A DNA-Based Archival Storage System

Luis Ceze

University of Washington
Microsoft Research

joint work with Karin Strauss, Doug Carmean, Georg Seelig, James Bornholt, Randolph Lopez, Lee Organick, Yuan Chen, Chris Takahashi, Bichlien Nguyen, Sergey Yekhanin, Siena Dumas Ang.

Life evolved a fantastic storage medium…

DNA

Gene

Protein

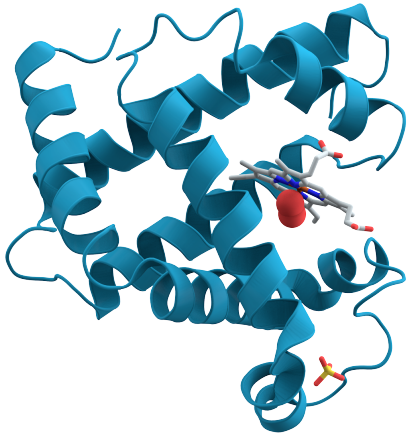

Function

100101010

But why?

DNA molecules for digital data

Extremely dense
Theory: 1 exabyte in 1 mm^3

Extremely durable
Half life > 500 years

And readers never become obsolete!

Comparing Storage Technologies

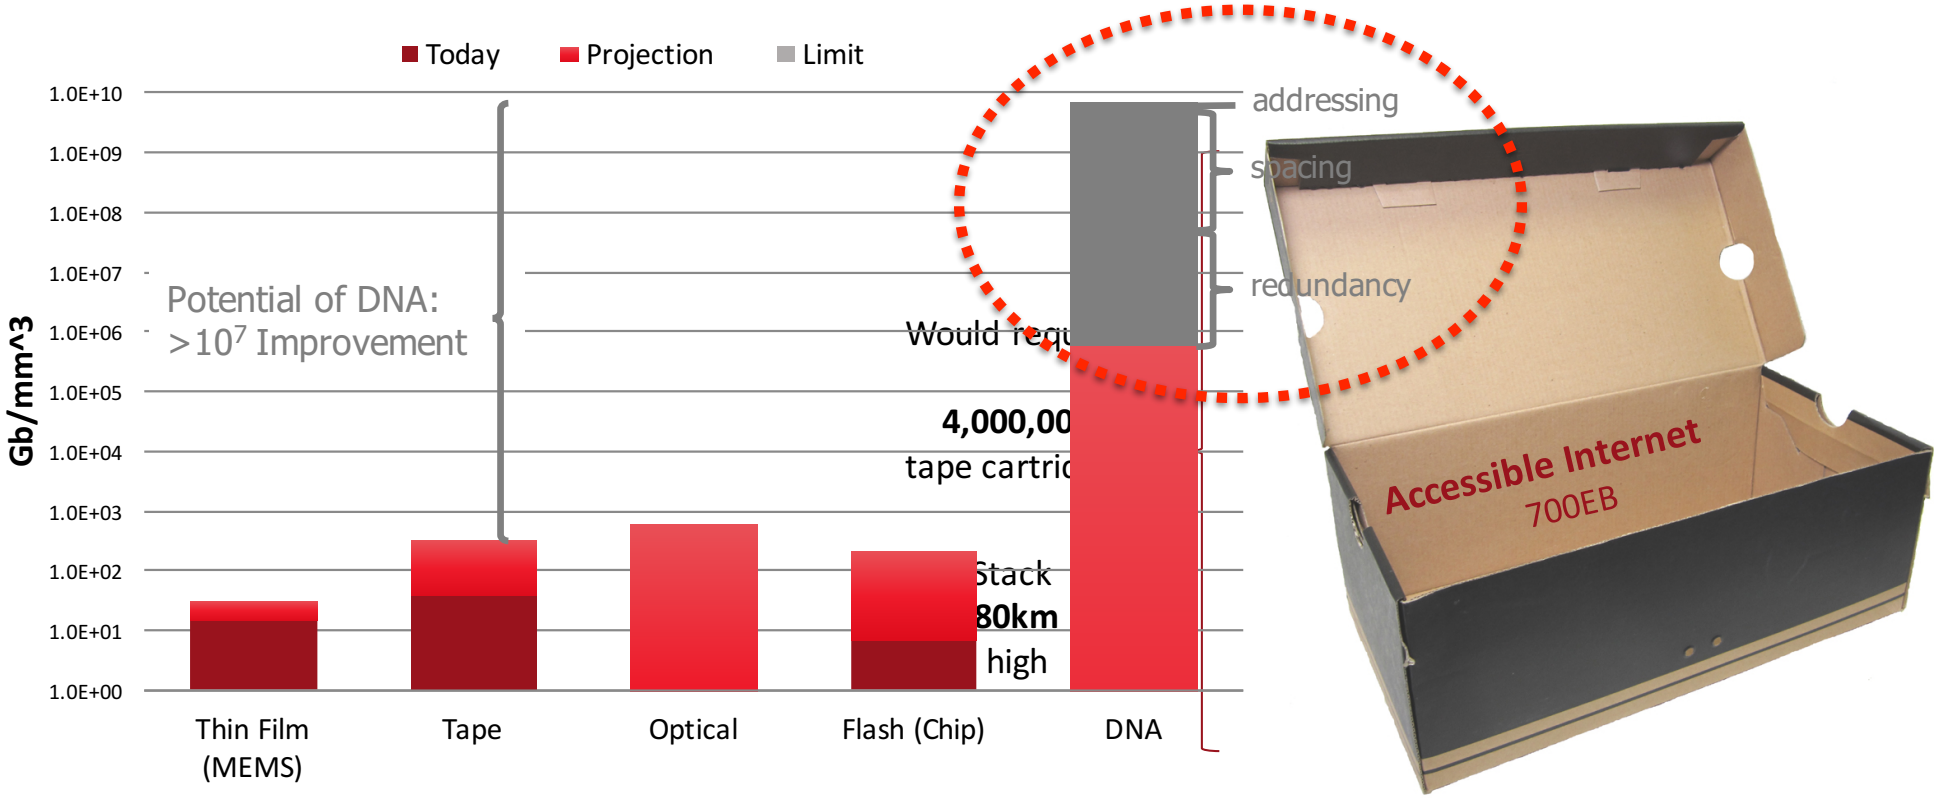

The ultimate hierarchy

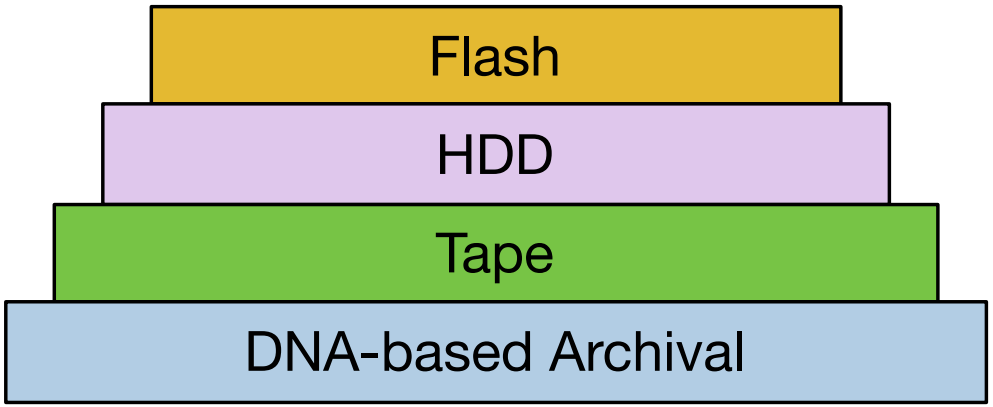

	Access Time	*Durability*
Flash	ms	~5 yrs
HDD	10s ms	~5 yrs
Tape	minutes	~15-30 yrs
DNA-based Archival	hrs	centuries

A DNA-based archival storage system

[ASPLOS'16]

11010101…

11010101…

Redundancy and density

Efficient retrieval

Write

Read

Wet lab experiments

Store

ACATCG…

MISL

Microsoft Research

DNA molecules

Four nucleotides:

- A Adenine
- C Cytosine
- G Guanine
- T Thymine

DNA strand (oligonucleotide) is a linear sequence of these nucleotides

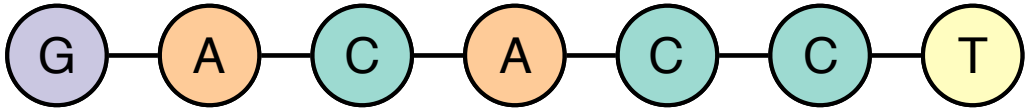

DNA molecules

Four nucleotides:

- A Adenine
- C Cytosine
- G Guanine
- T Thymine

DNA strand (oligonucleotide) is a linear sequence of these nucleotides

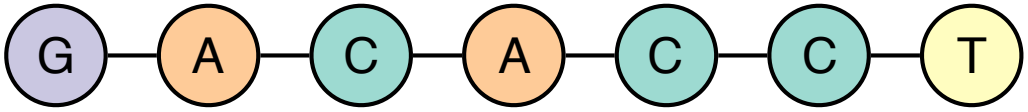

Two strands can bind to each other if they are complementary:

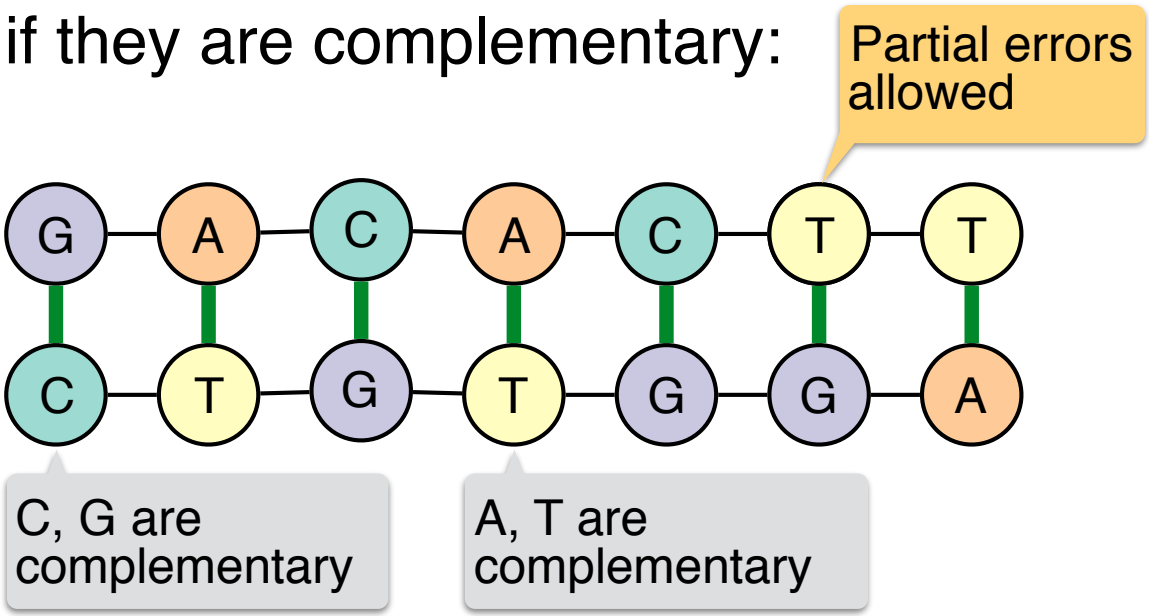

DNA data storage at 30,000 feet

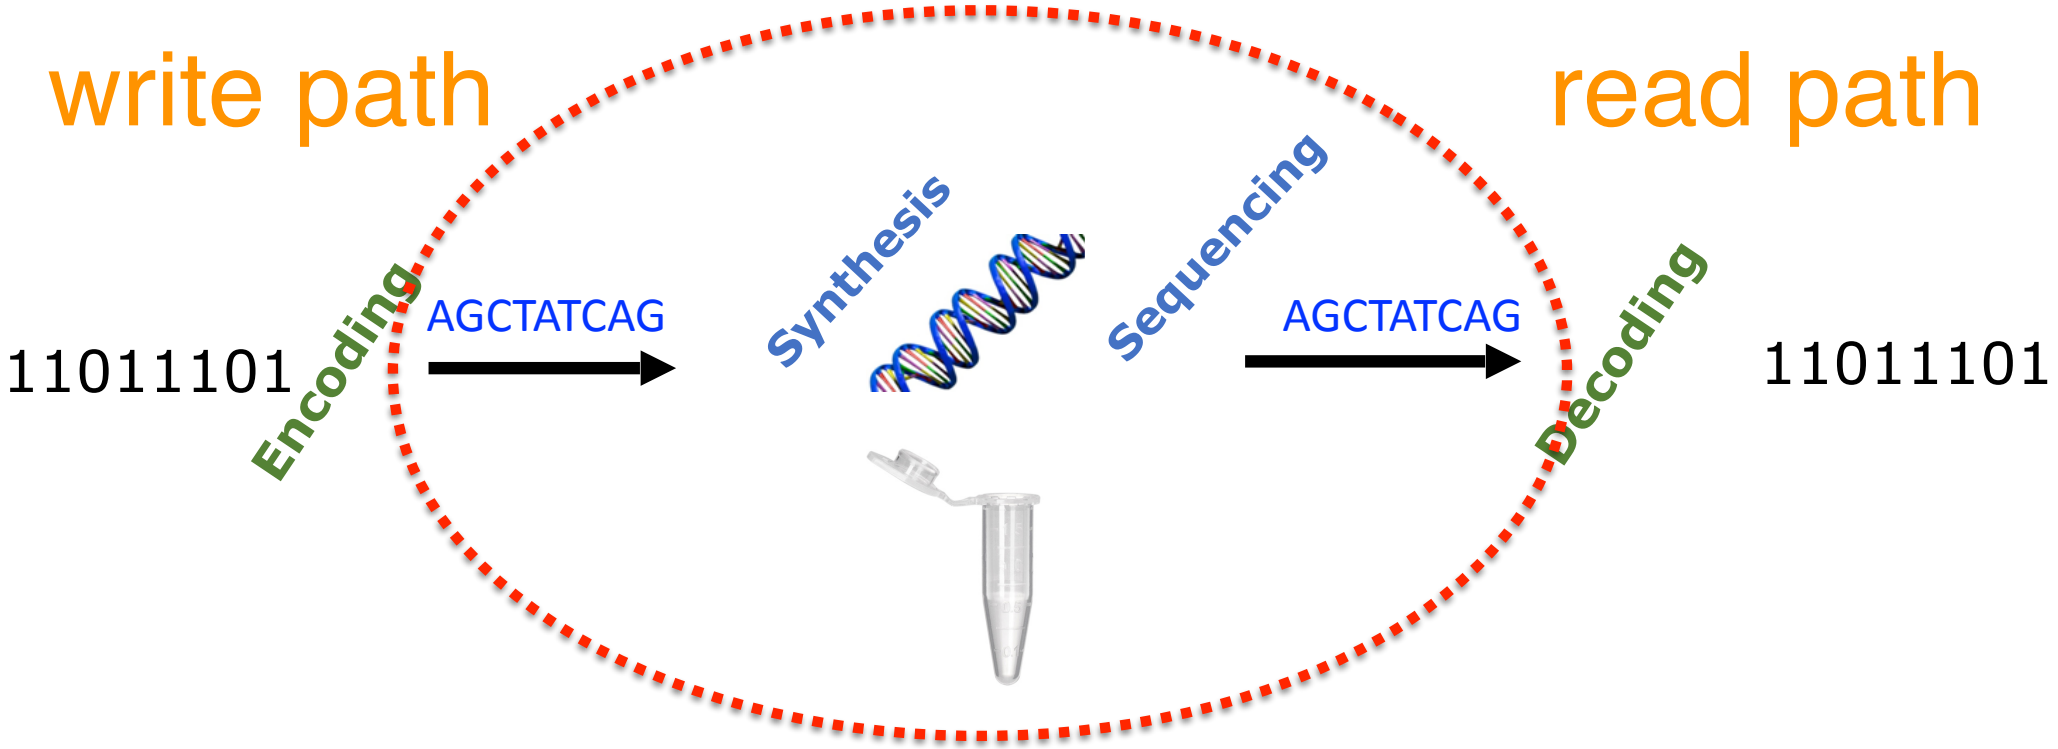

DNA Manipulation: Synthesis

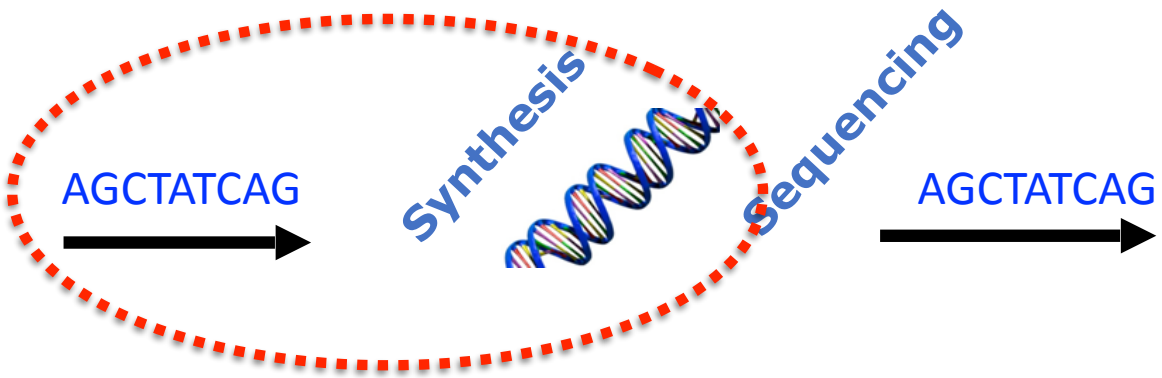

DNA Synthesis: manufacturing DNA strands

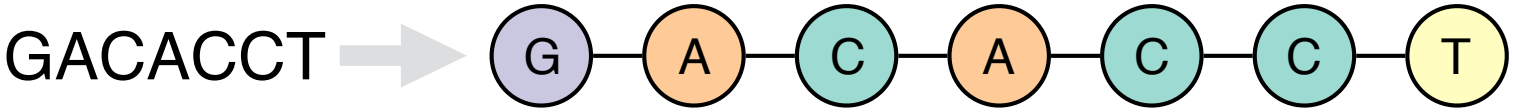

- Synthesis process appends one letter at a time
- Maximum practical sequence length of ~100s of letters
- Produces millions of copies of each sequence
- Can make many different sequences in parallel
- Normally used for genomics and genetic engineering

DNA Manipulation: Sequencing

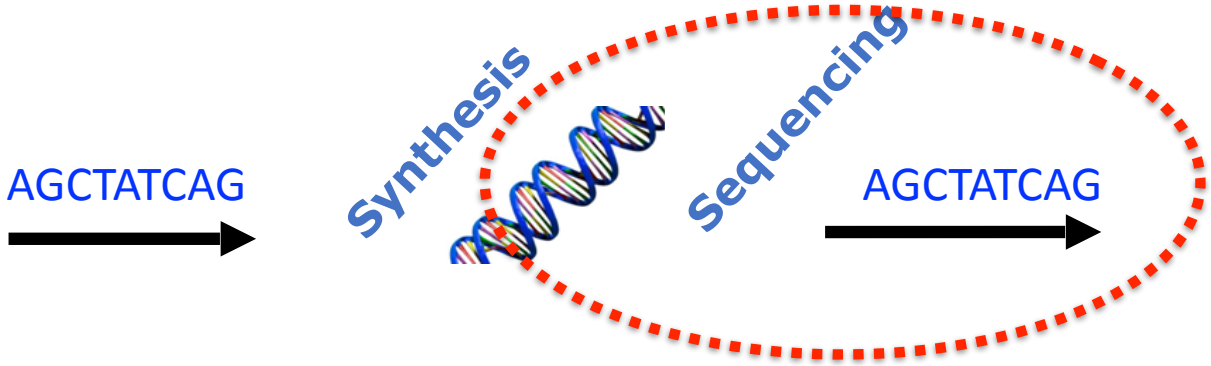

DNA Sequencing: reading DNA strands

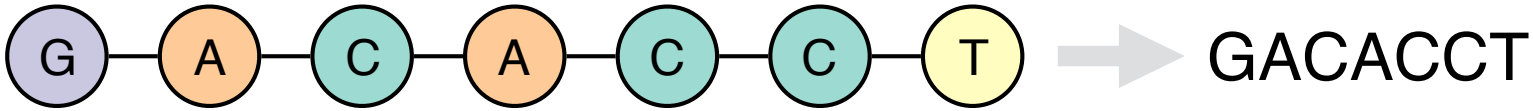

- Produces *many* reads of a strand
- Normally used for research and diagnostics
- Currently much higher throughput than synthesis

DNA data storage at 30,000 feet

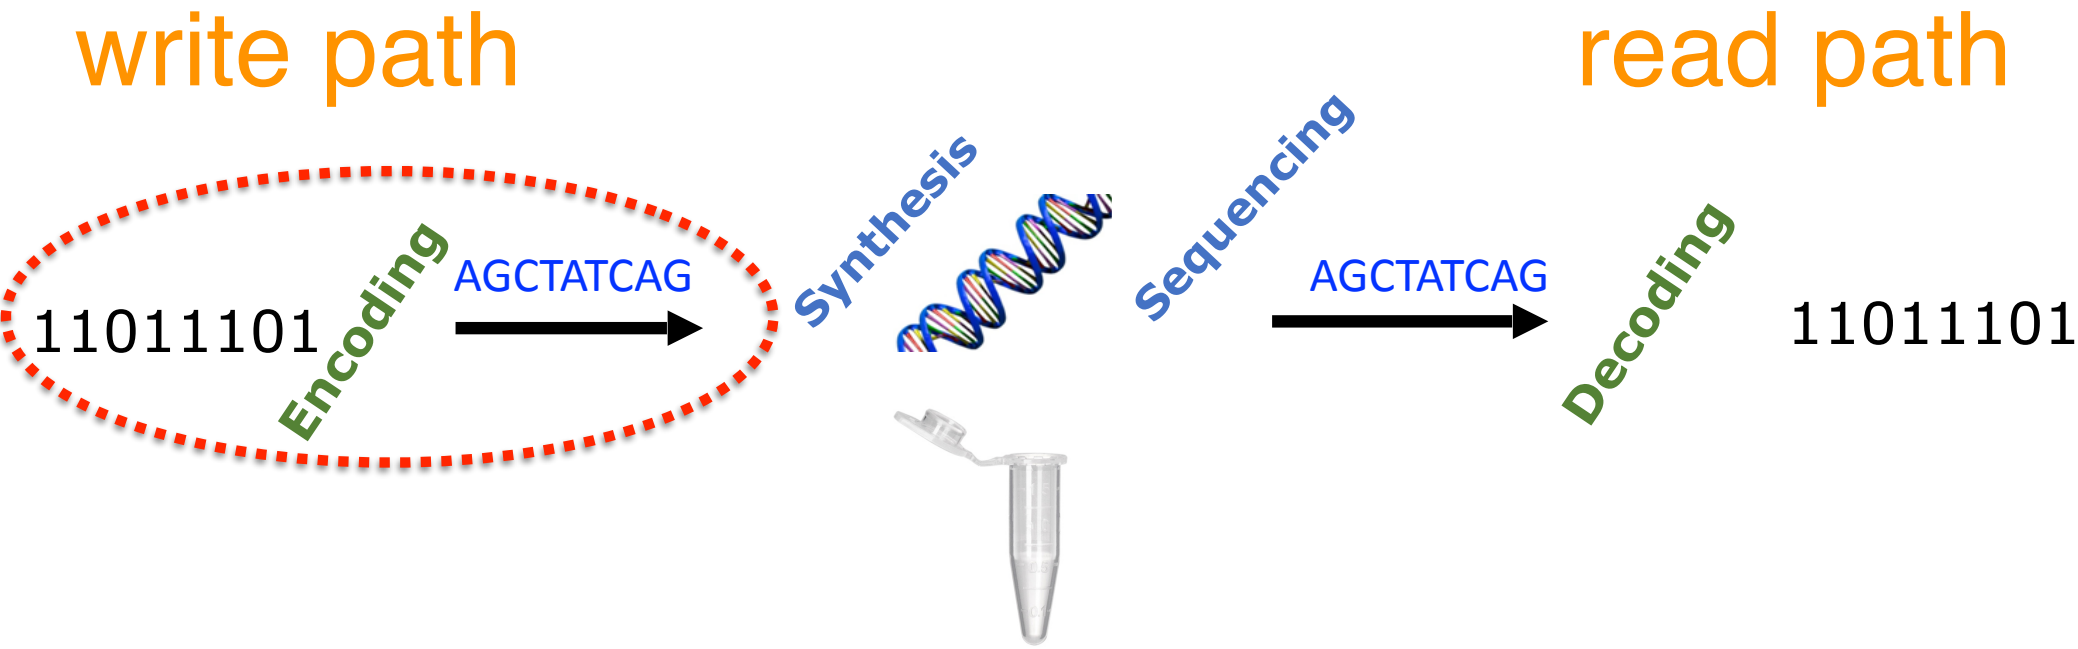

Writing digital data to DNA

The easy way: convert base 2 to base 4

10100011 10010001 11100111 11000101 10010100 10111101

2 2 0 3 2 1 0 1 3 2 1 3 3 0 1 1 2 1 1 0 2 3 3 1

G G A T G C A C T G C T T A C C G C C A G T T C

But this approach isn't feasible for more than a few bytes

P[Attach] = 99%

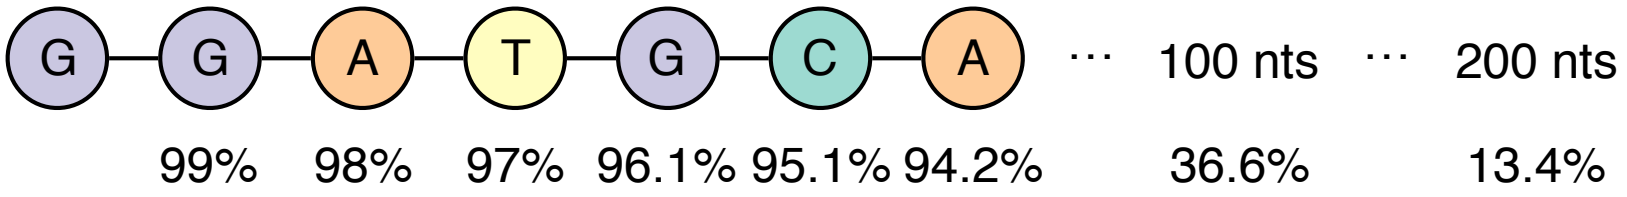

Chunking data

Break binary data into chunks stored in separate strands

10100011 10010001 11100111 11000101 10010100 10111101

2 2 0 3 2 1 0 1 3 2 1 3 3 0 1 1 2 1 1 0 2 3 3 1

G G A T G C A C T G C T T A C C G C C A G T T C

Chunking data

Break binary data into chunks stored in separate strands

10100011 10010001 11100111 11000101 10010100 10111101

2 2 0 3 2 1 0 1 3 2 1 3 3 0 1 1 2 1 1 0 2 3 3 1

A T G T T G G A T G C A C A A A A C A T C C

A T G T T T G C T T A C C A A A C C A T C C

A T G T T G C C A G T T C A A A G C A T C C

File identifiers ("primers")

Addresses within the file

ATGTT GGATGCAC AAAA CATCC

ATGTT TGCTTACC AAAC CATCC

ATGTT GCCAGTTC AAAG CATCC

File identifiers ("primers")

Addresses within the file

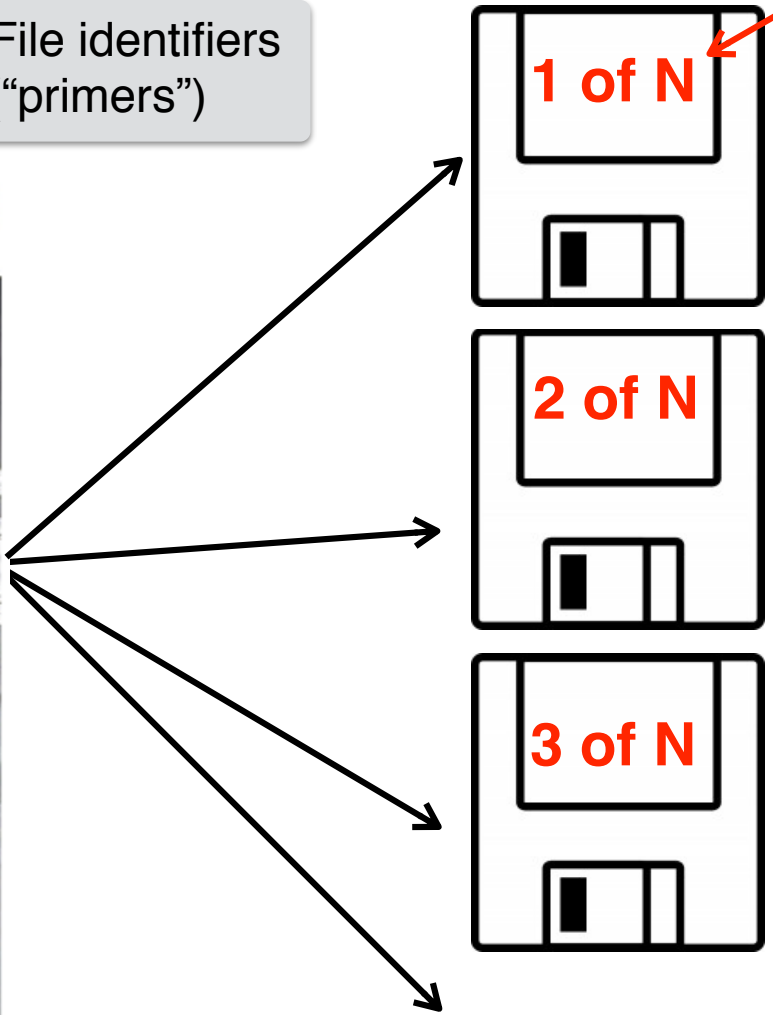

~ 10 bytes per DNA strand.

...

Other encoding considerations

Homopolymers are bad…

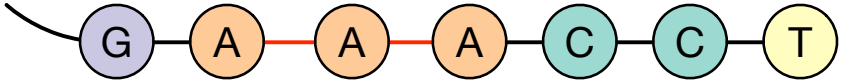

Secondary structures are problematic …

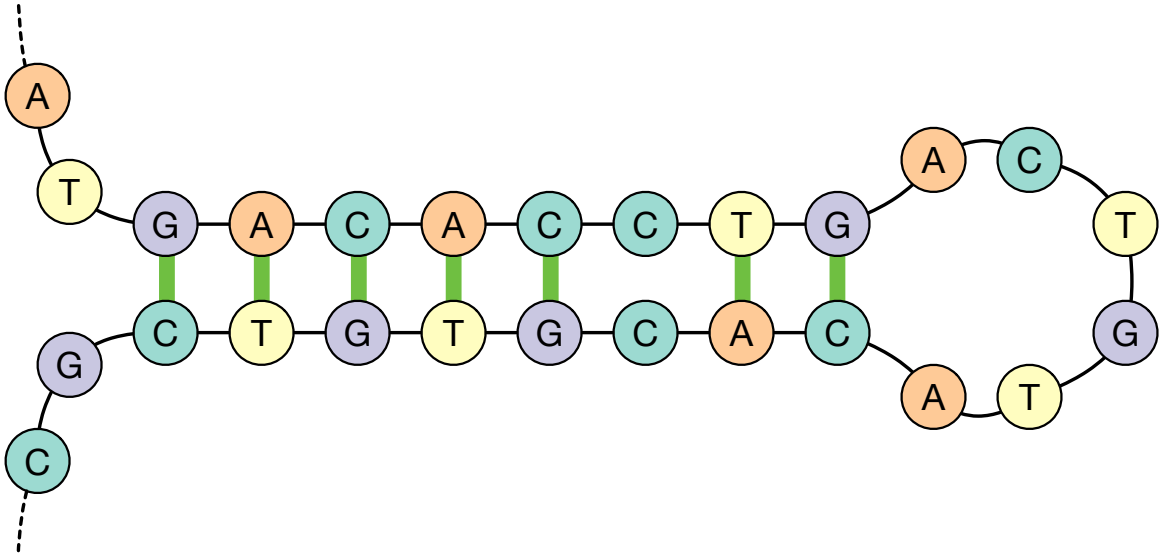

DNA data storage at 30,000 feet

write path

read path

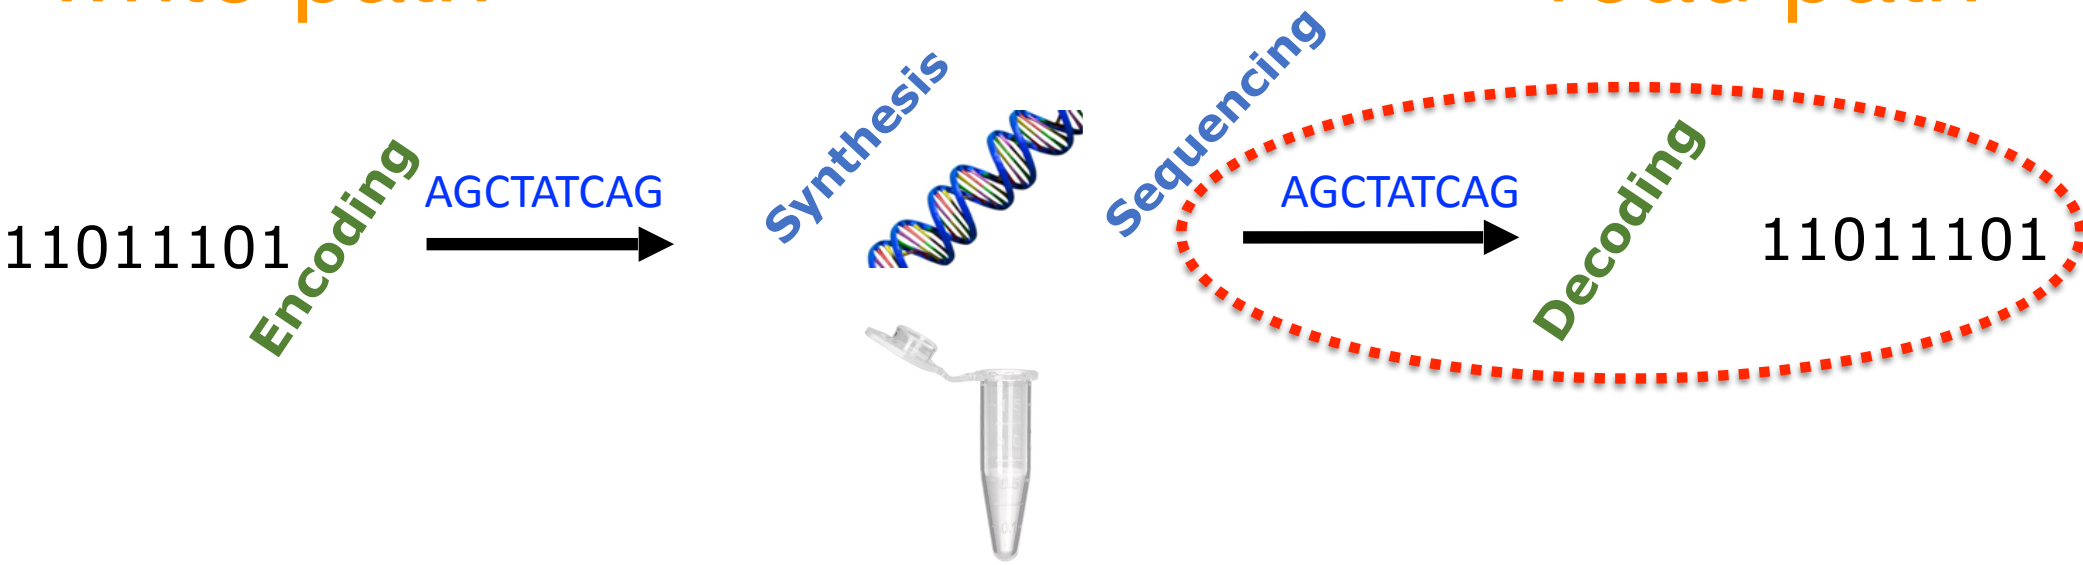

Random access?

Random access!

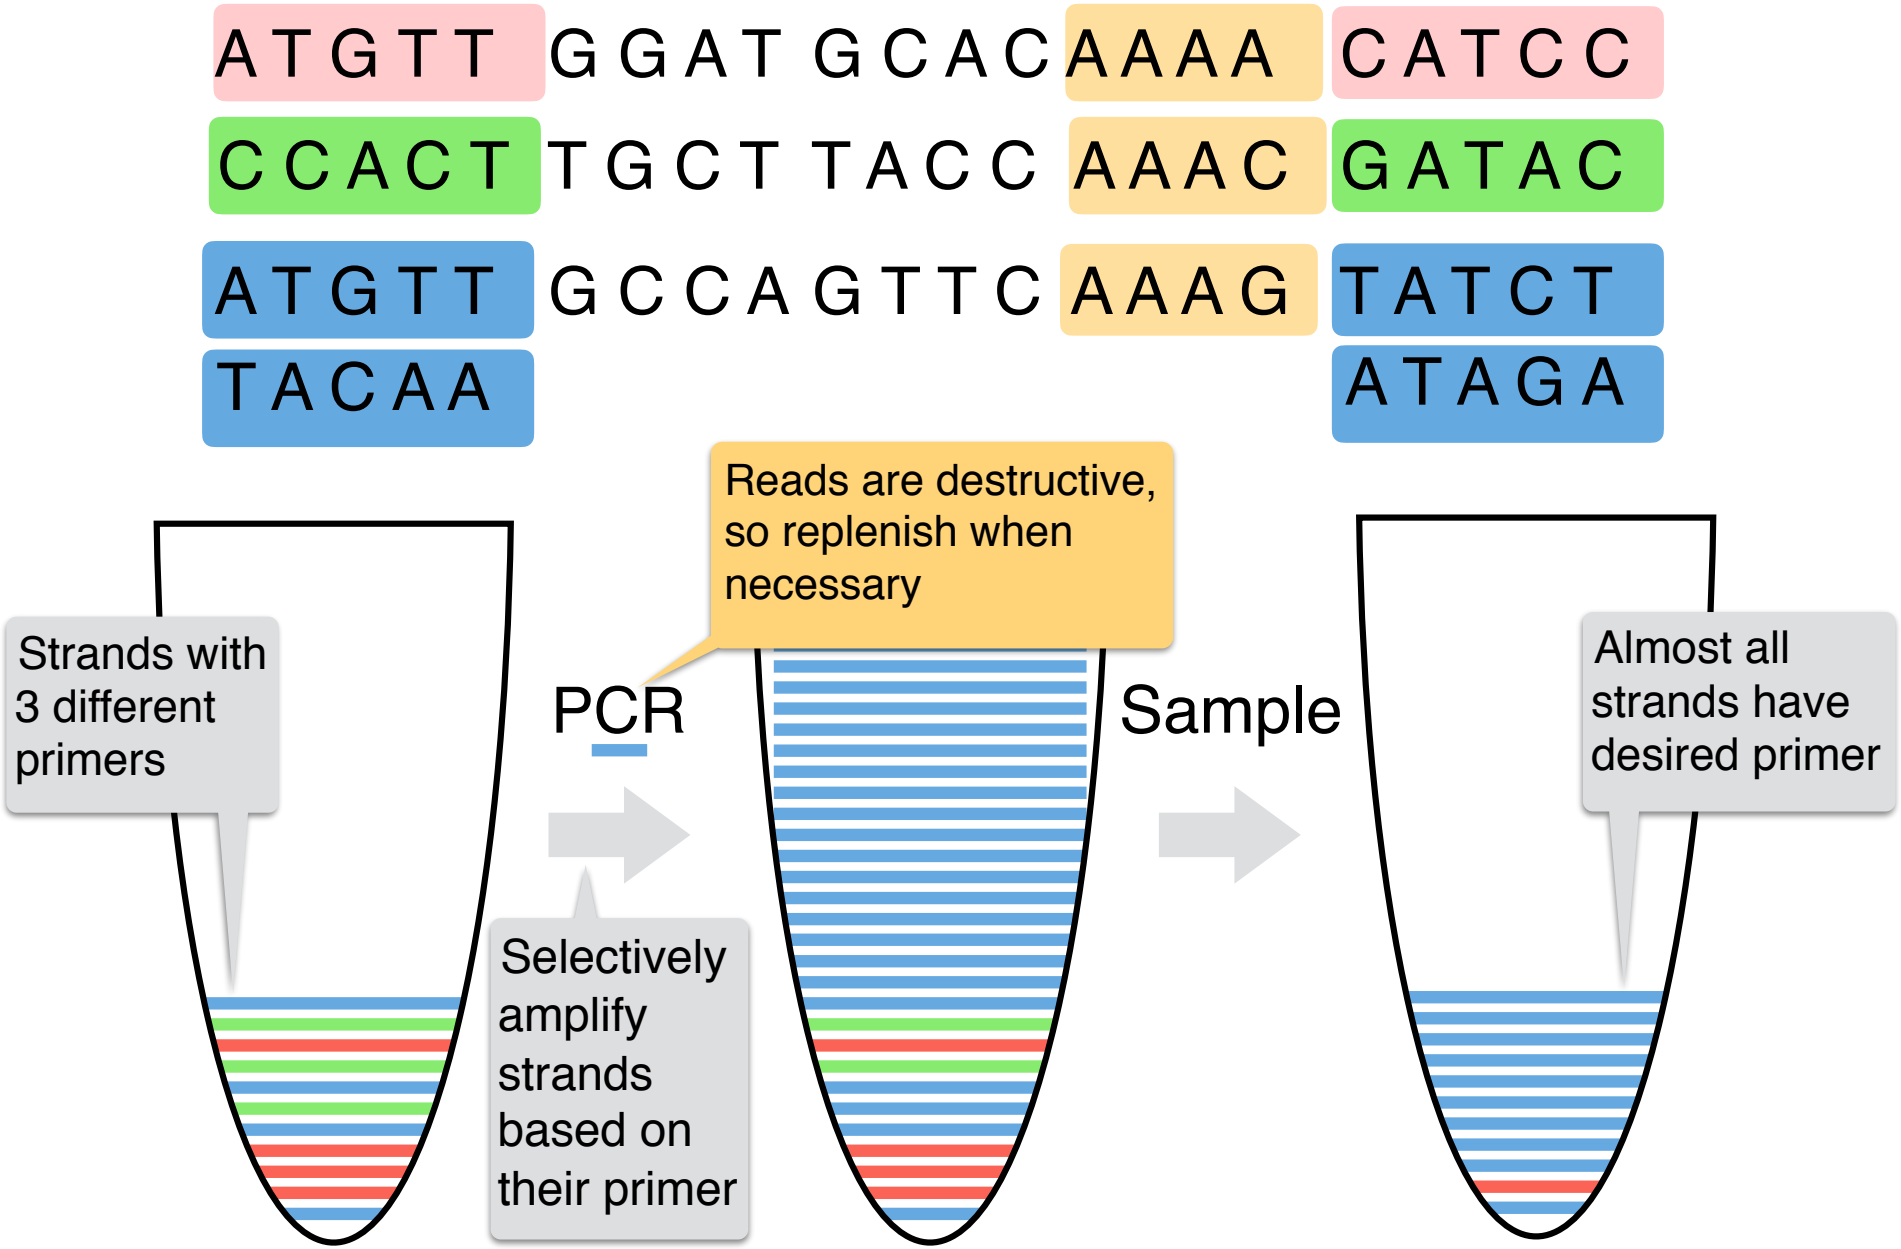

Going back to digital: Decoding

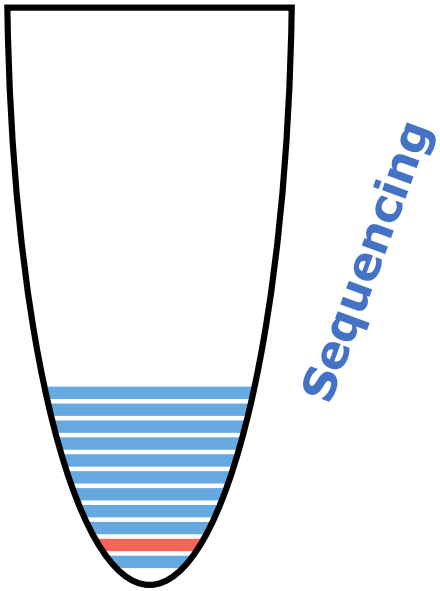

ATGTT GGAT GCAC AAAA CATCC
ATGTT TGCT TACC AAAC CATGC
ATGTT GCCA GTTC ACAGCATCC
ATGTT GGAT GCAC AAGACATCC
ATGTT TGCT TACC CAACCATCC
ATGTT GCCA GTTC AAAGCATCC
ATGTT GGAT GCAC AAGACATCC
ATGTT TGCT TACC CAACCATCC
ATGTT GCCA GTTC AAAGCATCC
.....

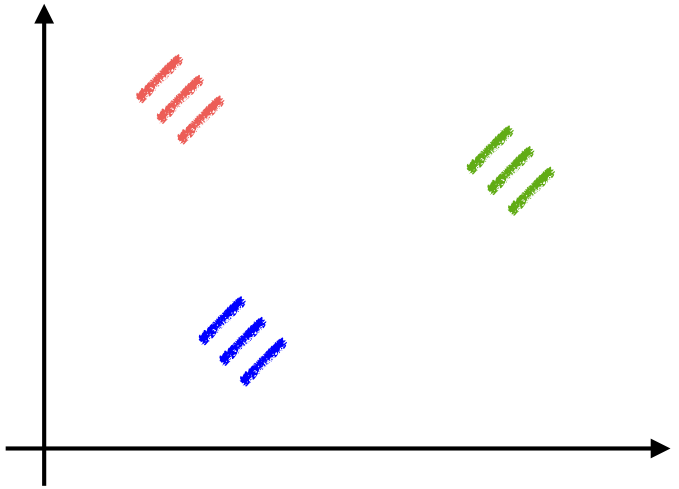

Error correction

…1101101011101…

Error correction

Both synthesis and sequencing are error prone:

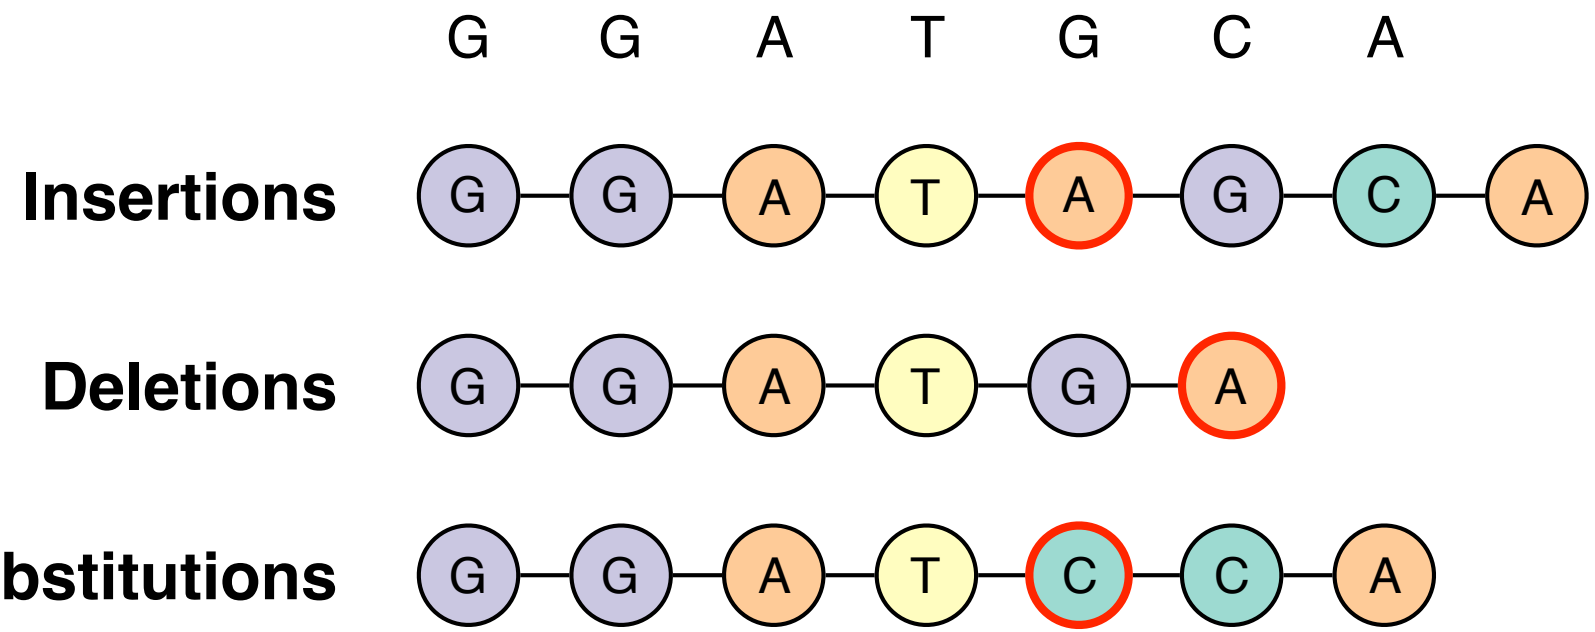

Aggregate error rates ~1% per position

Logical redundancy

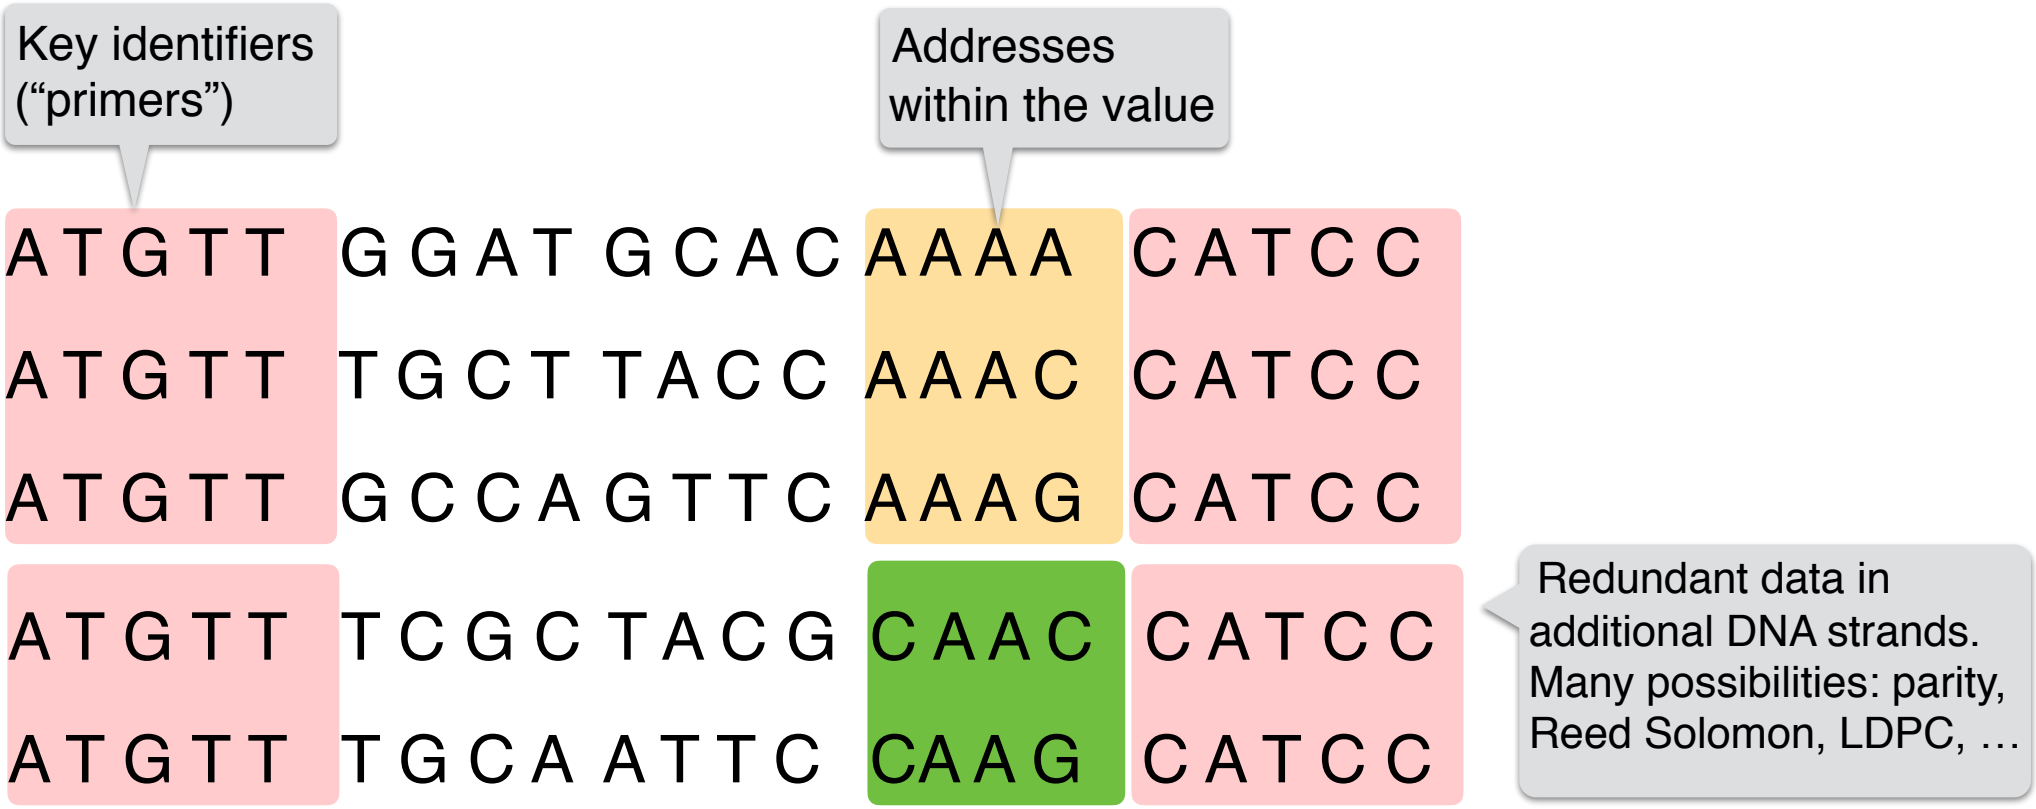

Putting it all together

SAS / SATA

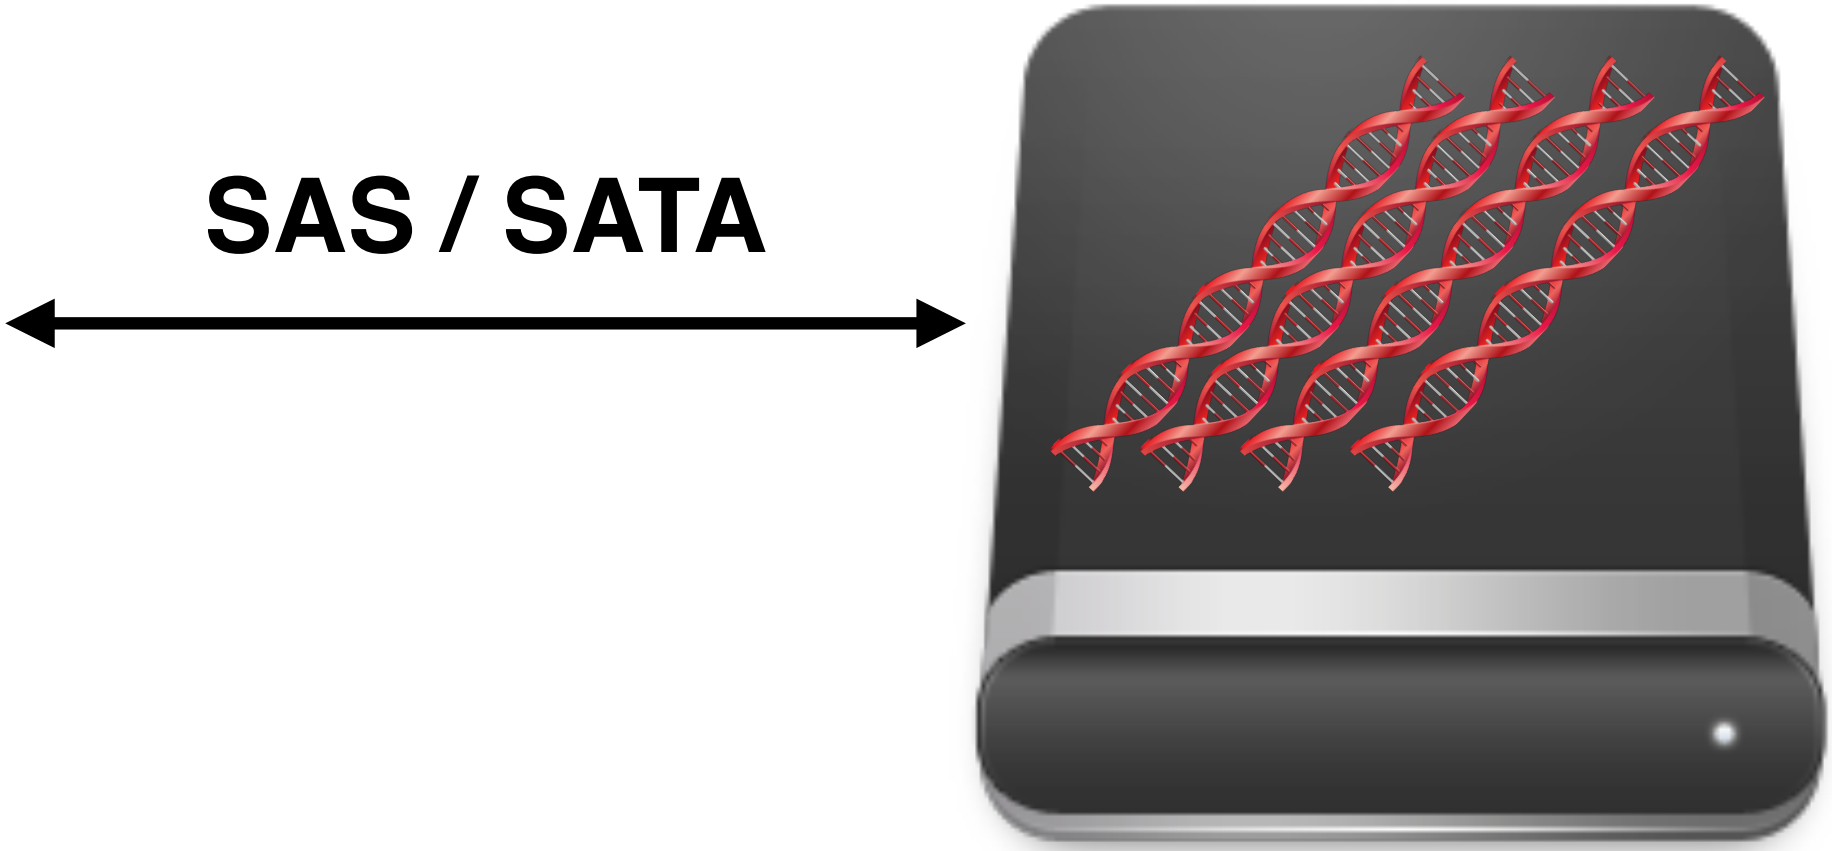

Putting it all together

Data address/key specifies physical location and primer for random access.

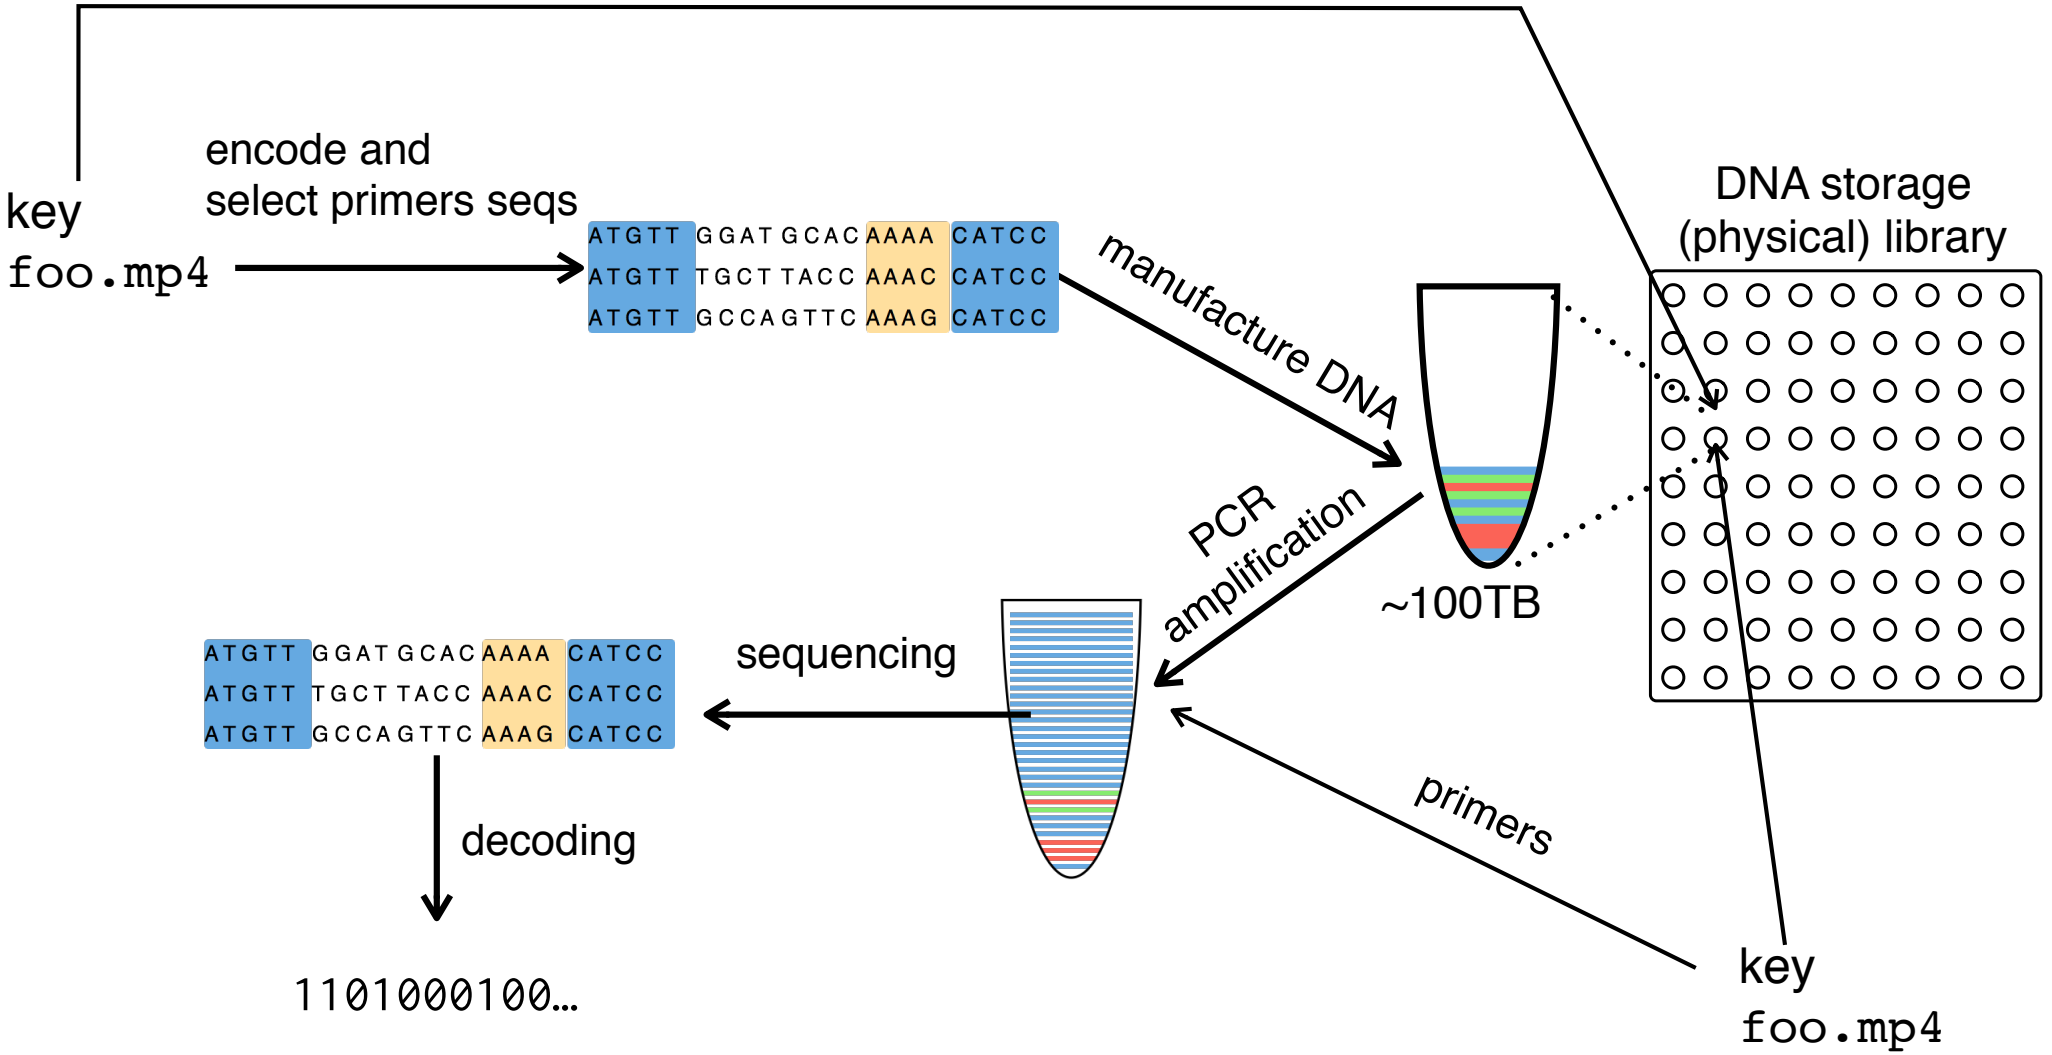

Wet lab results

Experiments

catcatgg

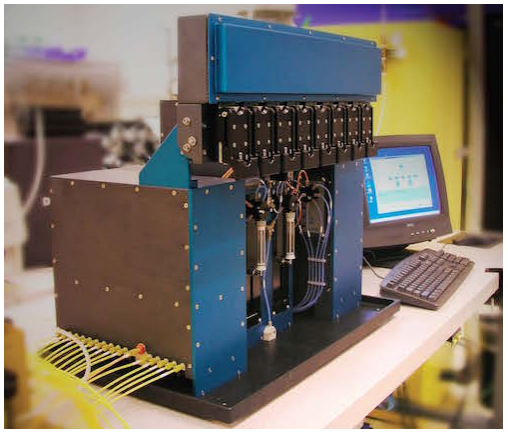

Throughput
MBs/week

catcatg**c**

FedEx

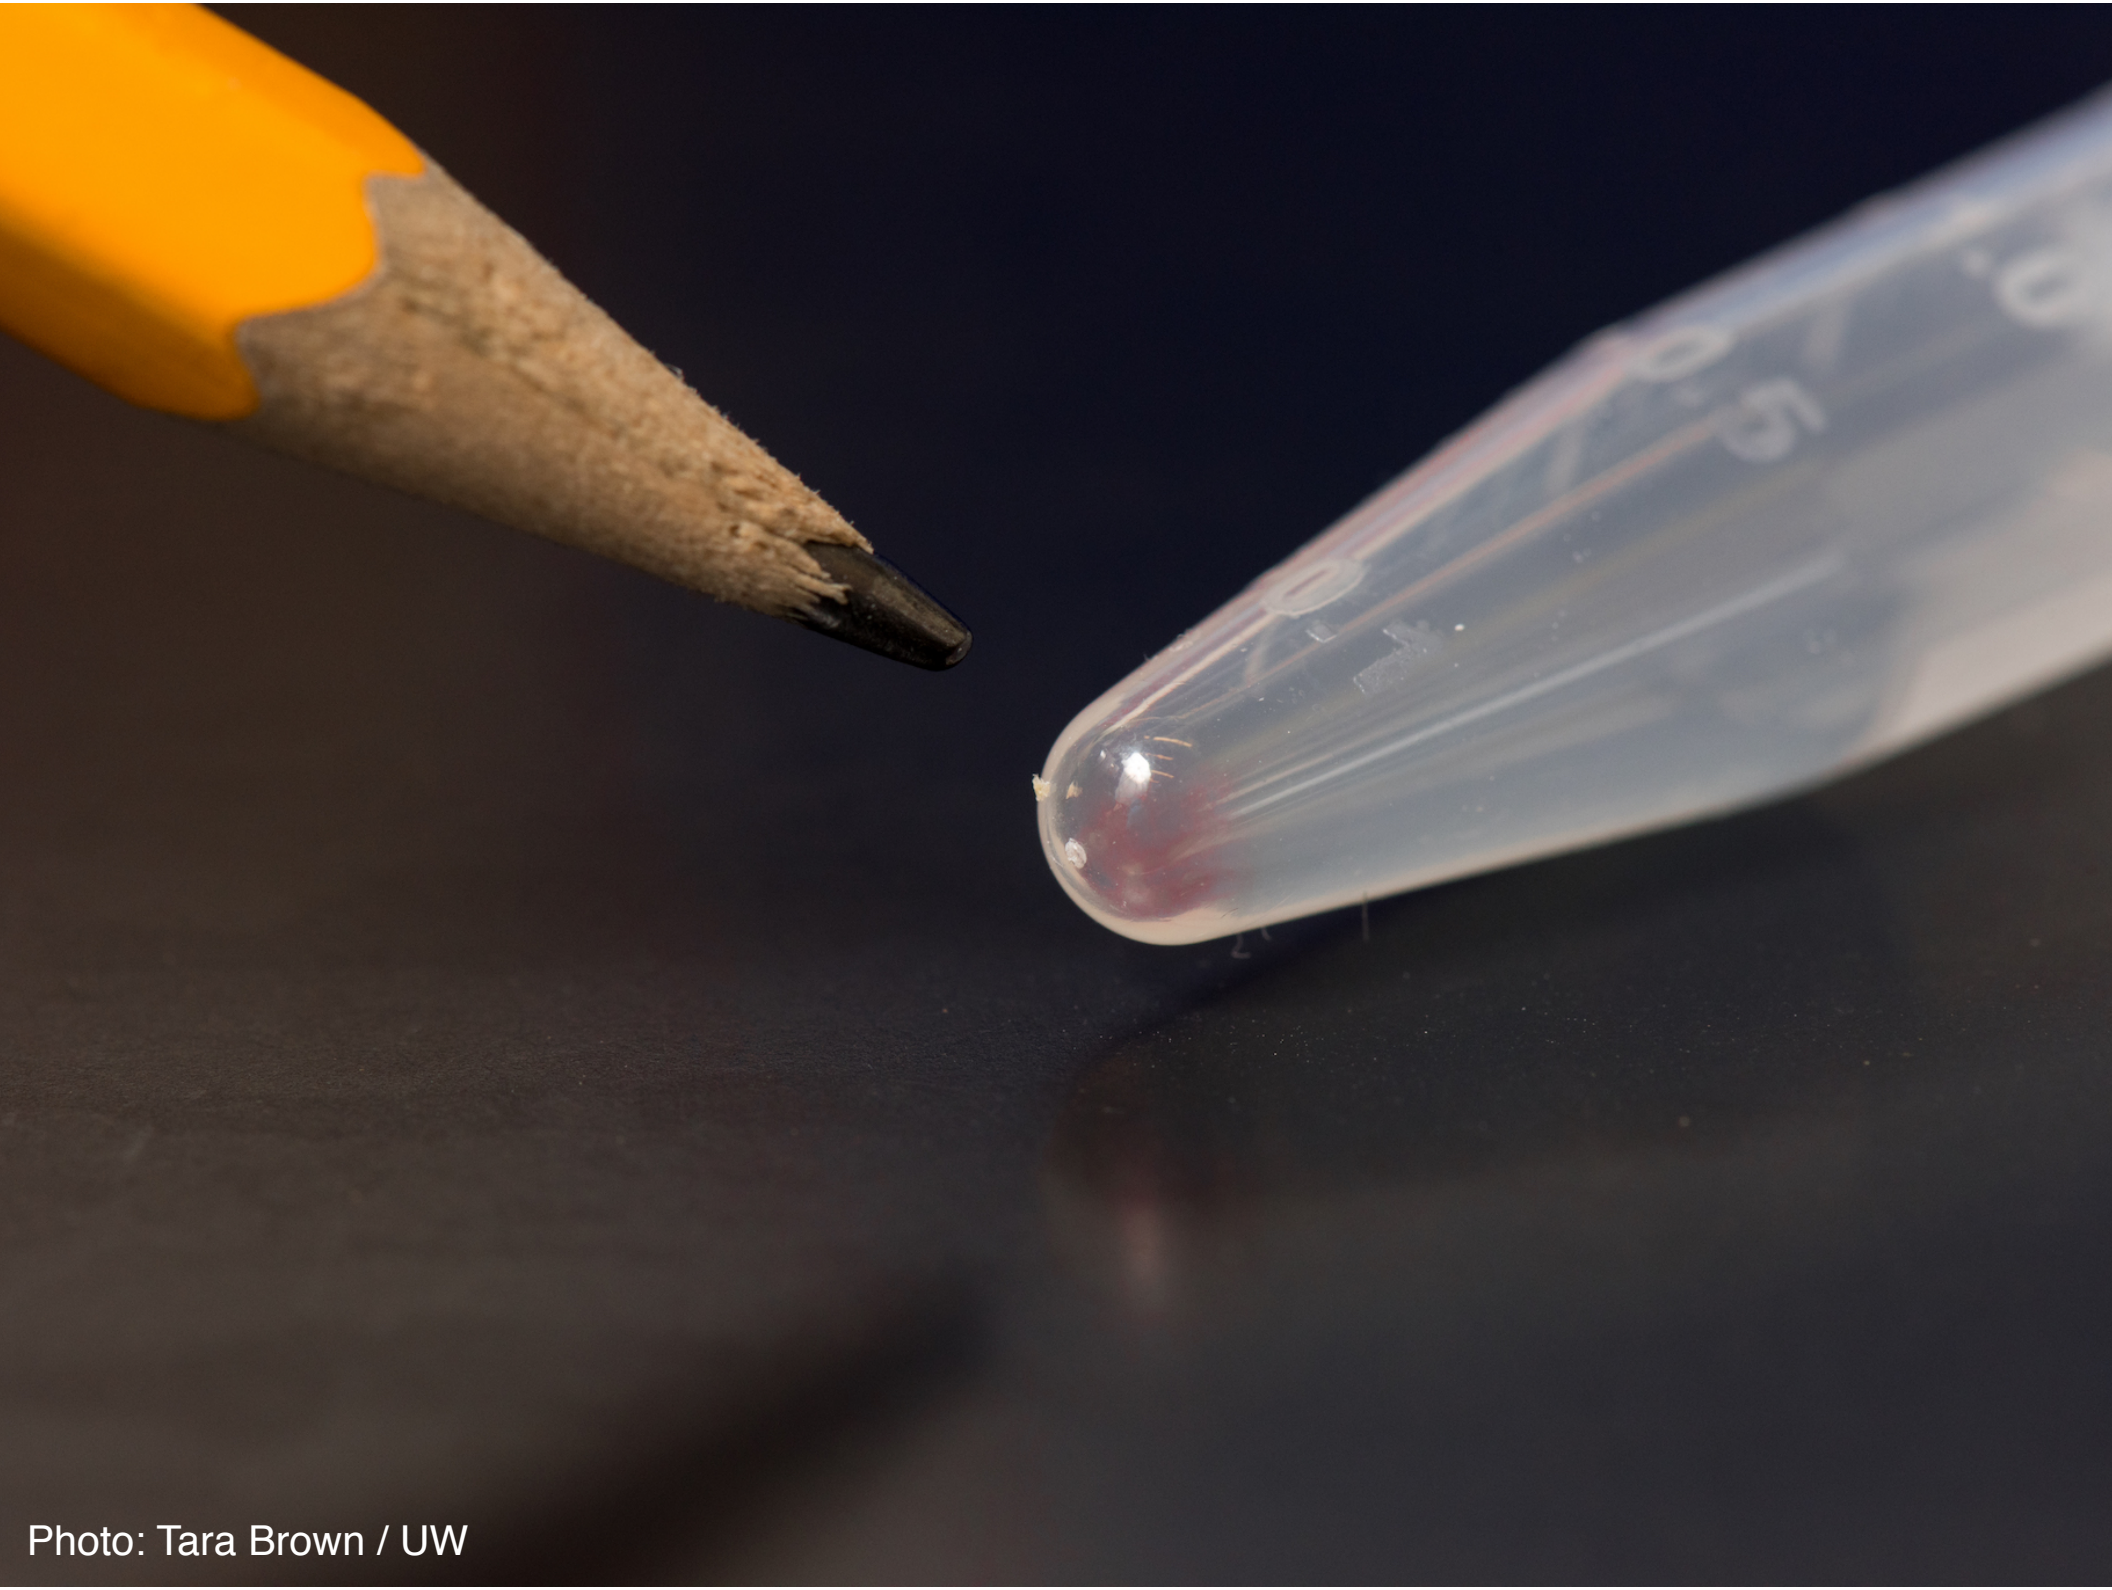

Photo: Tara Brown / UW

Decoding

Encoded and synthesized 3 files (151 kB):

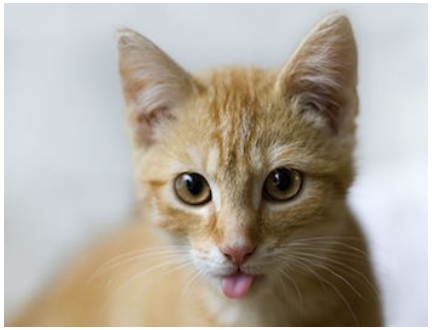

Selected and PCRed one file for random access (42 kB):

Sequenced and decoded the resulting amplified pool:

Recovered *every bit* despite errors in synthesis and sequencing

The importance of redundancy

Key identifiers ("primers")

Addresses within the value

A T G T T G G A T G C A C A A A A C A T C C

A T G T T

A T G T T

A T G T T

A T G T T

If we ignore redundancy data, we cannot recover the file.

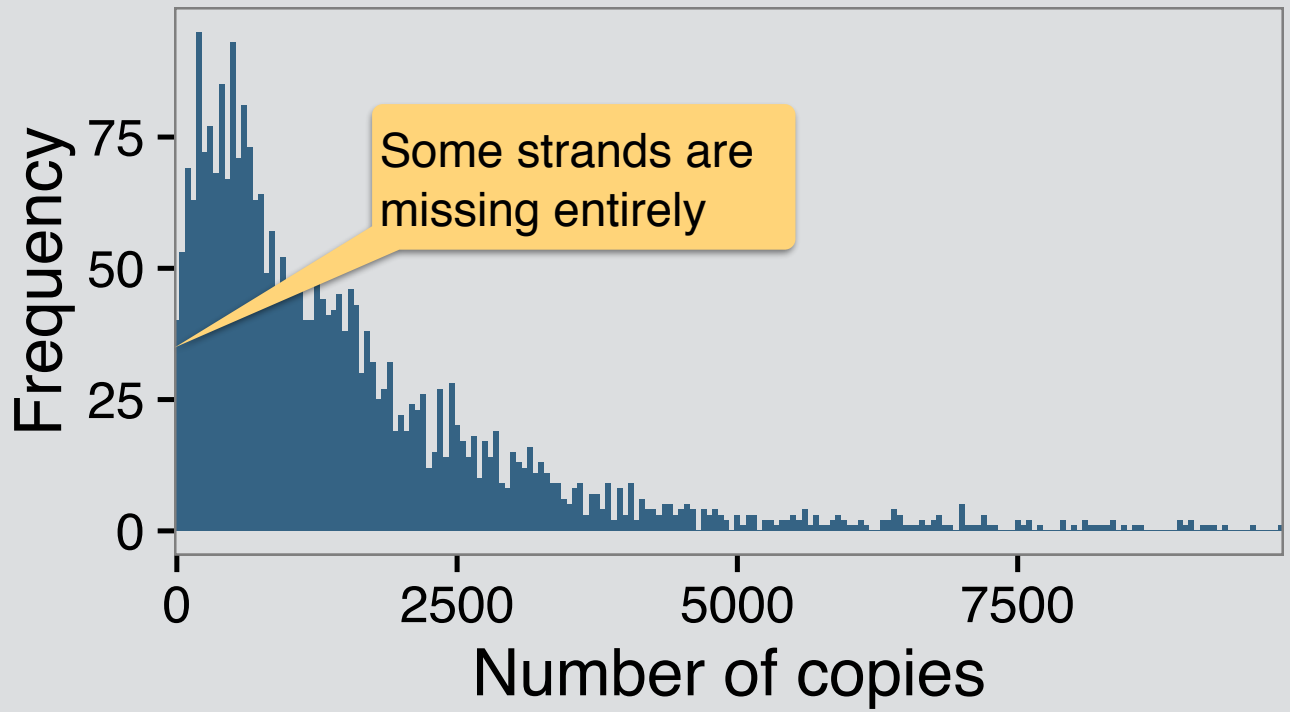

ıdant data in
nal DNA strands.
ossibilities: parity,
Solomon, LDPC, …

Error Analysis: Synthesis and Sequencing

Both synthesis and sequencing are error prone, but how much?

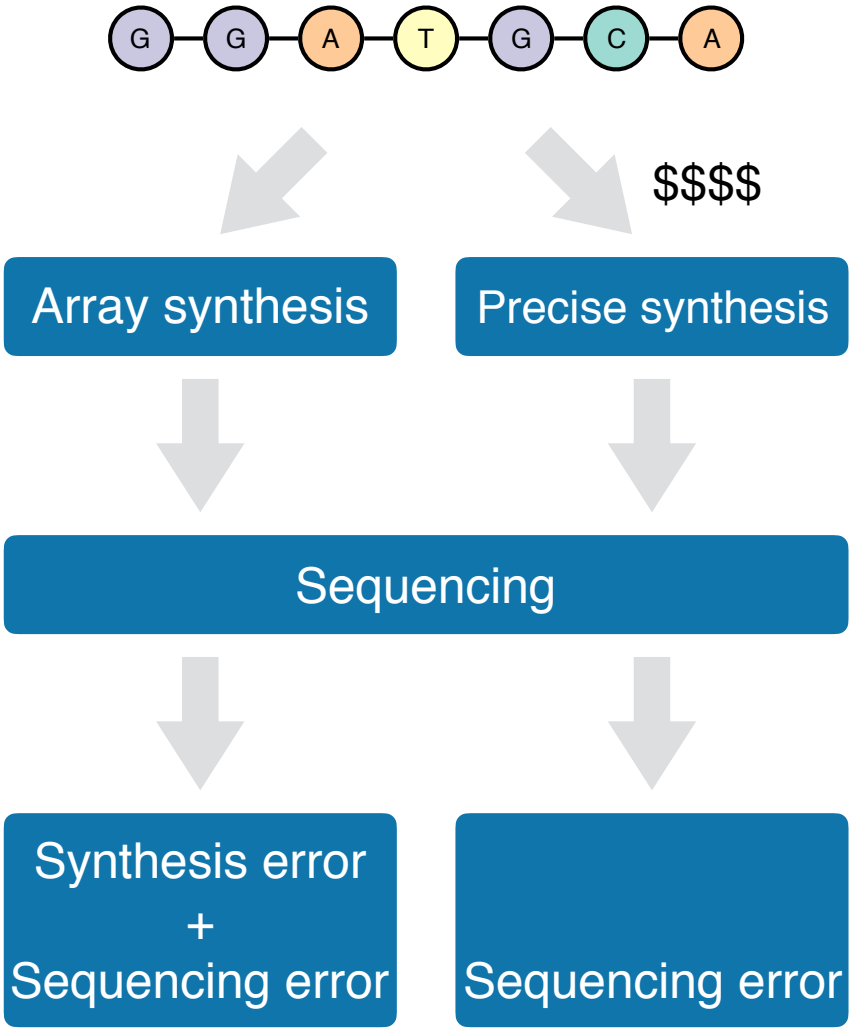

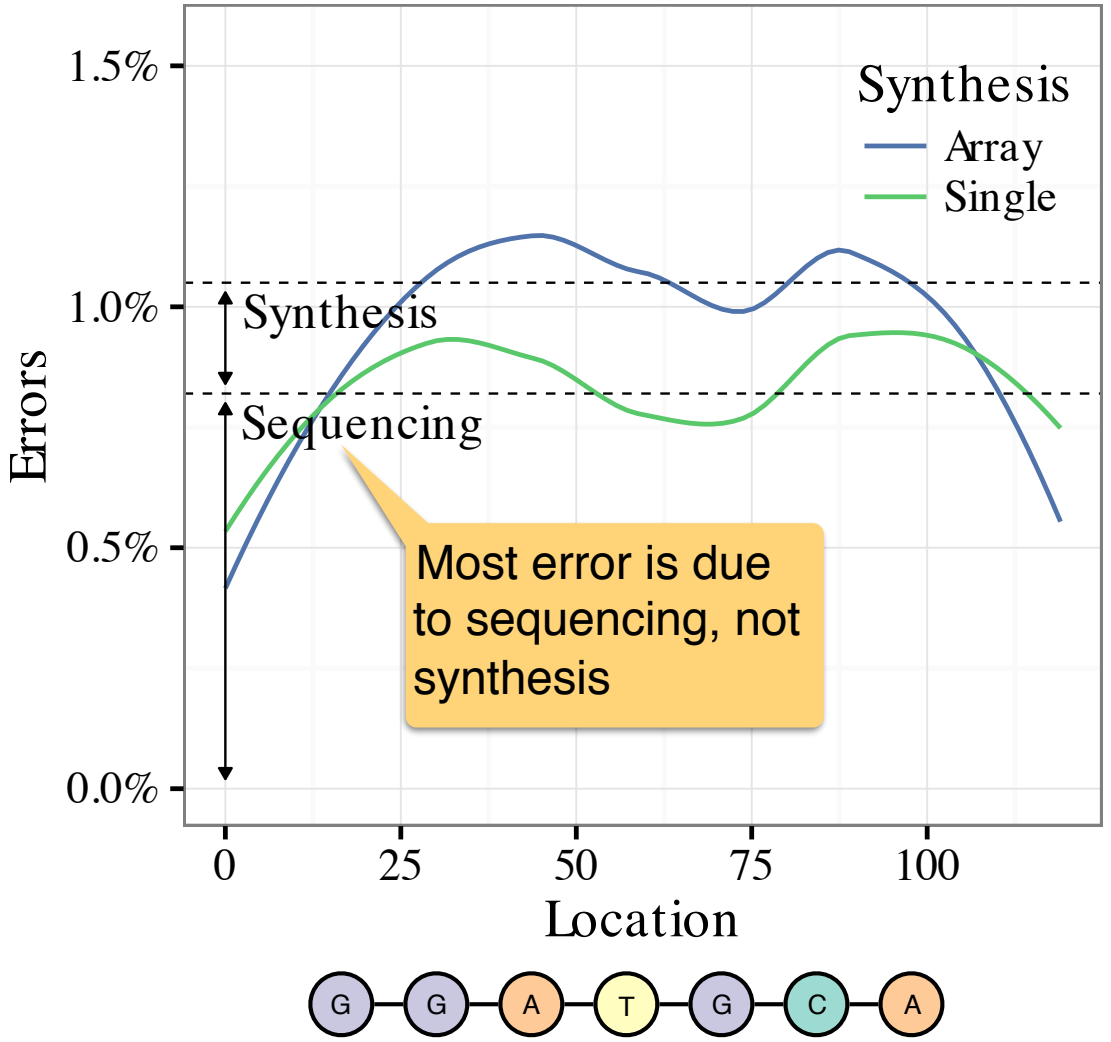

Durability

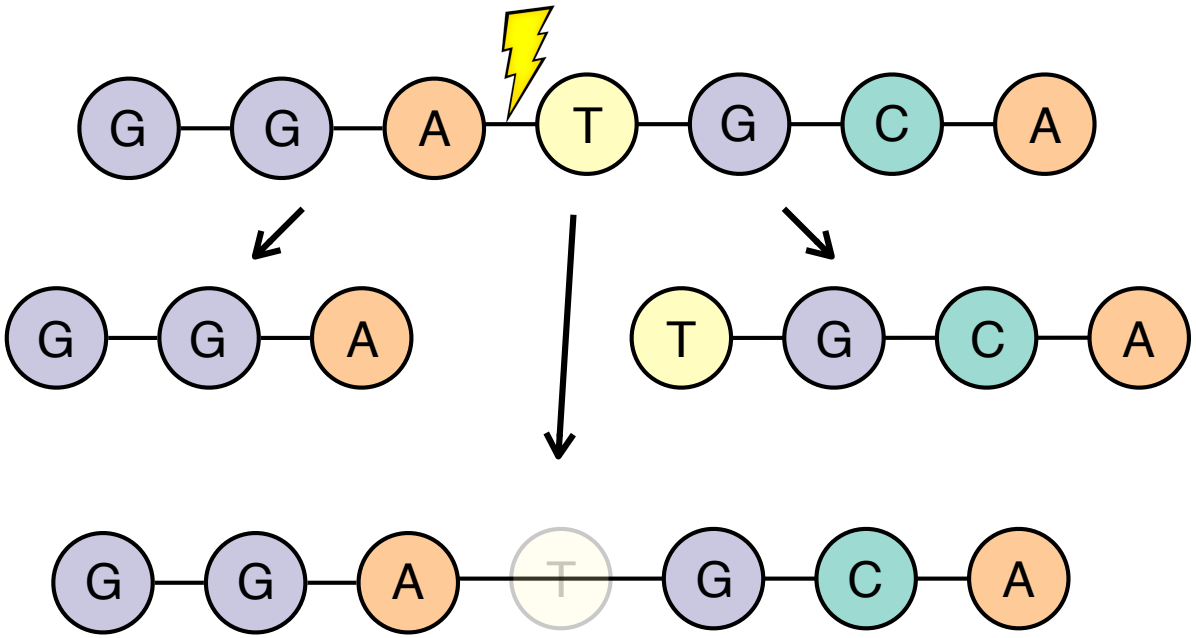

Water
Radiation
…

Durability

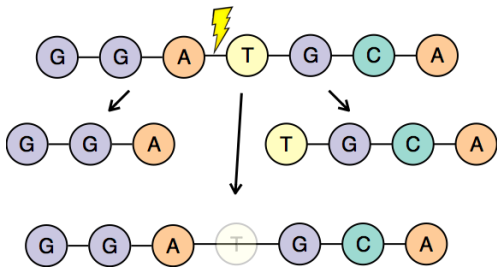

The half-life of DNA in bone: measuring decay kinetics in 158 dated fossils

Morten E. Allentoft, Matthew Collins, David Harker, James Haile, Charlotte L. Oskam, Marie L. Hale, Paula F. Campos, Jose A. Samaniego, M. Thomas P. Gilbert, Eske Willerslev, Guojie Zhang, R. Paul Scofield, Richard N. Holdaway, Michael Bunce

Half-life in bone: ~521 years.

ETH zürich

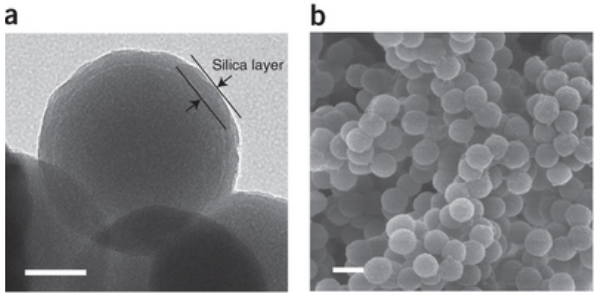

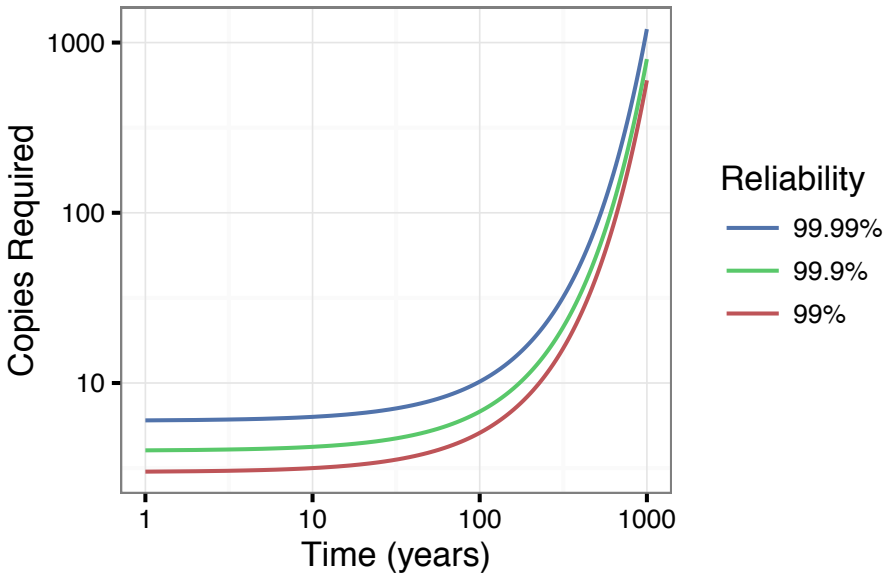

DNA synthetic fossils survive:

Time	Temperature
1 week	70˚C
2,000 years	10˚C
2,000,000 years	-18˚C

Source: Grass et al. Robust Chemical Preservation of Digital Information on DNA in Silica with Error-Correcting Codes

MBs/week → GBs/second

DNA manipulation productivity is growing

And cost is decreasing…

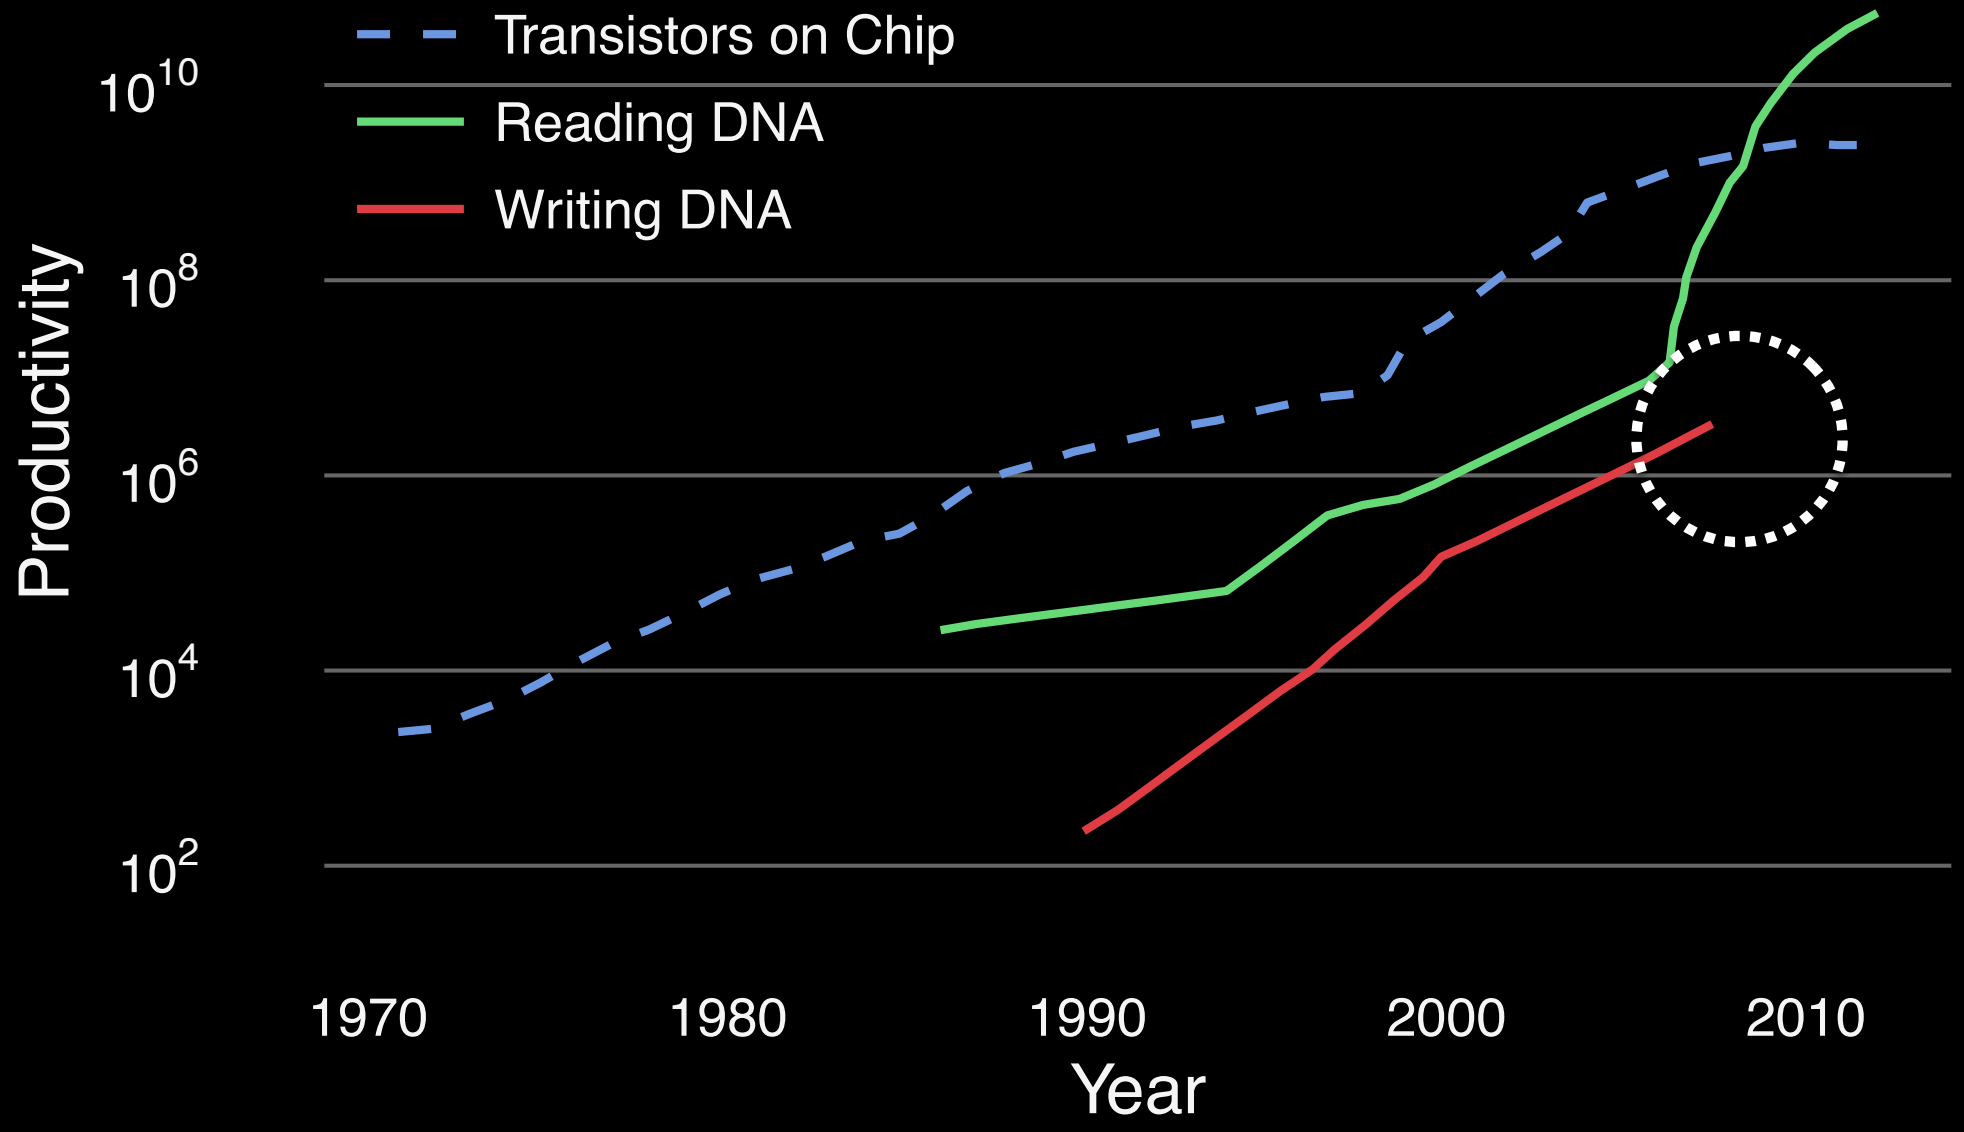

Source: Robert Carlson

Thanks!

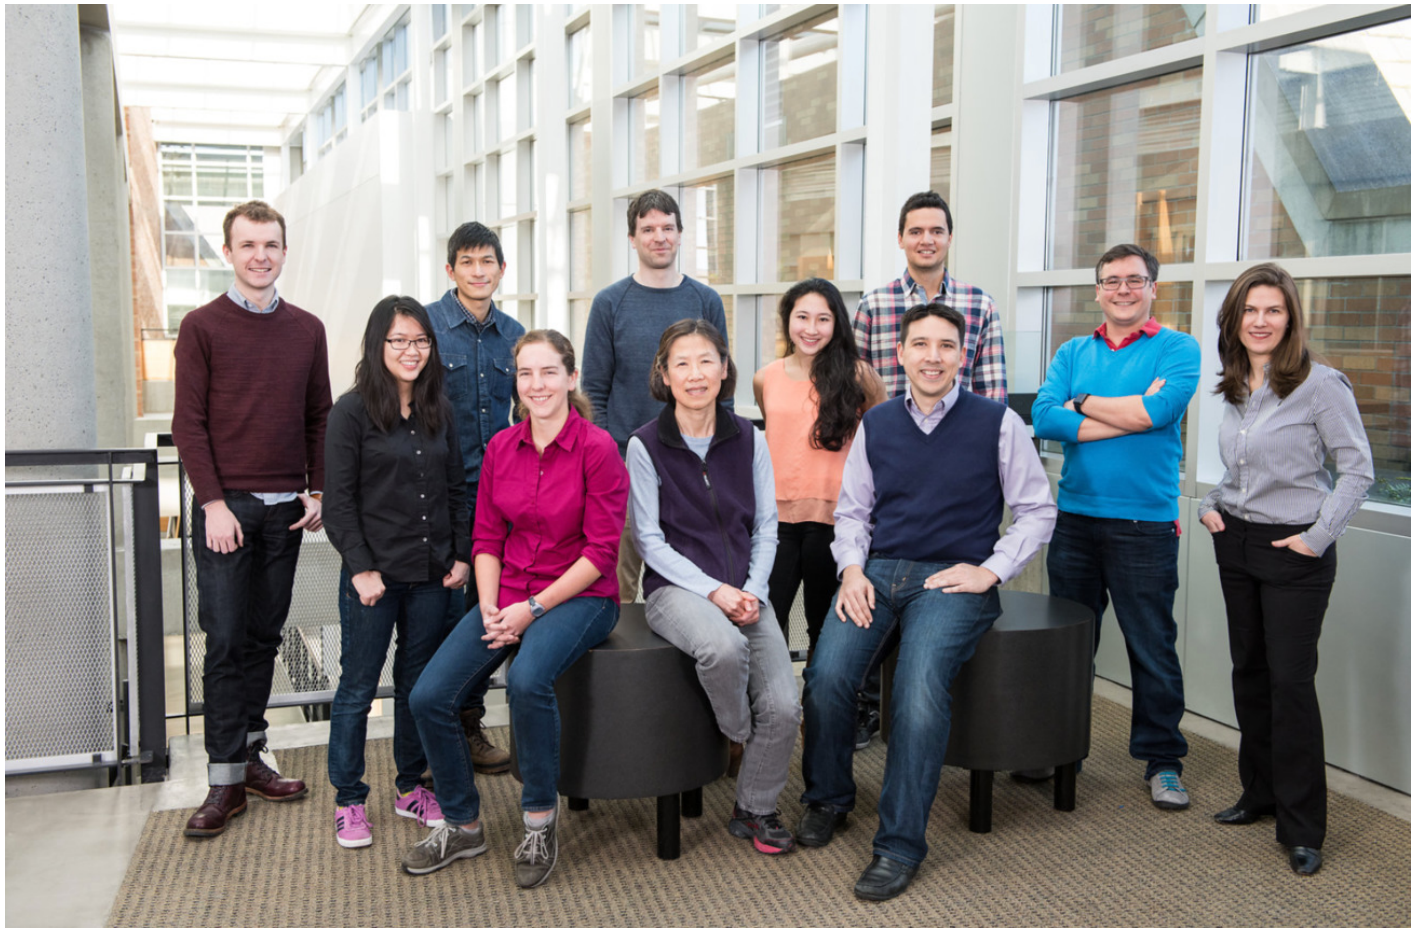

MISL

Molecular Information Systems Lab

Microsoft® Research