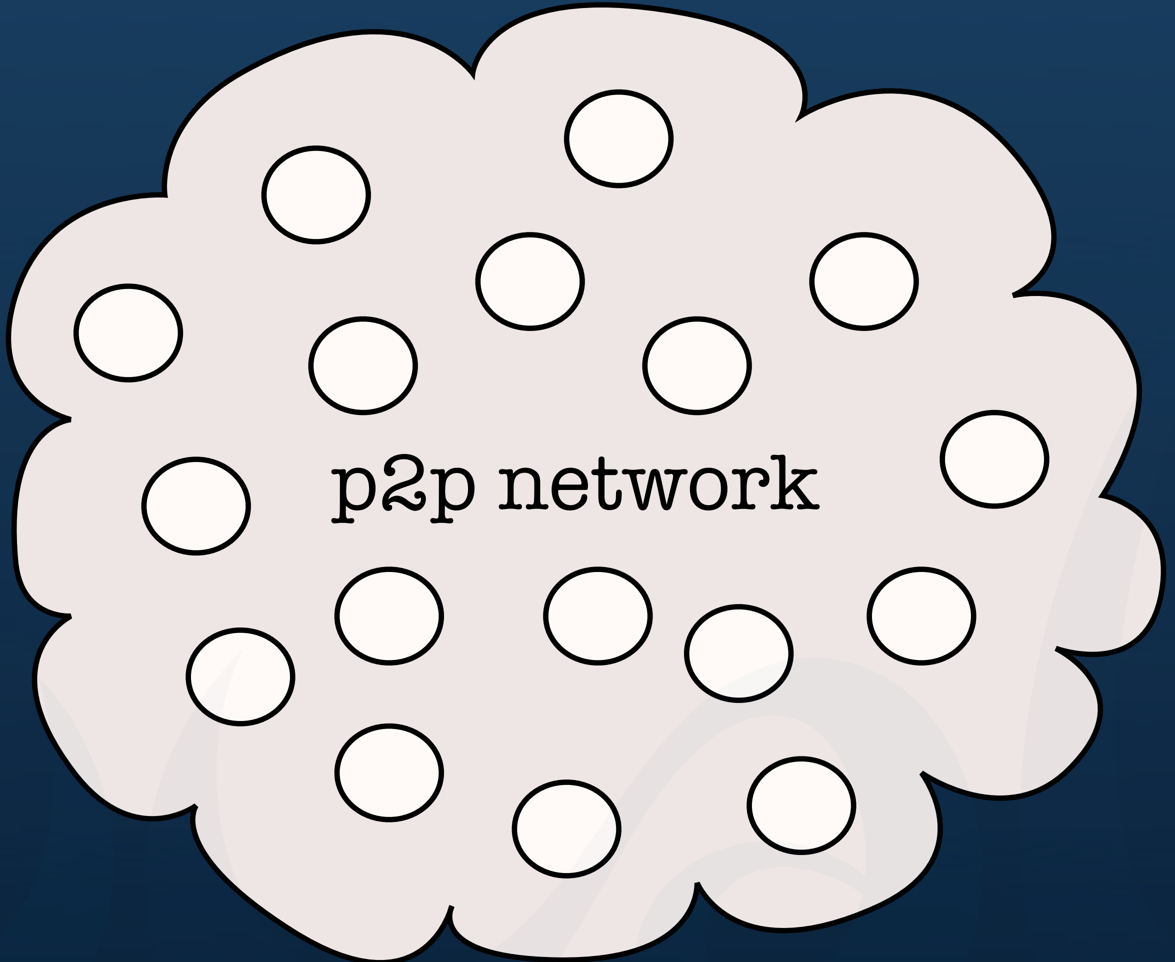
p2p network

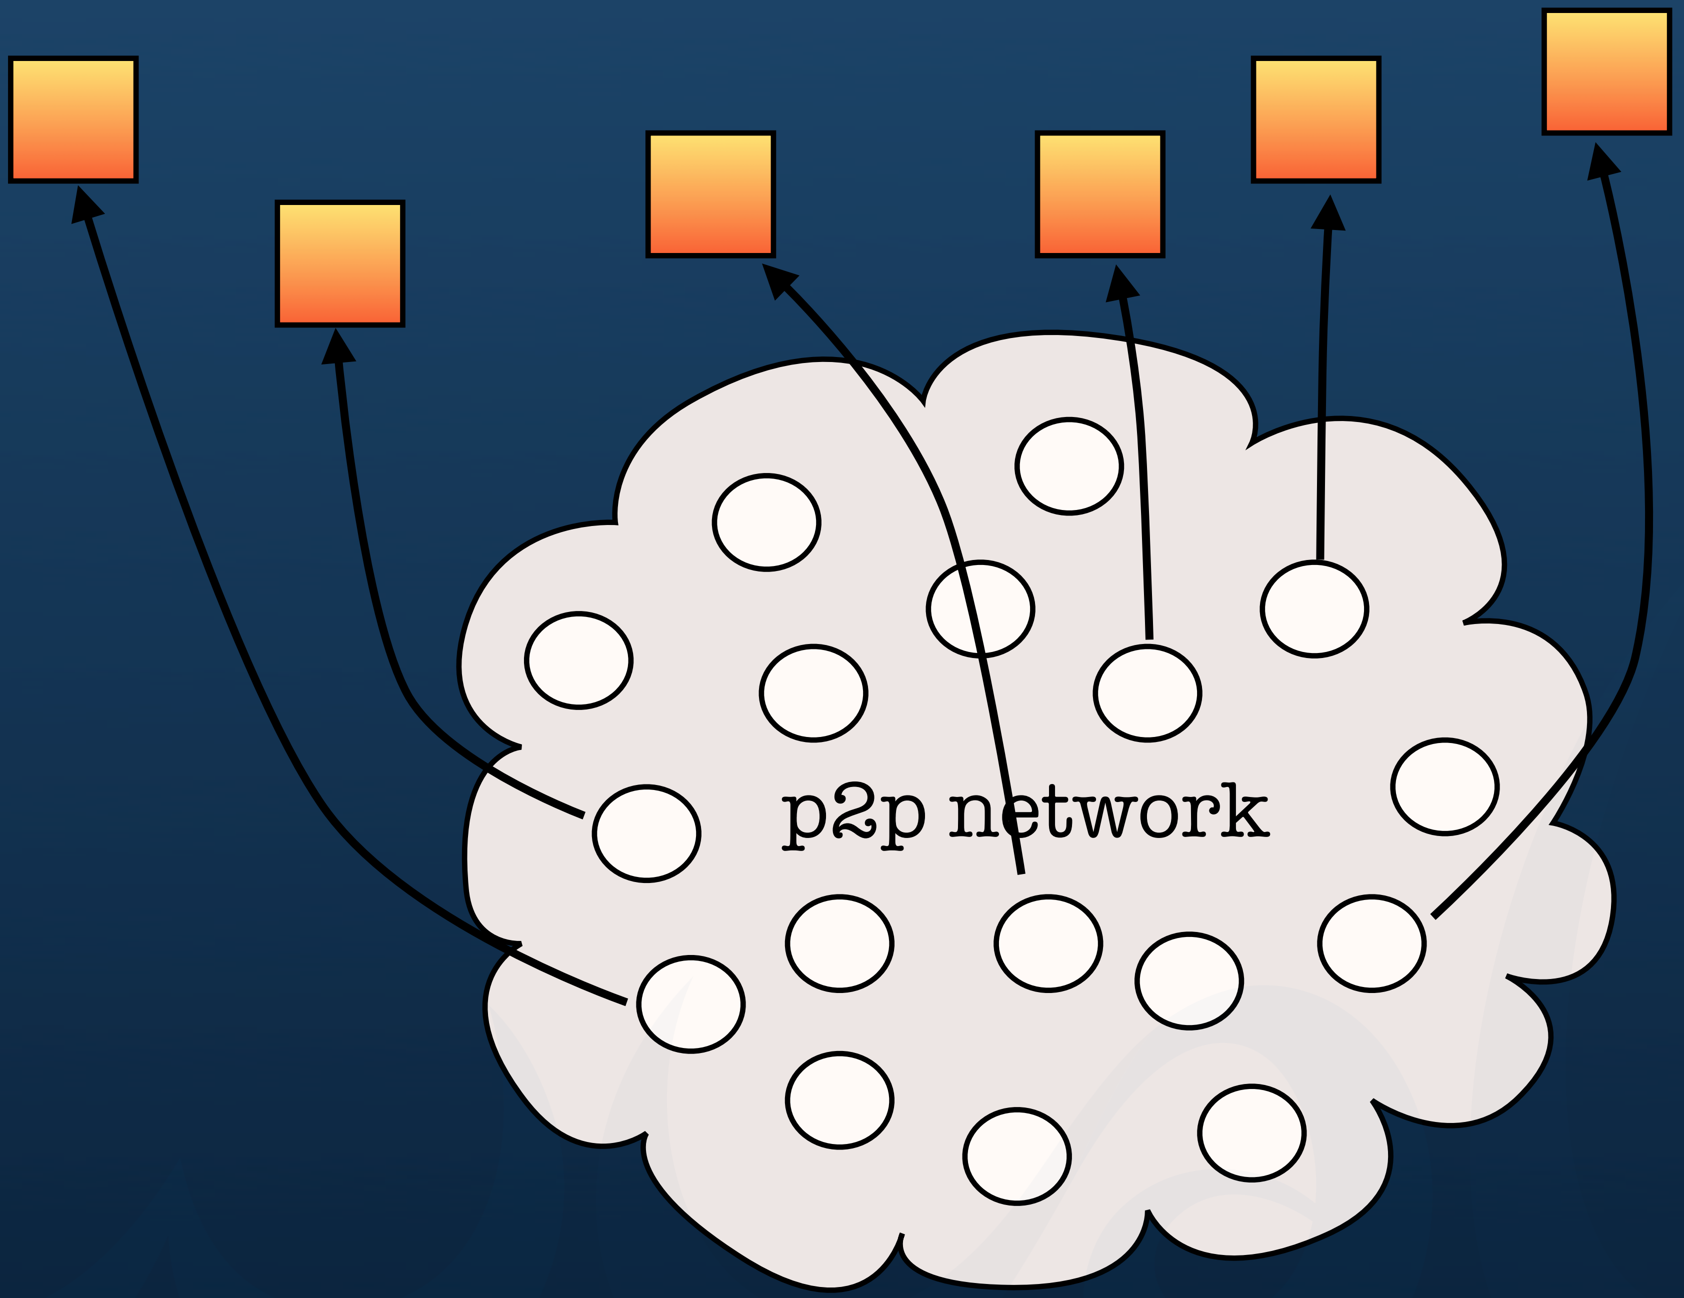

1. download subset of fragments (m)

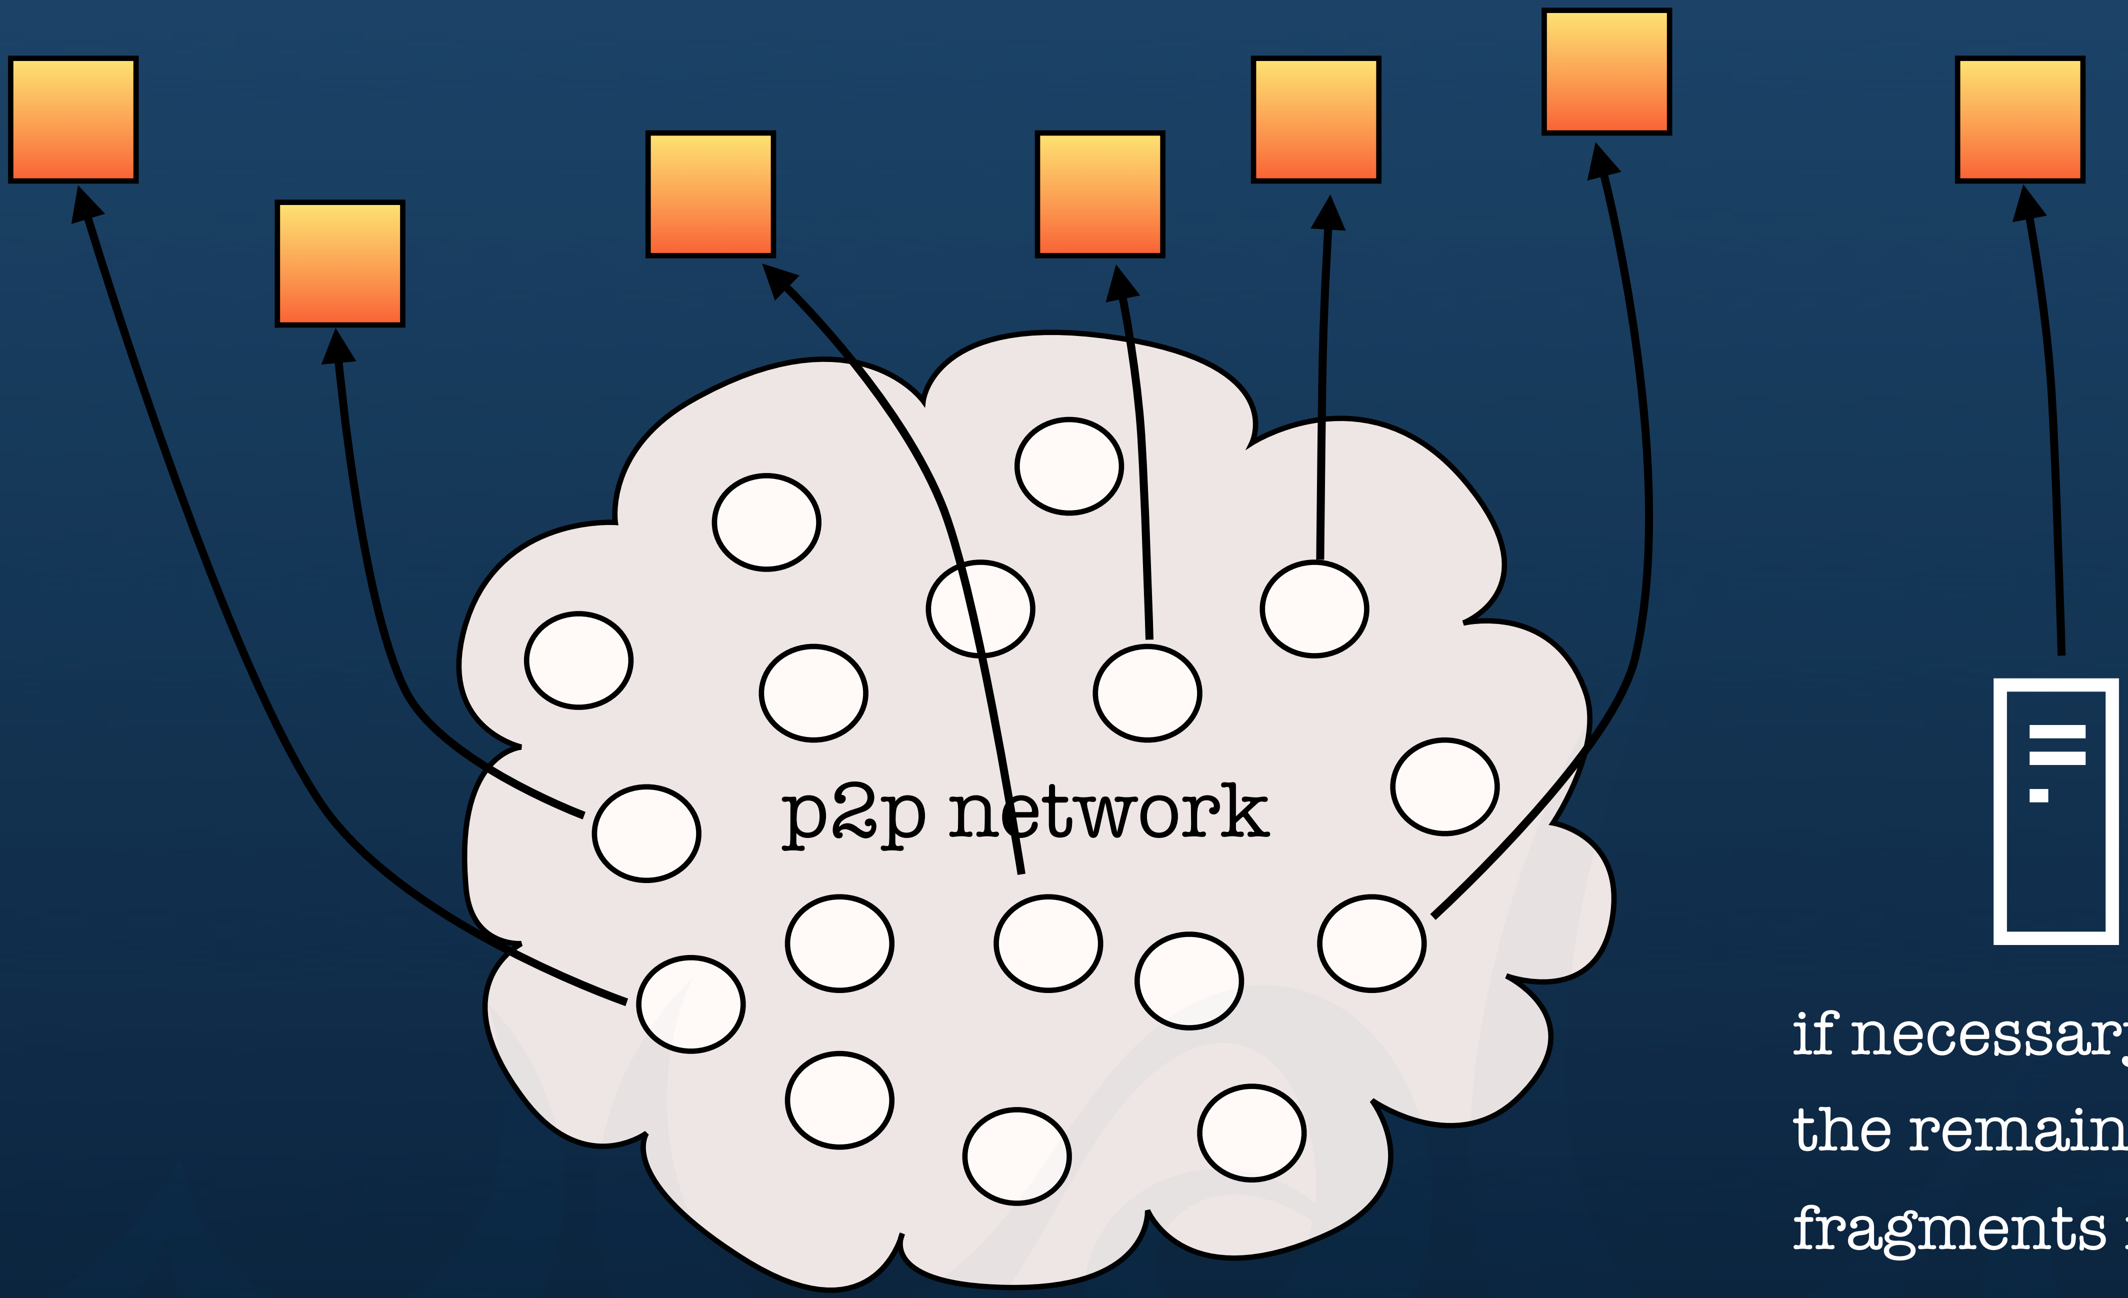

if necessary, get
the remaining
fragments from
our servers

1. download subset of fragments (m)

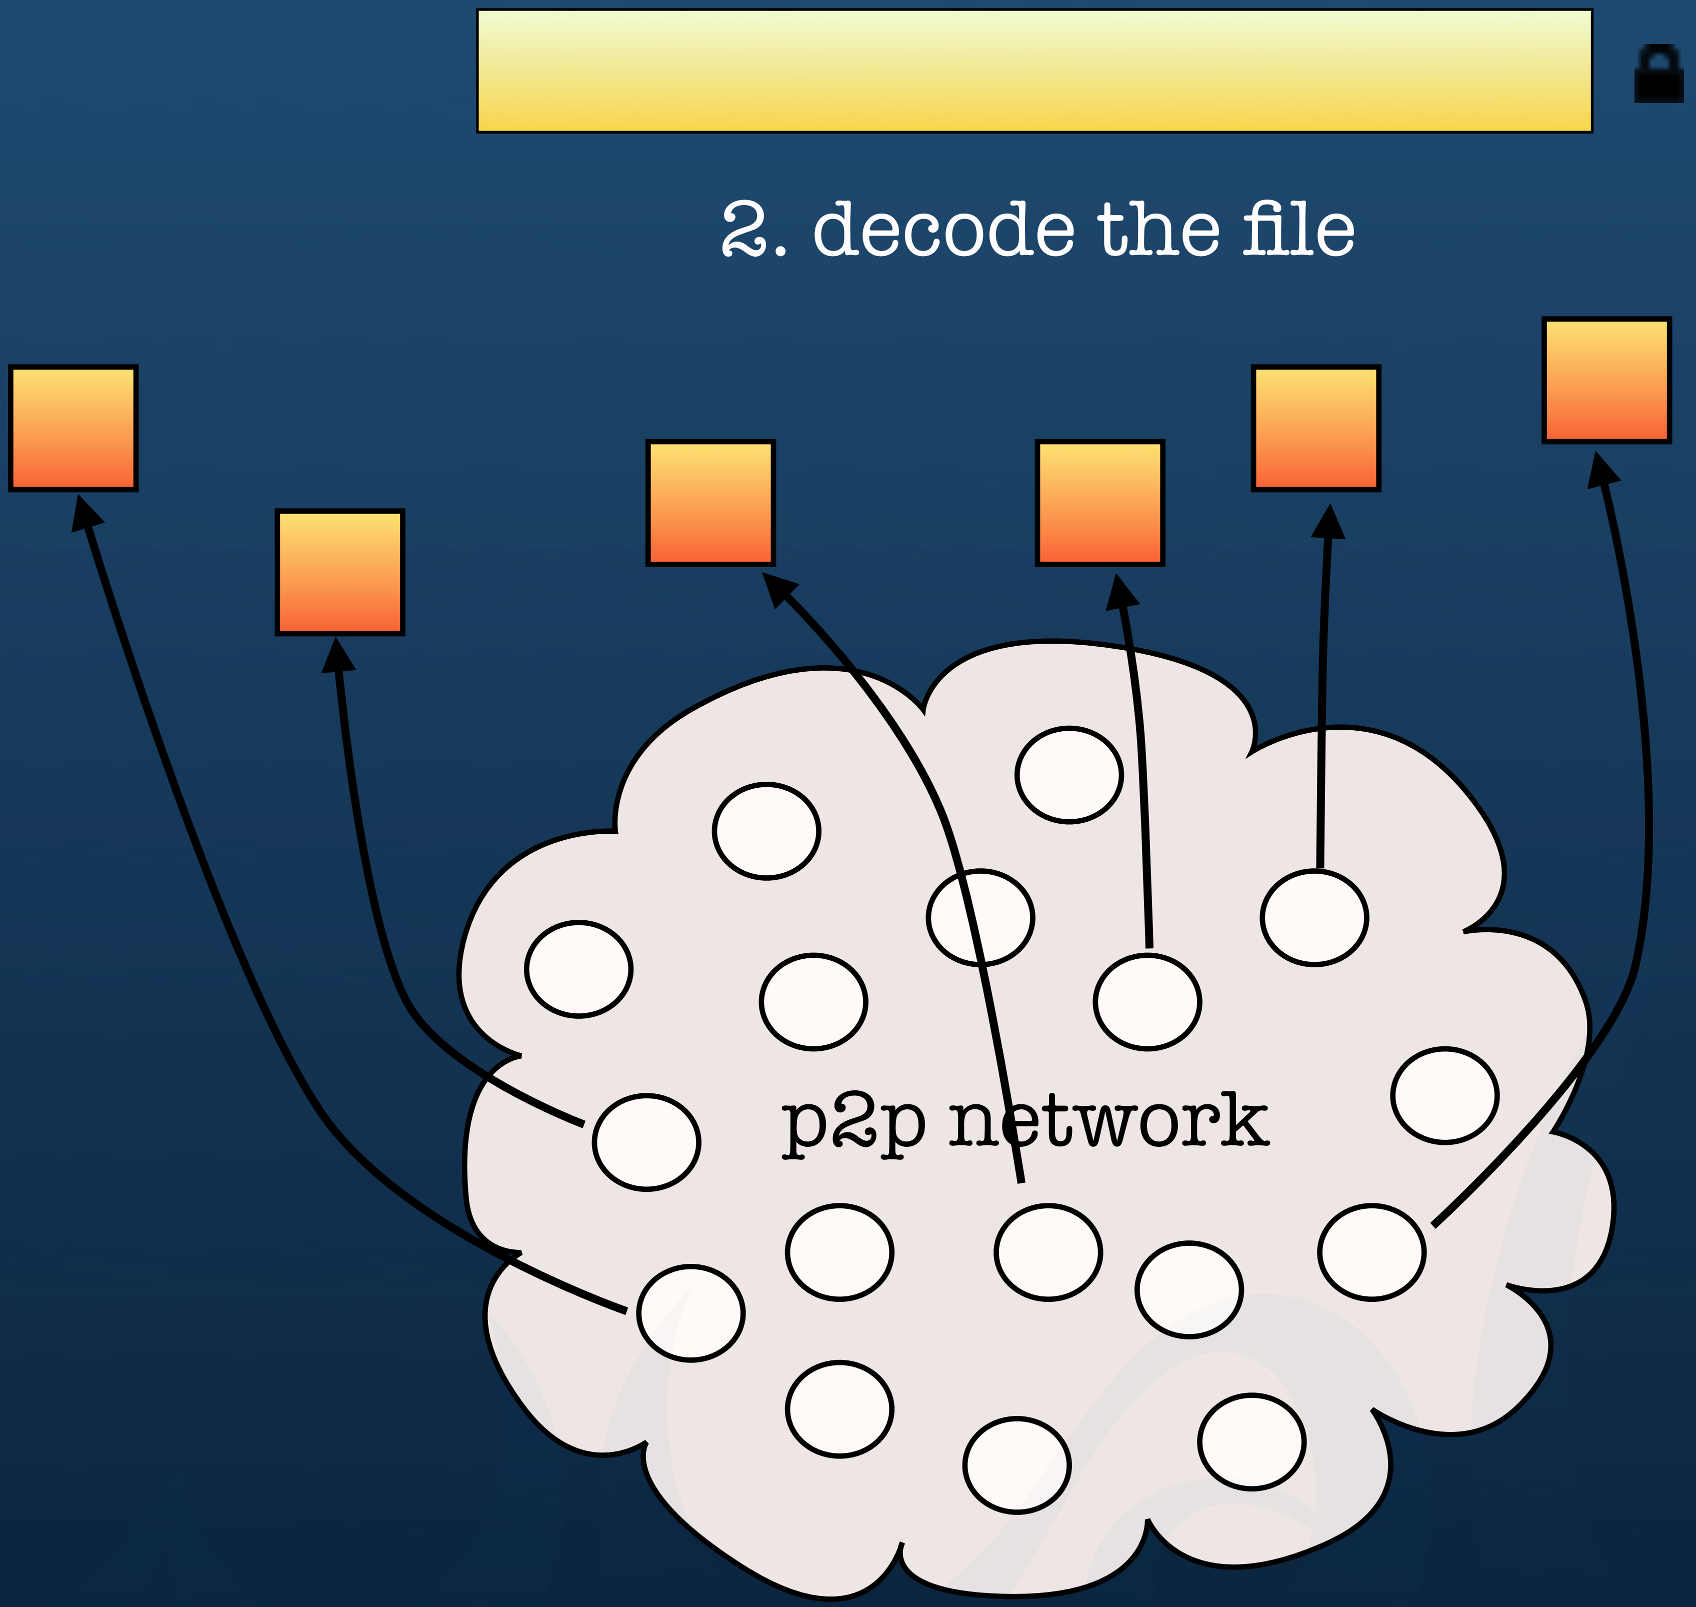

2. decode the file

1. download subset of fragments (m)

3. decrypt the file

2. decode the file

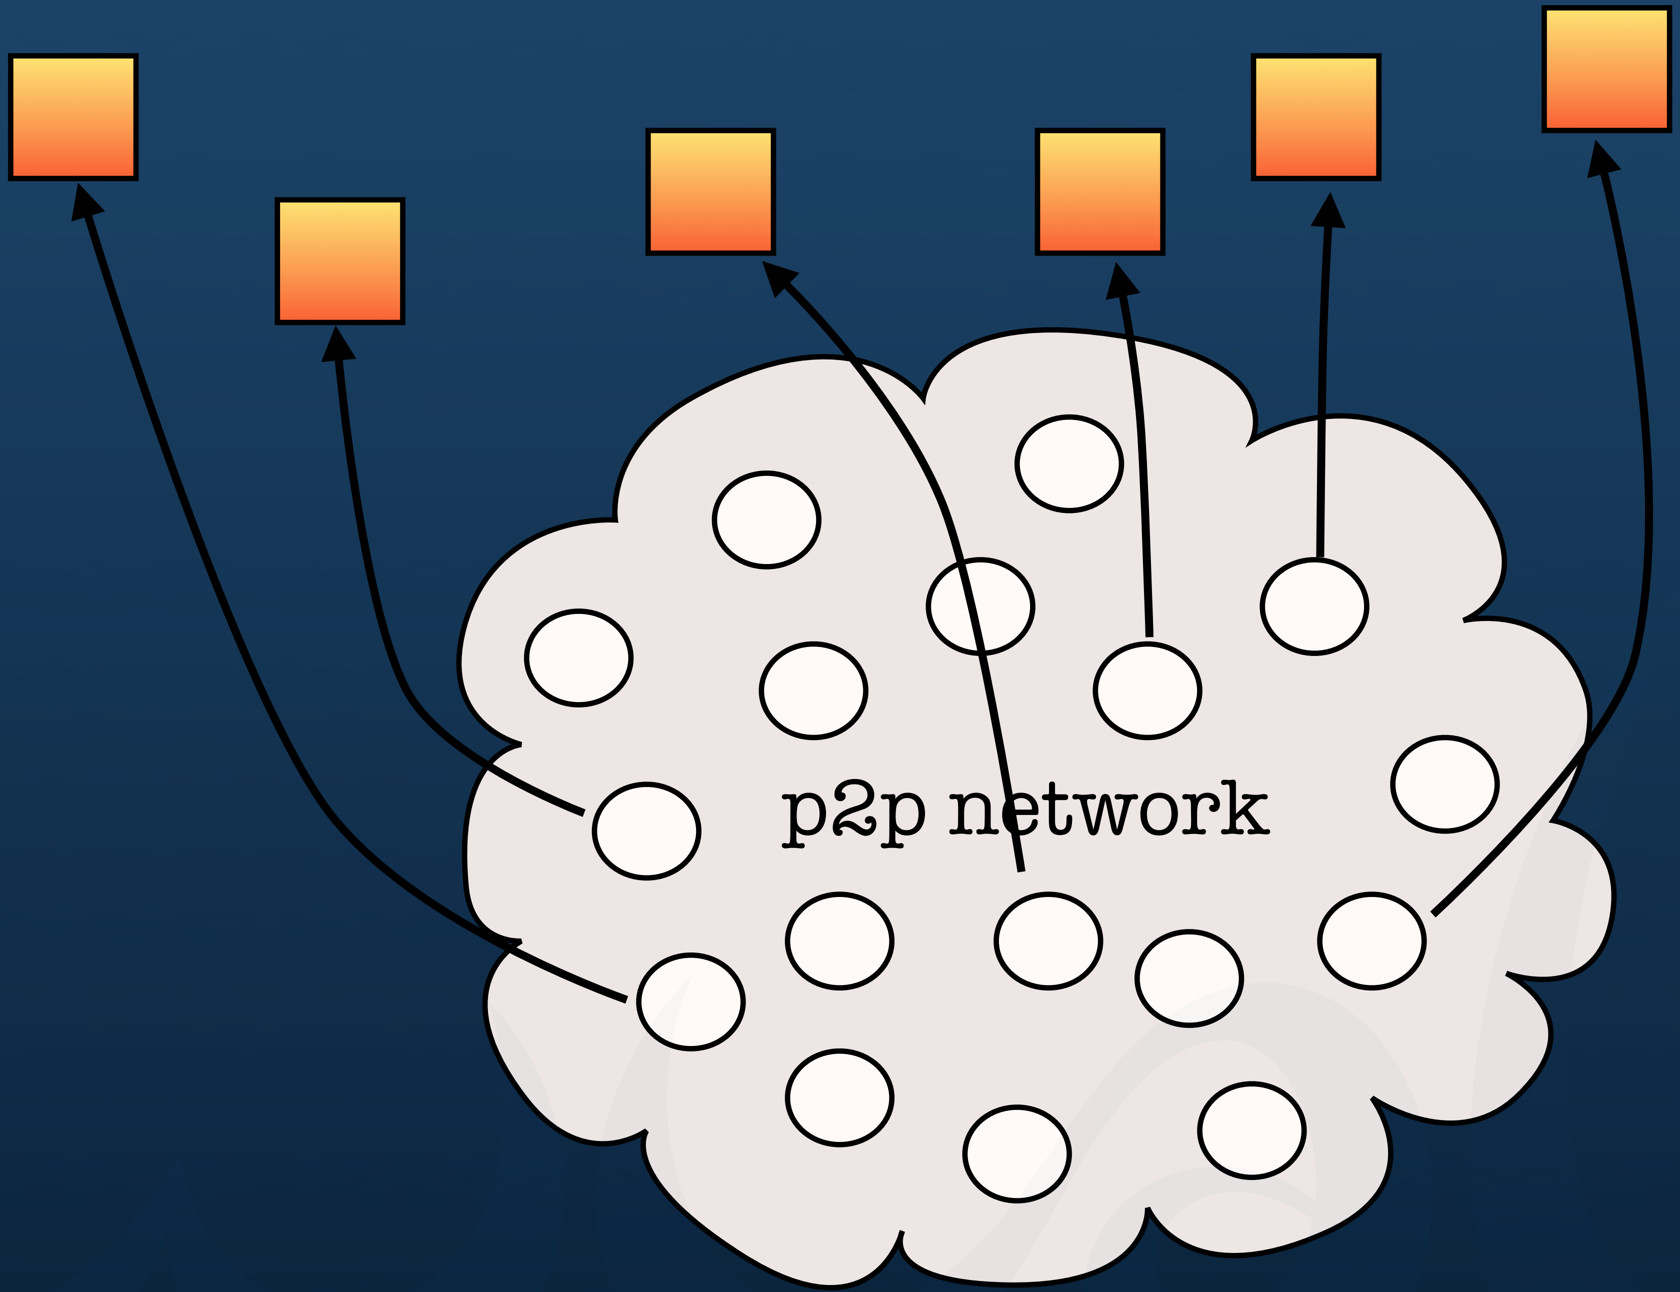

1. download subset of fragments (m)

bob plays roadtrip.mpg

2. decode the file

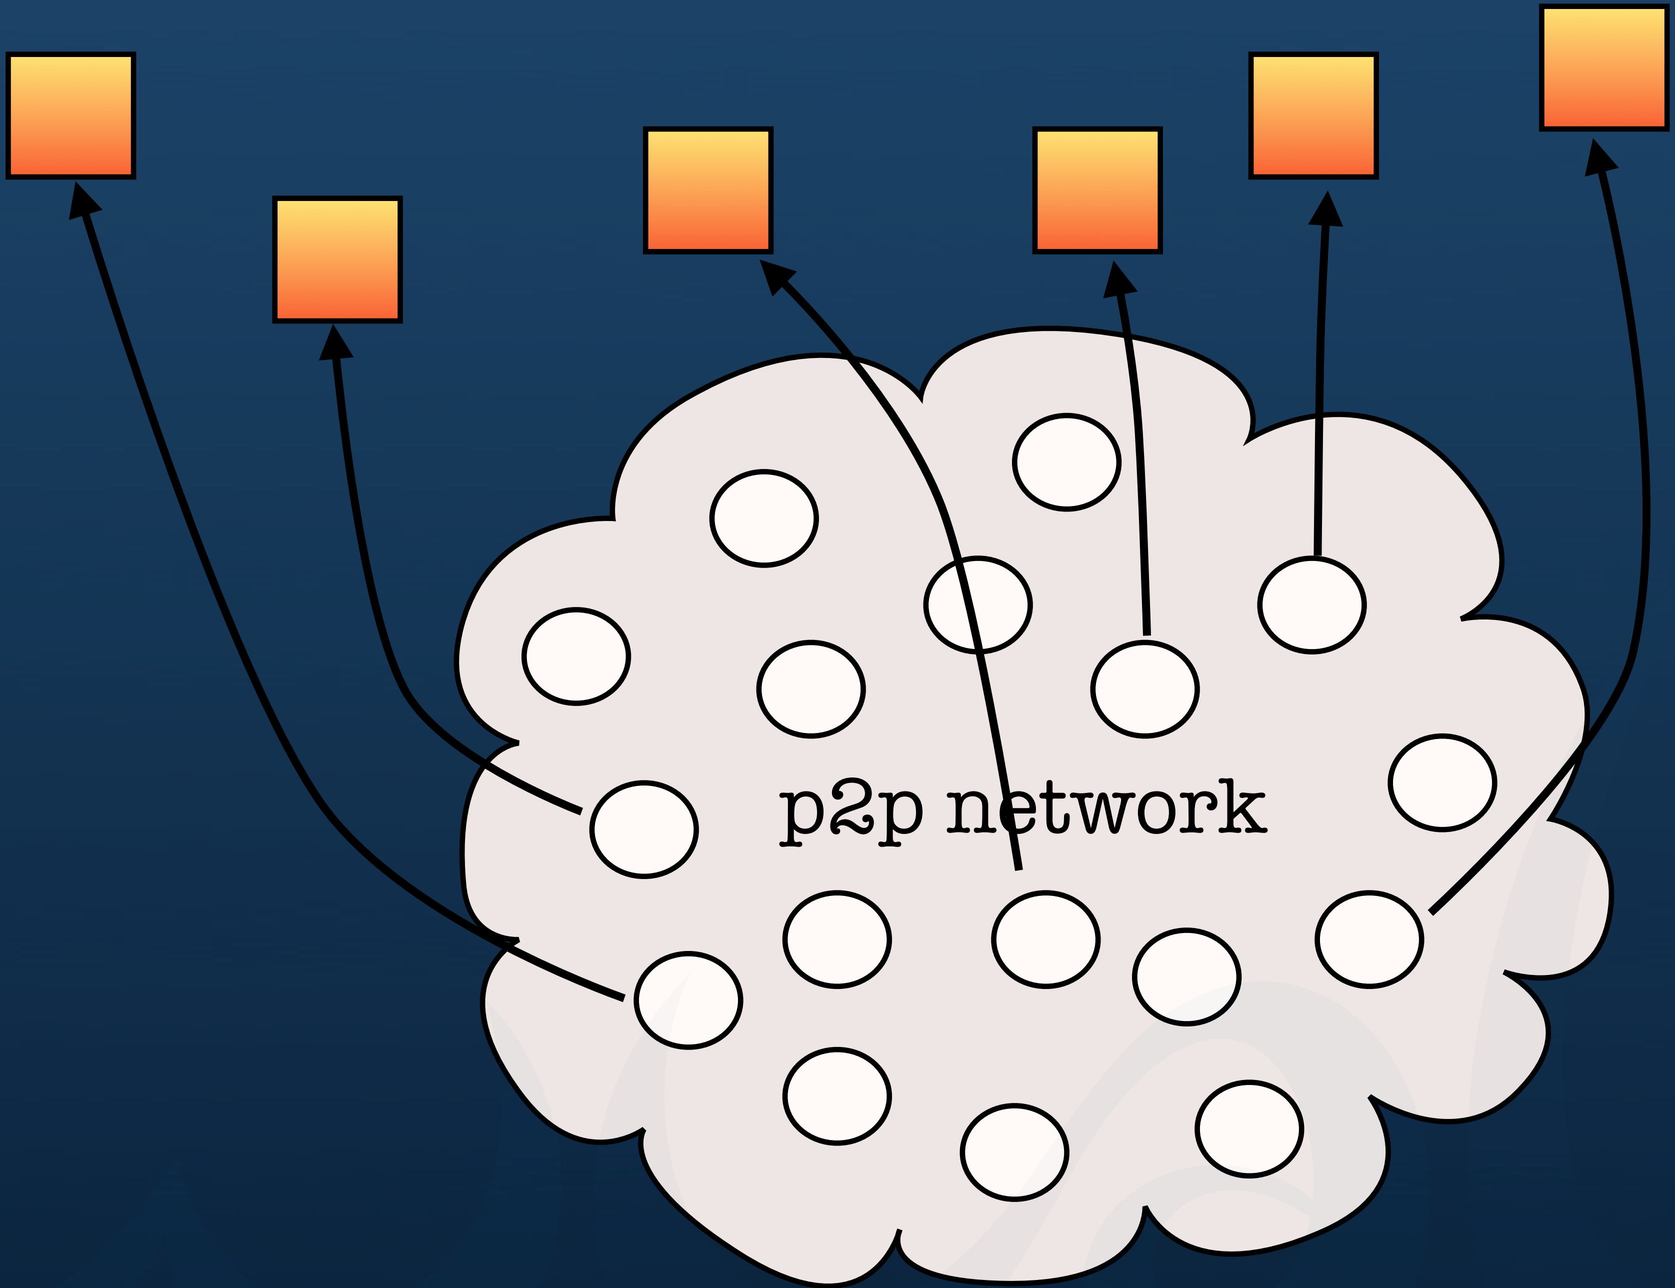

1. download subset of fragments (m)

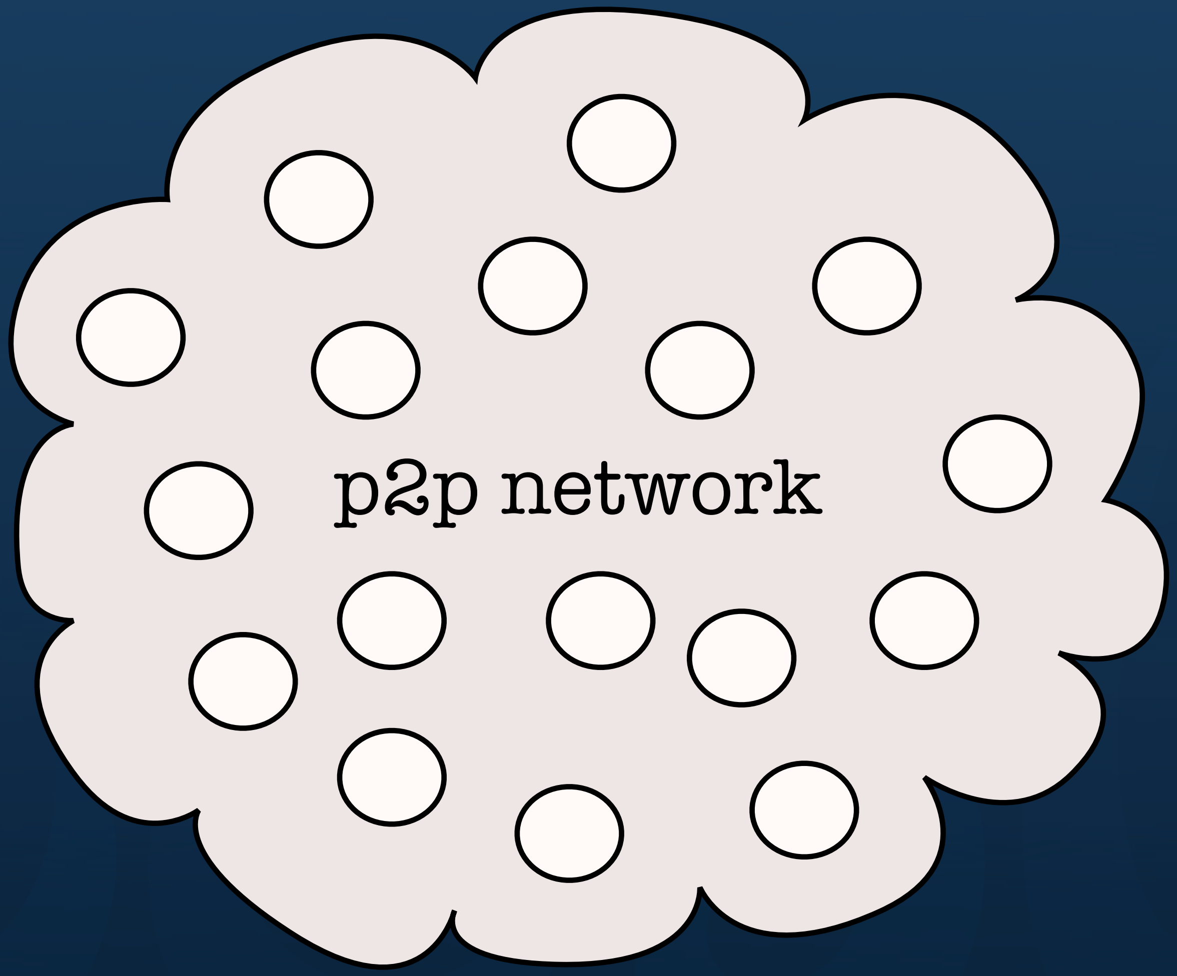

p2p network

maintenance

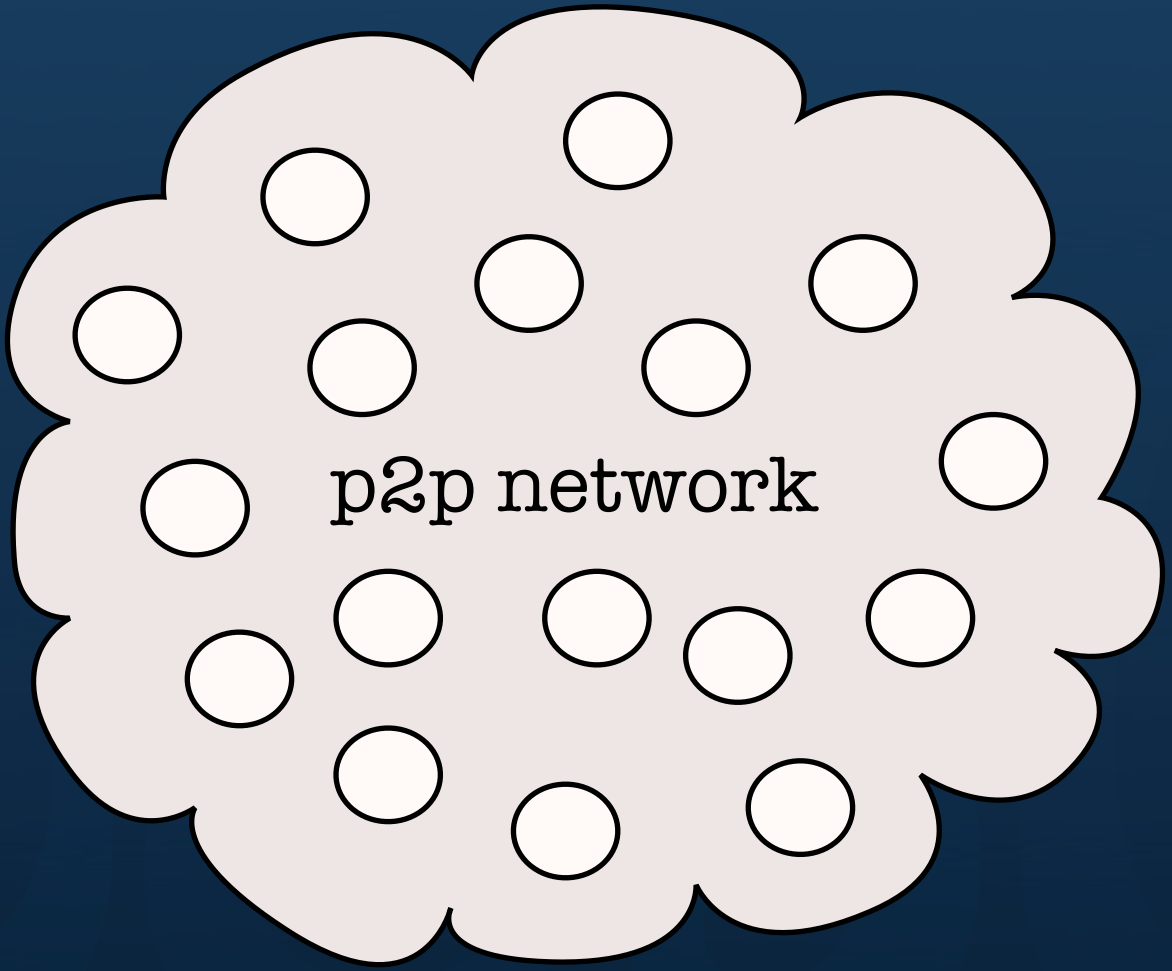

maintenance

alice’s computer checks and maintains her files

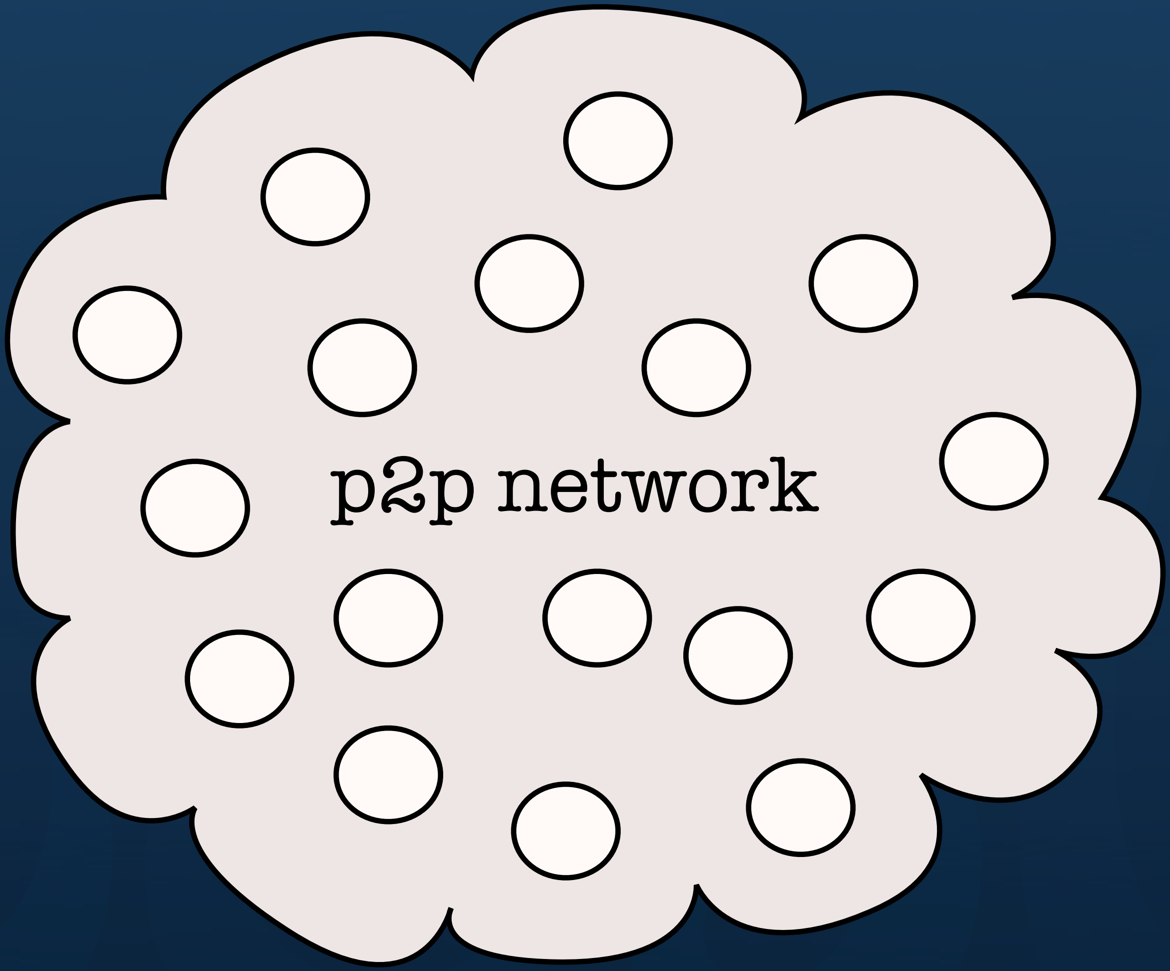

maintenance

alice’s computer checks and maintains her files

if necessary, it constructs new fragments and uploads them

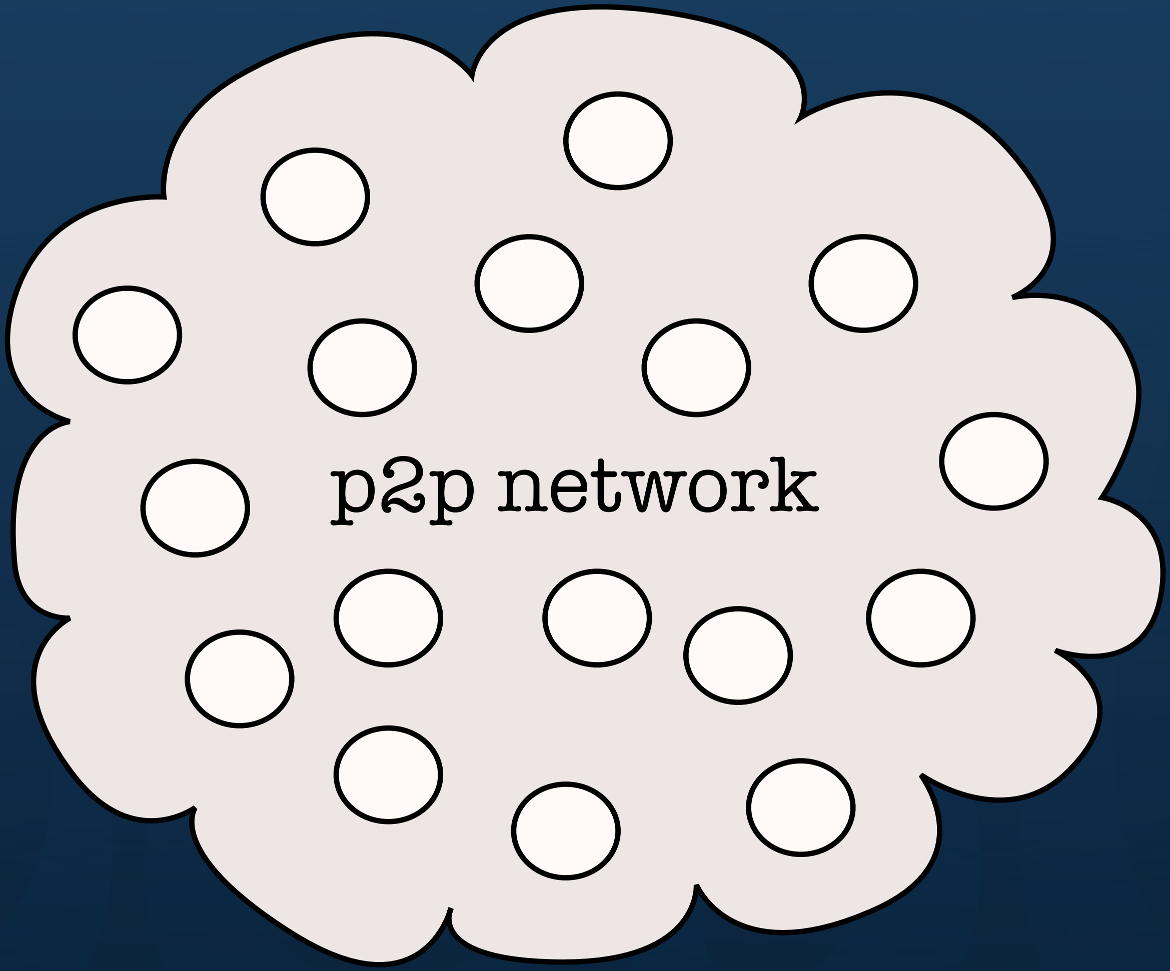

maintenance

alice’s computer checks and maintains her files

if necessary, it constructs new fragments and uploads them

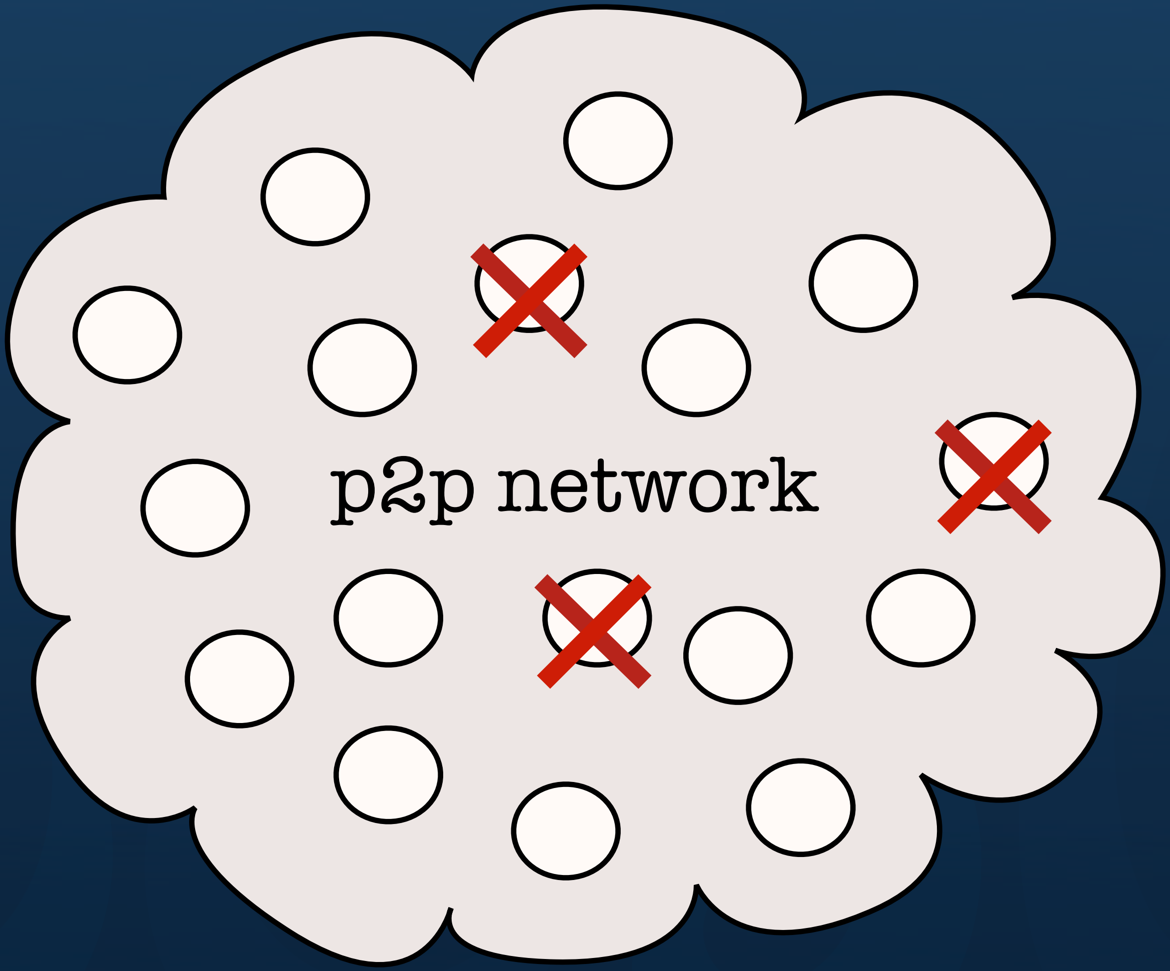

maintenance

alice's computer checks and maintains her files

if necessary, it constructs new fragments and uploads them

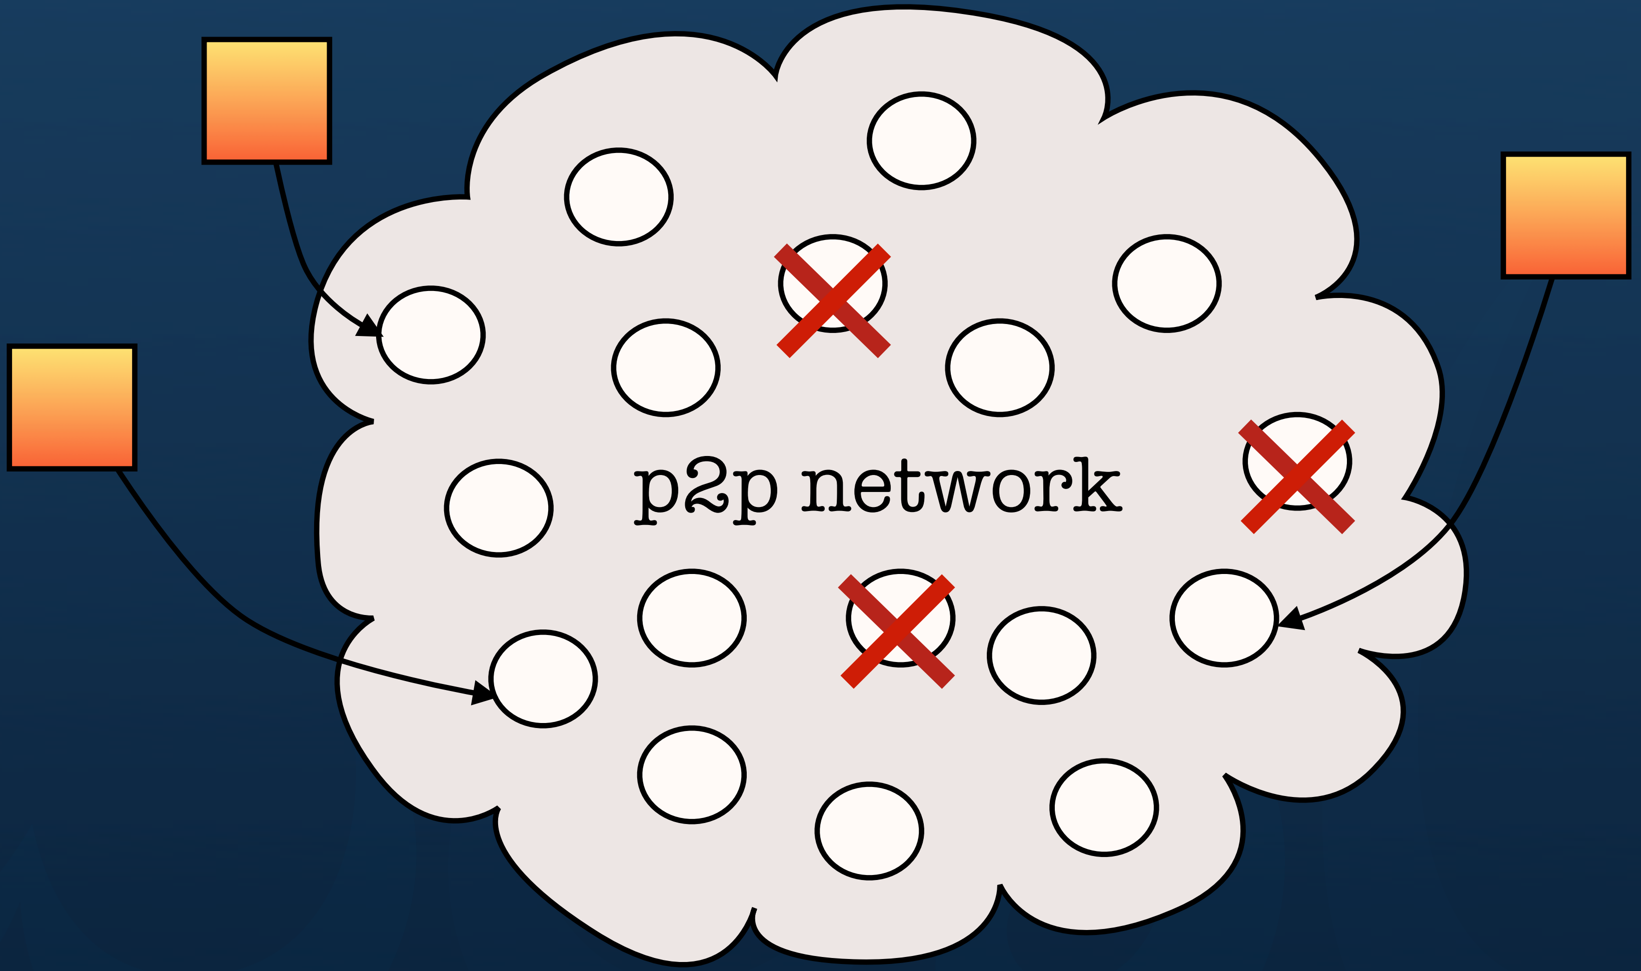

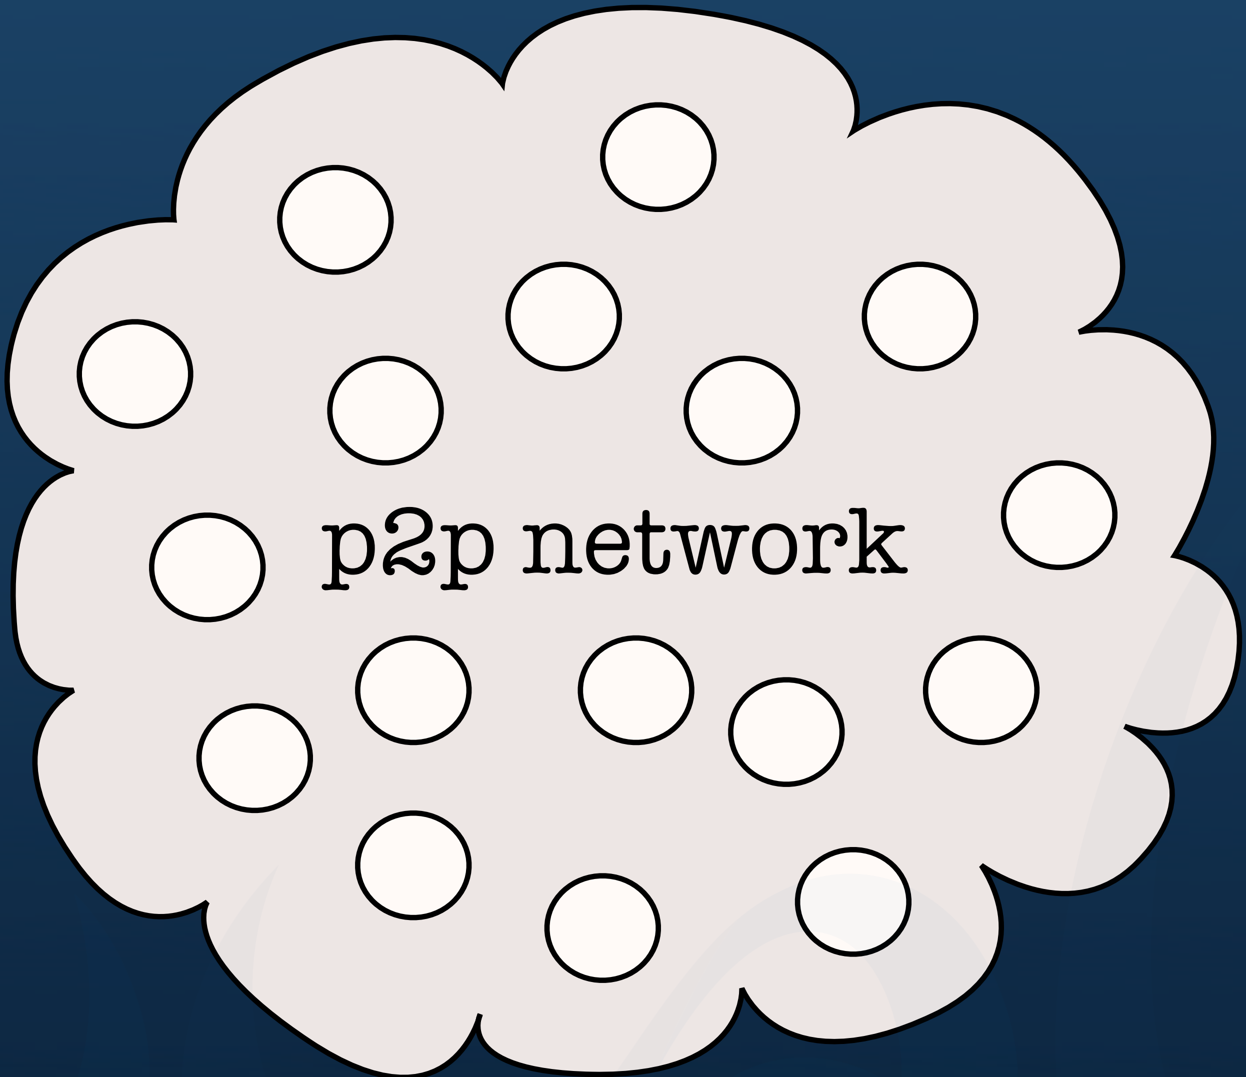
p2p network

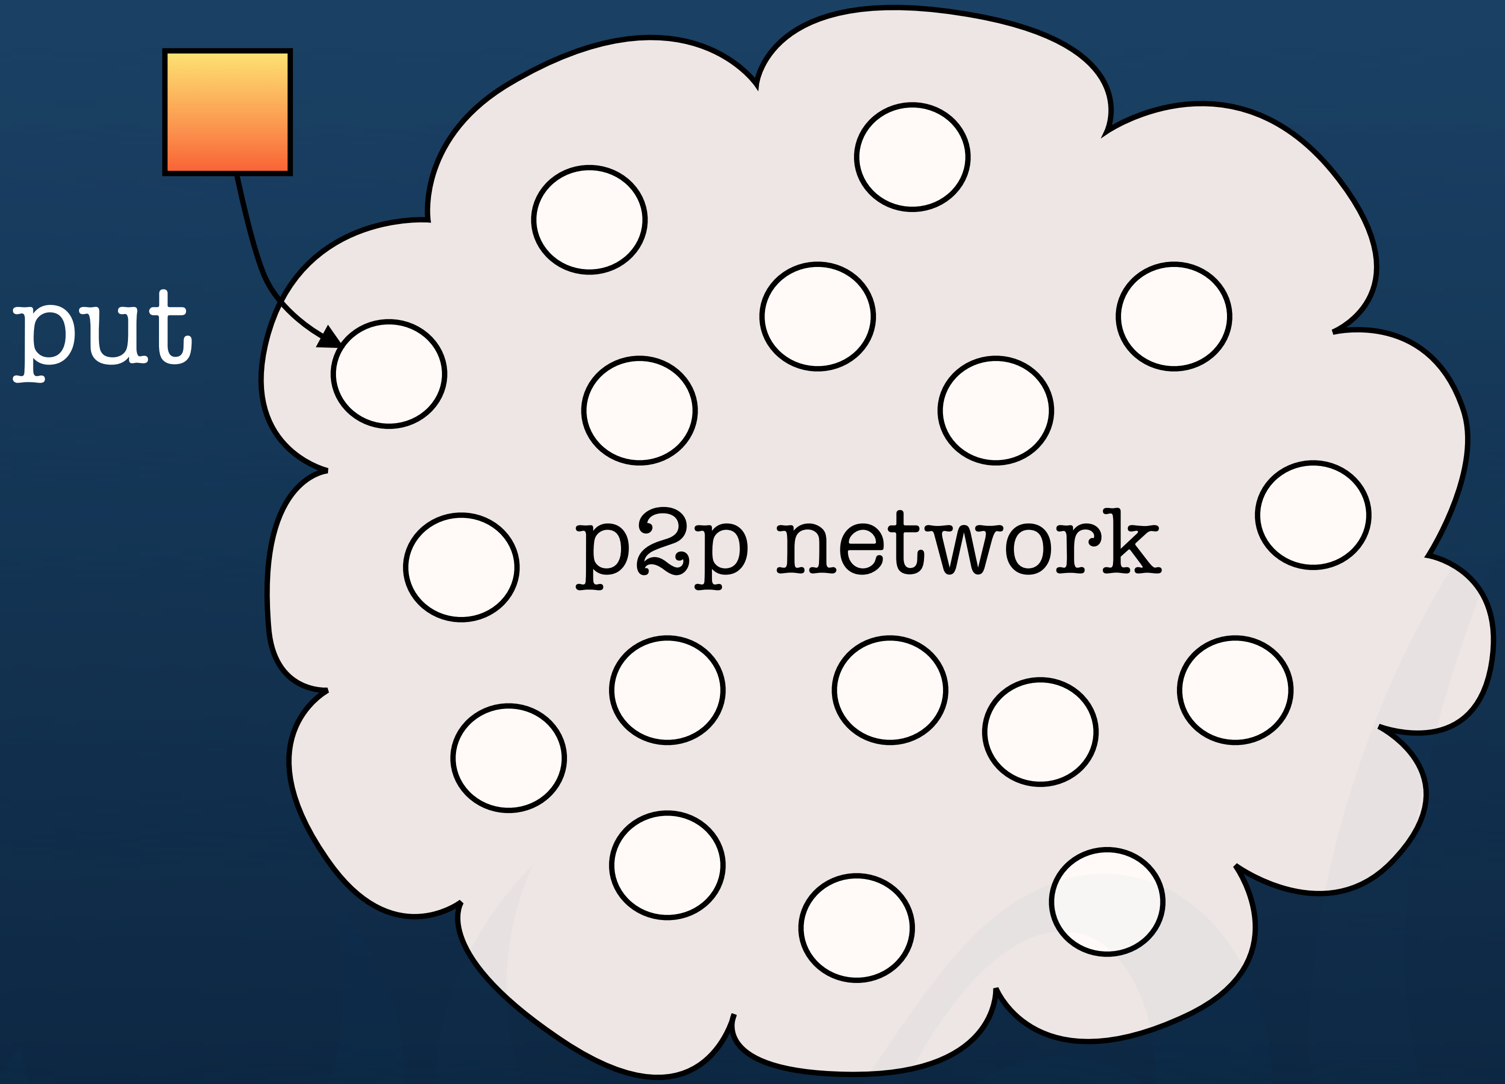
put
p2p network

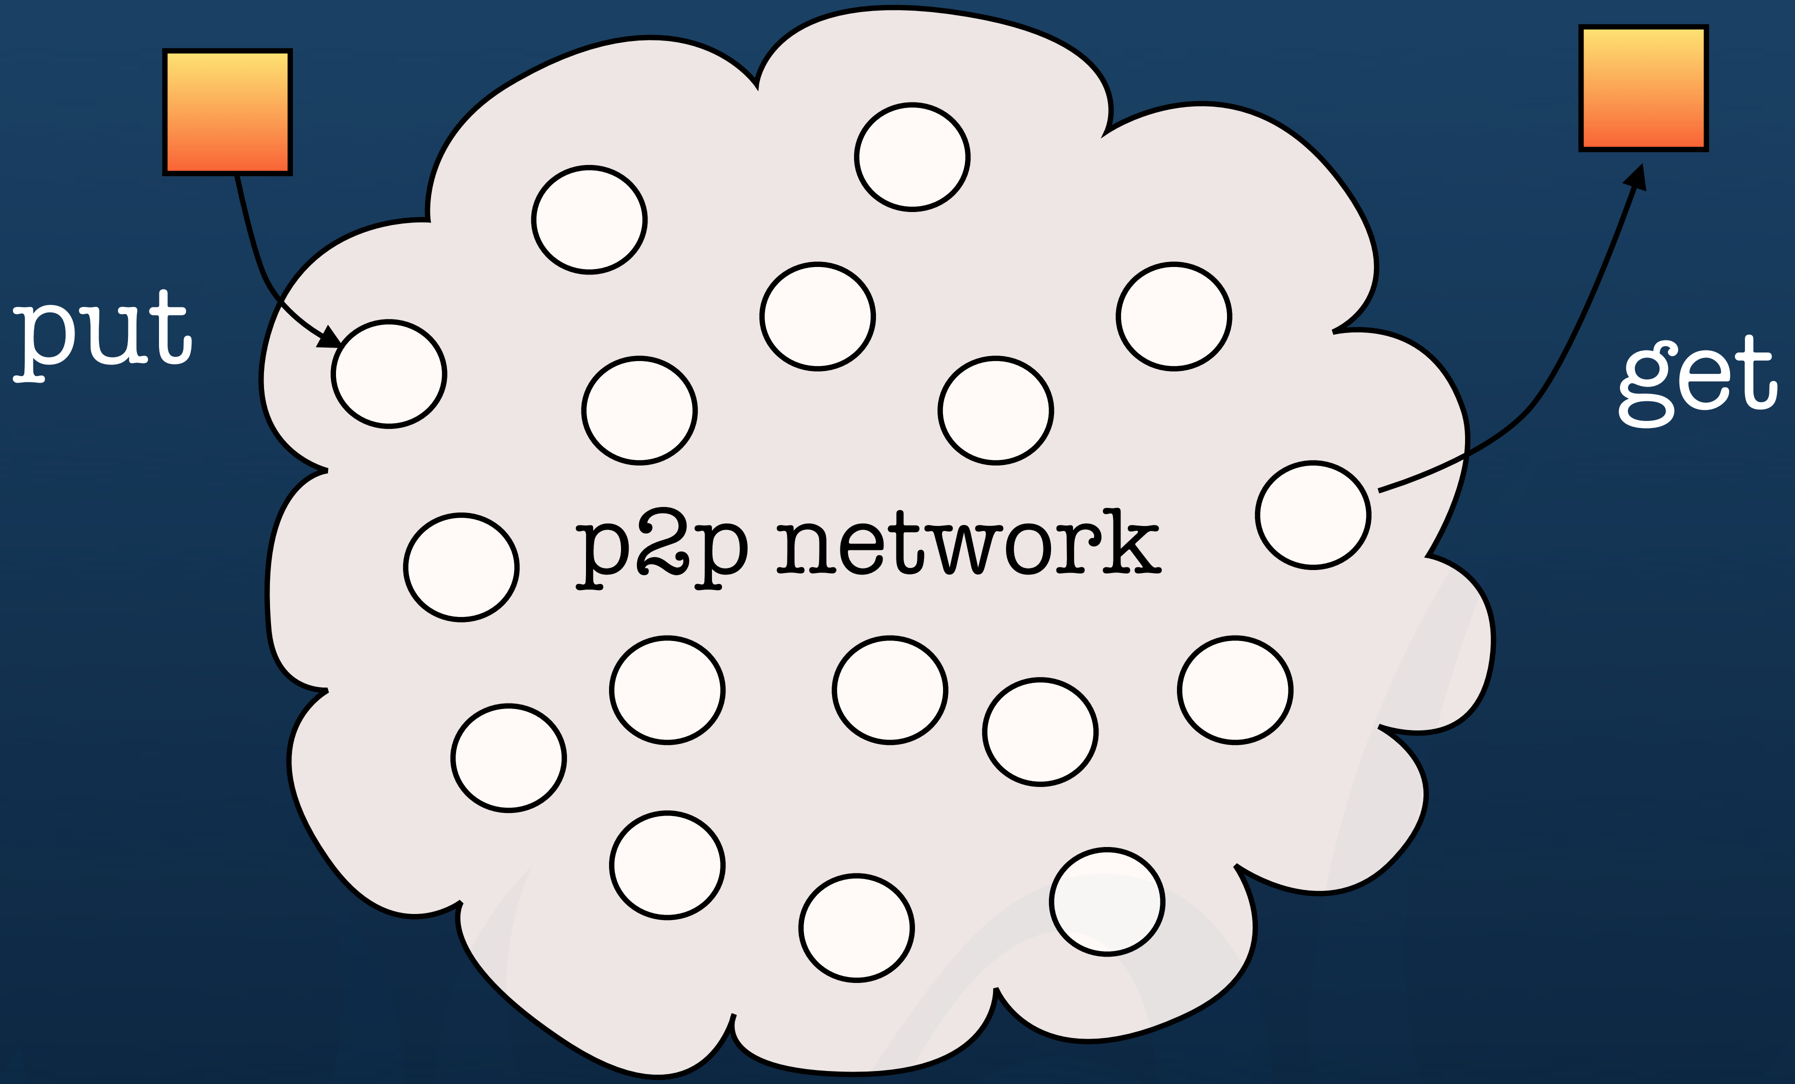
put
p2p network
get

distributed hash table (DHT)

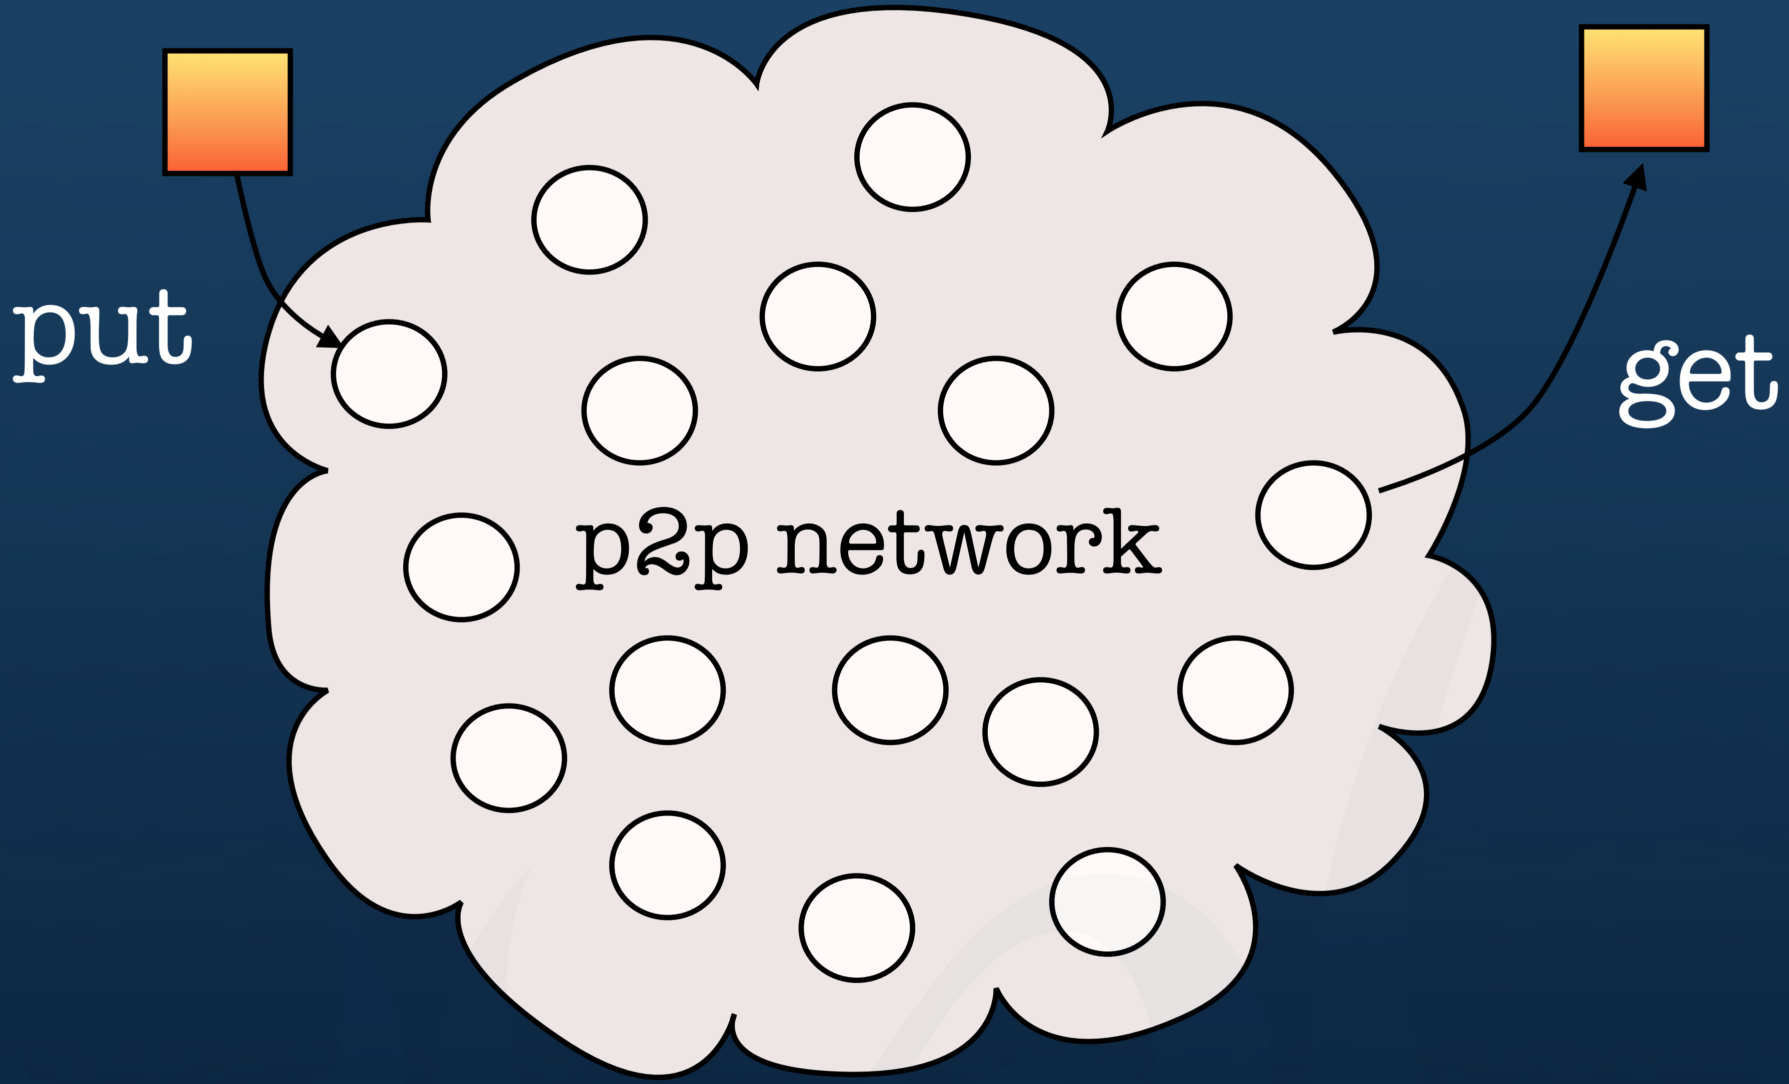

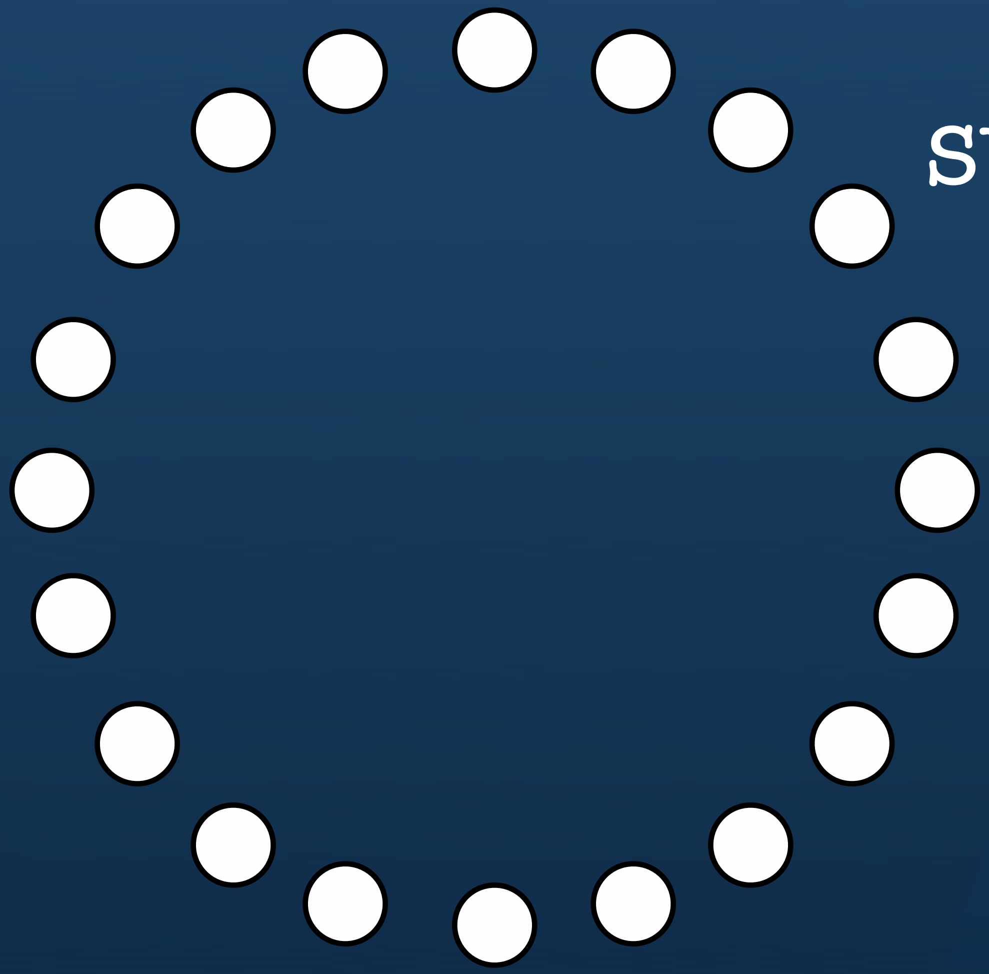

super nodes

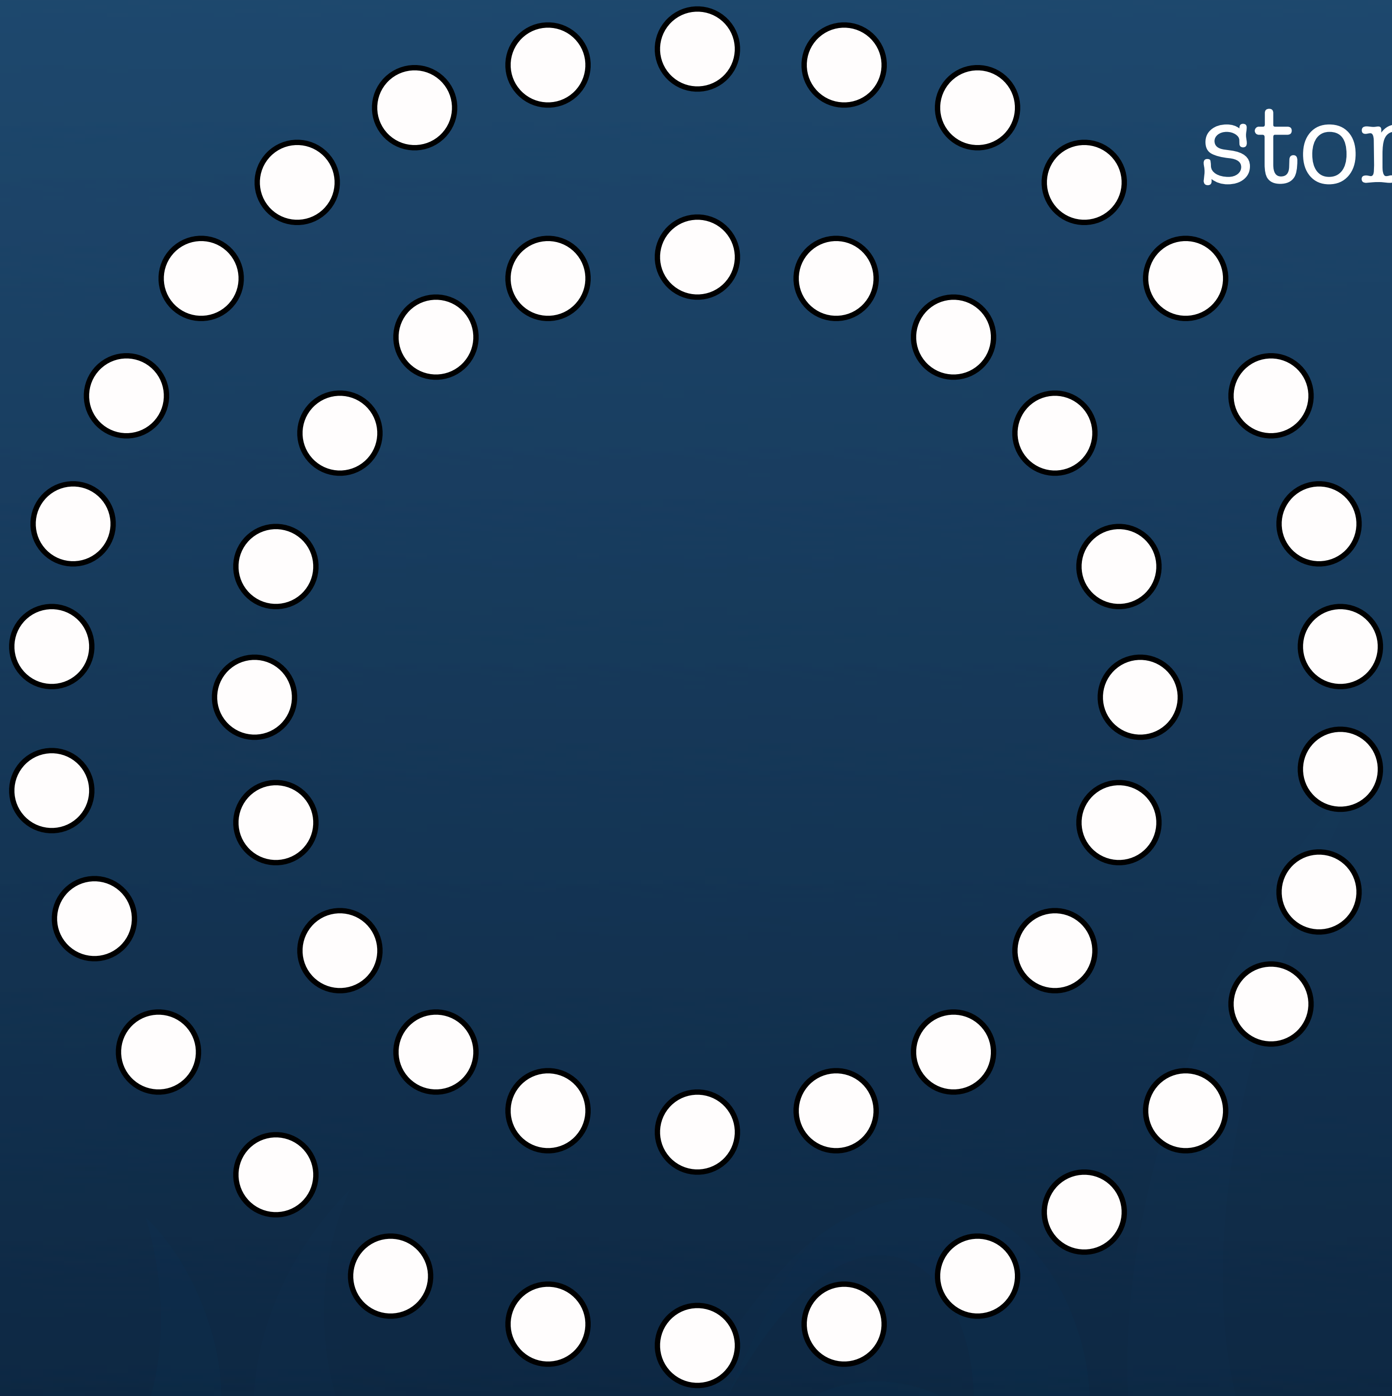

storage nodes

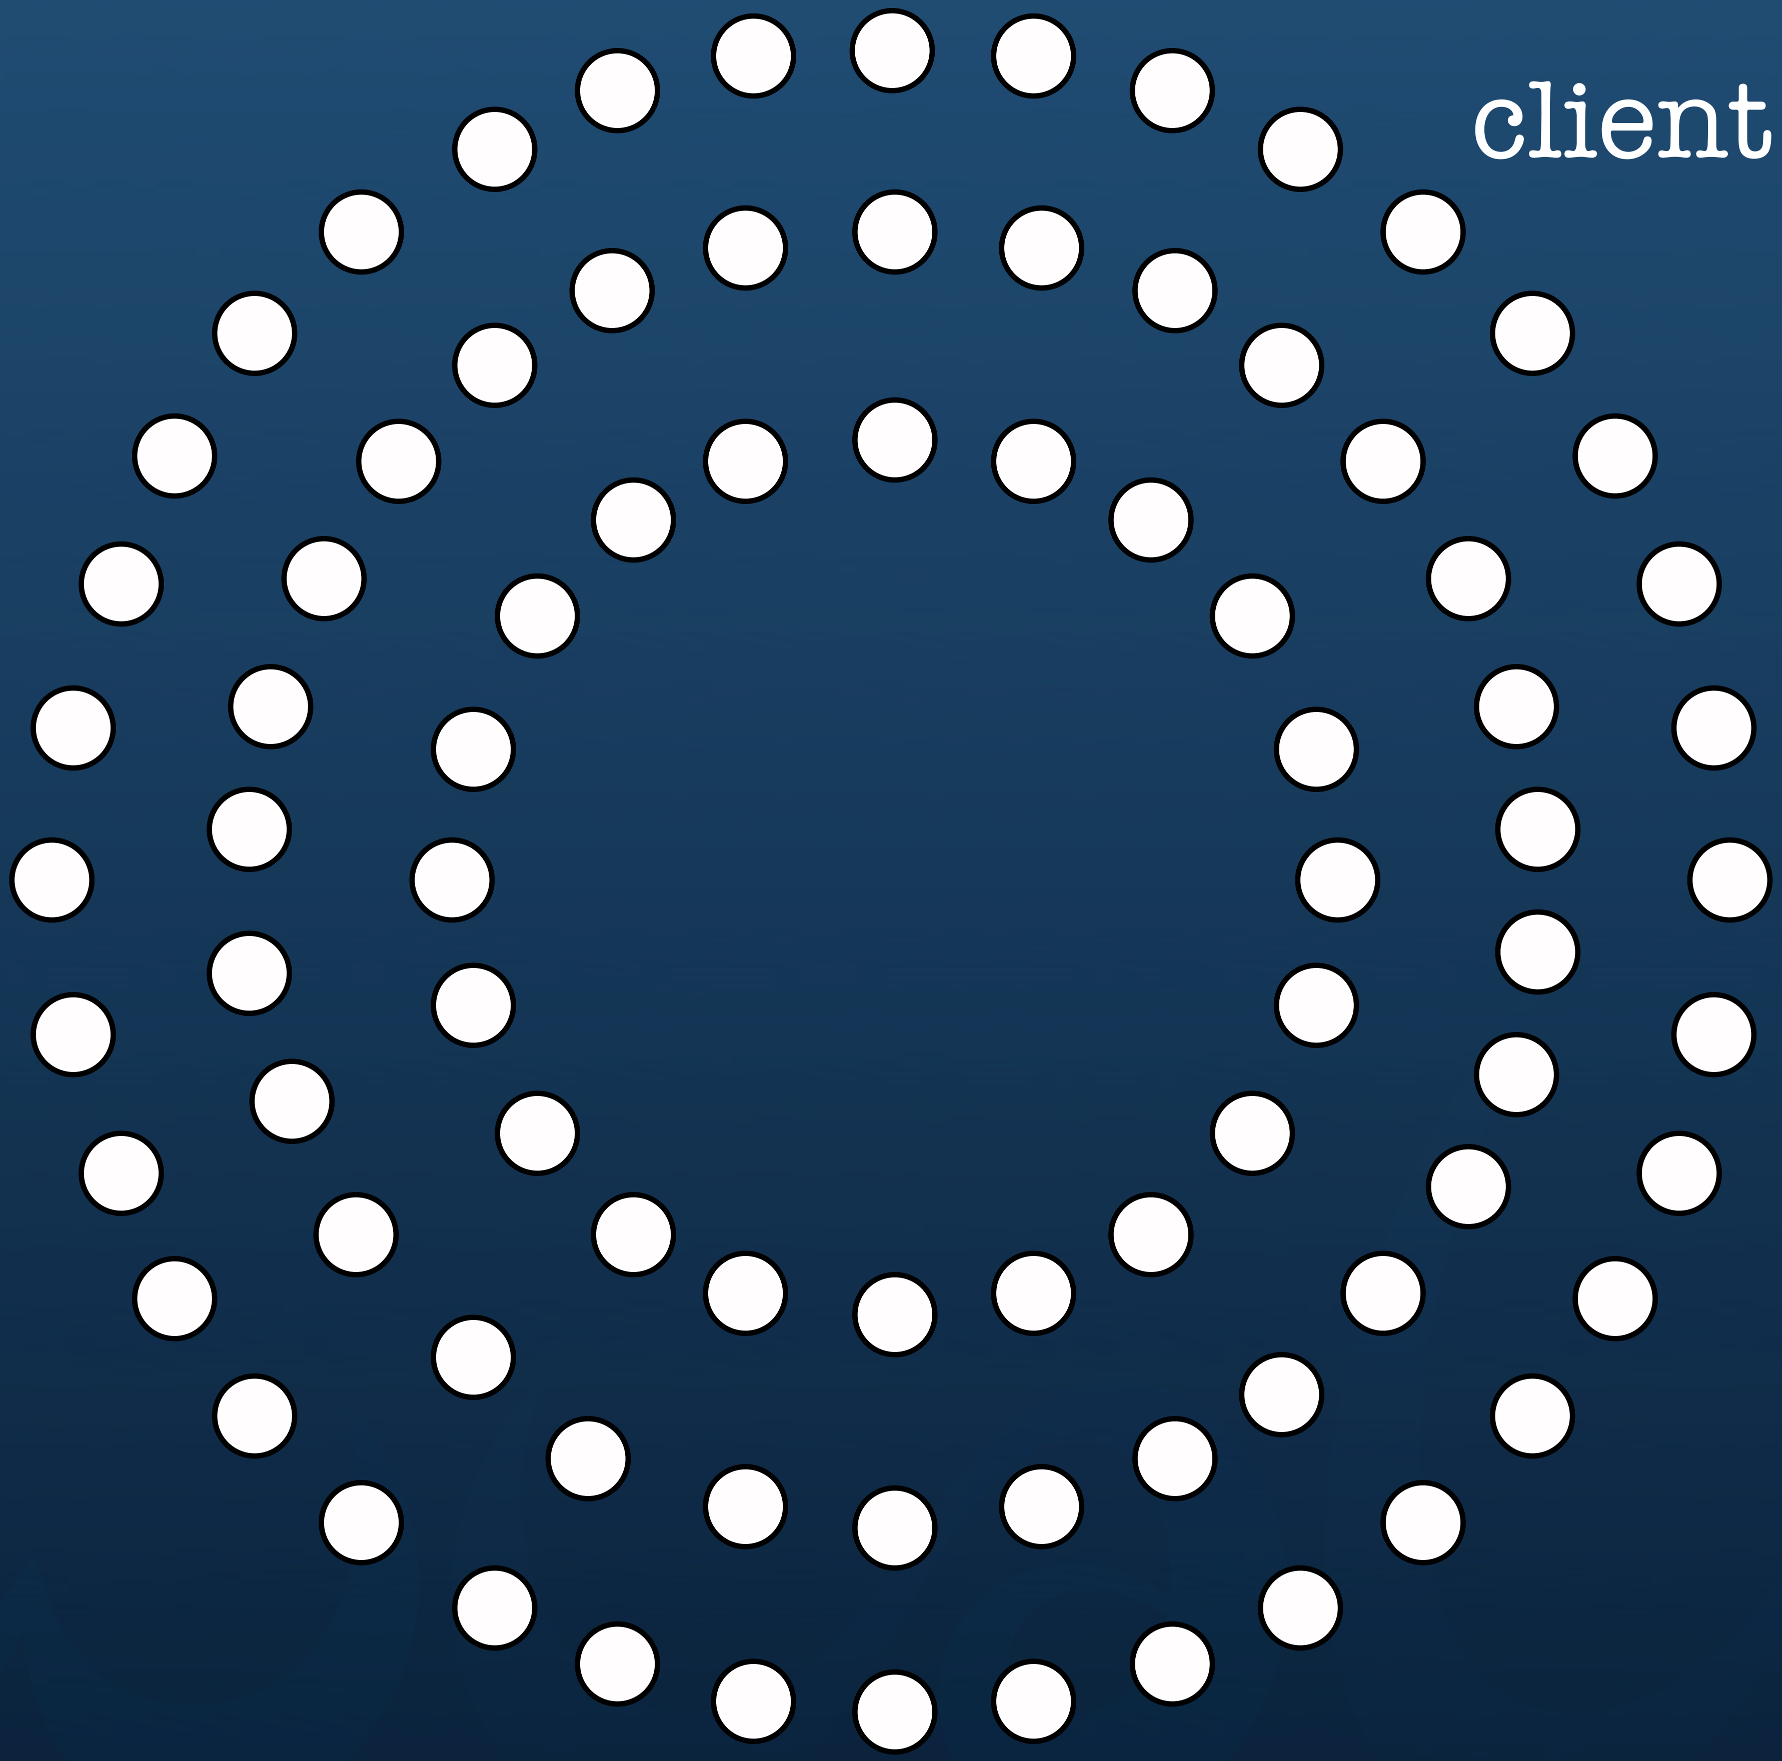

client nodes

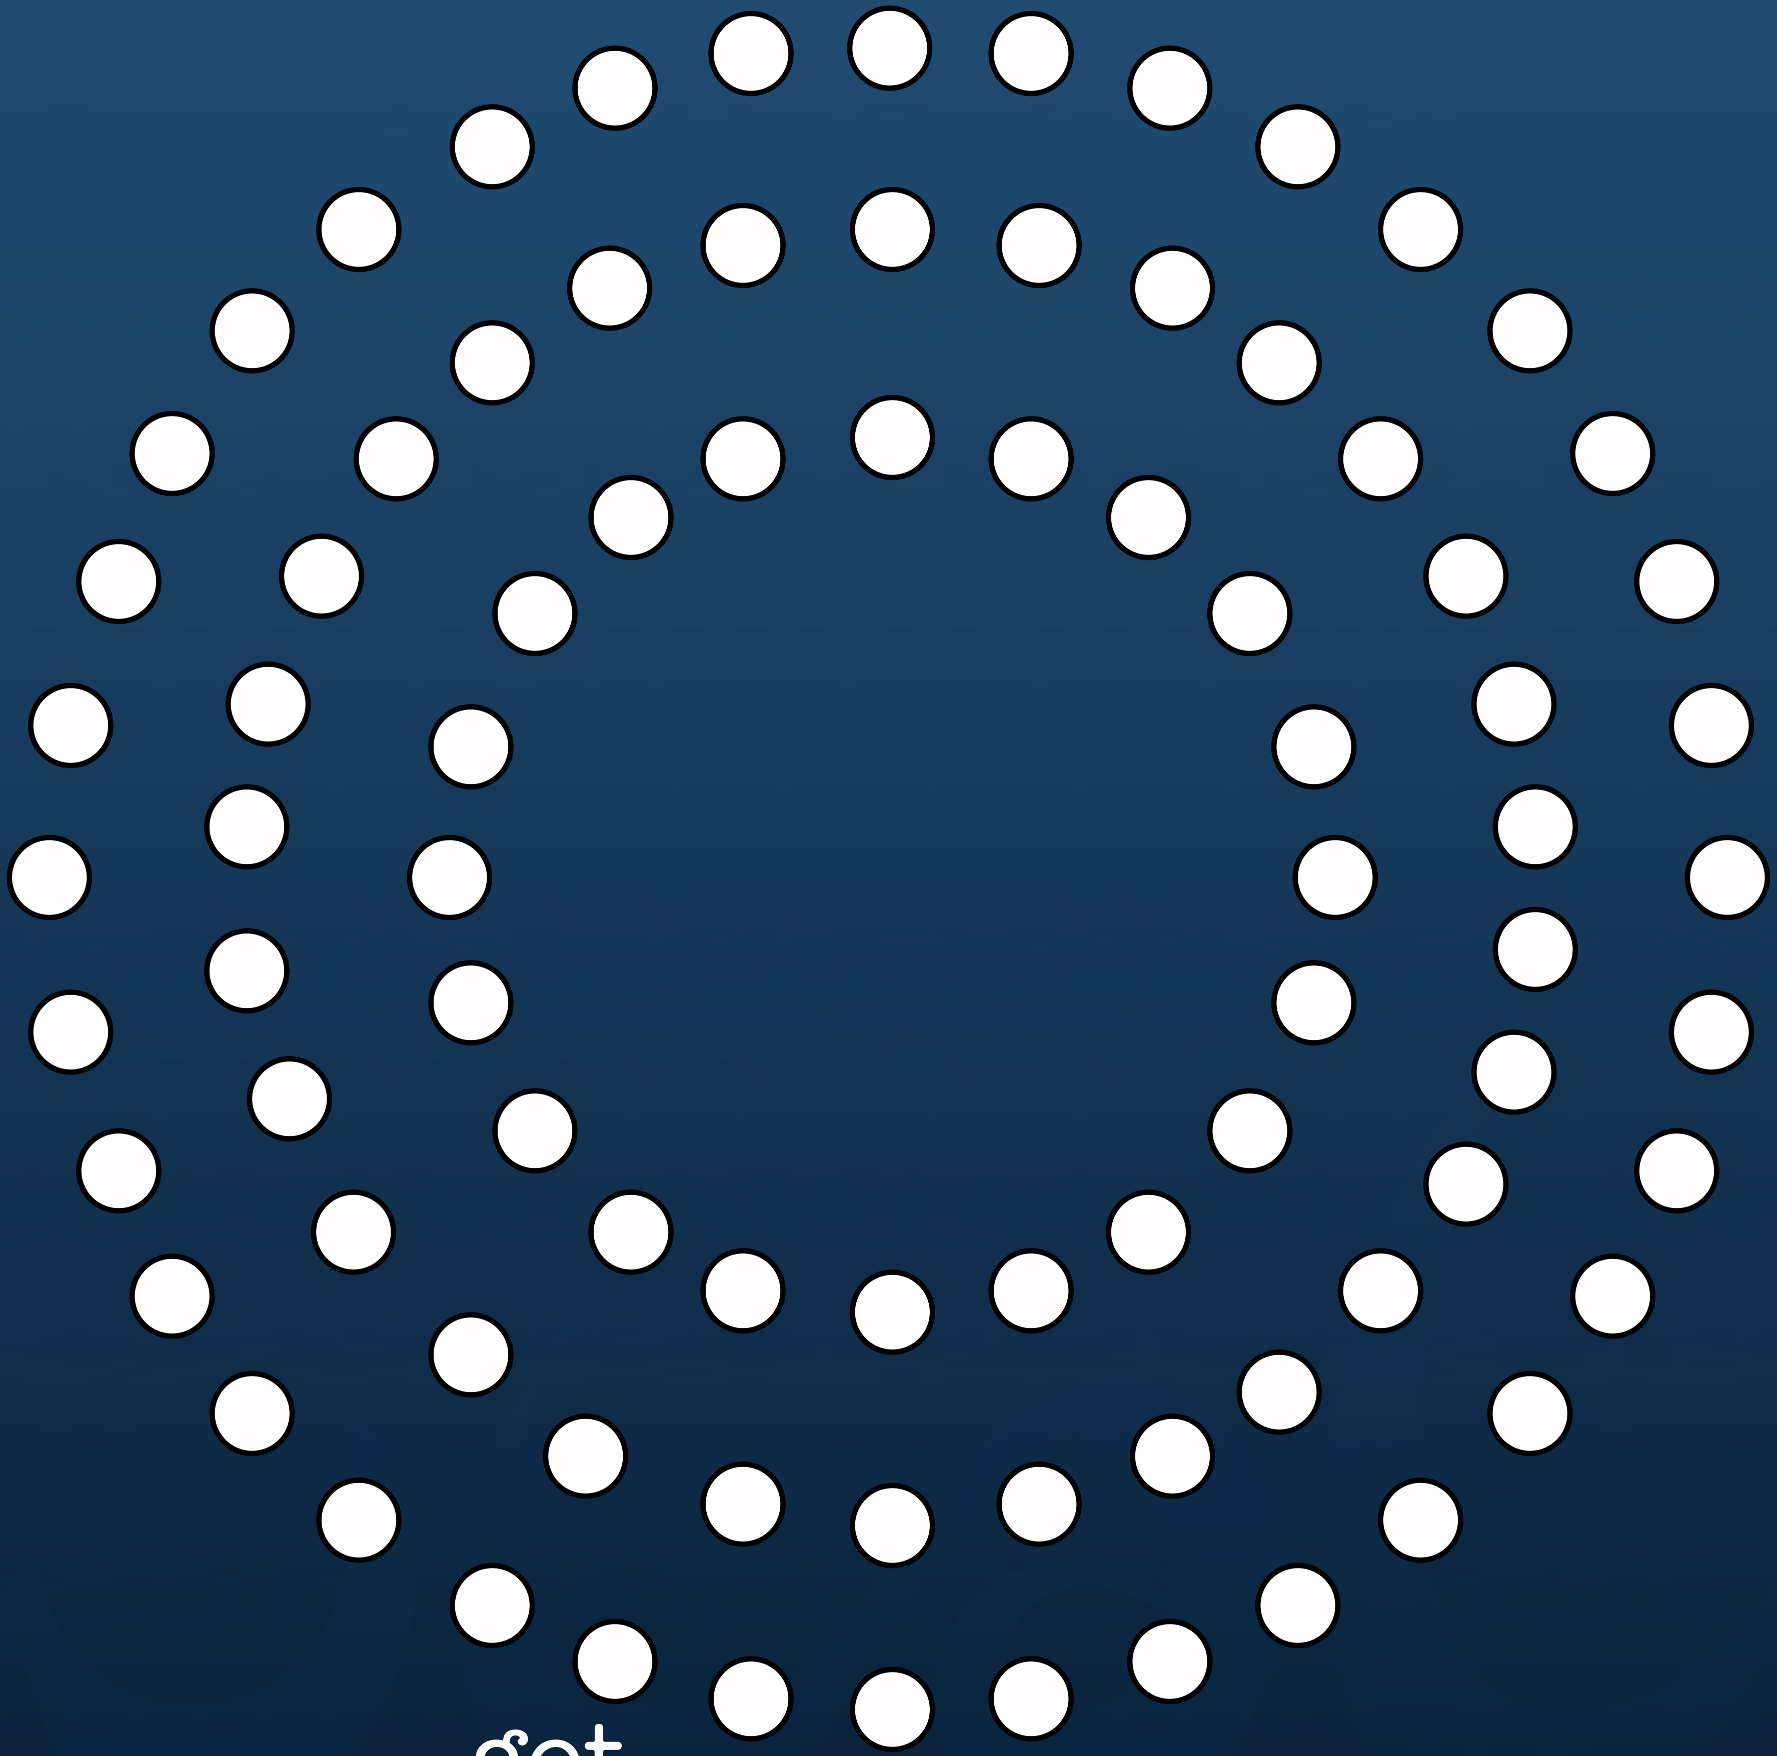

get

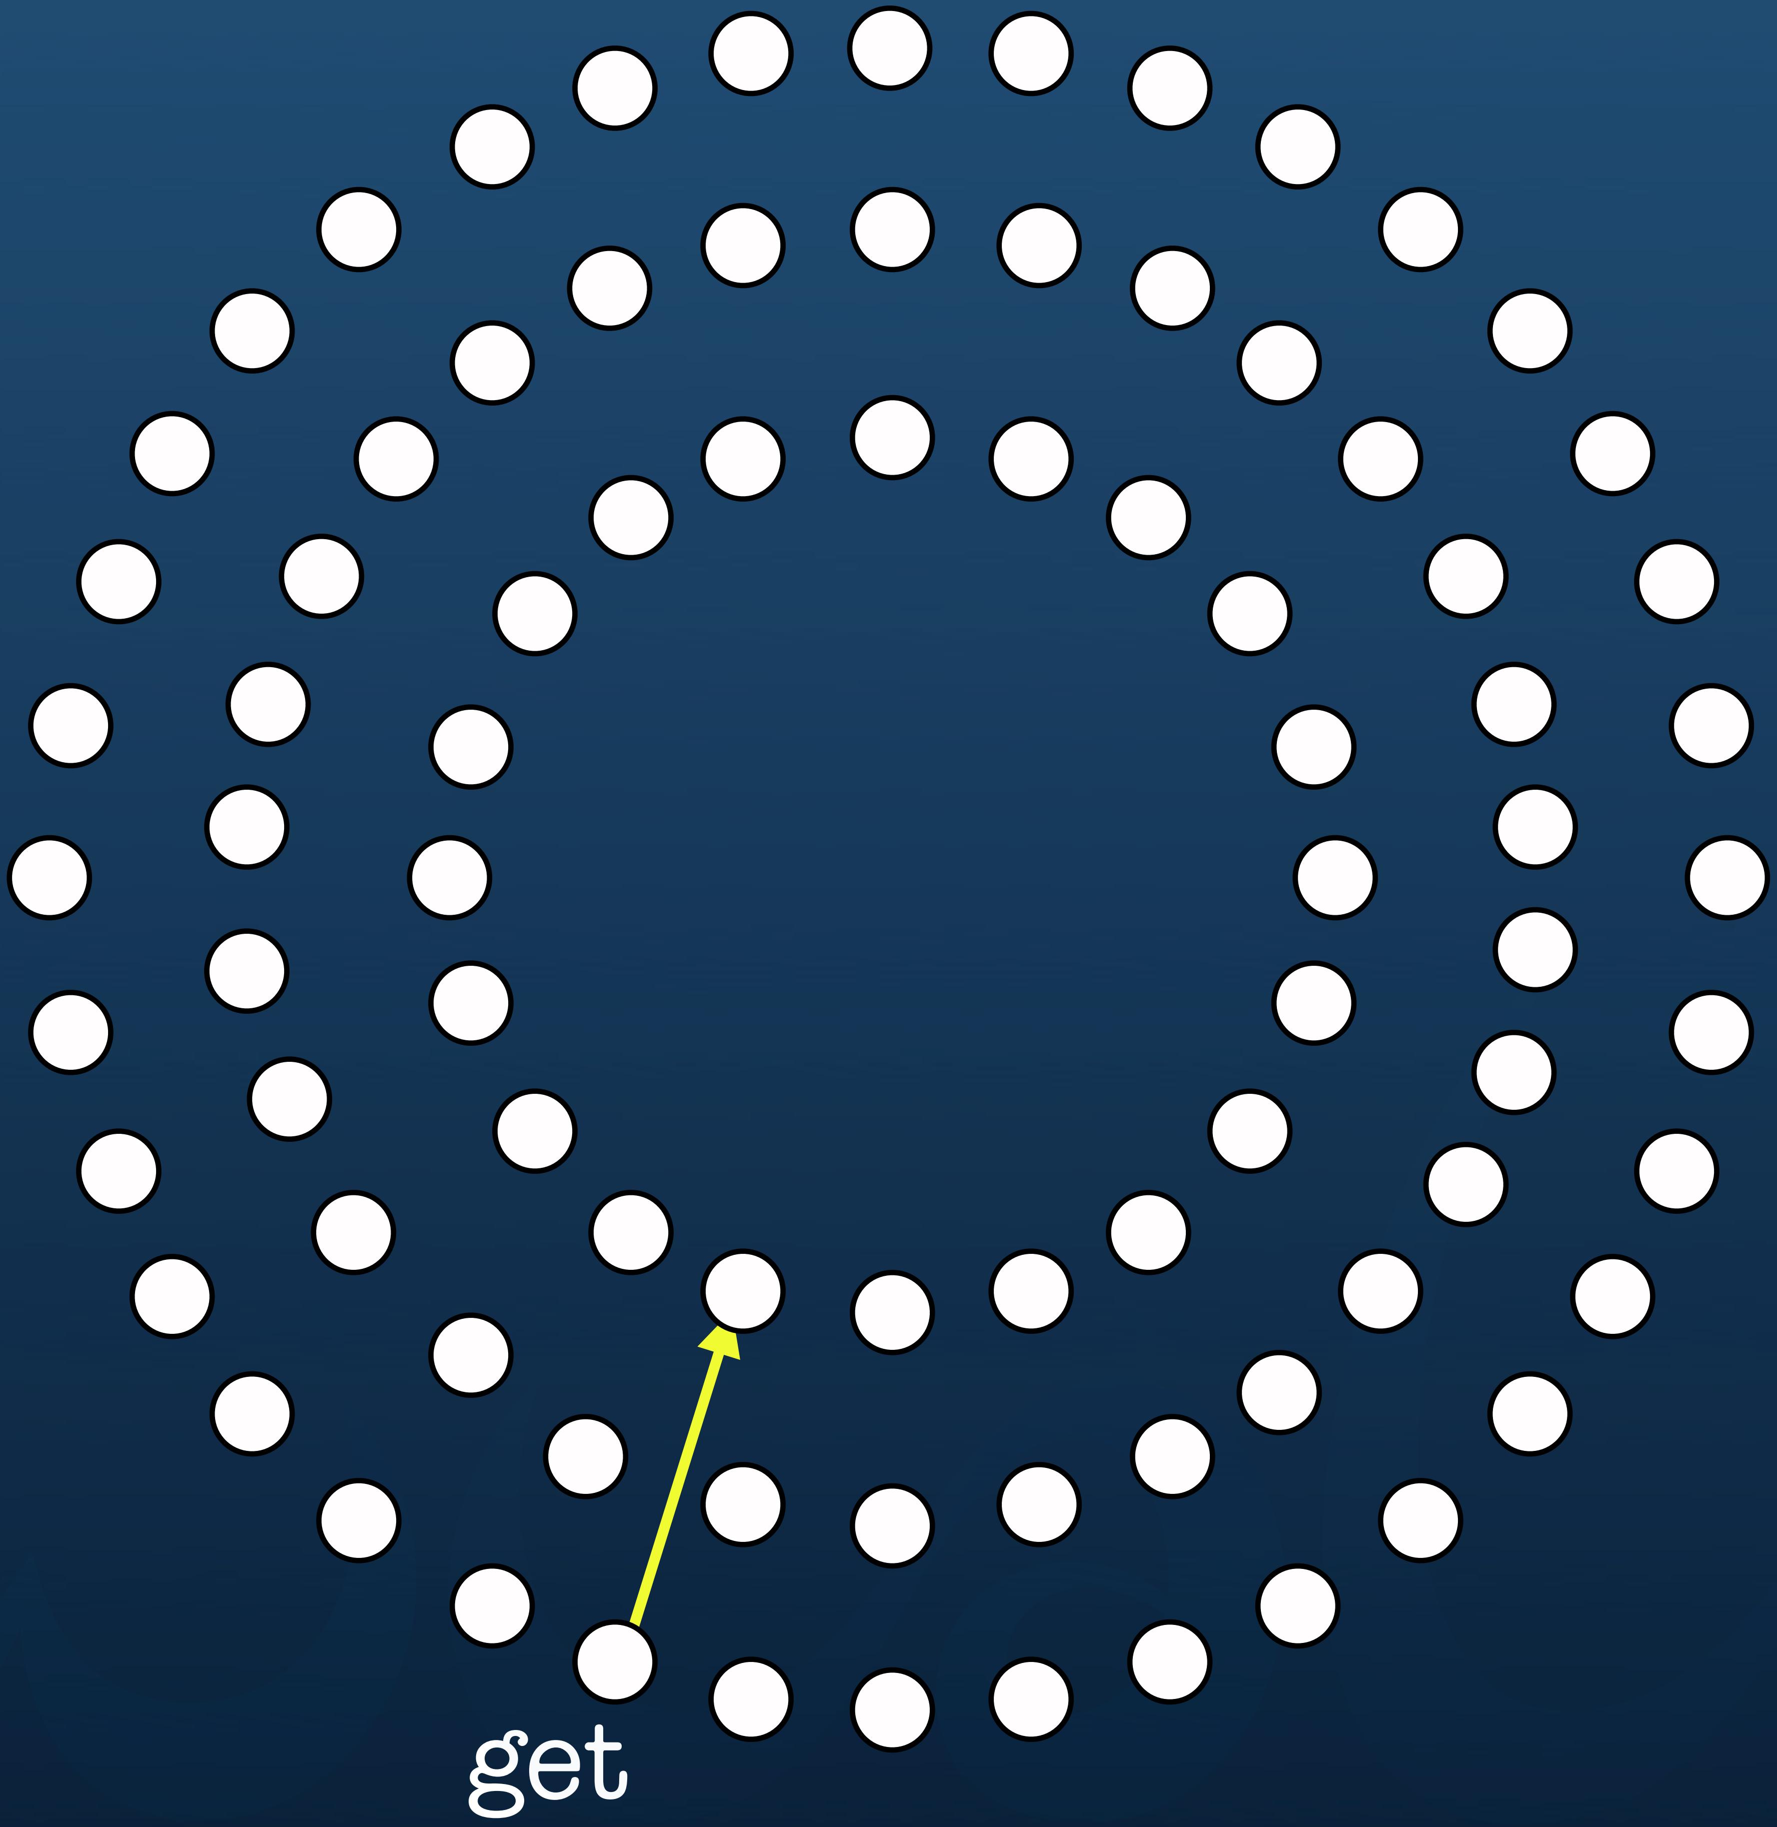
get

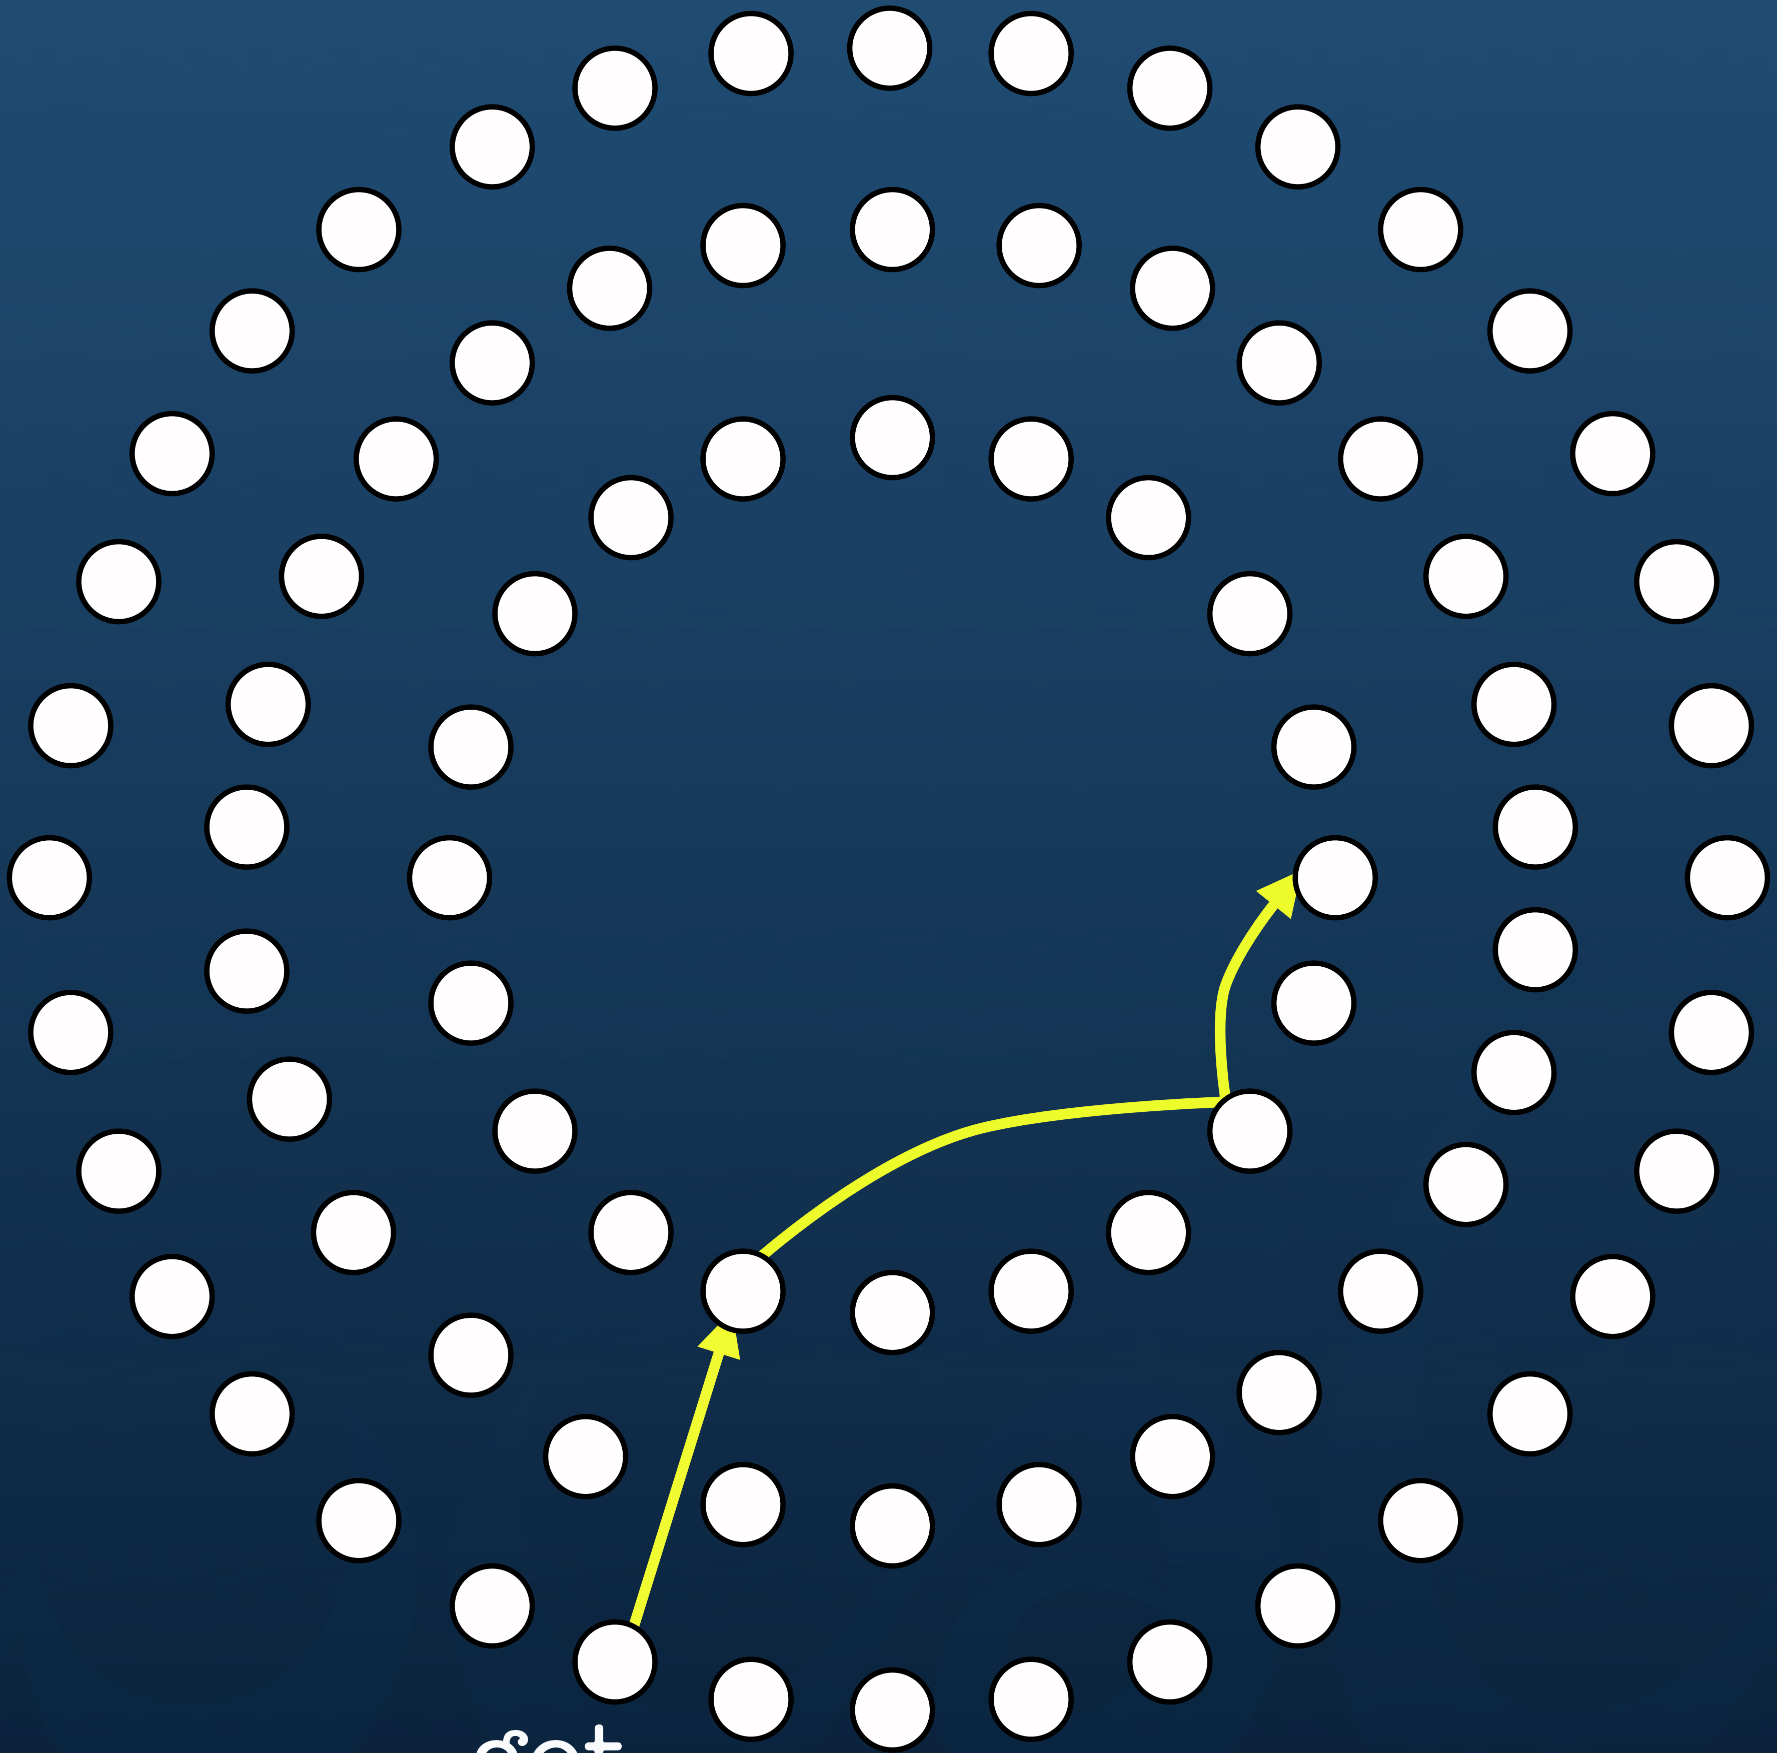

get

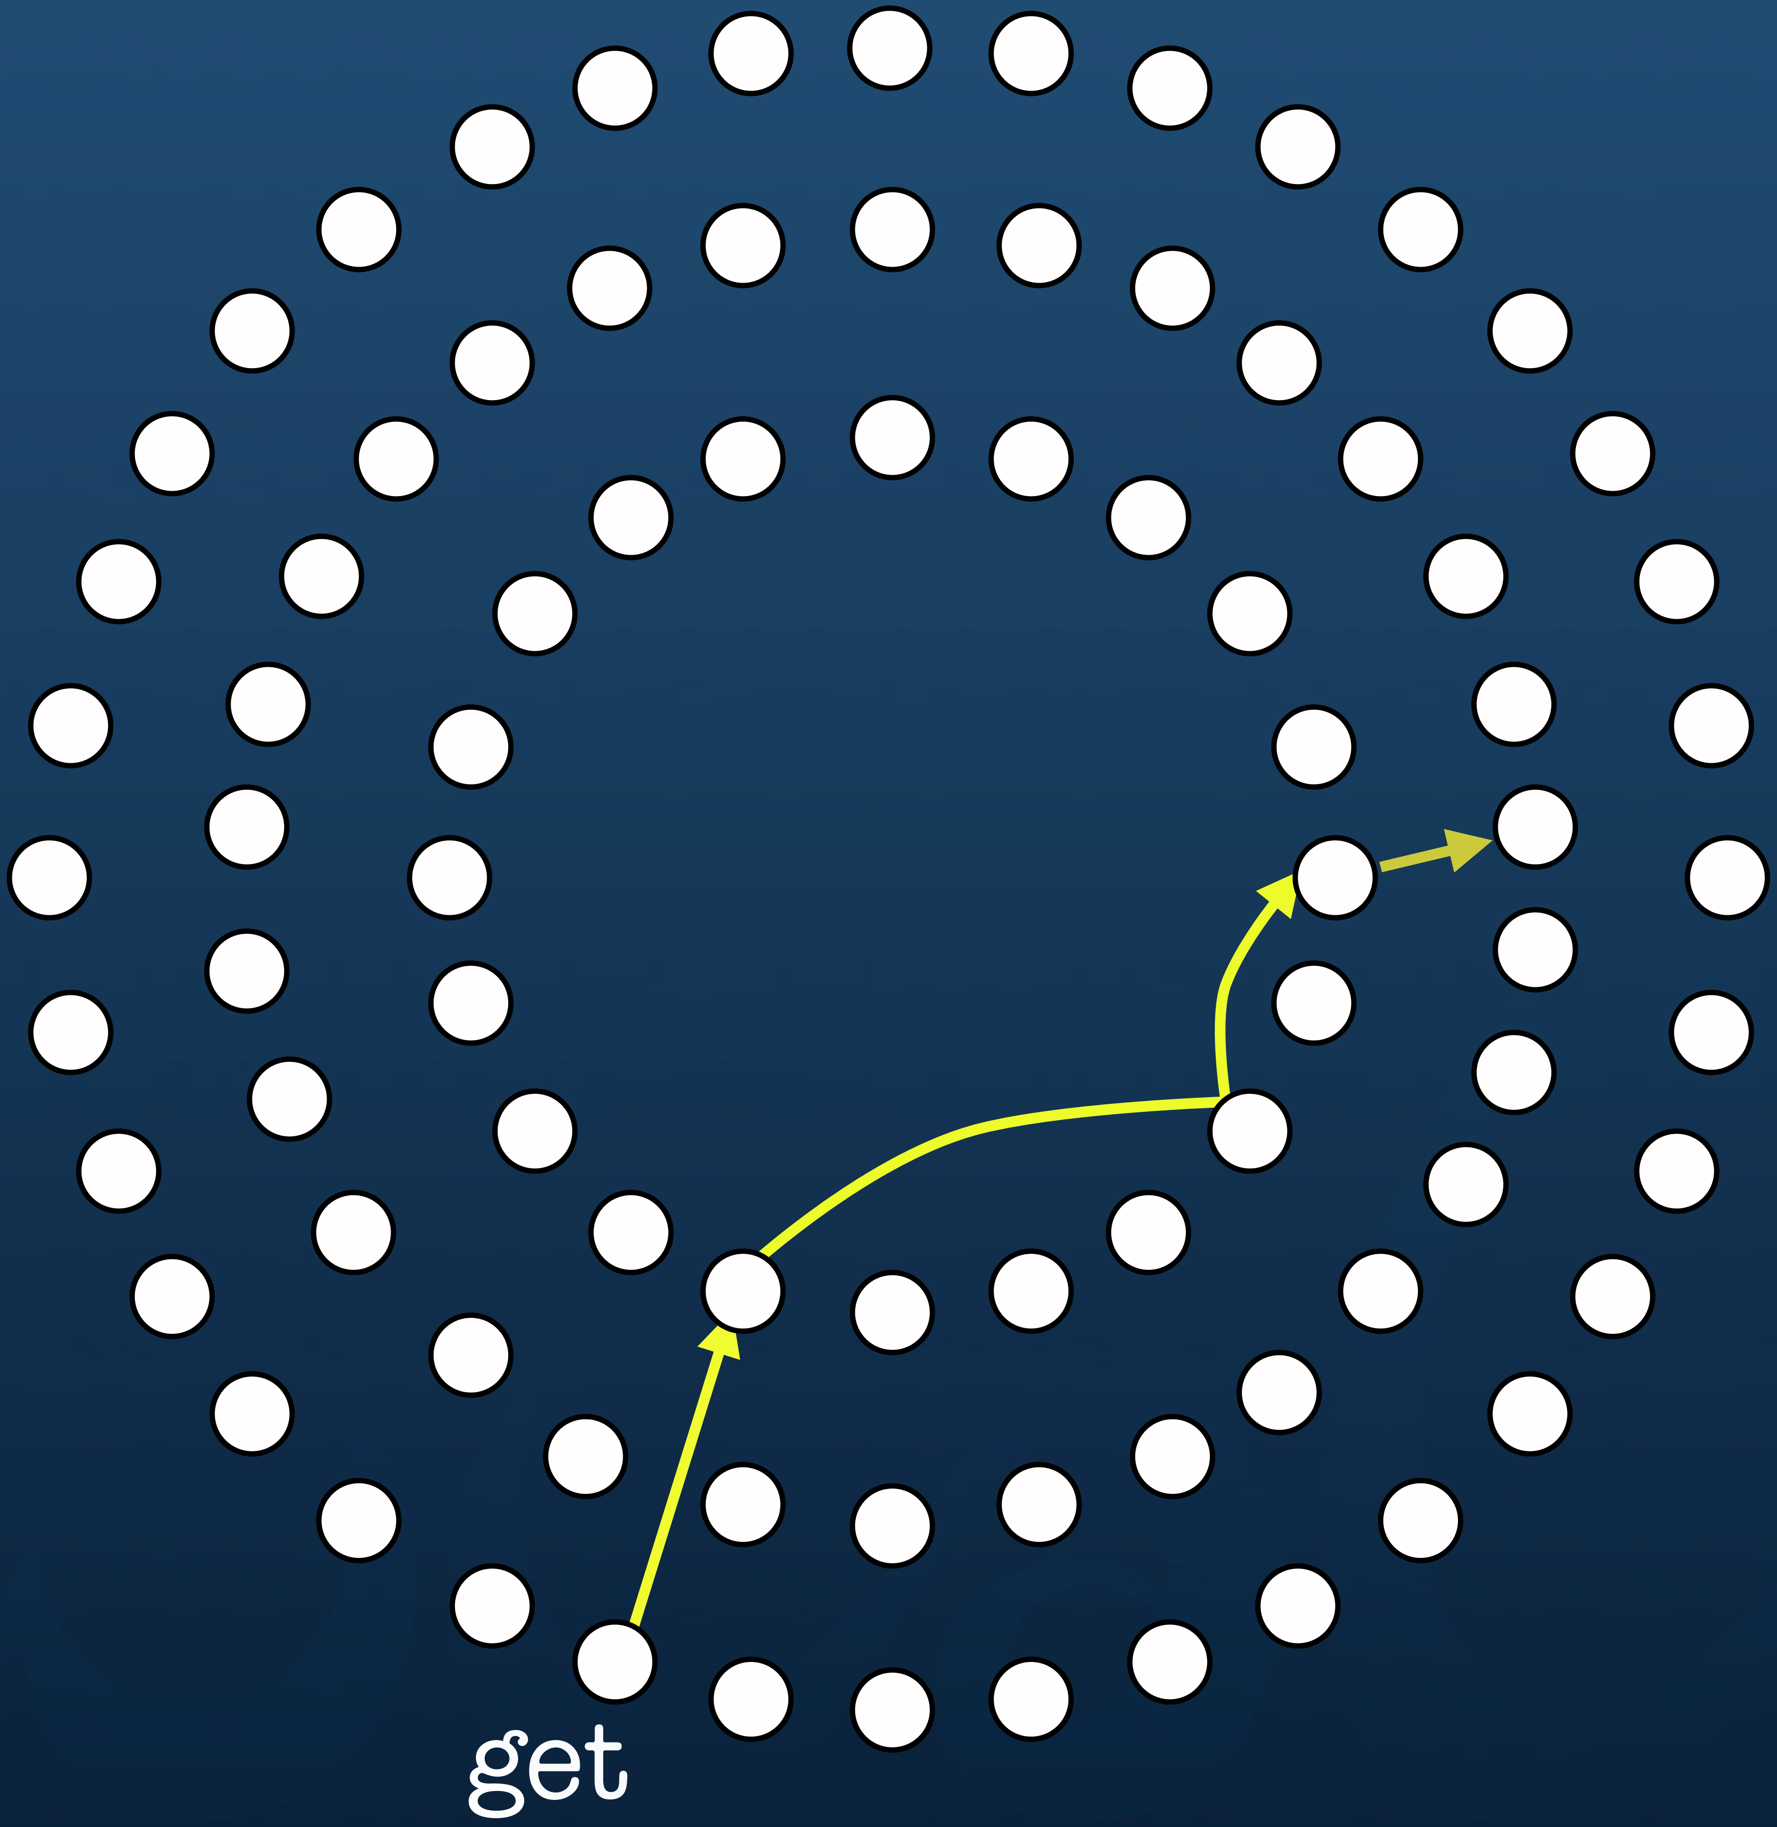
get

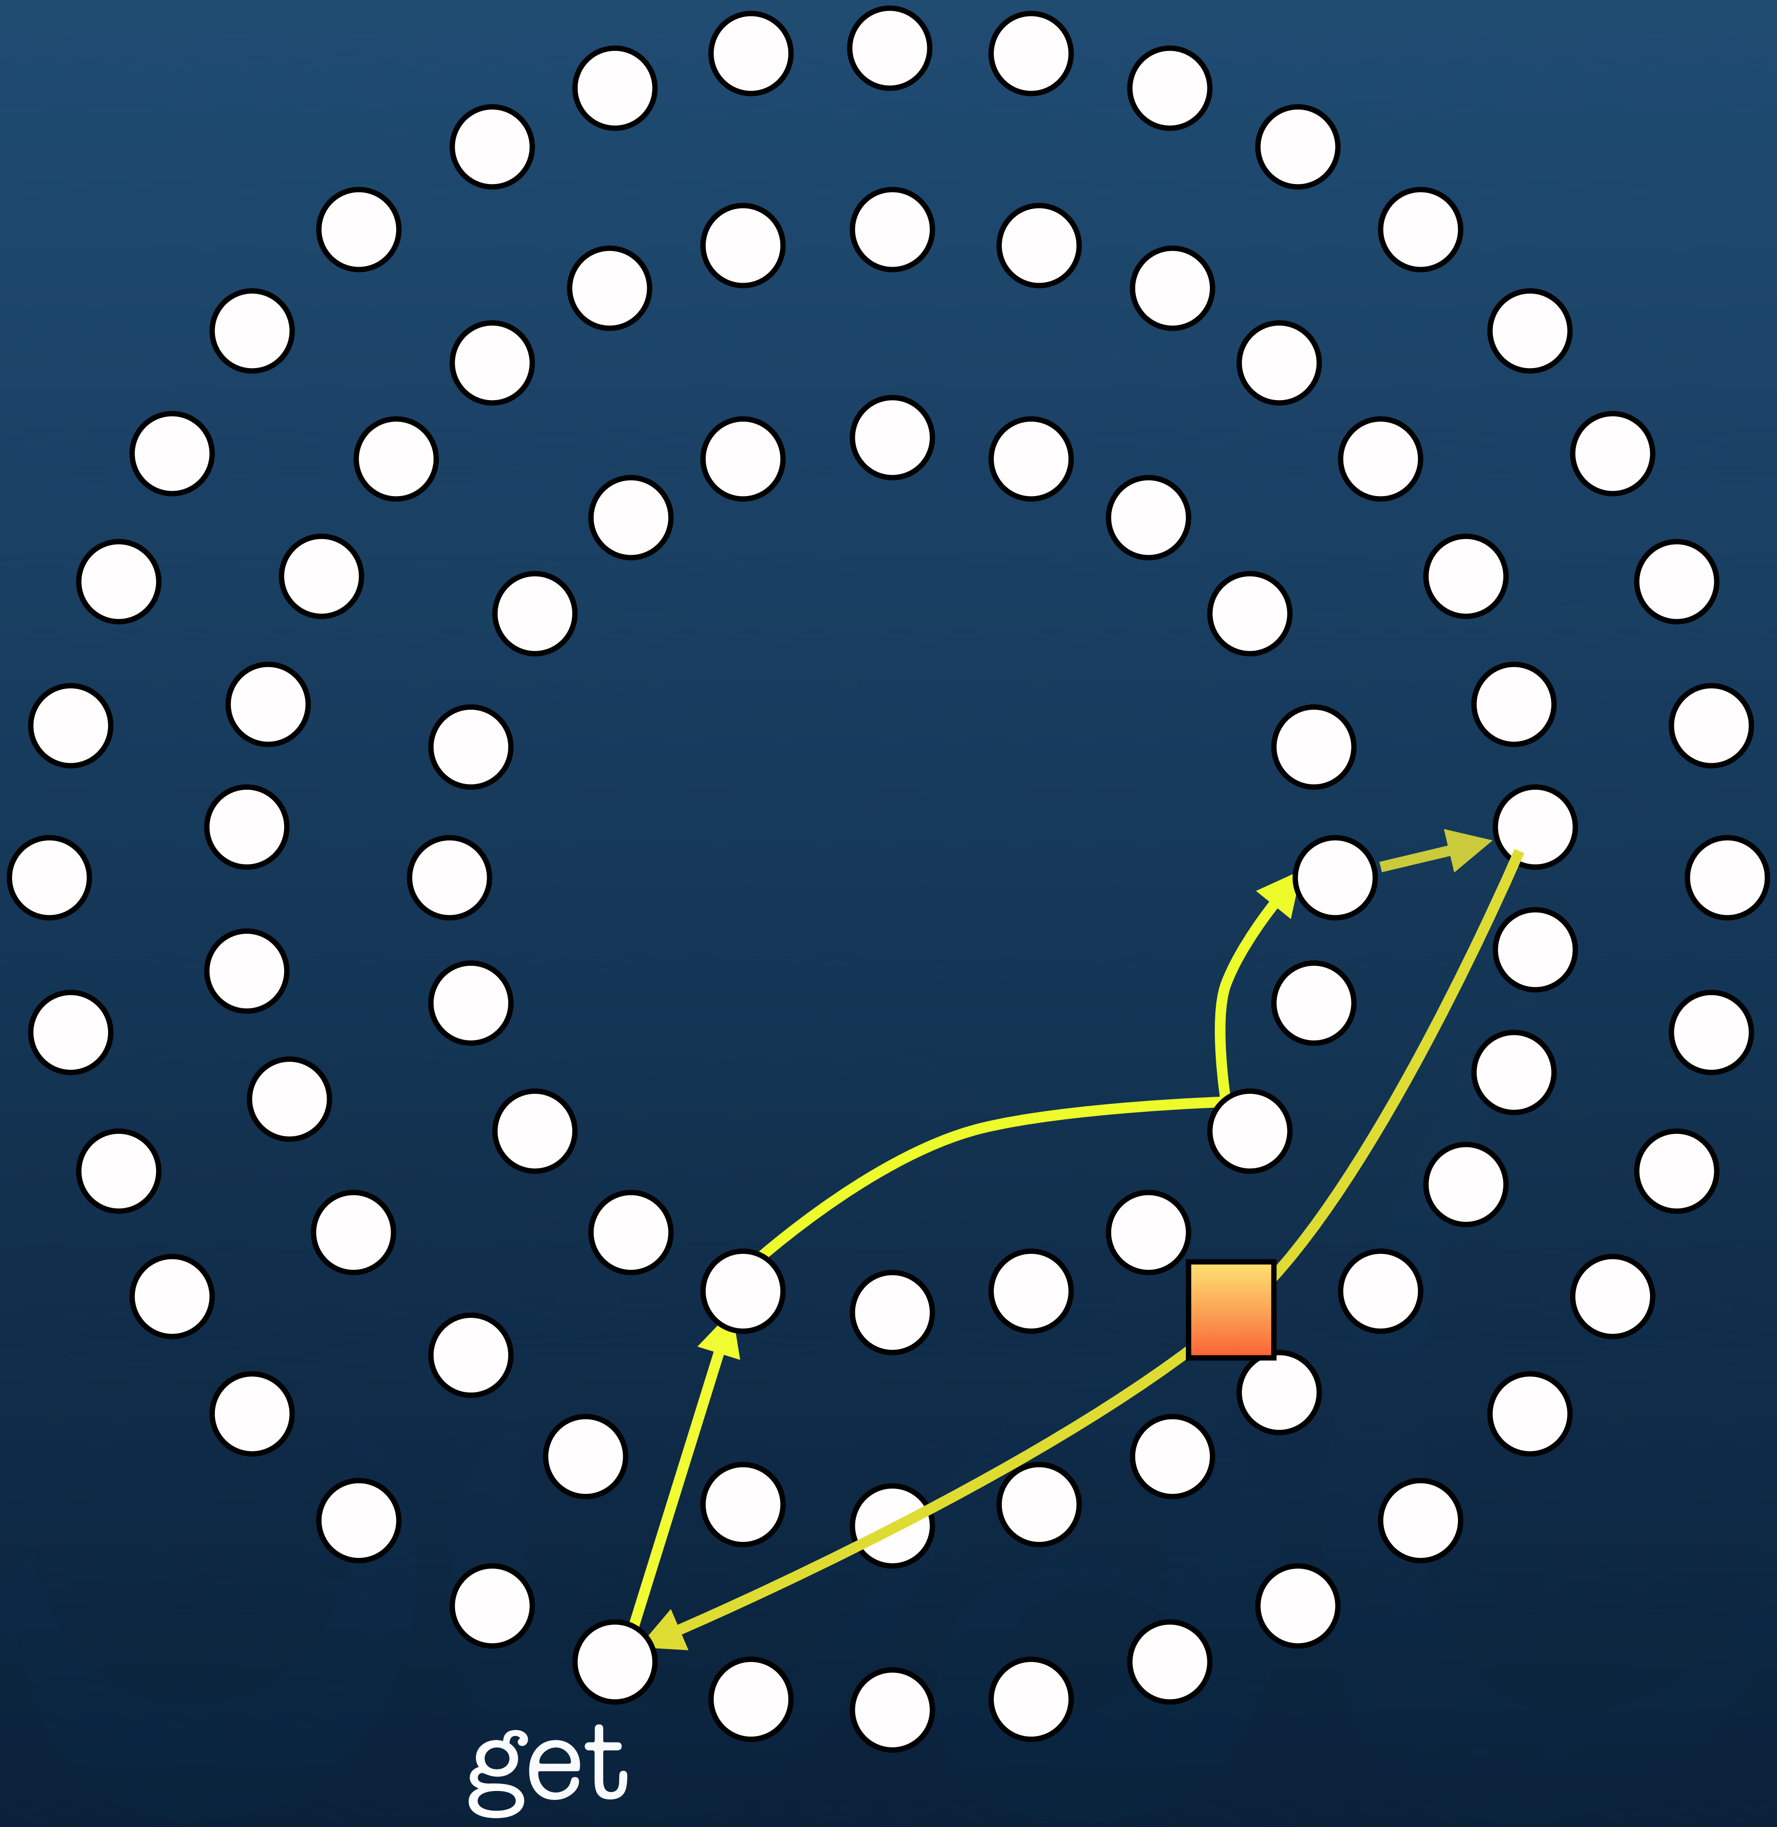

get

download of fragments (in parallel)

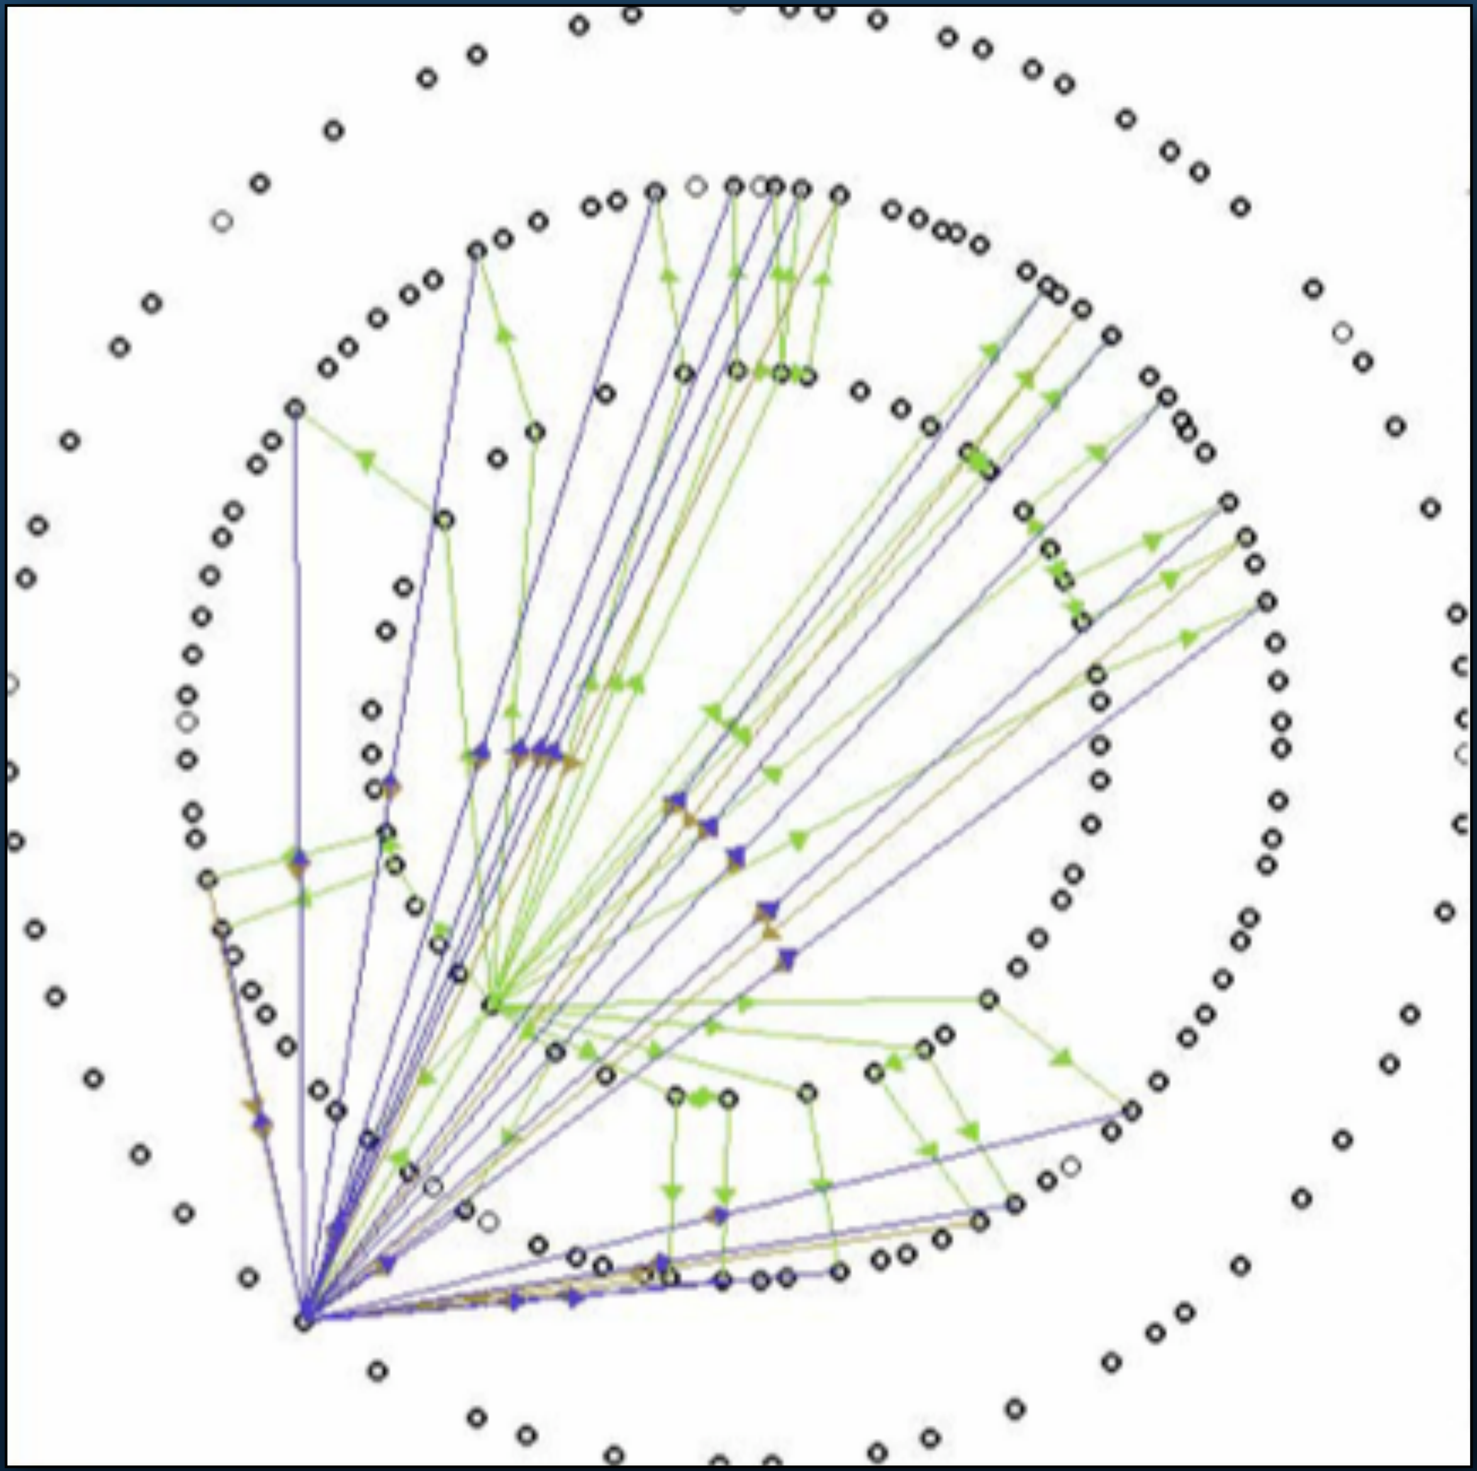

routing

napster: centralized :-(
gnutella: flooding :-(

chord, tapestry: structured overlay networks
O(log n) hops :-)
n = # super nodes

vulnerable to attacks (partitioning) :-(

super node
connected to direct neighbors
plus some random links

random links?
piggy-pack routing information

number of hops depends on

size of the network (n)
size of the routing table (R)

which itself depends on the traffic
we have lots of traffic due to erasure coding

simulation results

$n = 10^6$

$R = 1{,}000$: < 3 hops

$R = 100$: ~5 hops

reasonable already with moderate traffic

small world effects

(see milgram, watts & strogatz, kleinberg)

regular graph

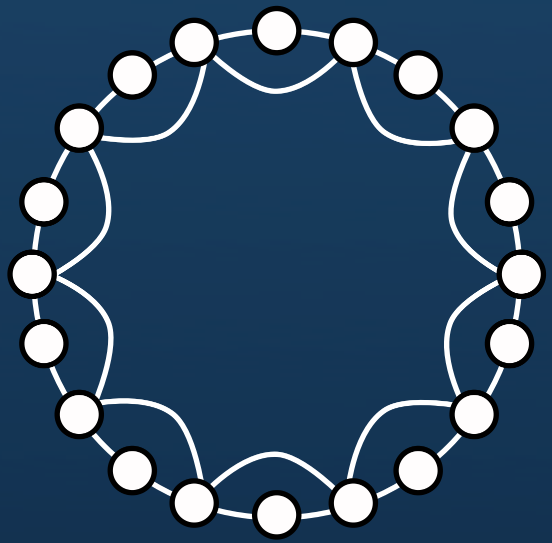

high diameter :-(
high clustering :-)

small world effects

(see milgram, watts & strogatz, kleinberg)

regular graph

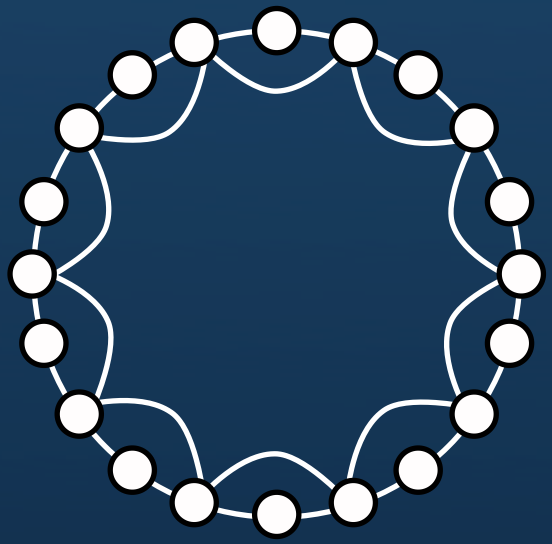

random graph

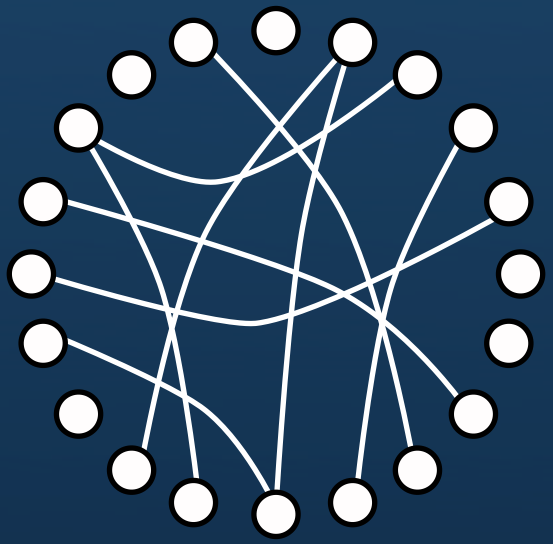

high diameter :-(
high clustering :-)

low diameter :-)
low clustering :-(

small world effects

(see milgram, watts & strogatz, kleinberg)

regular graph

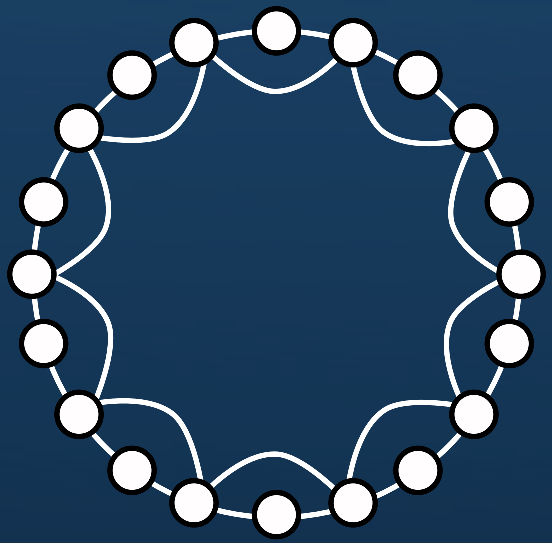

high diameter :-(
high clustering :-)

random graph

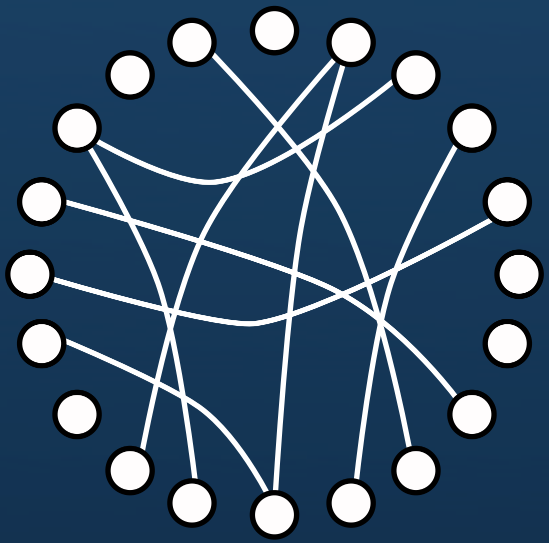

low diameter :-)
low clustering :-(

mix

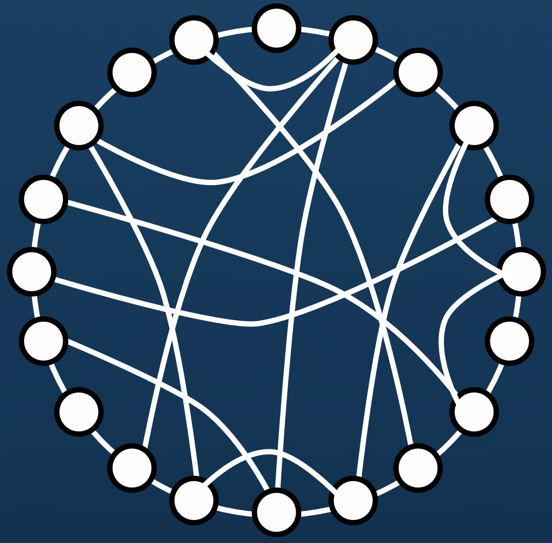

low diameter :-)
high clustering :-)

routing table

$n = 10^9$, $R = 10,000$

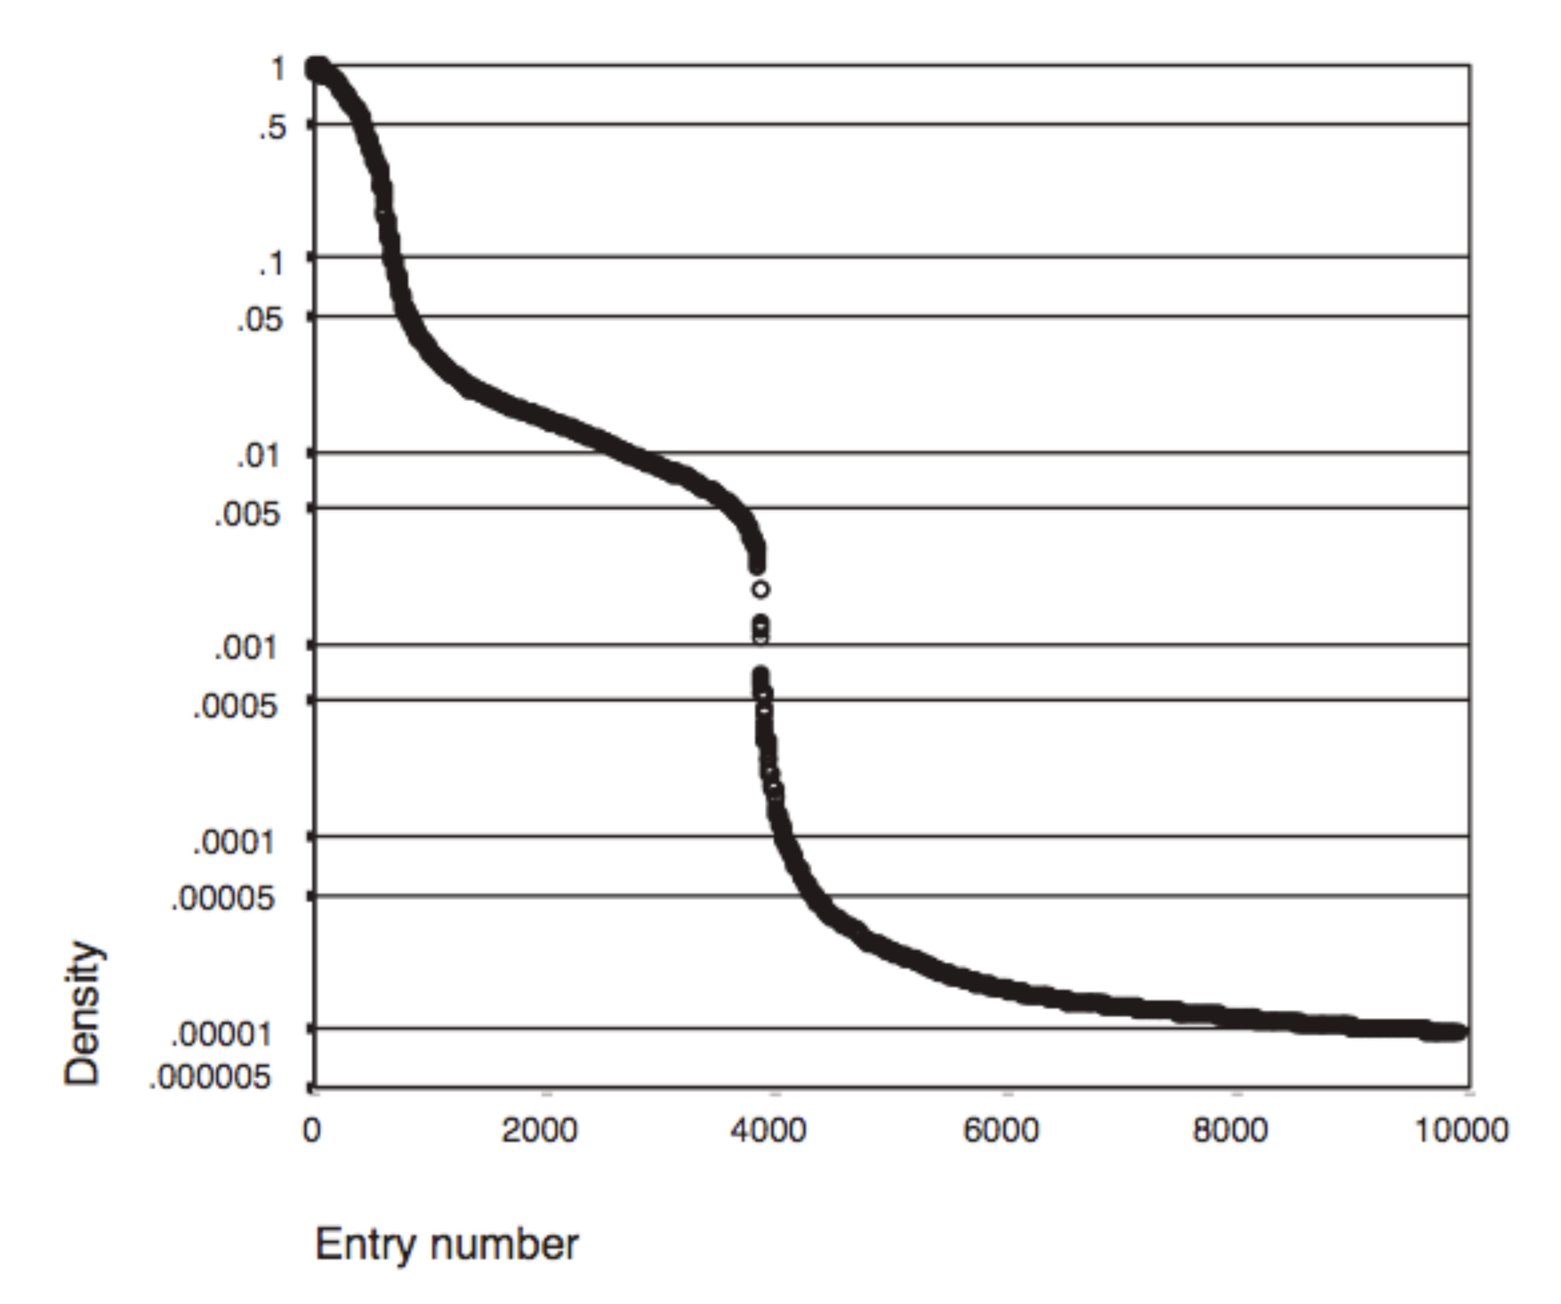

incentives, fairness
prevent free-riding

local disk space
online time

upload bandwidth

online storage = local disk space * online time

example: 10 GB disk space, 70% online --> 7 GB

we have different mechanisms to measure
and check these two variables

trading storage

only if you want to (you start with 1 GB)

you must be online at least 17% of the time

(≈ 4 hours a day, running average)

storage can be earned on multiple computers

upload bandwidth

the more upload bandwidth you provide,
the more download bandwidth you get

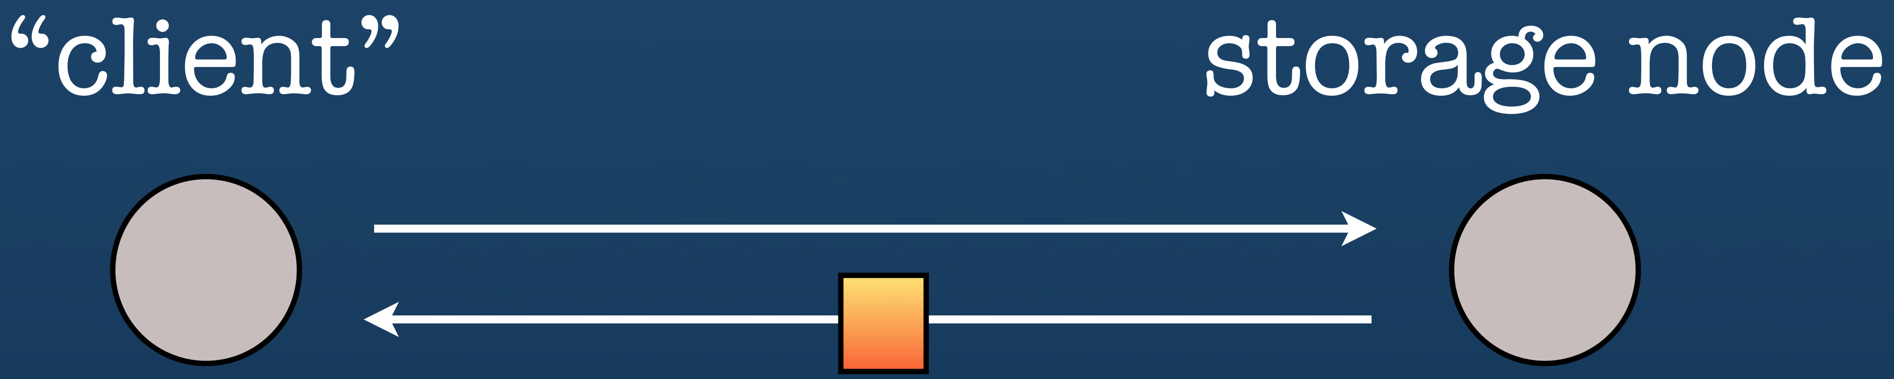

asymmetric interest

tit-for-tat doesn't work :-(

believe the software? hack it (kazaa lite) :-(

distributed reputation system
that is not susceptible to false reports
and other forms of cheating

must scale well with number of transactions
we have lots of small transactions due to erasure coding

🗎 Havelaar, NetEcon 2006

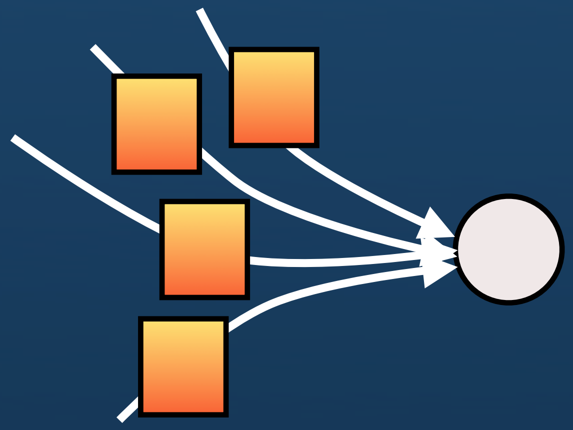

1. lots of transactions
"observations"

Havelaar, NetEcon 2006

2. every round (e.g., a week) send observations to pre-determined neighbors (hash code)

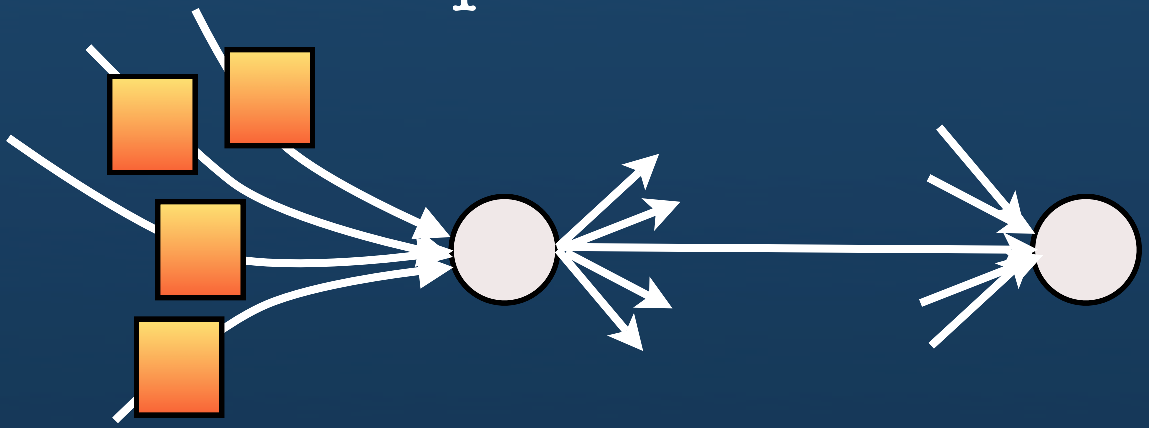

1. lots of transactions “observations”

Havelaar, NetEcon 2006

2. every round (e.g., a week)
send observations to
pre-determined neighbors (hash code)

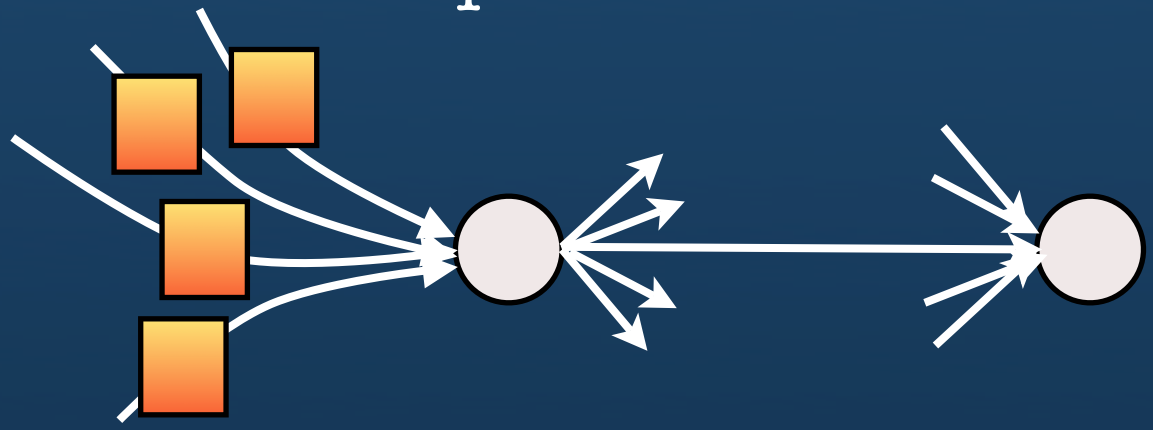

3. discard ego-reports,
median, etc.

1. lots of transactions
“observations”

Havelaar, NetEcon 2006

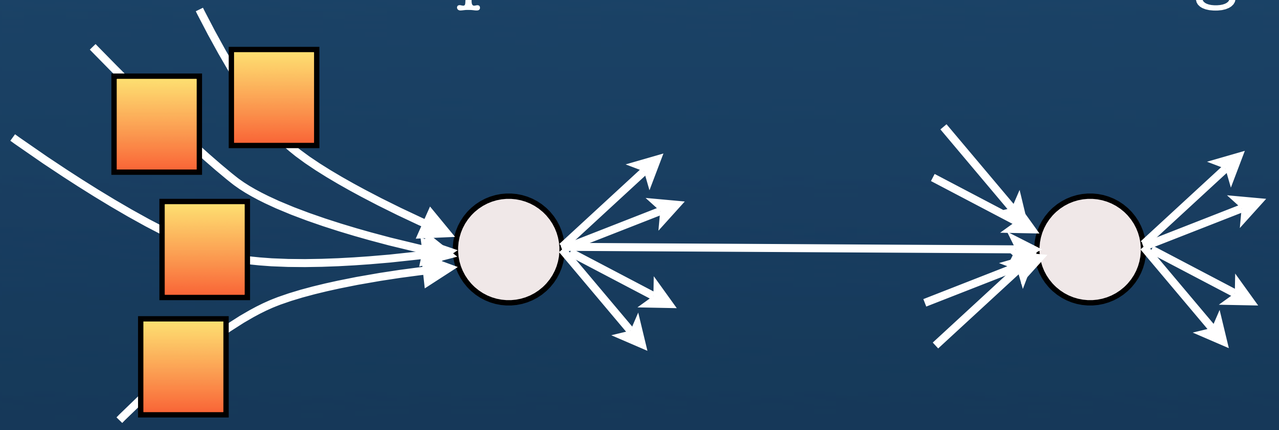

1. lots of transactions
“observations”

2. every round (e.g., a week)
send observations to
pre-determined neighbors (hash code)

3. discard ego-reports,
median, etc.

4. next round, aggregate

Havelaar, NetEcon 2006

2. every round (e.g., a week)
send observations to
pre-determined neighbors (hash code)

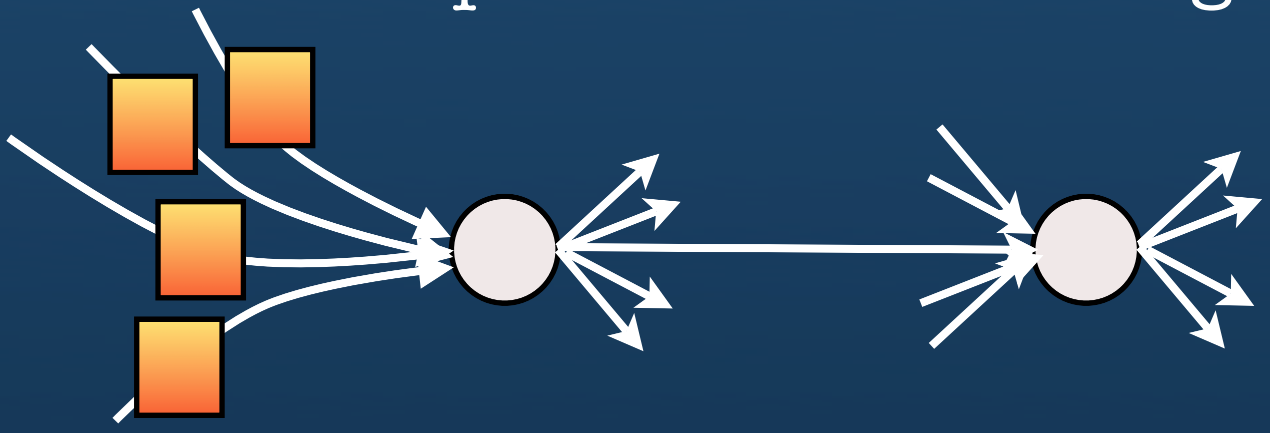

3. discard ego-reports,
median, etc.

4. next round, aggregate

1. lots of transactions
"observations"

5. update reputation
of storage nodes

Havelaar, NetEcon 2006

2. every round (e.g., a week)
send observations to
pre-determined neighbors (hash code)

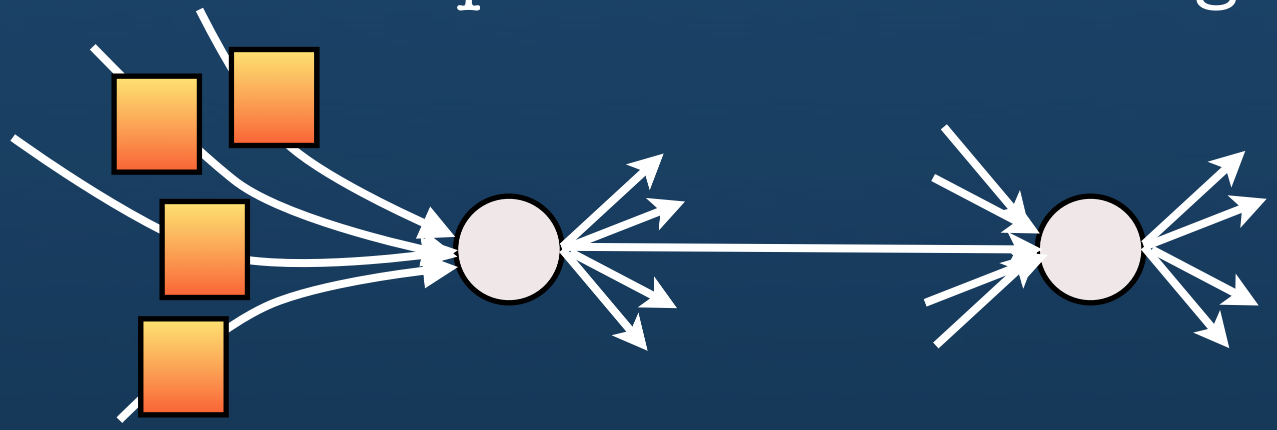

3. discard ego-reports,
median, etc.

4. next round, aggregate

1. lots of transactions
“observations”

5. update reputation
of storage nodes

rewarding:
upload bandwidth
proportional
to reputation

Havelaar, NetEcon 2006

local approximation of contribution

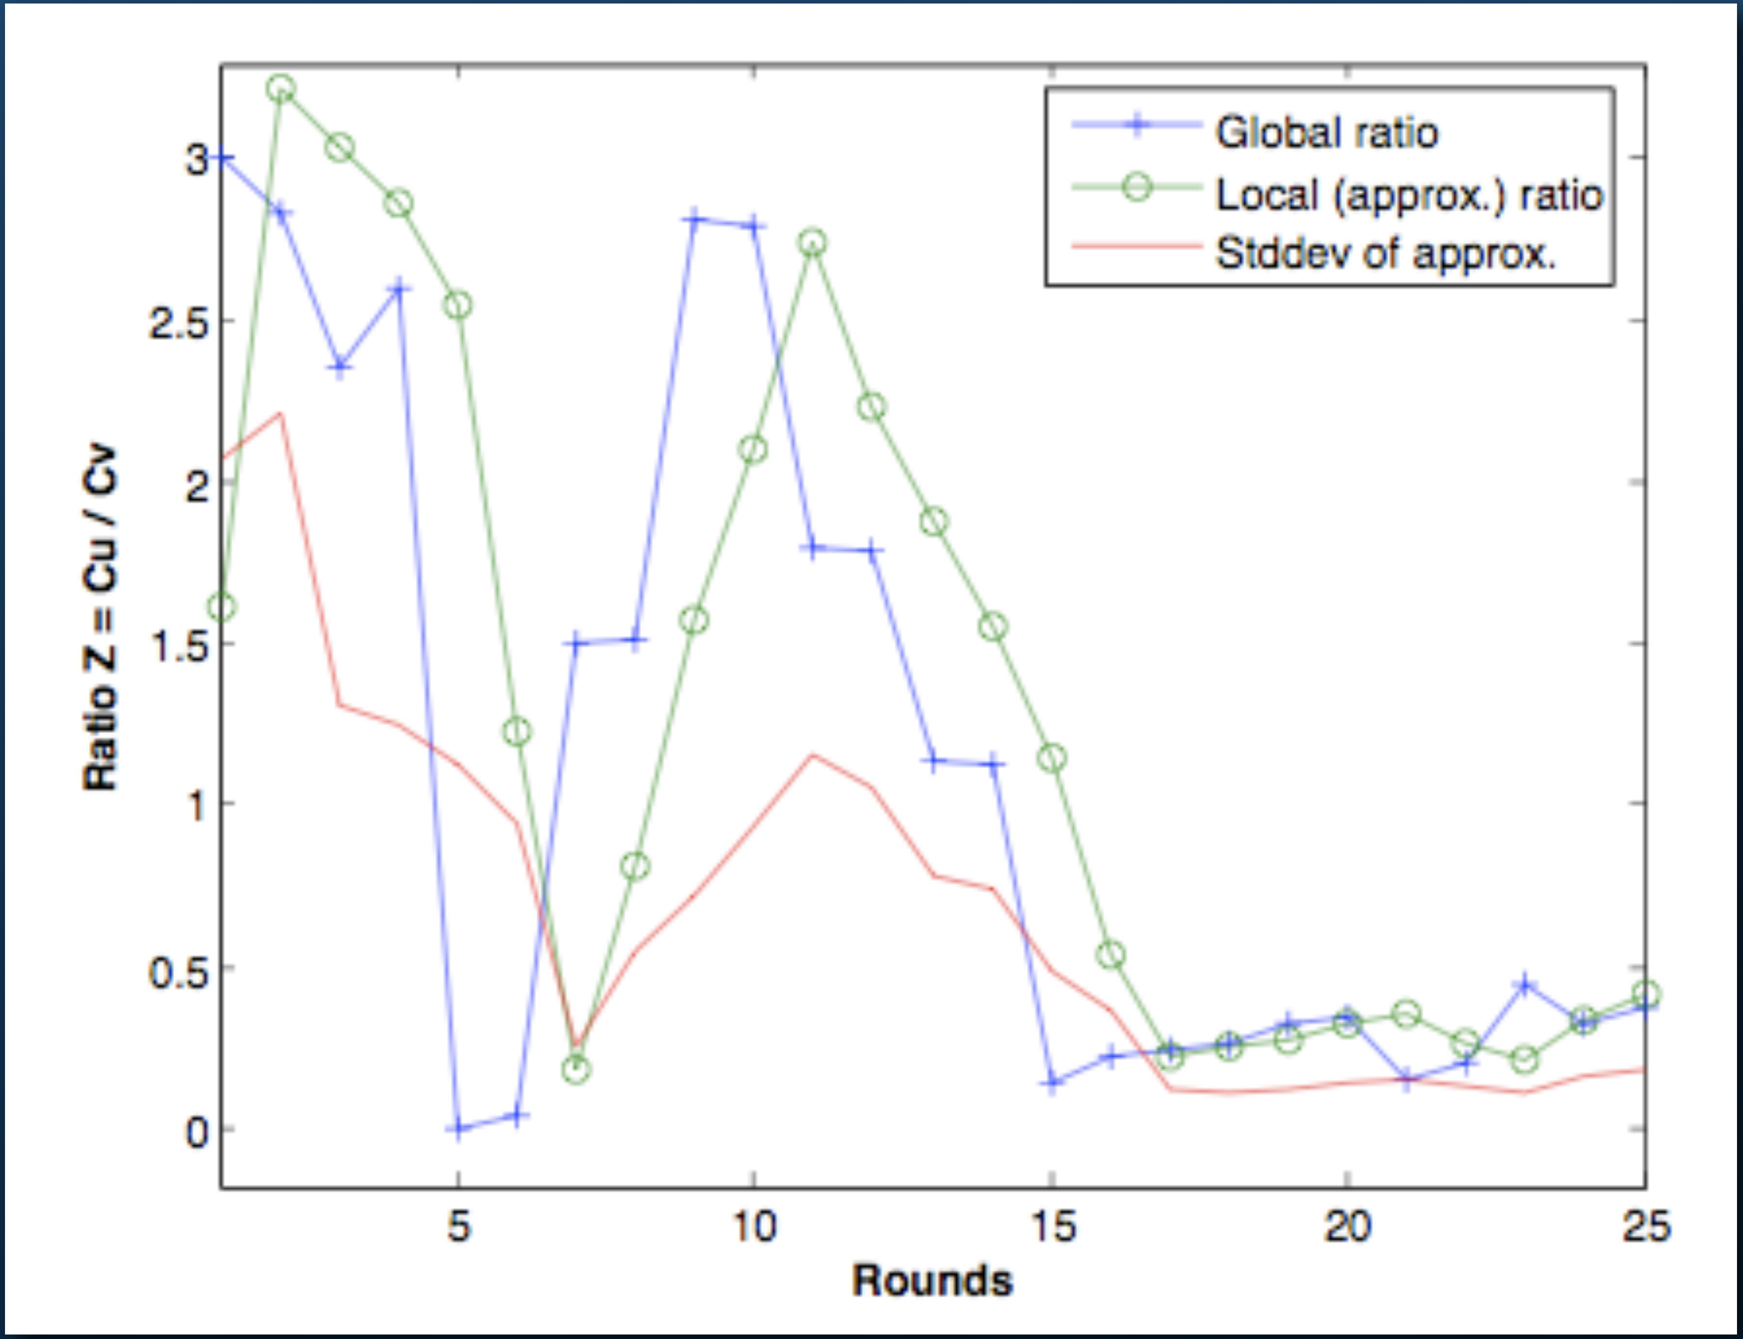

Havelaar, NetEcon 2006

“client”

storage node

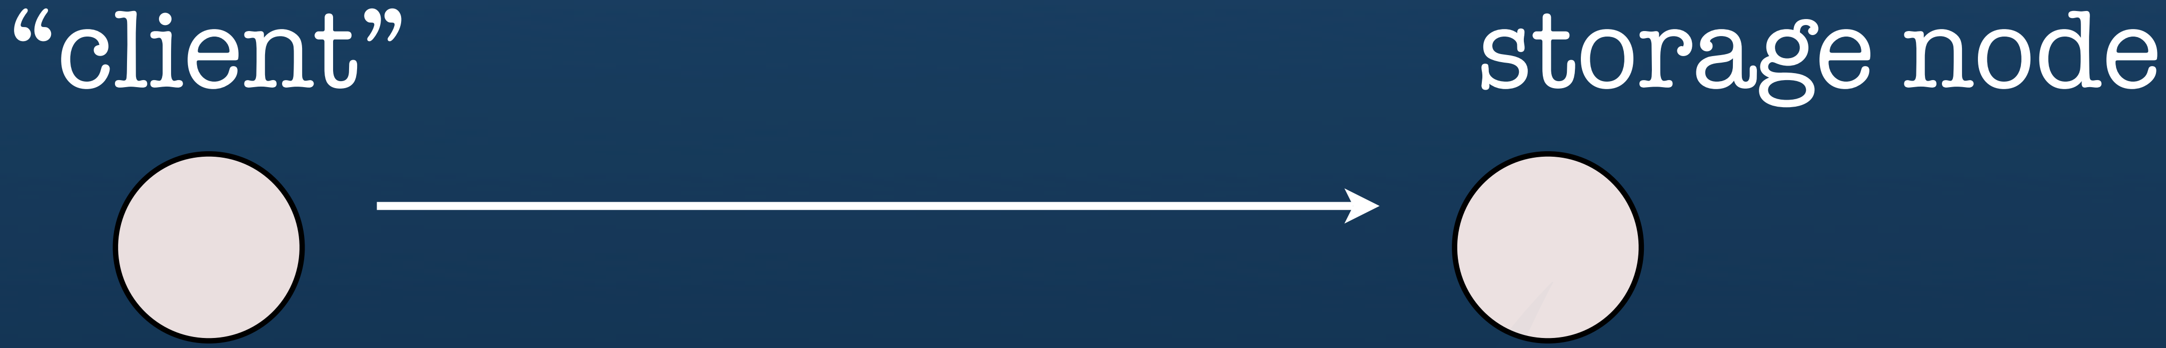

“client”

storage node

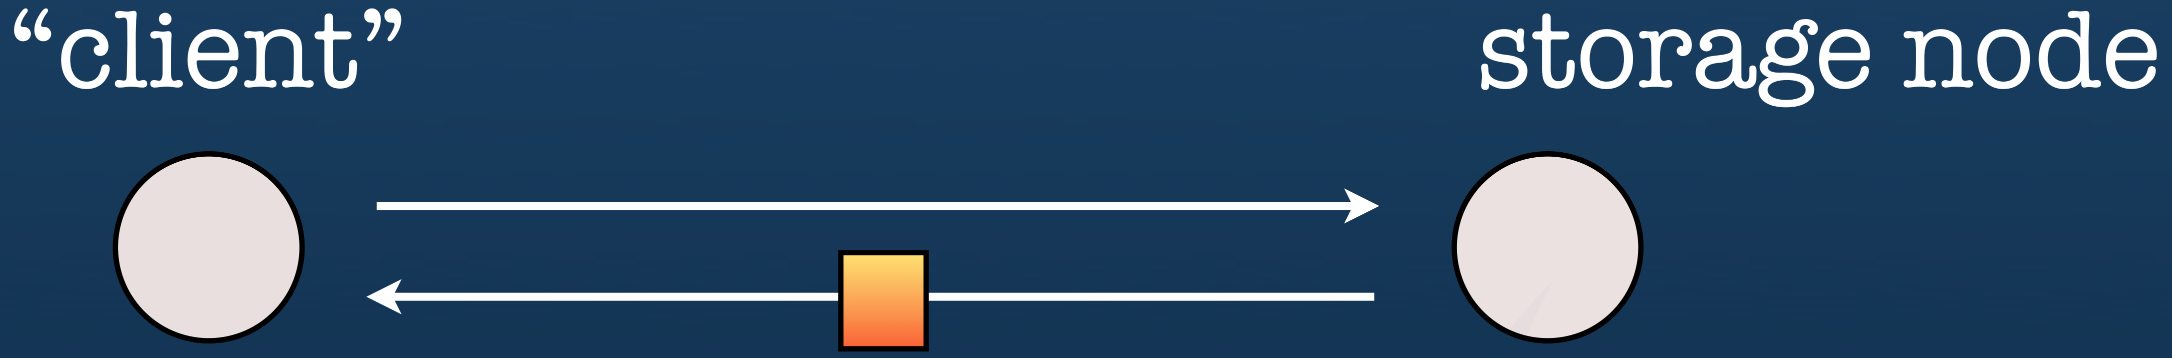
“client”
storage node

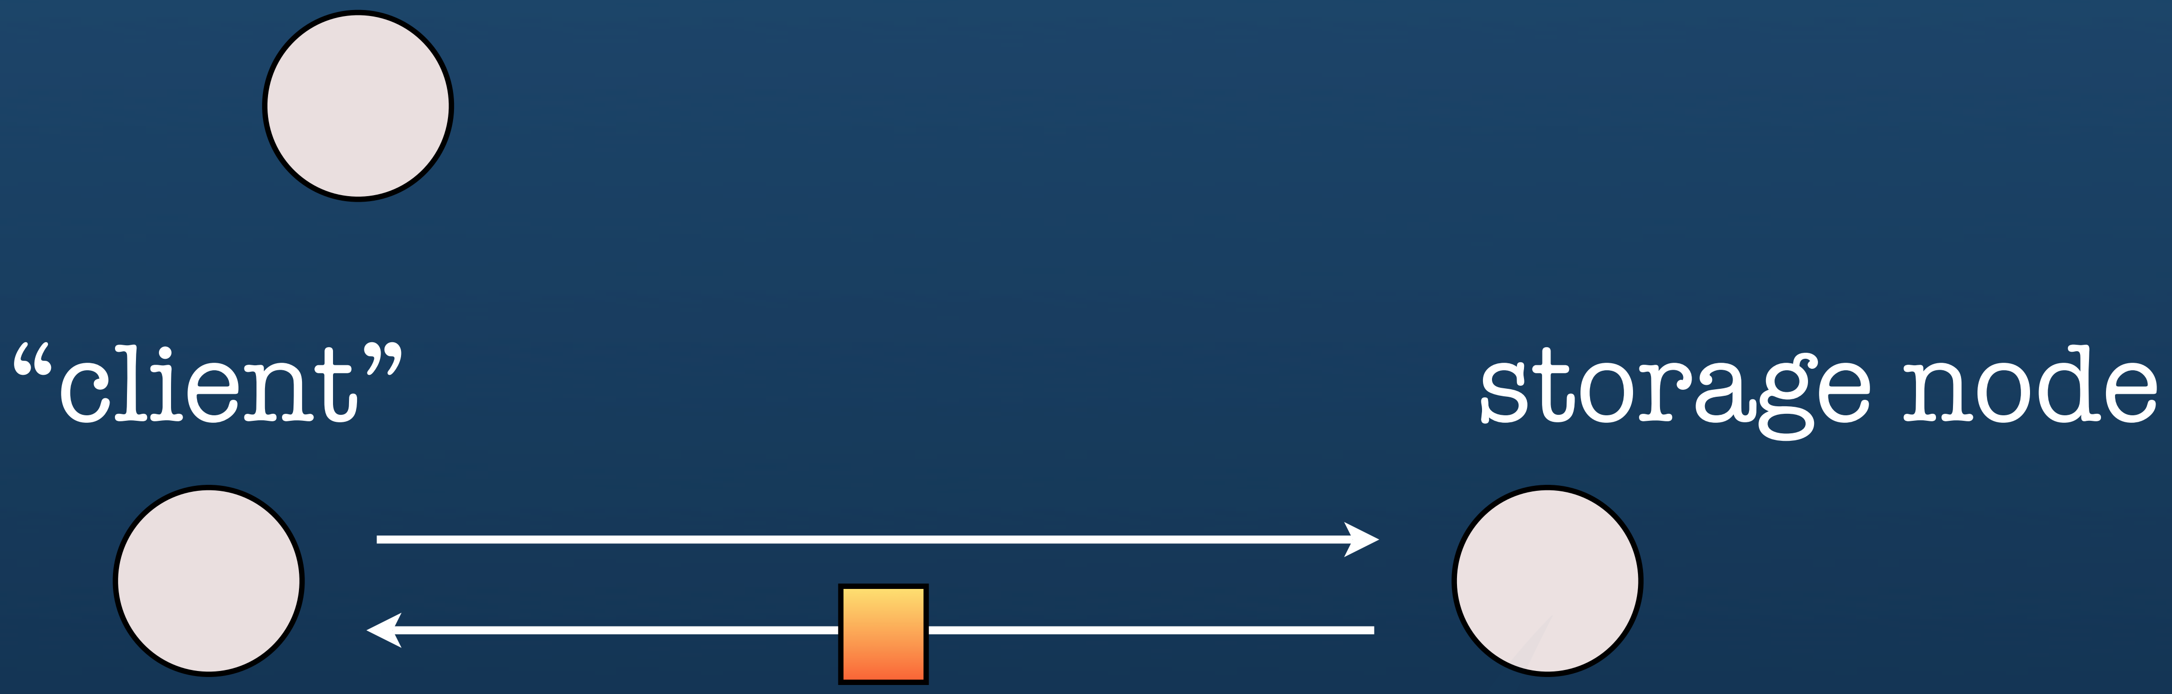
"client"
storage node

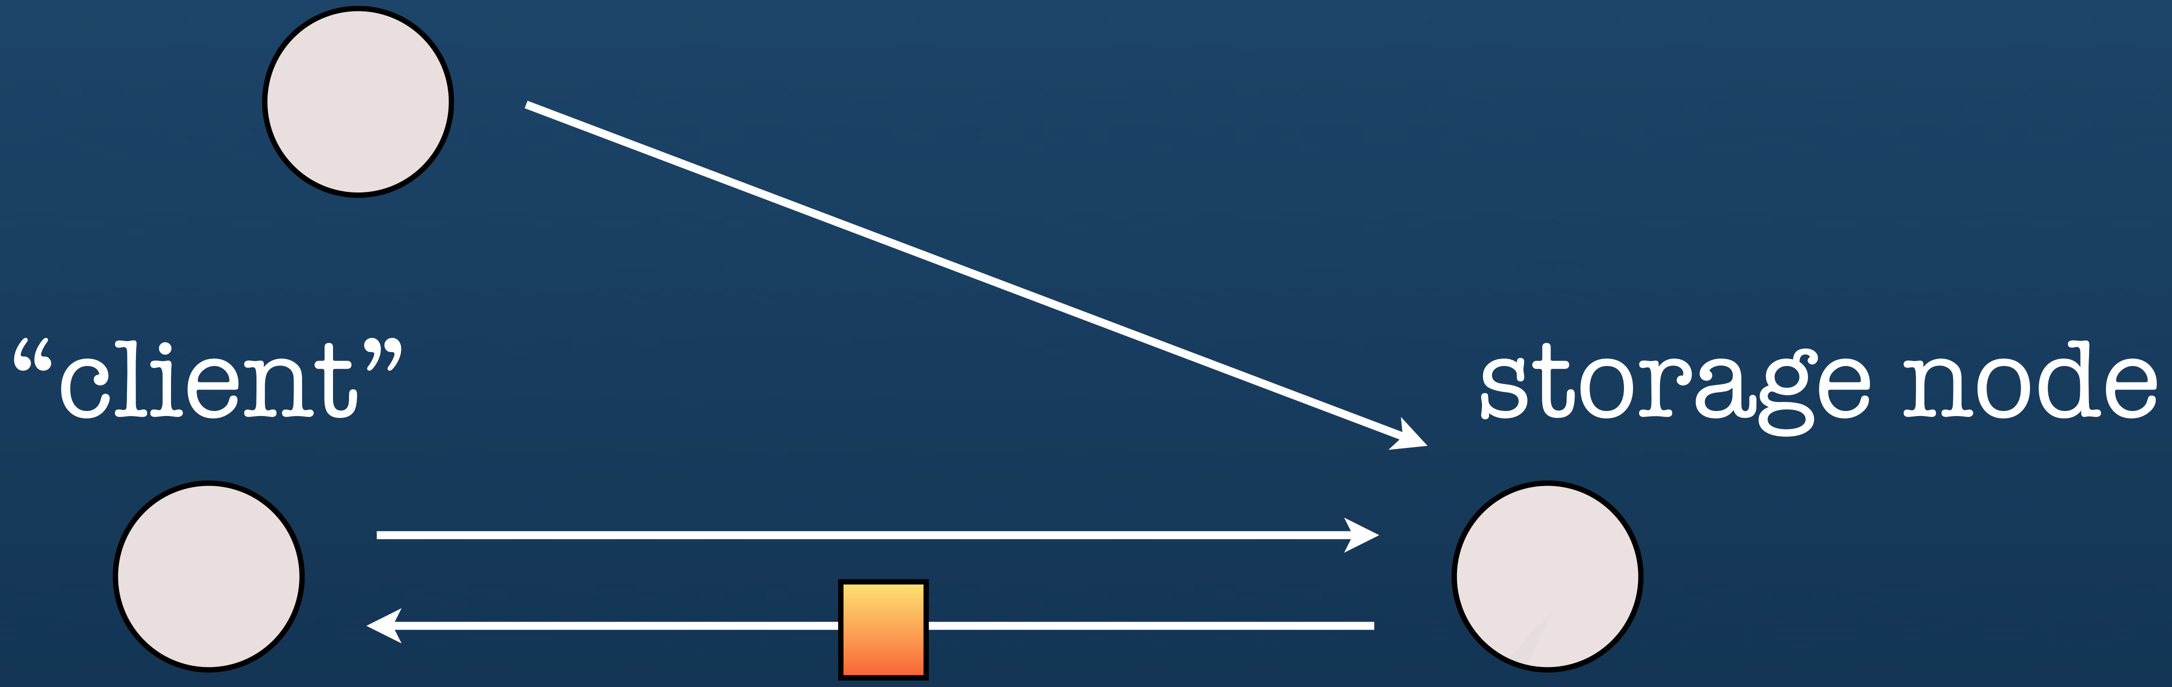
"client"
storage node

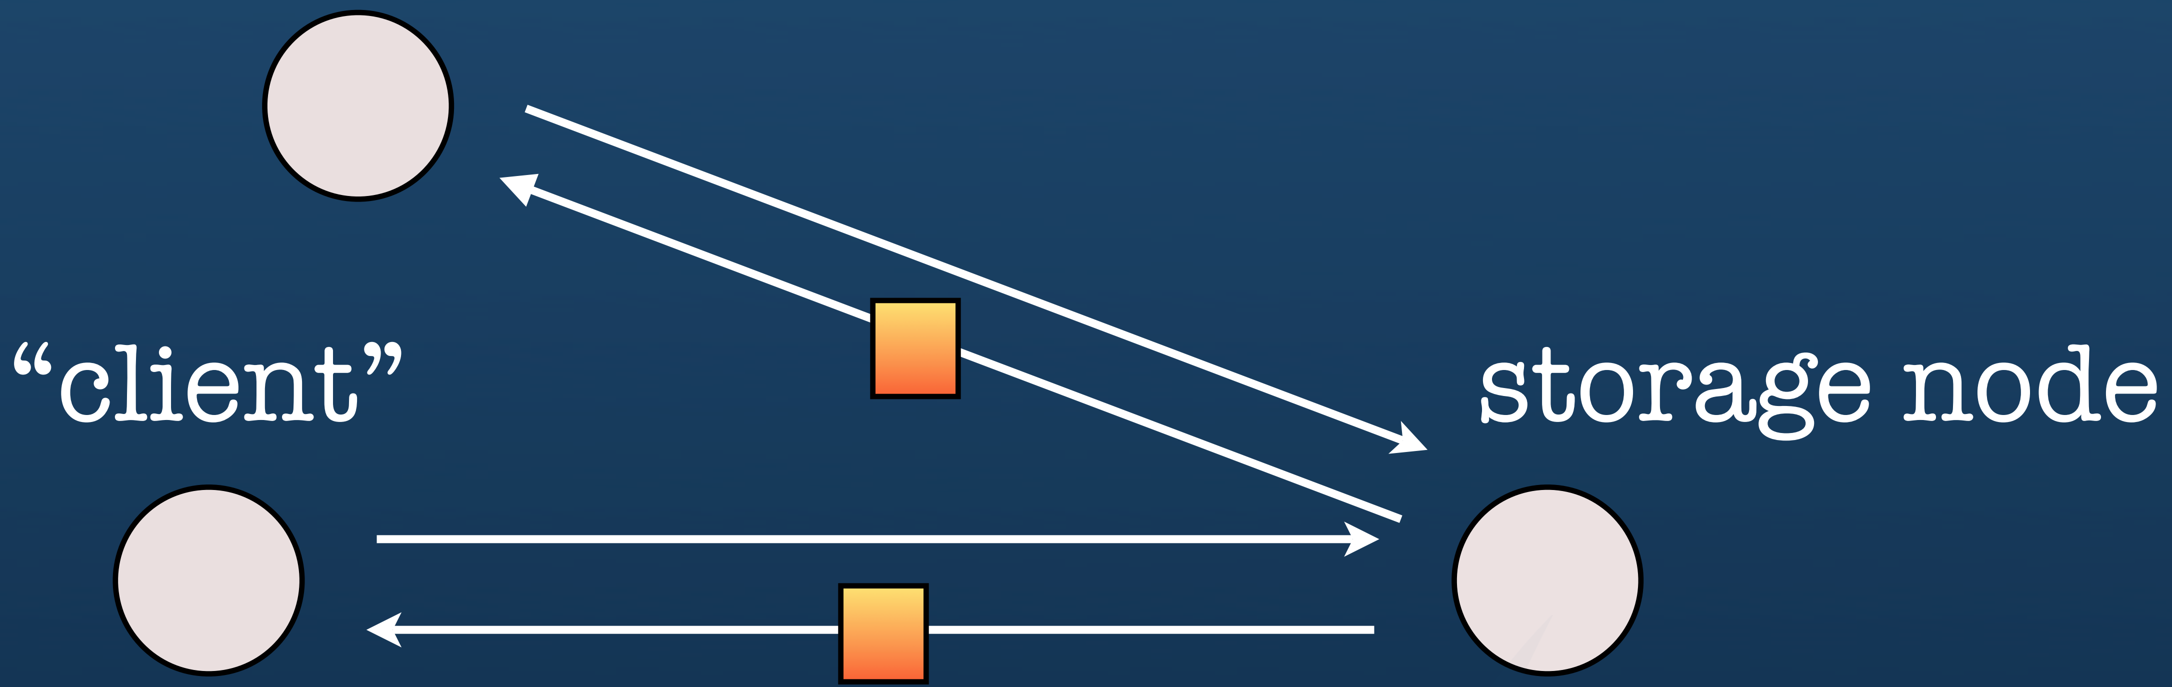
"client"
storage node

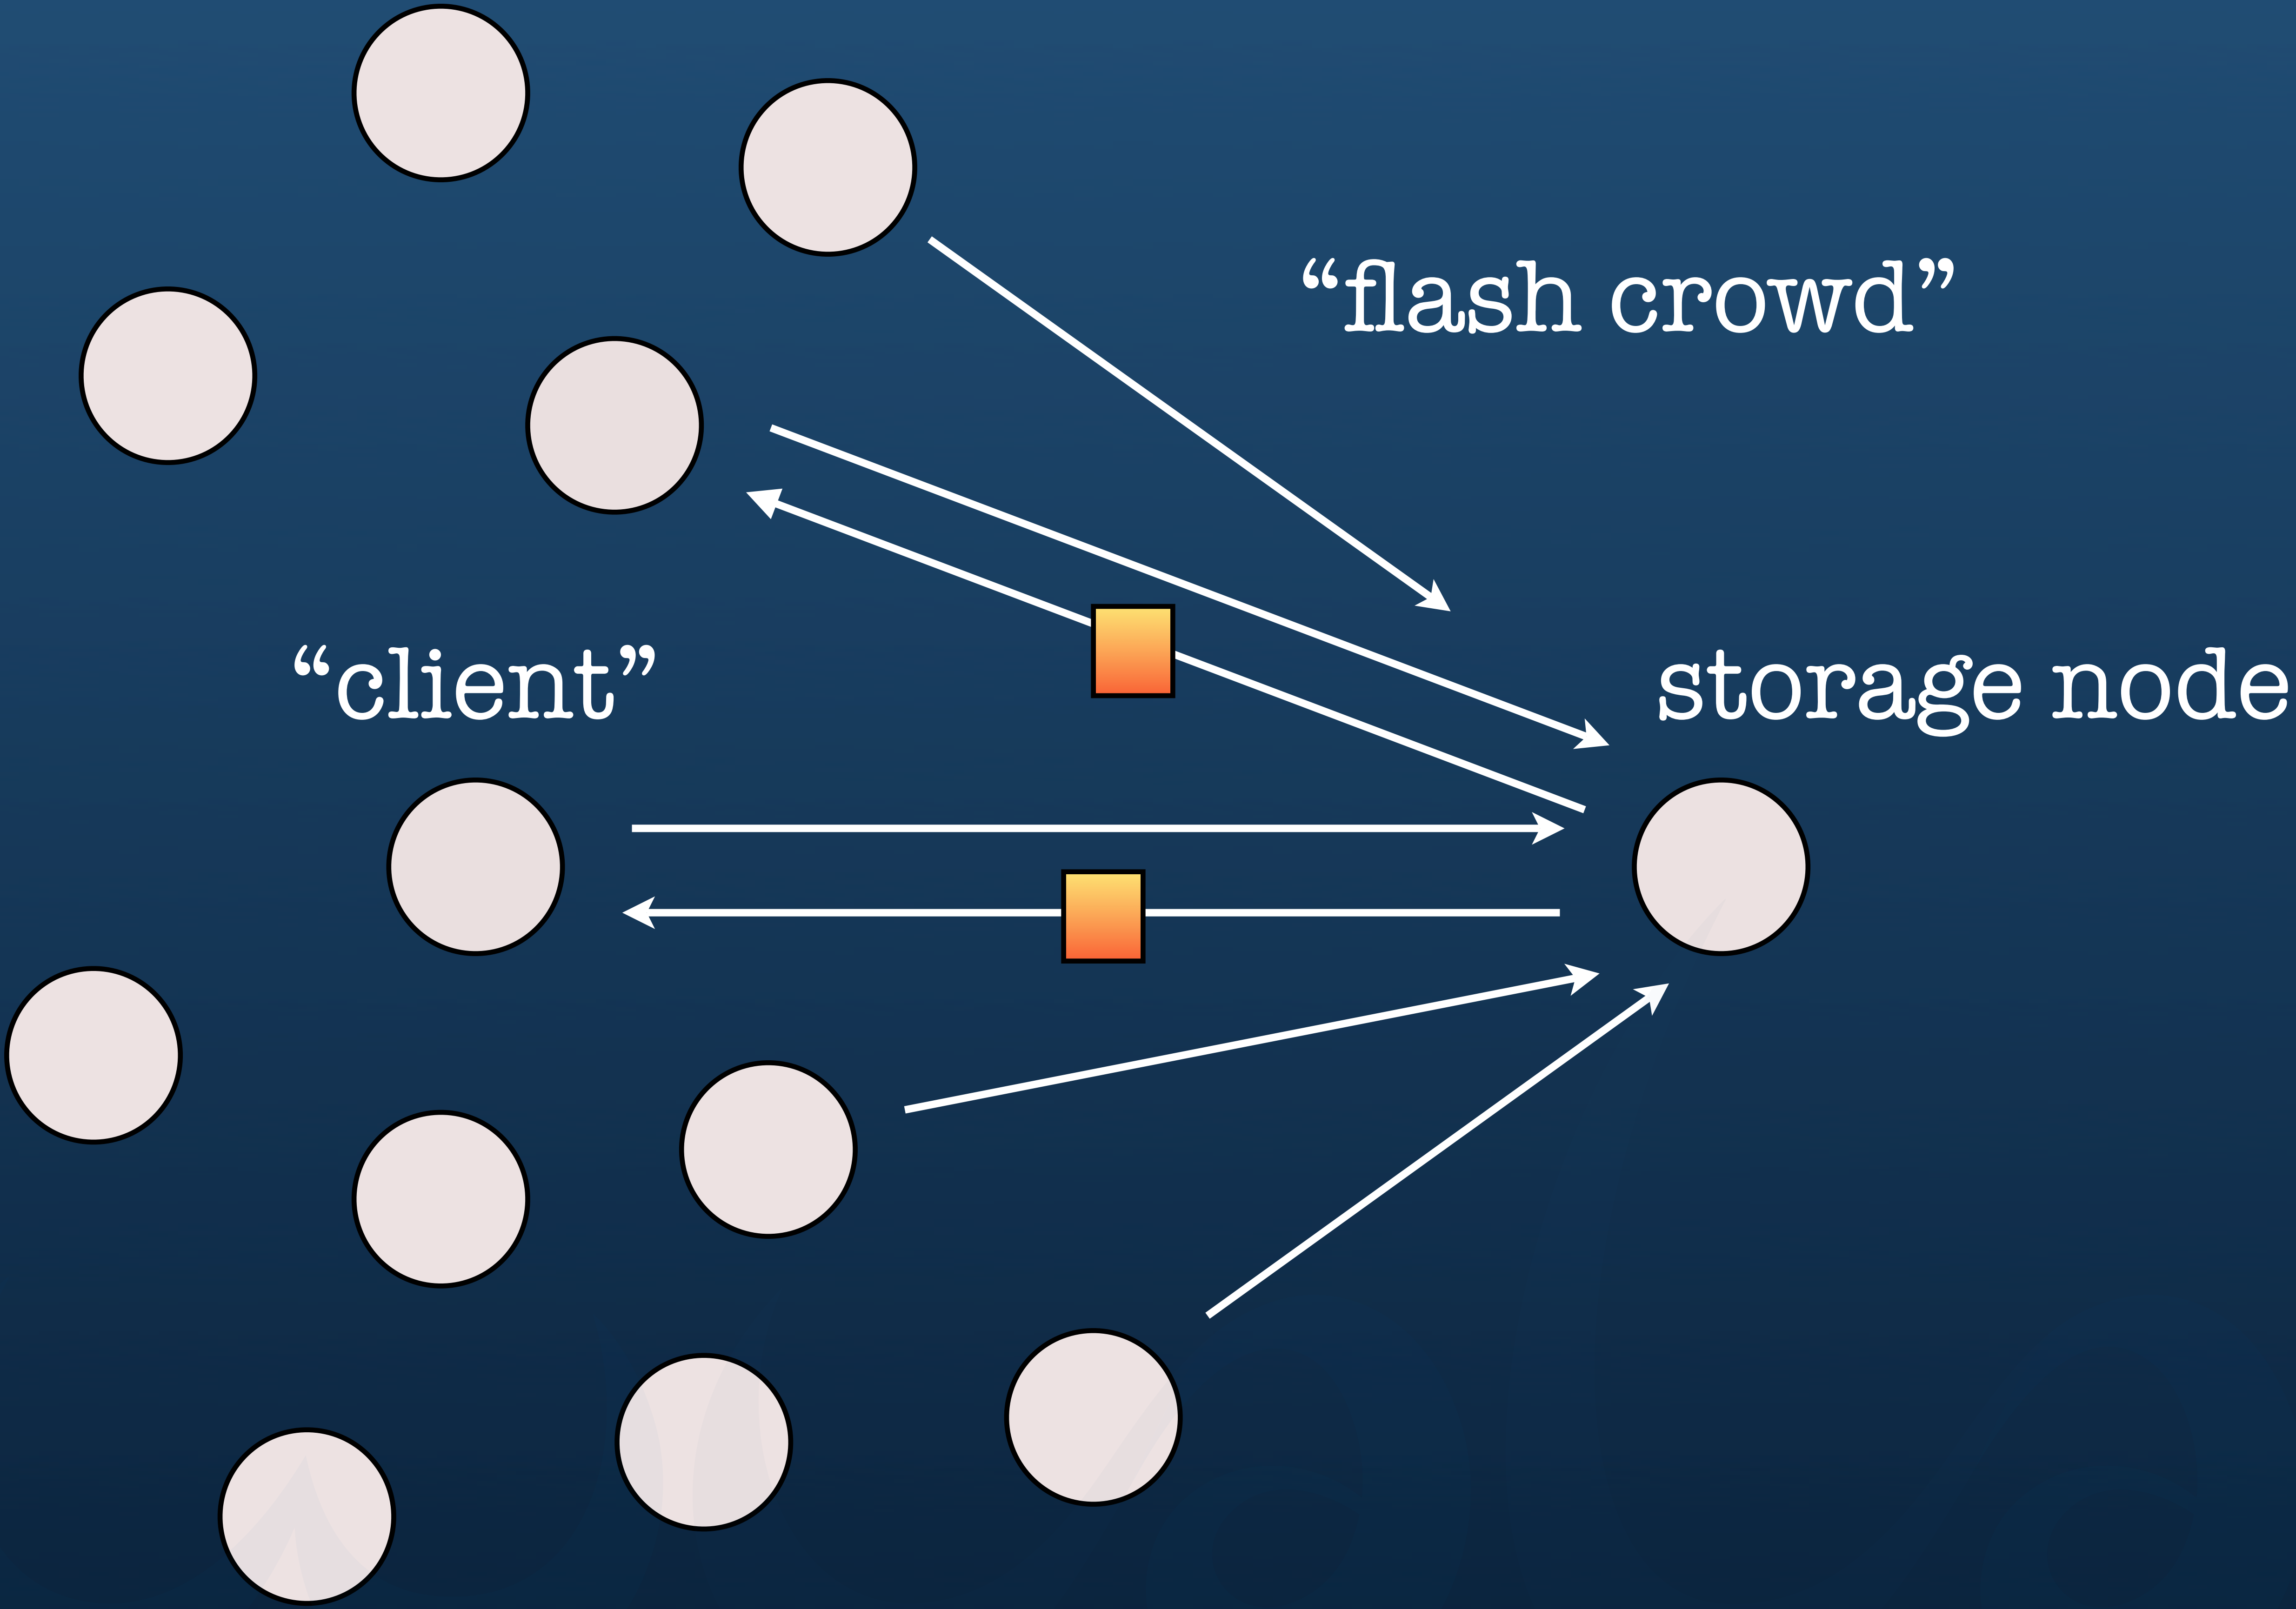
"flash crowd"
"client"
storage node

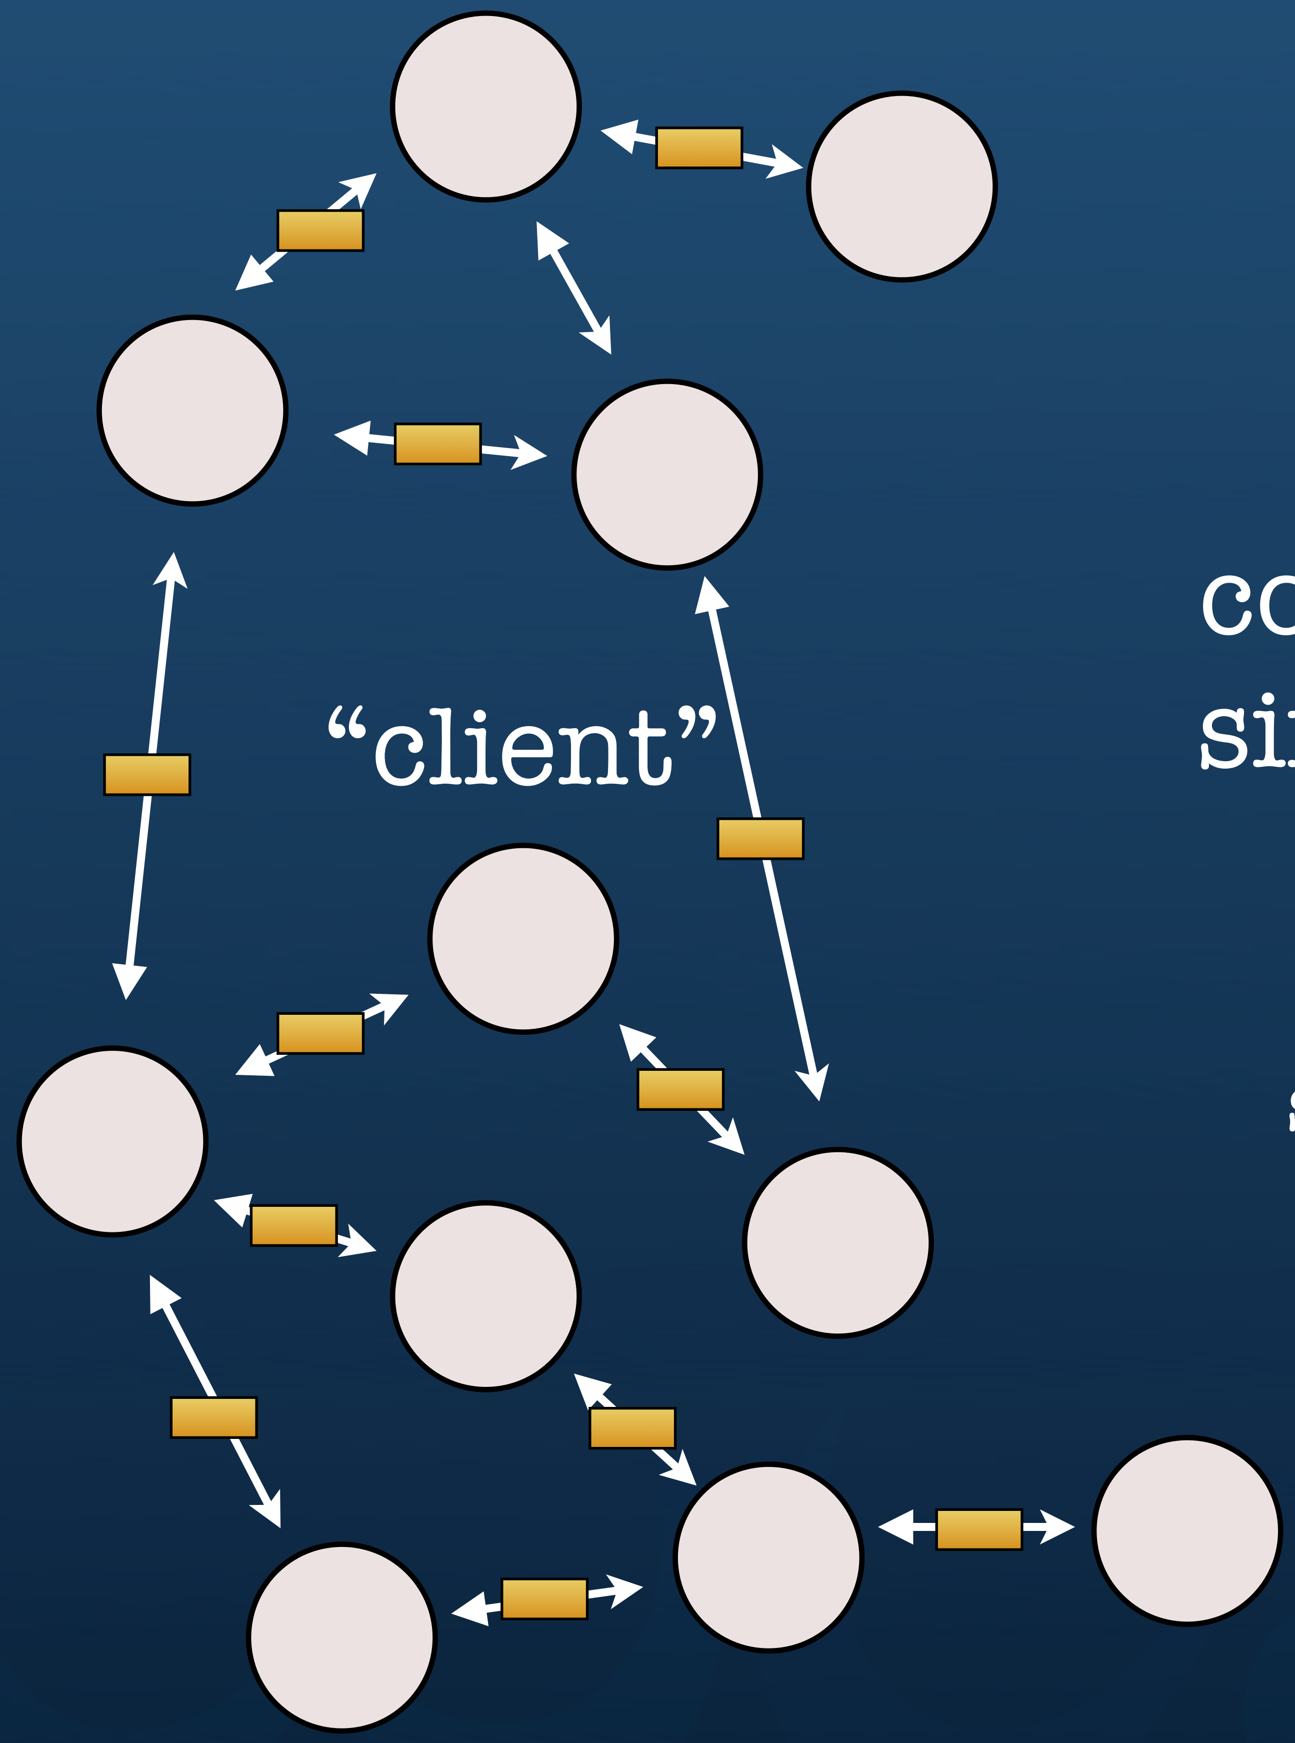

content distribution
similar to bittorrent
tit-for-tat

some differences due to
erasure codes

encryption

128 bit AES for encryption
2048 bit RSA for authentication

all data is encrypted (file + meta data)
all cryptographic operations performed locally
(i.e., on your computer)

access control

cryptographic tree structure
untrusted storage
doesn’t reveal who has access
very efficient for typical operations
(grant access, move, etc.)

📄 Cryptree, SRDS 2006

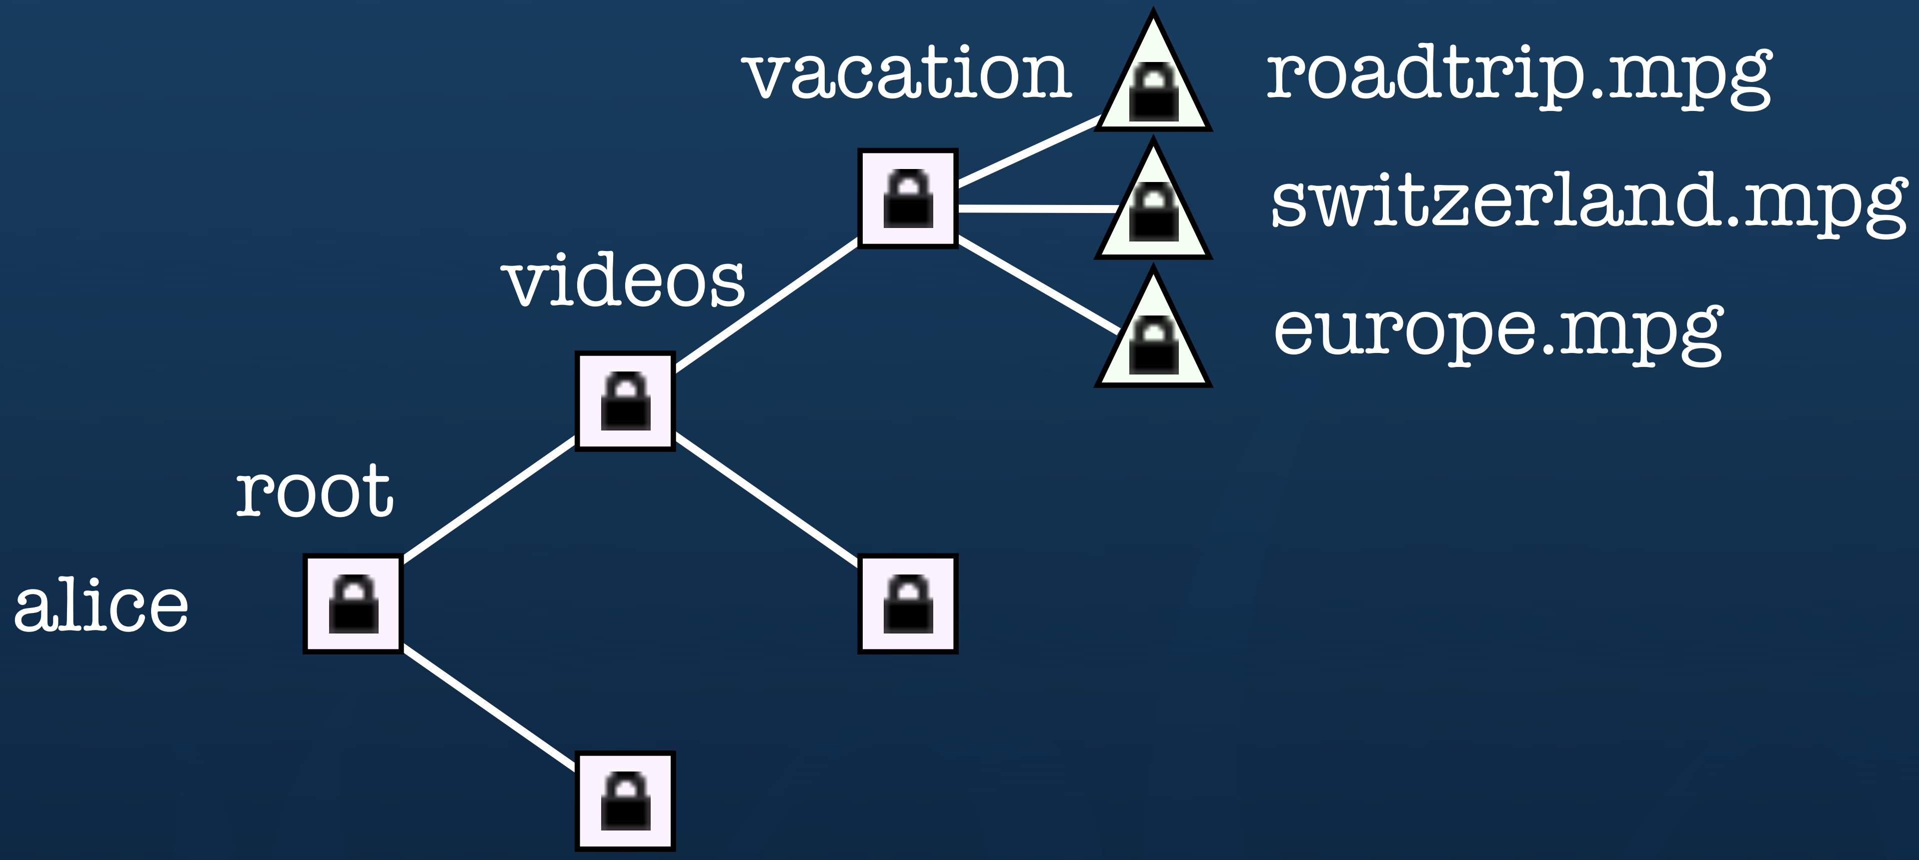

Cryptree, SRDS 2006

bob doesn't see that
claire has also access
and vice versa

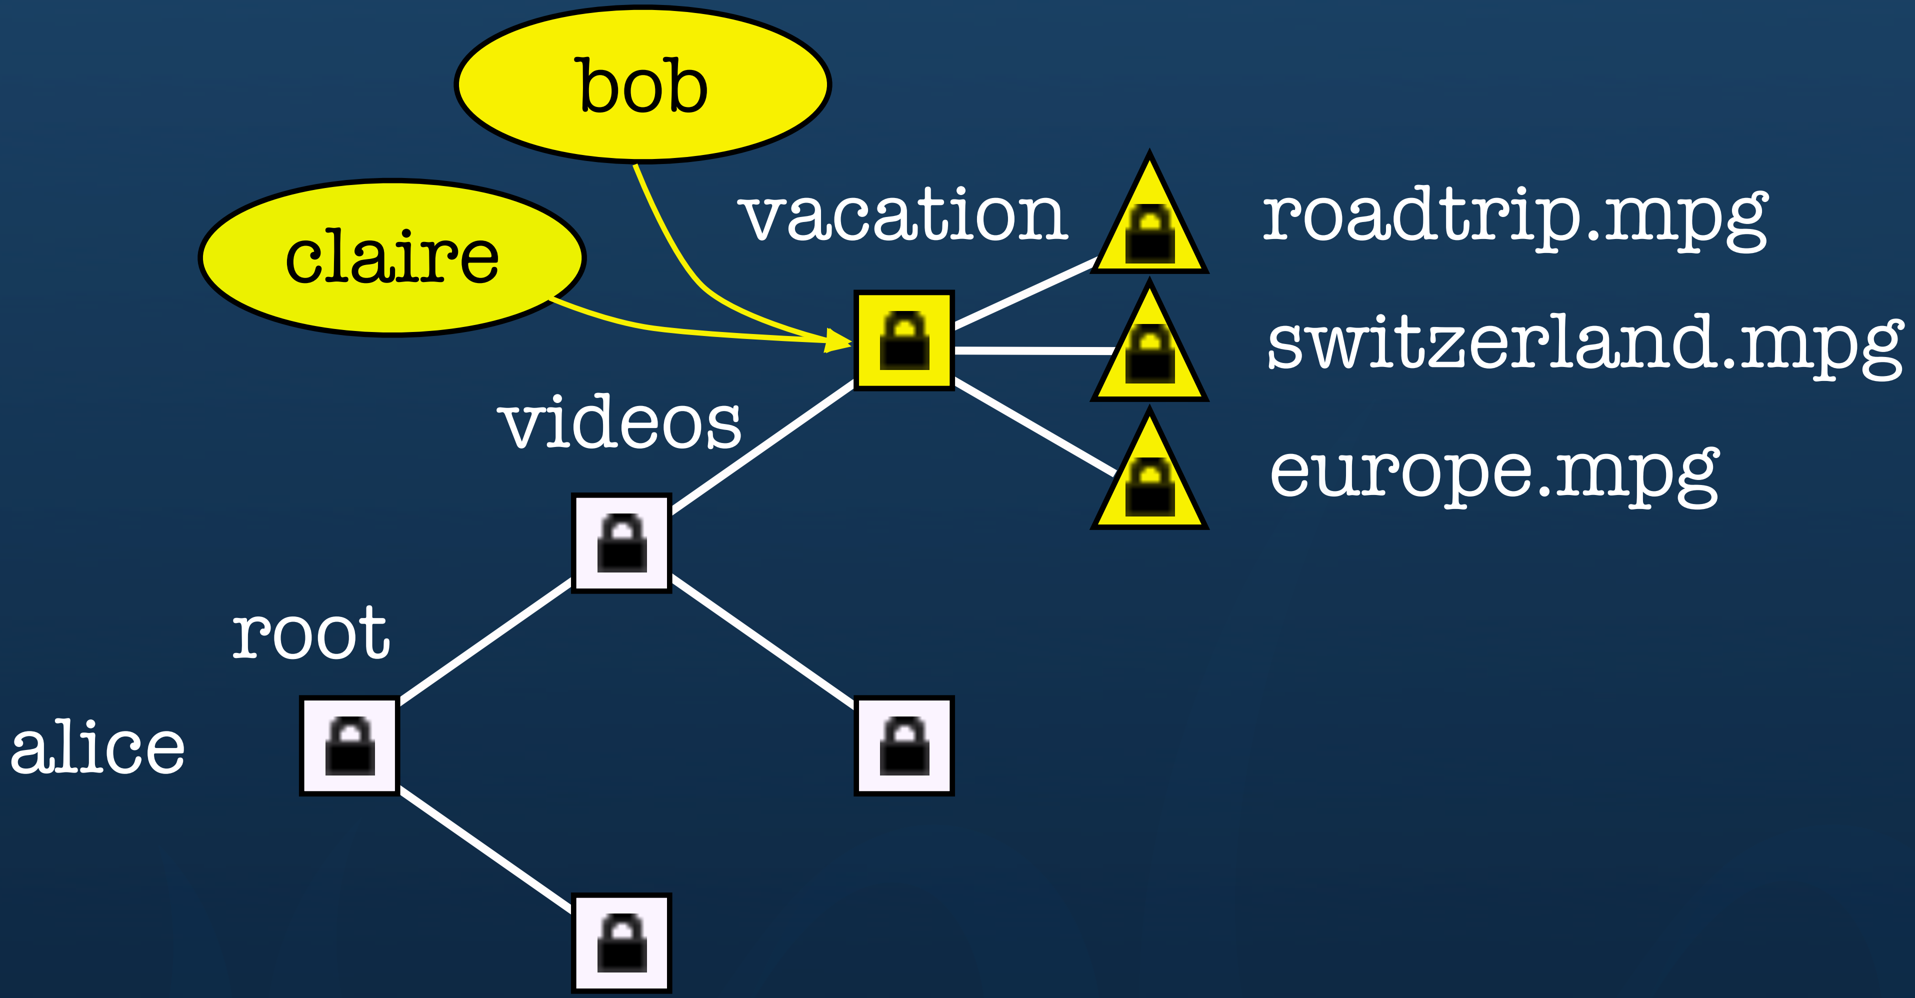

Cryptree, SRDS 2006

granting access to this and all subfolders takes just one operation all subkeys can be derived from that parent key

bob doesn't see that claire has also access and vice versa

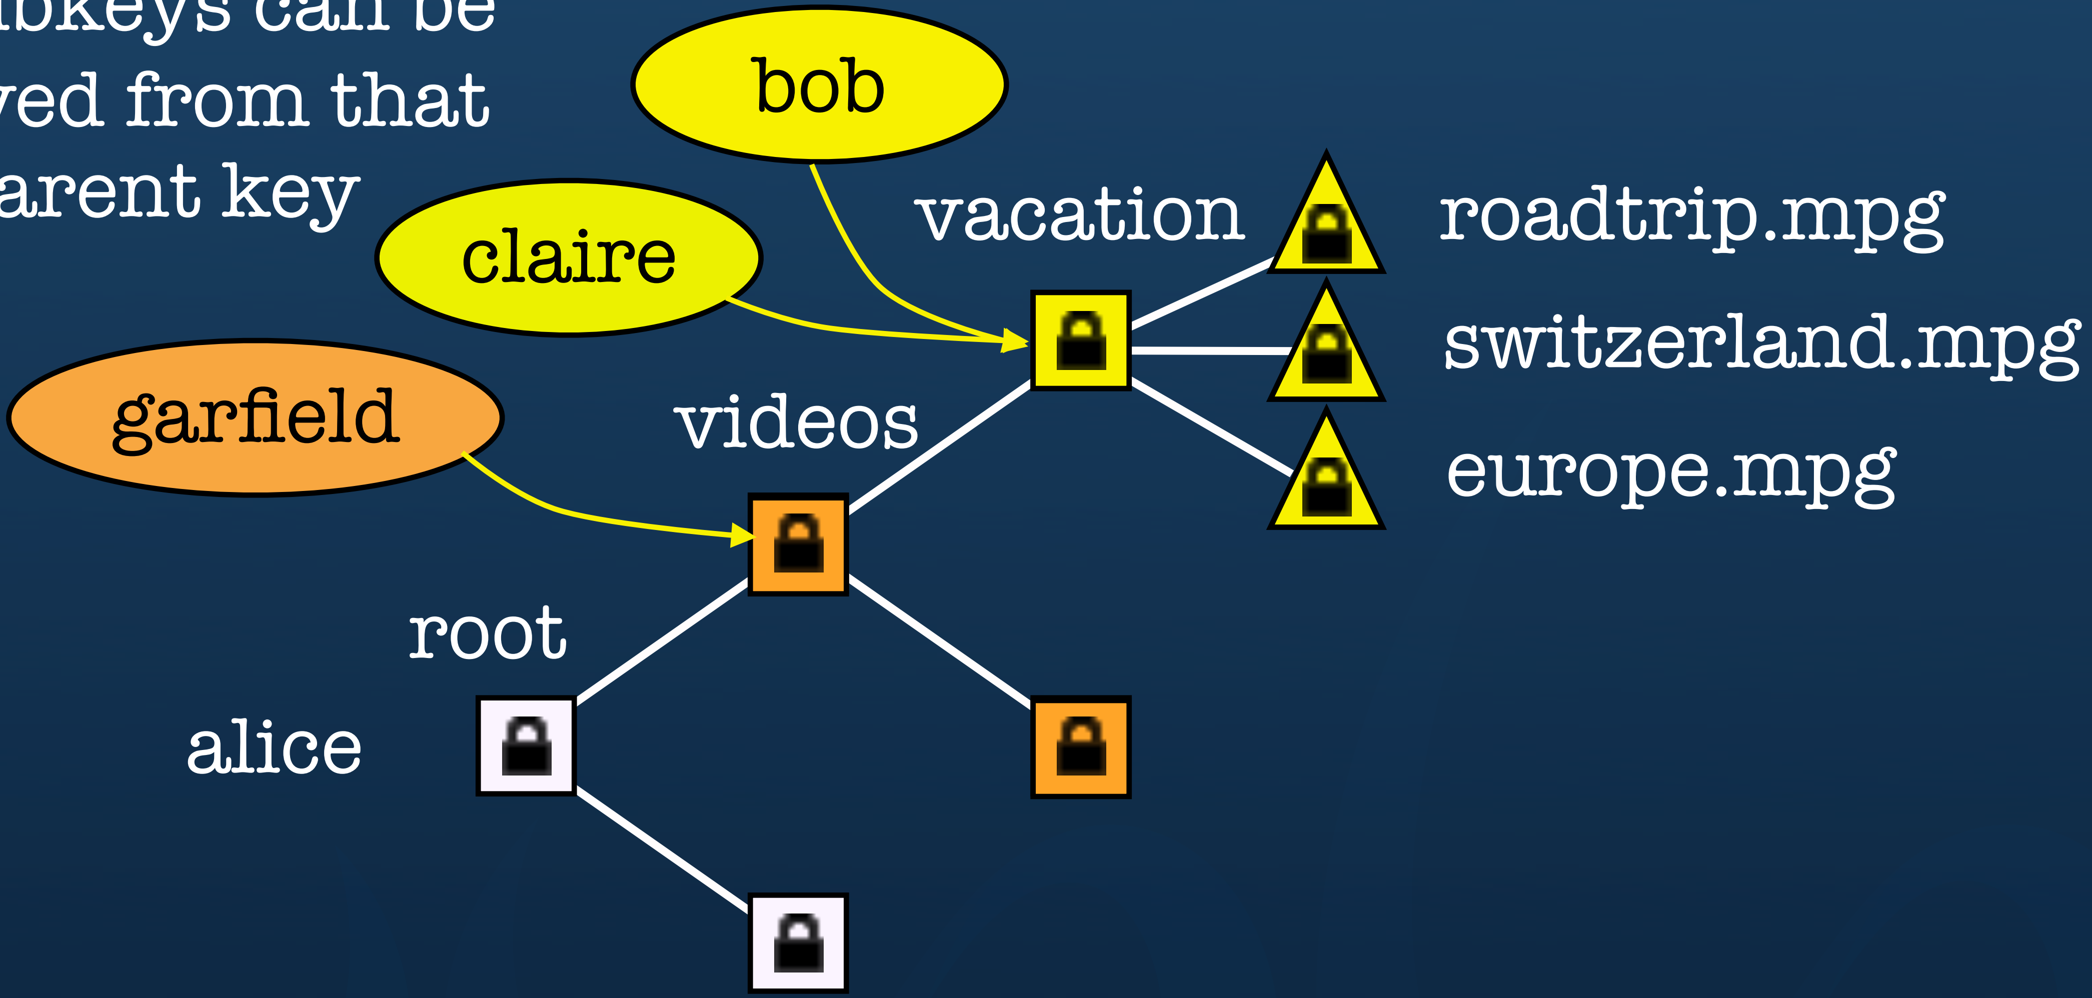

Cryptree, SRDS 2006

demo

thank you!

Invitation for the closed alpha

1. http://download.wua.la
2. Run the installer
3. Enter your invitation code:

CERN