

Large-Scale Physical Genome Mapping

Anthony Bonner

Dept. of Computer Science
University of Toronto

Joint work with the Whitehead
Institute/MIT Center for Genome Research

Outline

- Introduction to Molecular Biology
- Genome Mapping
- Noisy Data
- Integrating many forms of data.
- Our approach
 - Condense mapping data into a graph.
 - Graph algorithms
 - Graph visualization
- Summary

**CLINTON
MISSION TO EUROPE**



Genetics

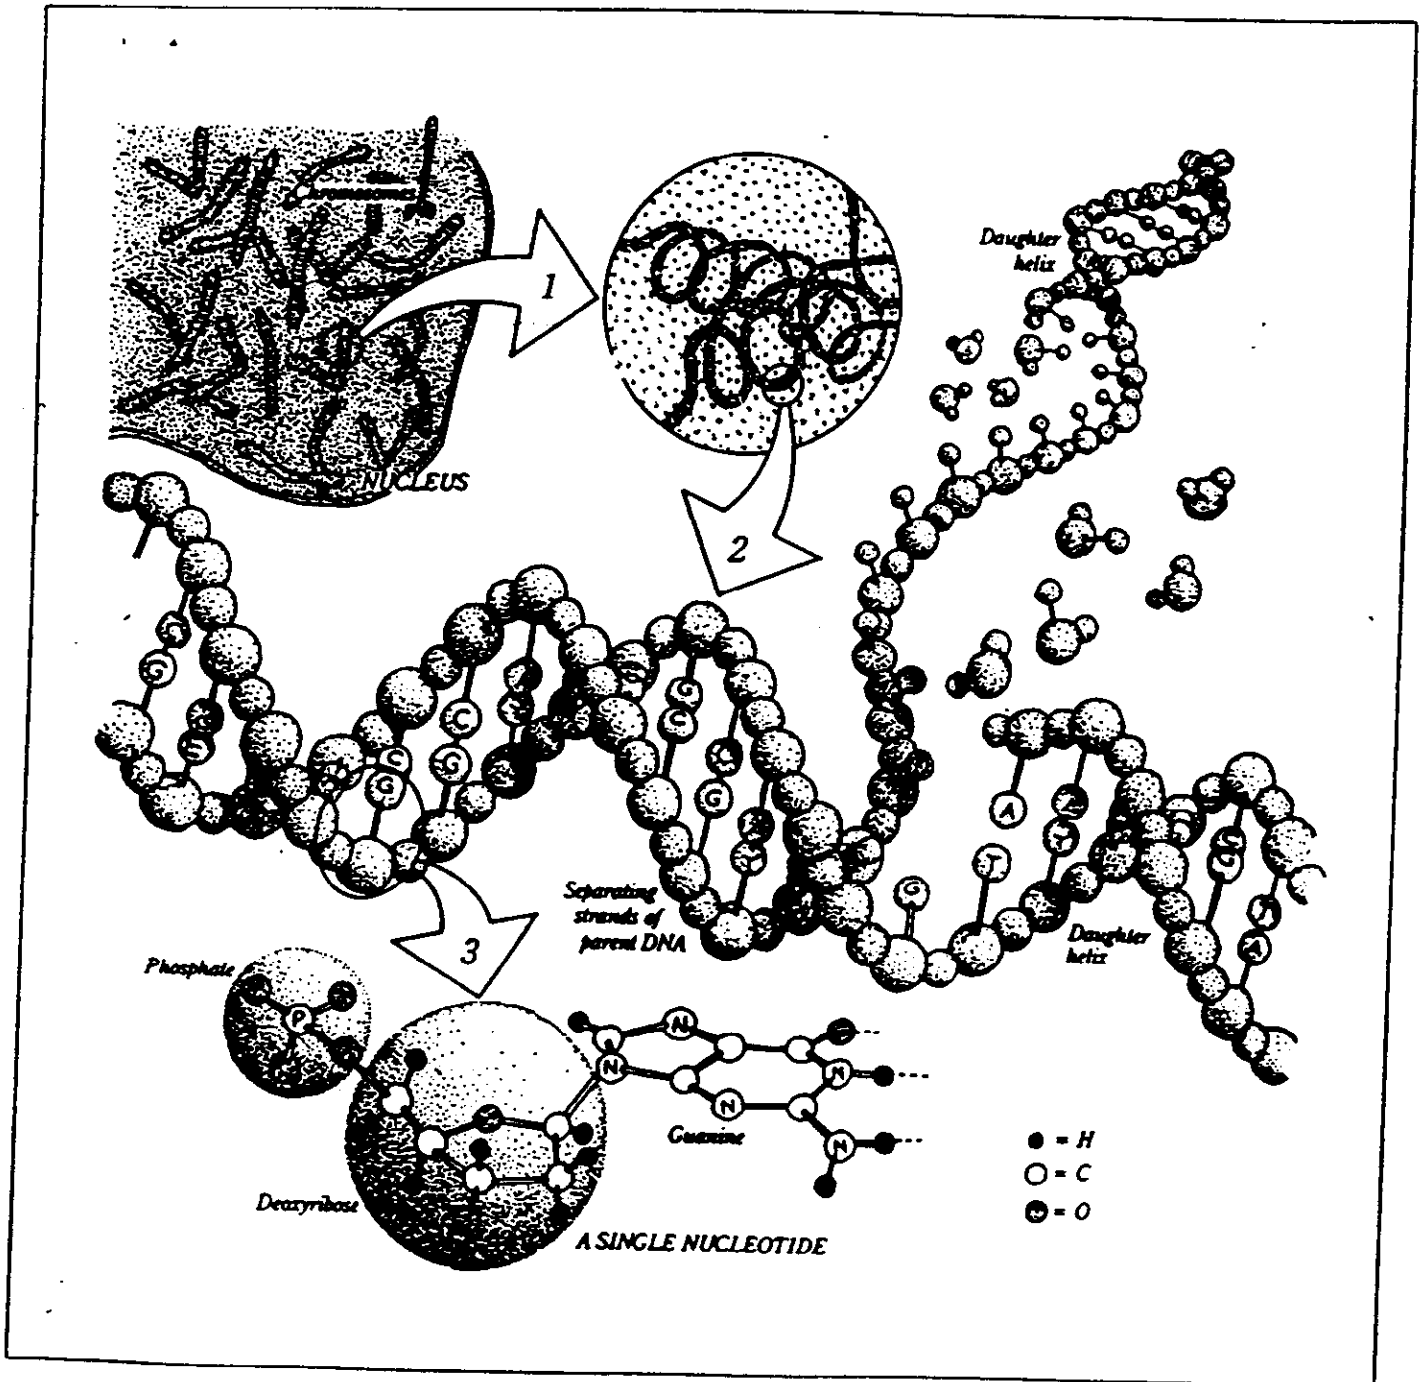
THE FUTURE IS NOW

**New breakthroughs can cure diseases and save lives,
but how much should nature be engineered?**

The Human Genome at Four Levels of Detail.

Primer on Molecular Genetics

Appendix A



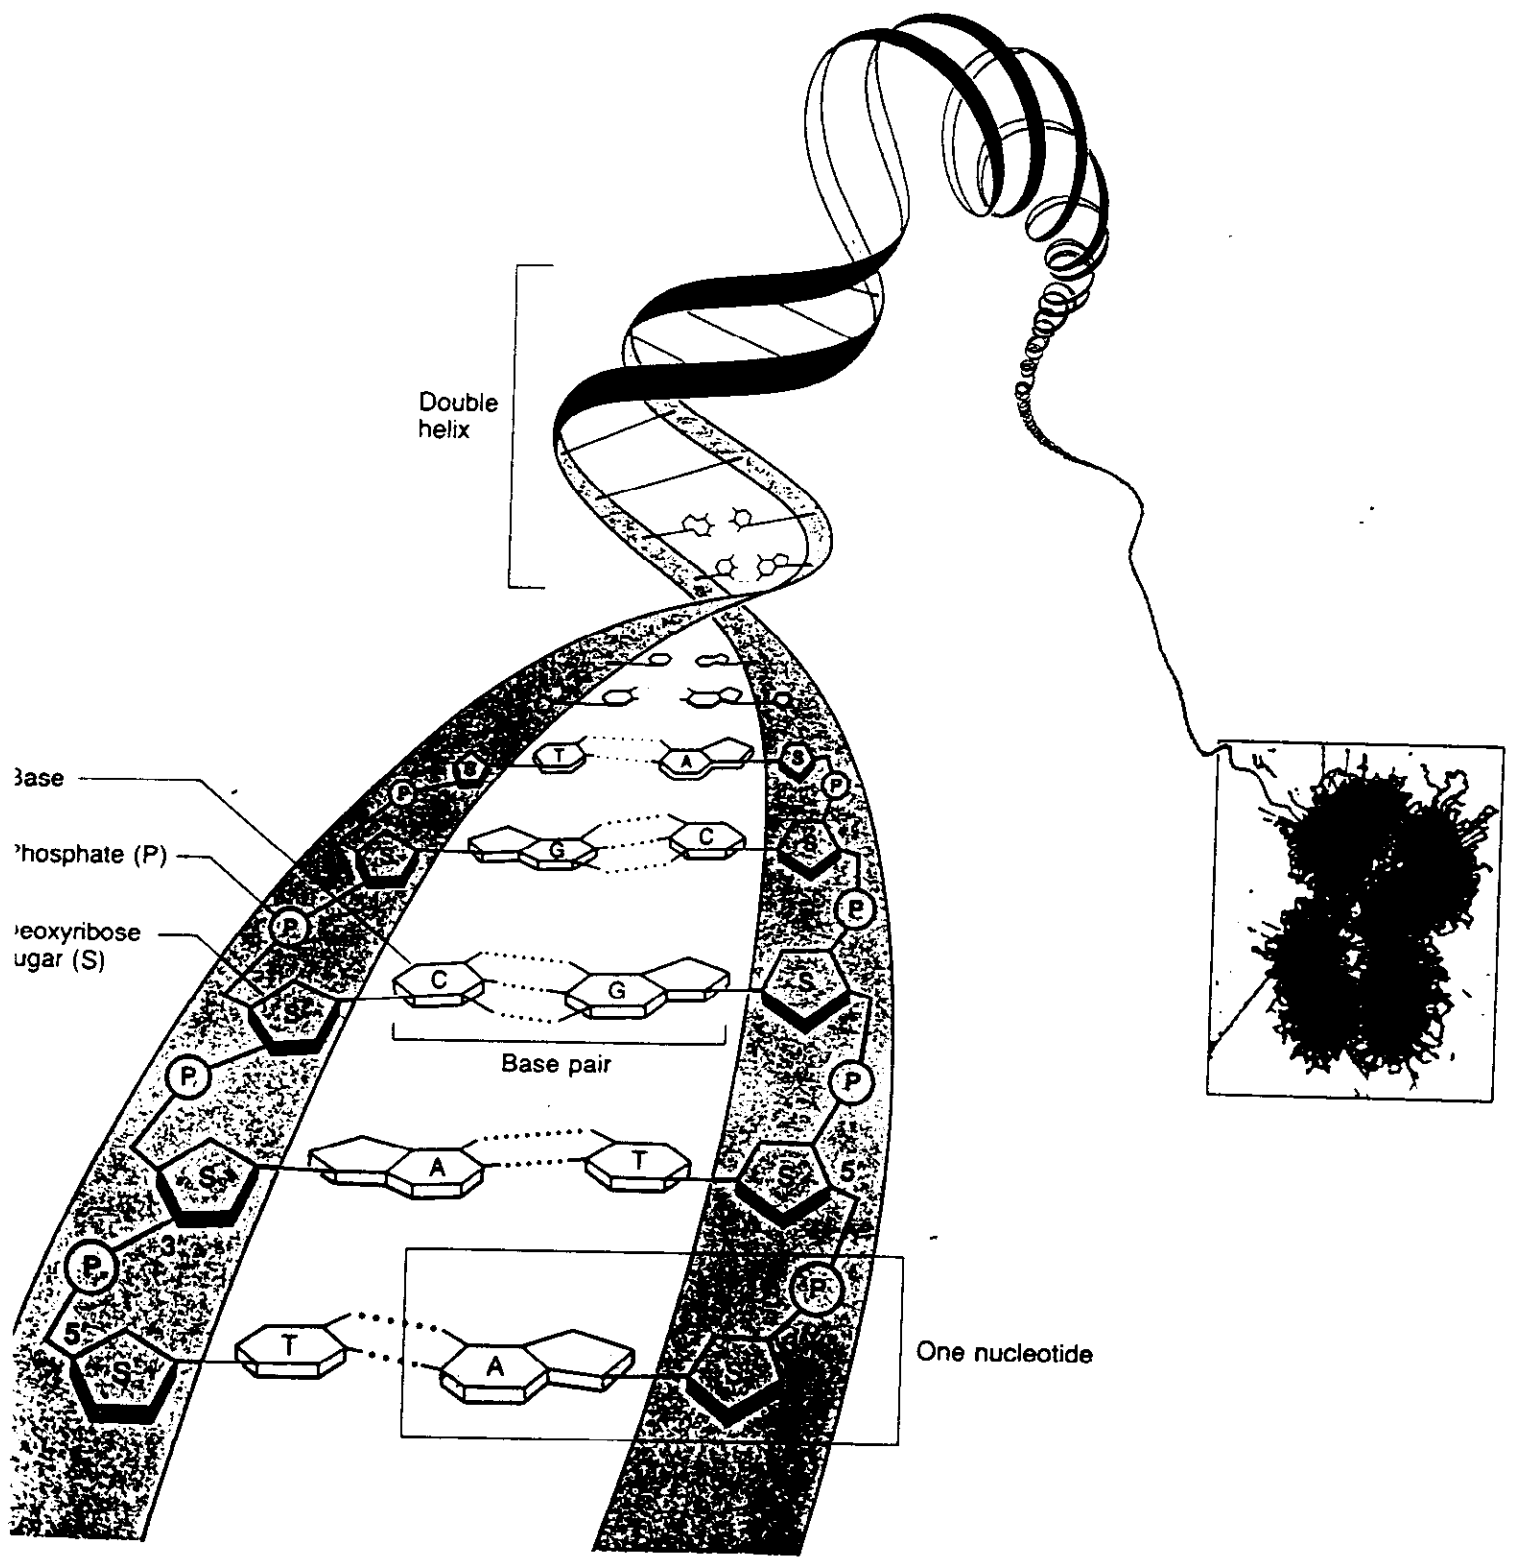
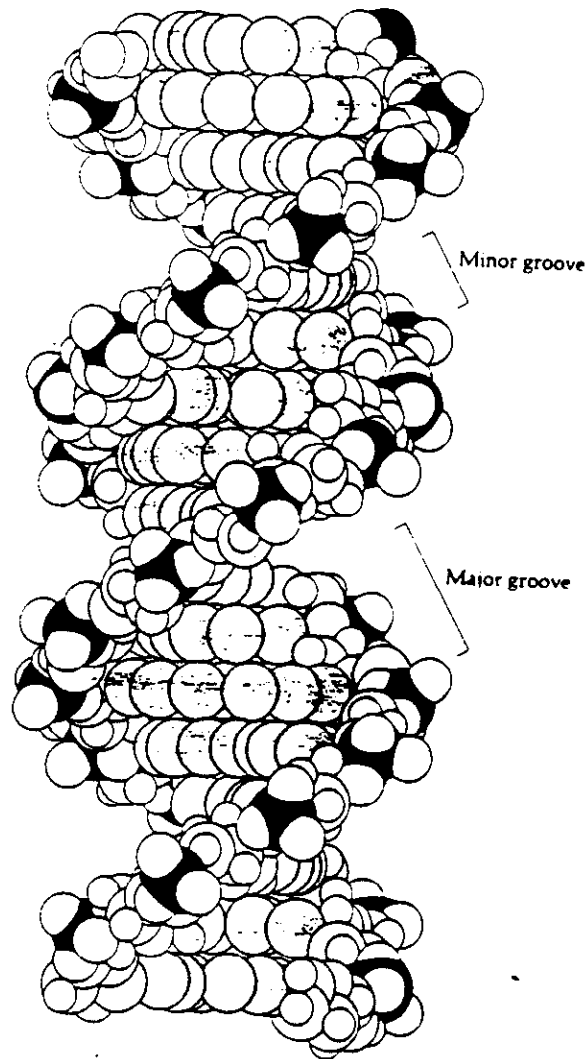


Figure 2.5 The arrangement and association of nucleotides in the DNA double helix.

DNA is a complex biological polymer:



which may be represented by a string of symbols from the set

{ A G C T }

each of which represents a chemical group (nucleotide)

A Portion of a Human Hemoglobin Gene:

1 GTAAGCAGGT TGTGGTTGAG AAAGGAAAGT GTGAAACAGG GACCCAGAGG
51 GAGAGGTGGG GGGATGGCGC TGCTCAGTTT GGTTTGAGGG ACTTGCTTCT
101 CTGACCAAGG TAGGAGGATA CTAACTTCTT CCCAAACTGC CATCACTGGA
151 GACATAGTAA GGGGTAAGAA AGTGTGTCCG GGCAACTGAT AAGGATTCCC
201 TGCACCTAGG GGAAGCACAA CCCAGCCCCA GAATCTCAGG GGCCCTAACA
251 AGTTTTACTG GGTAGAGCAA GCACAAACCA GCCAATGAGT AACTGCTCCA
301 AGGGCGTGTC CACCCTGCCT GGAGGACAGC CCTTGGAGGG CATATAAGTG
351 CTACTTGCTG CAGGTCCAAG ACACTTCTGA TTCTGACAGA CTCAGGAAGA
401 AACCATGGTG CTCTCTGGGG AAGACAAAAG CAACATCAAG GCTGCCTGGG
451 GGAAGATTGG TGGCCATGGT GCTGAATATG GAGCTGAAGC CCTGGAAAGG
501 TGAGAACAGG ACCTTGATCT GTAAGGATCA CAGGATCCAA TATGGACCTG
551 GCACTCGCTC AGTGGGCAGC TTCTAACTAT GCTTTTCTGT GACCTCAACT
601 TCTCTTCTCT CCTTCTCCCA GGATGTTTGC TAGCTTCCCC ACCACCAAGA
651 CCTACTTTCC TCACTTTGAT GTAAGCCACG GCTCTGCCCA GGTC AAGGGT
701 CACGGCAAGA AGGTCGCCGA TGCCTGGCC AGTGCTGCAG GCCACCTCGA
751 TGACCTGCCC GGTGCCTTGT CTGCTCTGAG CGACCTGCAT GCCACAAGC
801 TCGTGTGGA TCCCGTCAAC TTCAAGGTAT GCGCTGGGAC CTGGCAGGCG
851 GCATCTGGGA CCCCTAGGAA GGGCTTGGGG GTCCTCGTGC CCAAGGCAGG
901 GAACATAGTG GTCCCAGGAA GGGGAGCAGA GGCATCAGGG TGTCCACTTT
951 GTCTCCGCAG CTCCTGAGCC ACTGCCTGCT GGTGACCTTG GCTAGCCACC
1001 ACCCTGCCGA TTTCACCCCC GCGGTACATG CCTCTCTGGA CAAATTCCTT
1051 GCCTCTGTGA GCACCGTGCT GACCTCCAAG TACCGTTAAG CTGCCCTCTG
1101 CGGGGCTTGC CTTCTGGCCA TGCCCTTCTT CTCTCCCTTG CACCTGTACC
1151 TCTTGGTCTT TGAATAAAGC CTGAGTAGGA AGAAGCCTGC ATGCCTGGTT
1201 CTCTGCGTCT GCAAAGGTGT CATGTTTAGT GTGGGGATGC CGCAGCTCAT
1251 TTGCCATGGG GCAGTAAAGA CAAGGTTTCA AGCAAAAAGC ATAATTGGAT
1301 GCCTACACAC ACACACATAT GTCTTCTGAG TCTGGGCAGC AGTCCCTCCC
1351 AAGCCCTCCA CTGACAGCCA TGTGTCTTCT CCTCGAGCCA AAGAAGCCAA

Humans cells are estimated to contain 3×10^9 characters
of independent information

GENOME SIZES

- E. Coli
 - 4.5 Million base pairs.
 - 4000 genes.
- Fungi, Yeasts:
 - 10-100 Million bp.
- Human:
 - 3 Billion bp.
 - 100,000 genes.

THE HUMAN GENOME PROJECT

- An International Effort to map & sequence the human genome.
 - Made possible by genetic breakthroughs in 1970's & 80's.
 - Biologists are now overwhelmed with genome data.
- ⇒ Data Management Problems

Genome Research

- Gene Hunting
- Genome Mapping
- Genome Sequencing
- Related Activities:
 - Evolution
 - Protein Structure
 - Cellular Metabolism
 - Gene Regulation

Genome Maps

- Gives the location of important or easily identifiable sites on each chromosome.

Two Kinds

- Linkage Maps:
 - Course grained
 - Lots of statistics
- Physical Maps:
 - Fine grained
 - Lots of algorithms

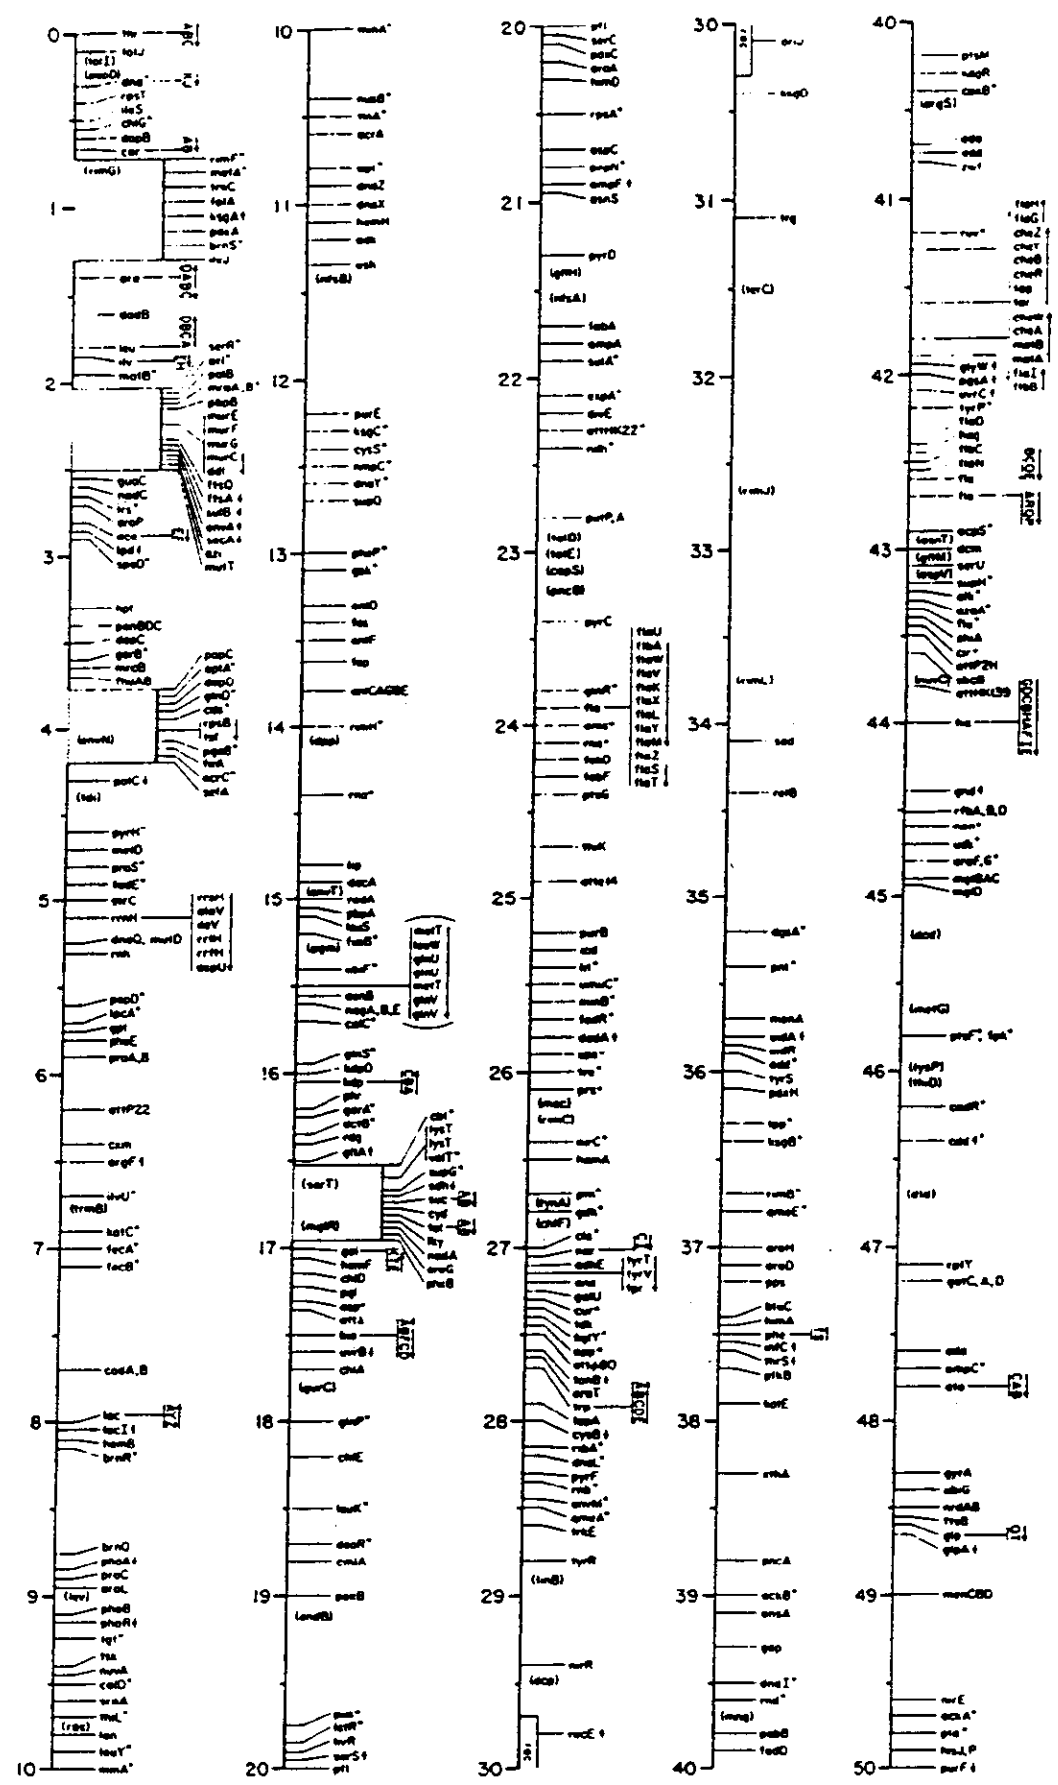
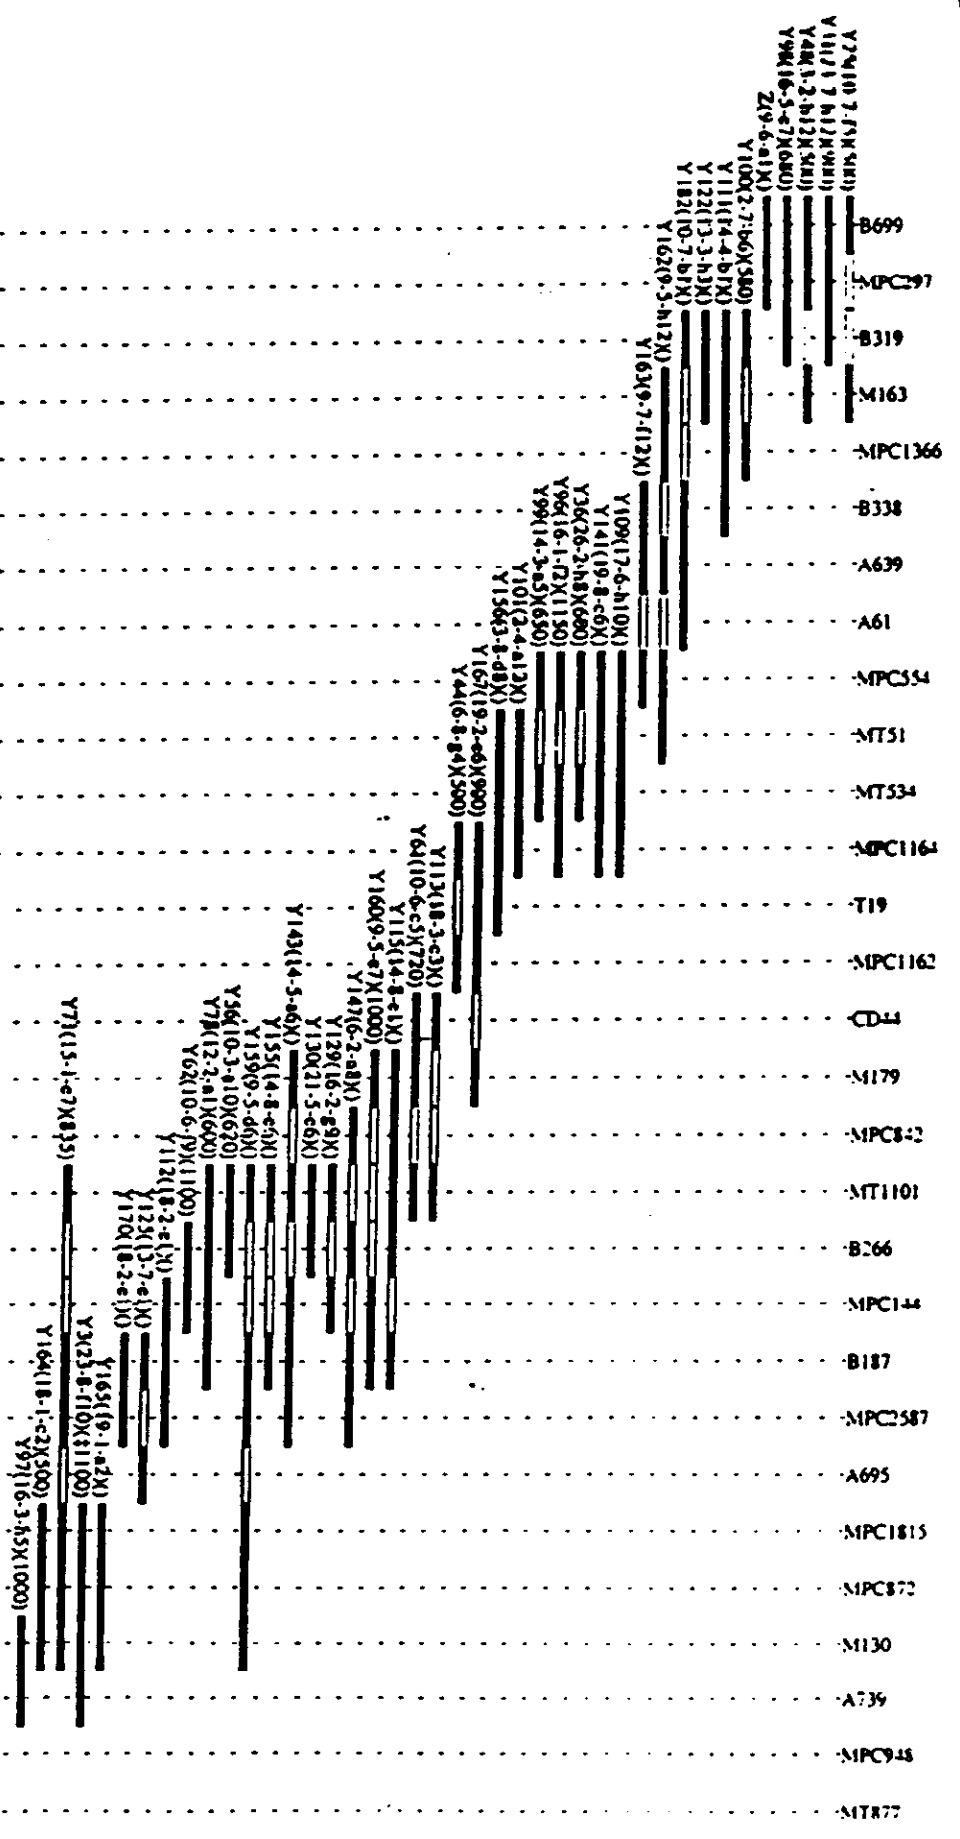


Figure 7-22
The complete genetic map of *E. coli*. [Courtesy of Barbara J. Bachman, *Microbiol. Rev.* 47 (1983):180]

Genetic Linkage Map

Tiny Portion of a Physical Genome Map



MOUSE CHR. 2

PETER GROOT

Building a Physical Genome Map

- Smash (many copies of) a genome into thousands of fragments (clones).
 - Experimentally, determine which pairs of clones overlap.
 - Computationally, use the overlap data to assemble a map ("jigsaw puzzle").
 - Main problems:
 - Noisy data
 - Anomalous data
 - Many forms of data
 - Subtle interrelationships
- } data cleansing
- } data integration

Overlap

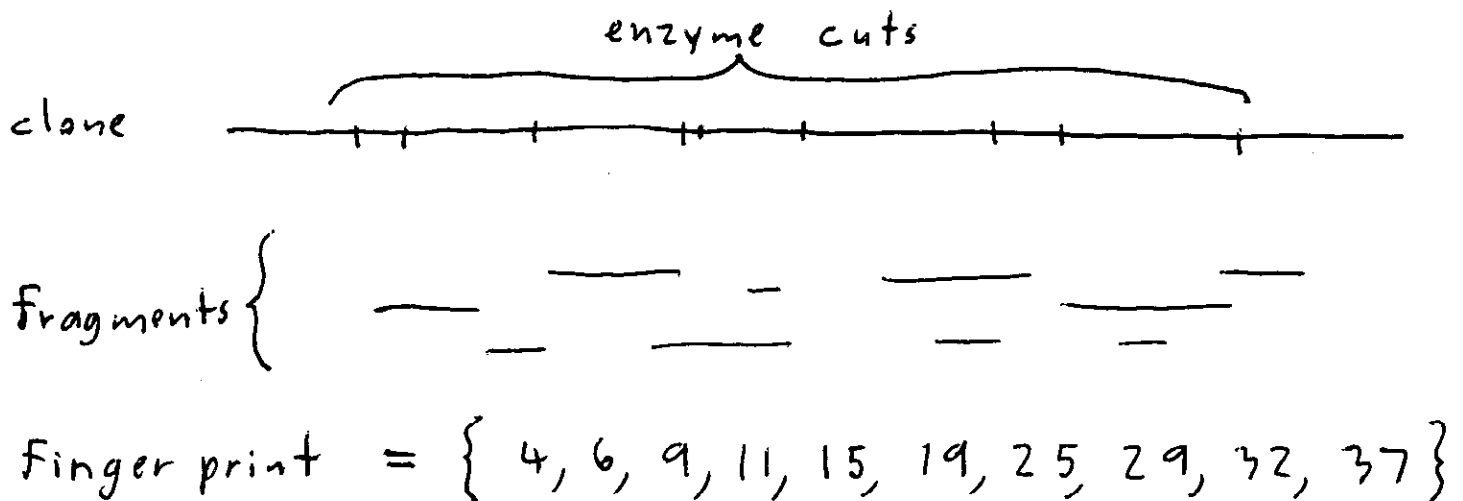
There are many ways to determine whether two clones overlap.

Some common methods:

- Finger prints
- STS content
- ALU PCR

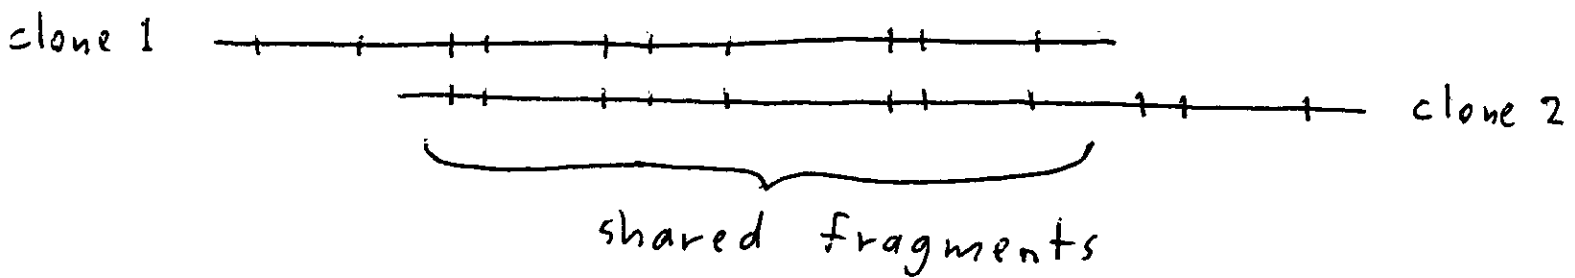
Finger Prints

- Use restriction enzymes to cut a clone wherever a given sequence occurs.
- eg. cut a clone at every occurrence of atcgat and gatc (complete digestion).
- Measure the lengths of the resulting fragments.
- The set of lengths is called a fingerprint.



Finger Print Overlap

IF the Fingerprints of two clones
have many lengths in common,
Then the clones probably overlap.



$$\text{Fingerprint 1} = \{4, 5, 6, 7, 8, 10, 11, 12, 15, 19, 21\}$$

$$\text{Fingerprint 2} = \{4, 5, 6, 7, 8, 9, 12, 15, 19, 21, 23, 25\}$$

$$\text{Fingerprint 1} \cap \text{Fingerprint 2}$$

$$= \{4, 5, 6, 7, 8, 12, 15, 19, 21\}$$

STS Content

- An STS (Sequence Tag Site) is a fragment of genomic DNA several hundred base pairs long.
- With high probability, an STS will appear only once in a random sequence of 3 Billion base pairs.
- With high probability, two clones overlap if they hybridize with (hit) the same STS.



Problems: Biological Anomalies

- A genome is not a random sequence.
- Many regions (subsequences) repeat.
- IF an STS hybridizes with a repeat region, then clones that are far apart may appear to overlap.



Problems: Experimental Error

False Negatives:

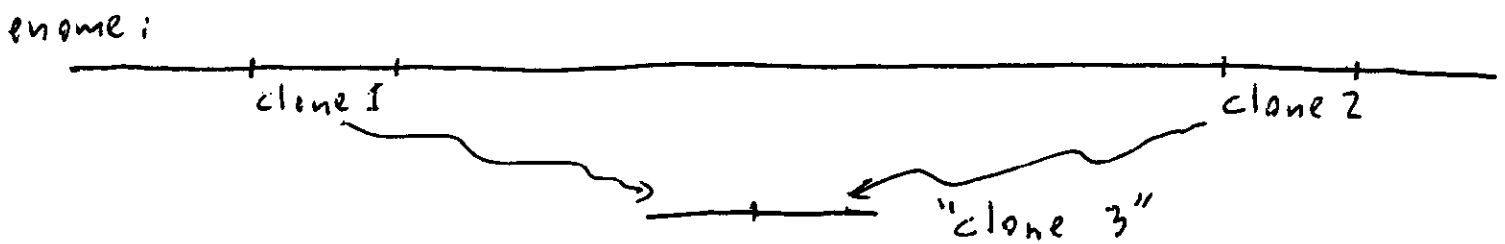
- An STS should hybridize with a clone, but fails to
- So, an overlap goes undetected.

False Positives:

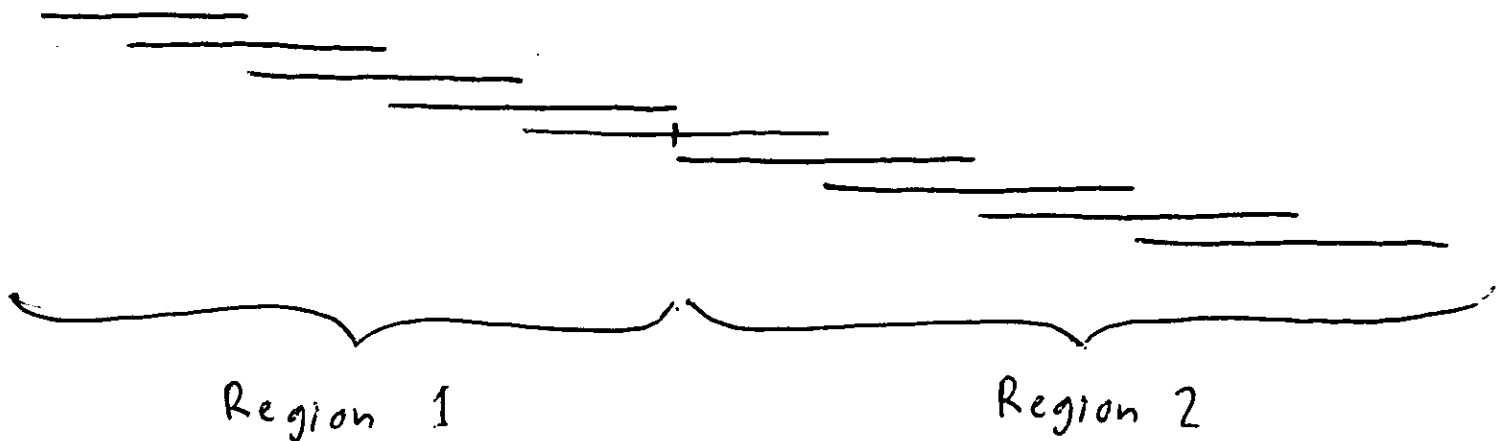
- An STS should not hybridize with a clone, but does.
- So, an overlap may erroneously be inferred.

Problems: Chimerism

- What appears to be a single clone, is actually two (or more) clones.
- One possible cause: When the genome is smashed into clones, some clones may fuse.

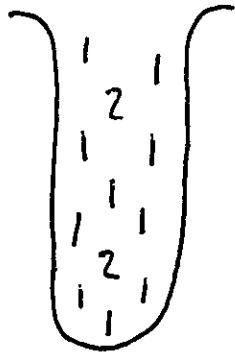


- Chimeric clones make discontinuous maps look contiguous:



Chimerism

- Another possible cause: Contamination
- A test tube containing (many copies of) clone 1, may also contain (some copies of) clone 2.



- Problem: Clones that overlap with either clone 1 or clone 2 will react positively with this test tube.

Summary of Problems

Biological Anomalies

- Repeat regions: The same sequence may be found in many places on a genome.

Overlap Errors

- False positives: Clones appear to overlap when actually they do not.
- False Negatives: Clones appear not to overlap when actually they do.

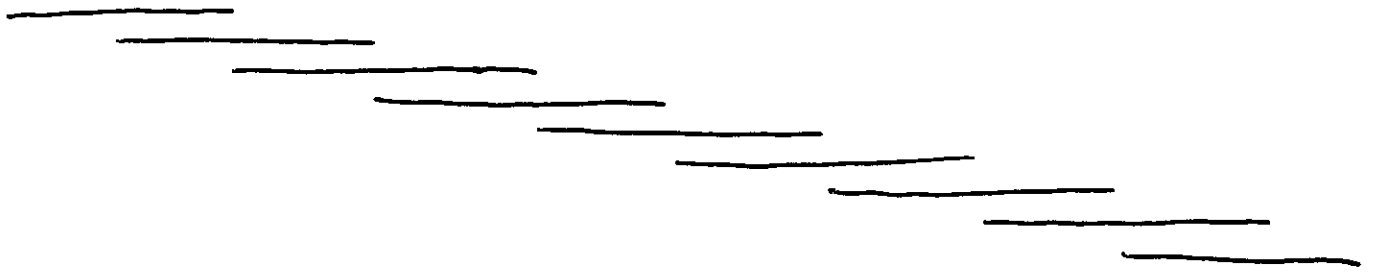
Identity Errors

- Chimerism: What appears to be one clone is actually two (or more) clones.

Assembling Large Physical Maps

- Given a collection of overlap data, how do we construct a map?

eg.



- Example: STS data
 - Ideal case (no noise)
 - False negatives
 - False positives
- Integrating many forms of overlap data
 - Detecting & removing errors & anomalies.

STS Data

(noise-free case)

Probe P "hits" DNA segment S.

data:

hits(p1, s1)

hits(p3, s2)

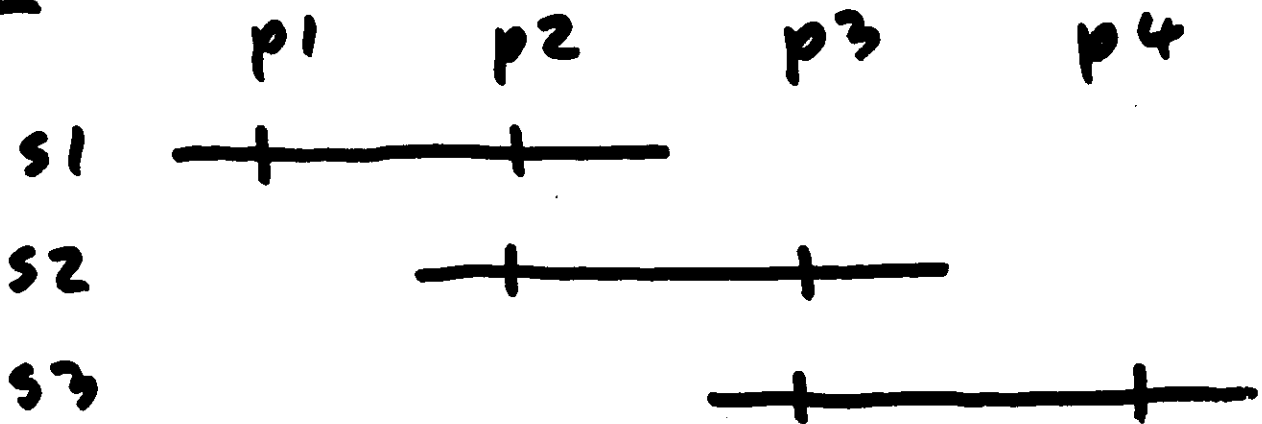
hits(p2, s1)

hits(p3, s3)

hits(p2, s2)

hits(p4, s3)

map:



False Negatives

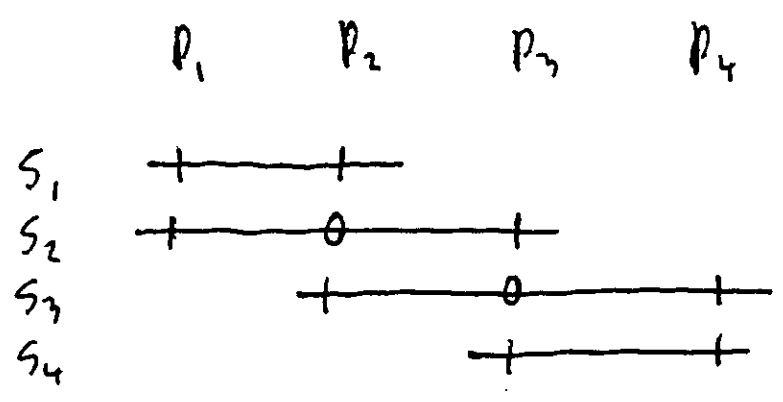
eg.

	P_1	P_2	P_3	P_4
S_1	x	x		
S_2	x		x	
S_3		x		x
S_4			x	x

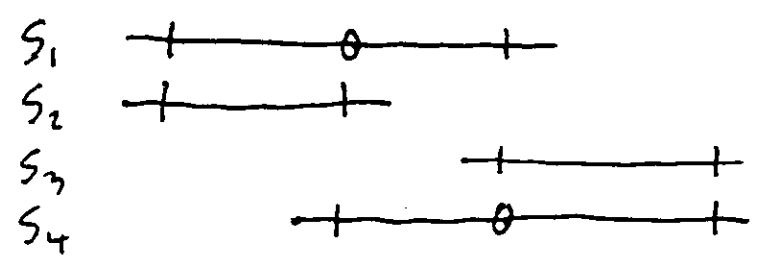
Date:

No map is consistent with this data.

eg.



Map 1:



Map 2:

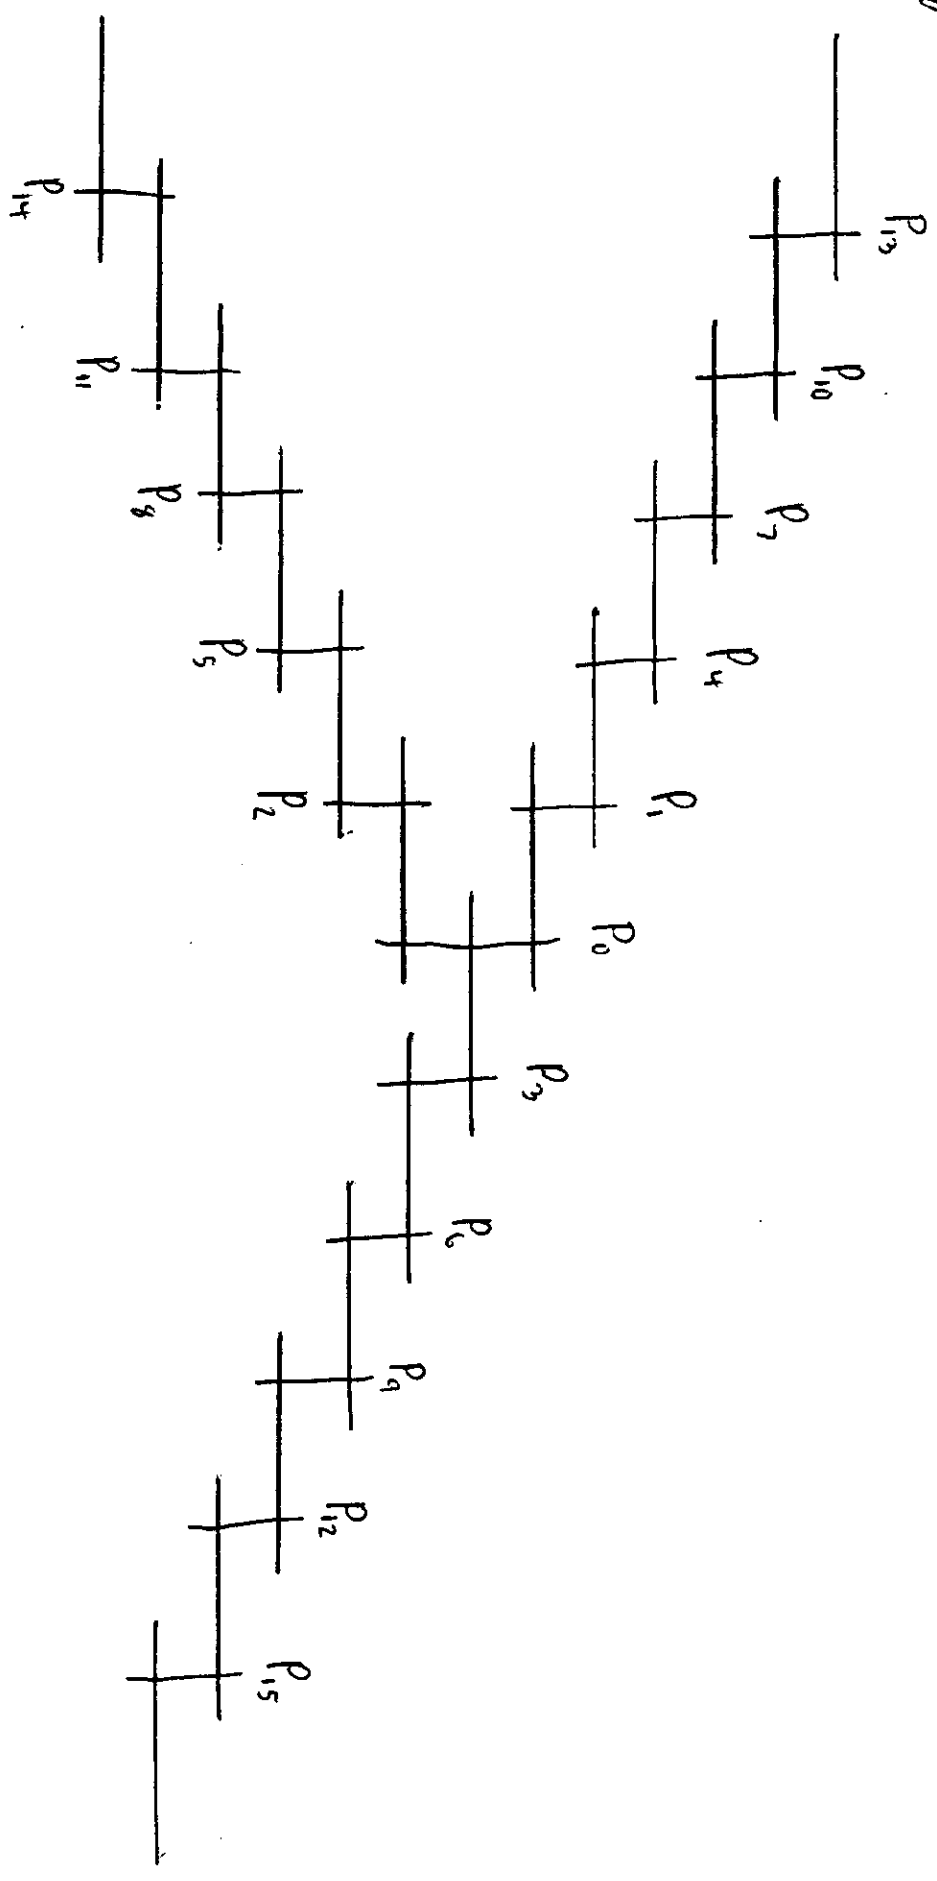
Y29(10-7-f1X581) _____ B699
 Y11(1-7-h12X981) _____ MPC1397
 Y44(1-2-h12X581) _____ B319
 Y99(16-3-e7X680) _____ M163
 Z(9-6-a1X) _____ MPC1366
 Y100(2-7-b6X580) _____ B334
 Y111(F4-4-b1X) _____ A639
 Y122(F3-3-b3X) _____ A61
 Y182(10-7-b1X) _____ MPC54
 Y162(9-3-h12X) _____ MT51
 Y163(9-7-f12X) _____ MT534
 Y100(17-6-h10X) _____ MPC1164
 Y141(19-8-e6X) _____ T19
 Y36(26-2-h8X680) _____ MPC1162
 Y96(16-1-f2X1150) _____ CDL
 Y99(14-3-e5X650) _____ M179
 Y101(2-4-a12X) _____ MPC8-2
 Y154(3-8-d8X) _____ MT1101
 Y167(19-2-e6X900) _____ B266
 Y44(6-8-g4X500) _____ MPC144
 Y113(18-3-e3X) _____ B187
 Y64(10-6-e3X720) _____ MPC2587
 Y115(14-8-e1X) _____ A695
 Y160(9-5-e7X1000) _____ MPC1815
 Y147(6-2-a8X) _____ MPC872
 Y129(16-2-e9X) _____ M130
 Y130(21-3-e6X) _____ A739
 Y143(14-5-e6X) _____ MPC948
 Y153(14-8-e1X) _____ MT872
 Y159(9-5-e6X) _____
 Y56(10-3-a10X670) _____
 Y78(12-2-a1X680) _____
 Y62(10-6-p9X1100) _____
 Y112(18-2-e1X) _____
 Y123(13-7-e1X) _____
 Y170(18-2-e1X) _____
 Y165(19-1-e2X) _____
 Y3(23-8-f10X1100) _____
 Y164(18-1-e2X500) _____
 Y87(16-3-f5X1000) _____

MOUSE CHR.:2
 PETER GROOT

Tiny Portion of a Physical Genome Map

False Positives

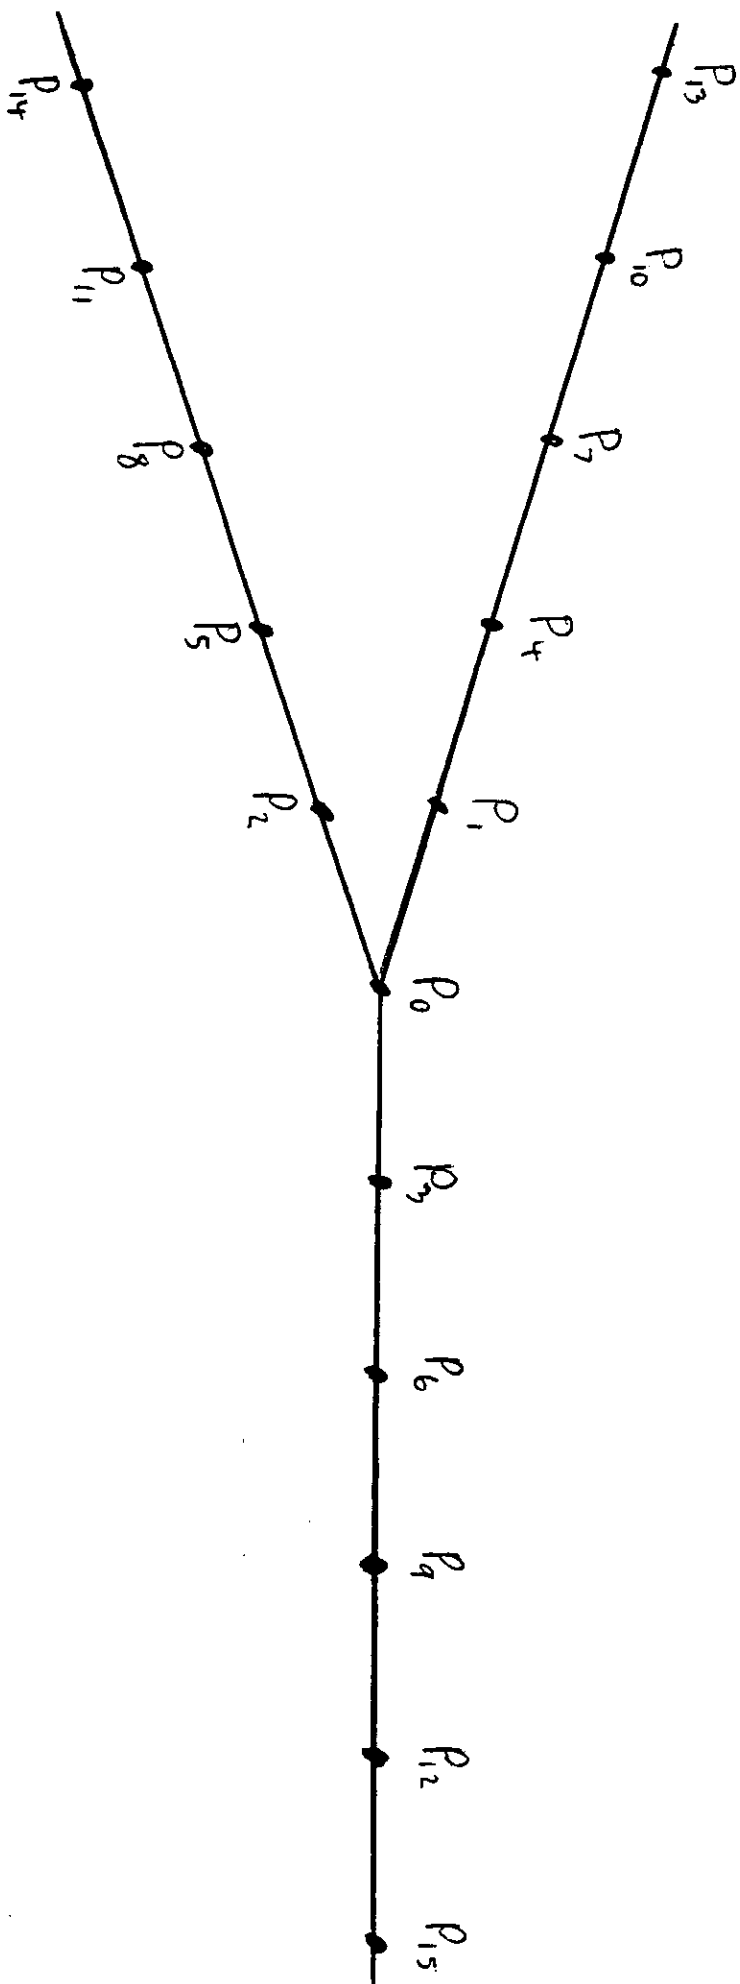
eg.



Unfortunately, no linear order of the probes is even approximately correct.

Observation

The data can be summarized as a graph in which each node is a probe.



Integrated Physical Genome Maps

- Integrating many forms of physical data into a single map.
- Problems: Noise & Complexity.
- Our Approach
 - Clusters: Overlap, Linkage, Ordering
 - Graphs: Algorithms & Visualization
 - Examples
- Summary

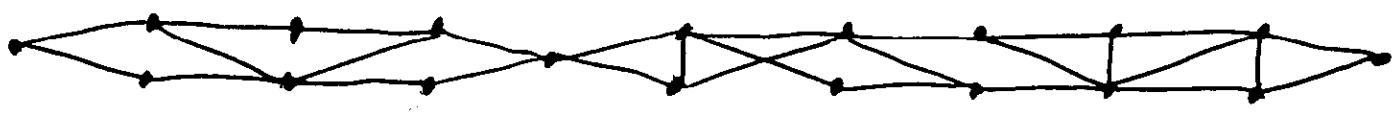
Assembling Integrated Maps: Problems

- Complex Data:
 - Many kinds of data,
 - Data from many labs,
 - Uneven data quality,
 - Noisy data: errors, ambiguities, contradictions, anomalies.
 - Subtle relationships between various forms of data and noise.
- Large and increasing volume of data.
- Algorithms are inflexible and limited to a few forms of data and noise.
- Requires much human intervention and biological expertise.
- Maps are full of errors.

Our Approach

- Abstract the genome data as a graph.
- Ideally, the graphs are long & thin, i.e., approximately "linear".

eg.

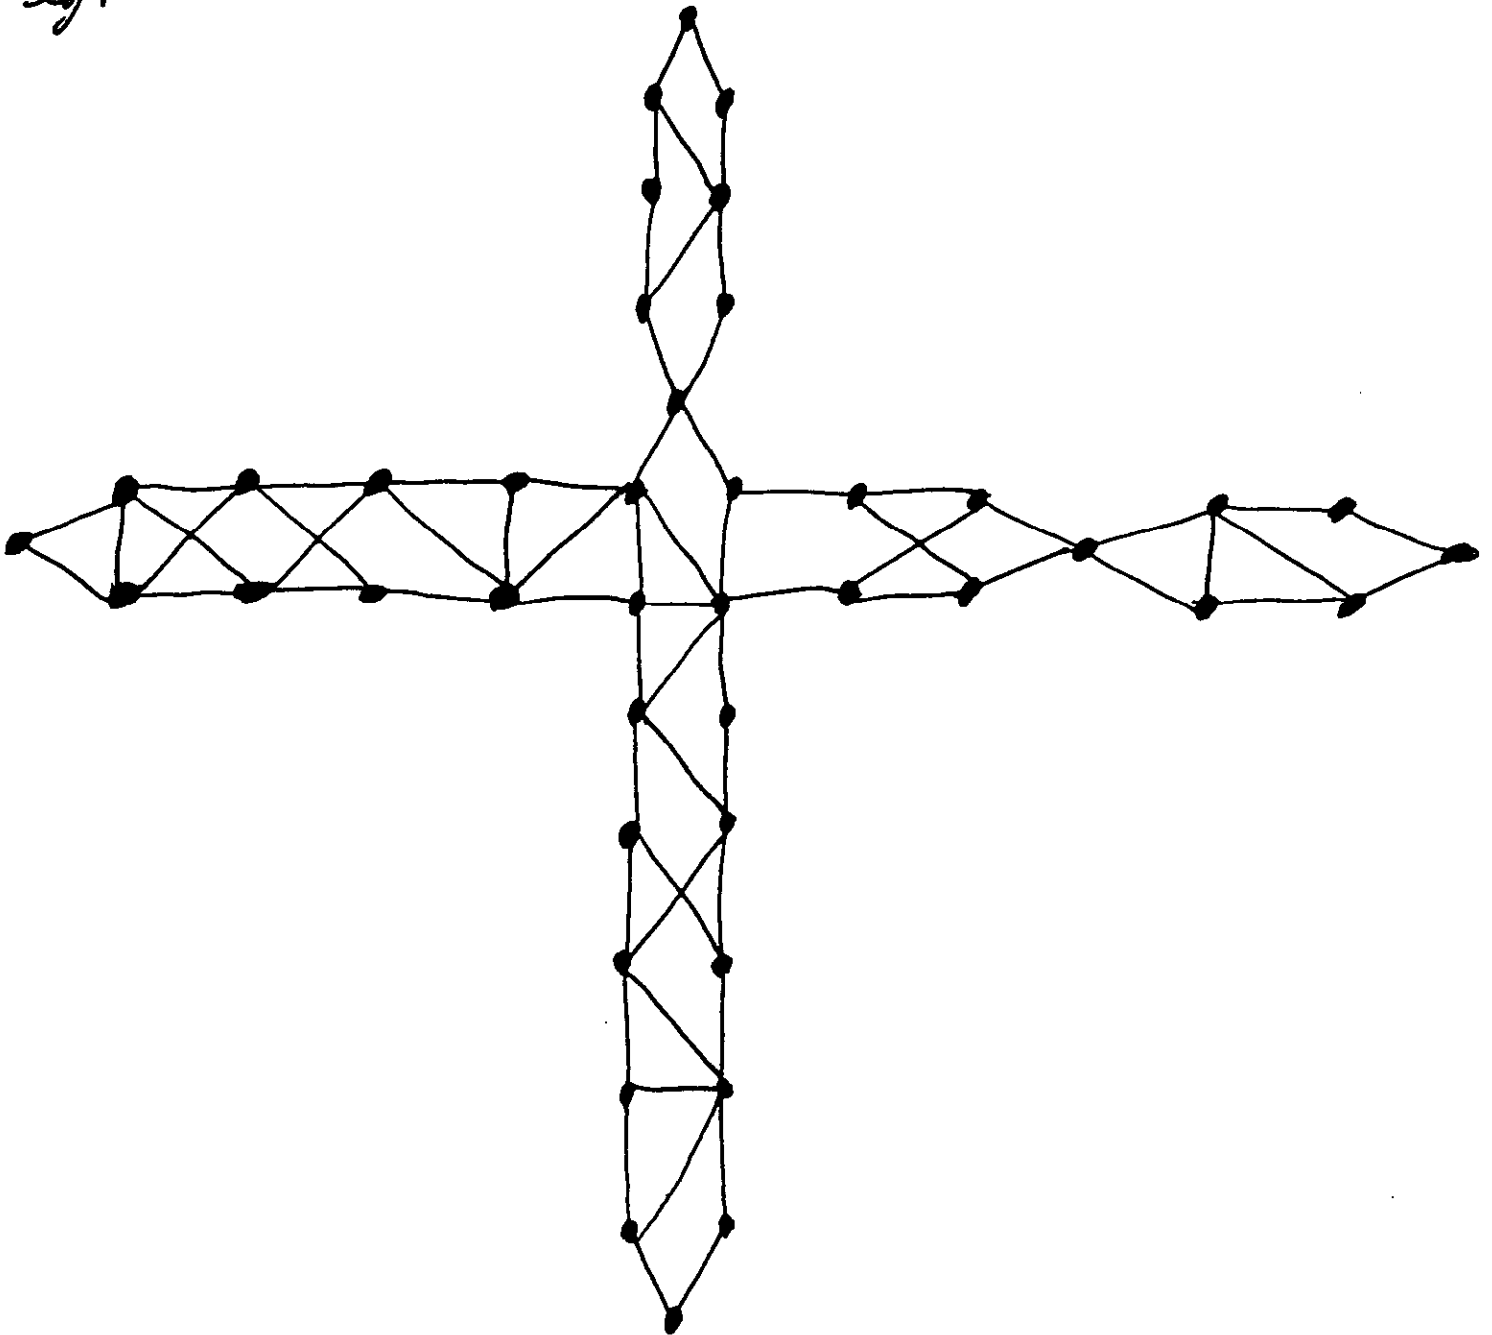


- Each node represents a point on the genome.
- An edge means that two nodes are "close" together on the genome.

Problems

Noise in the data distorts the linear structure of the graph:

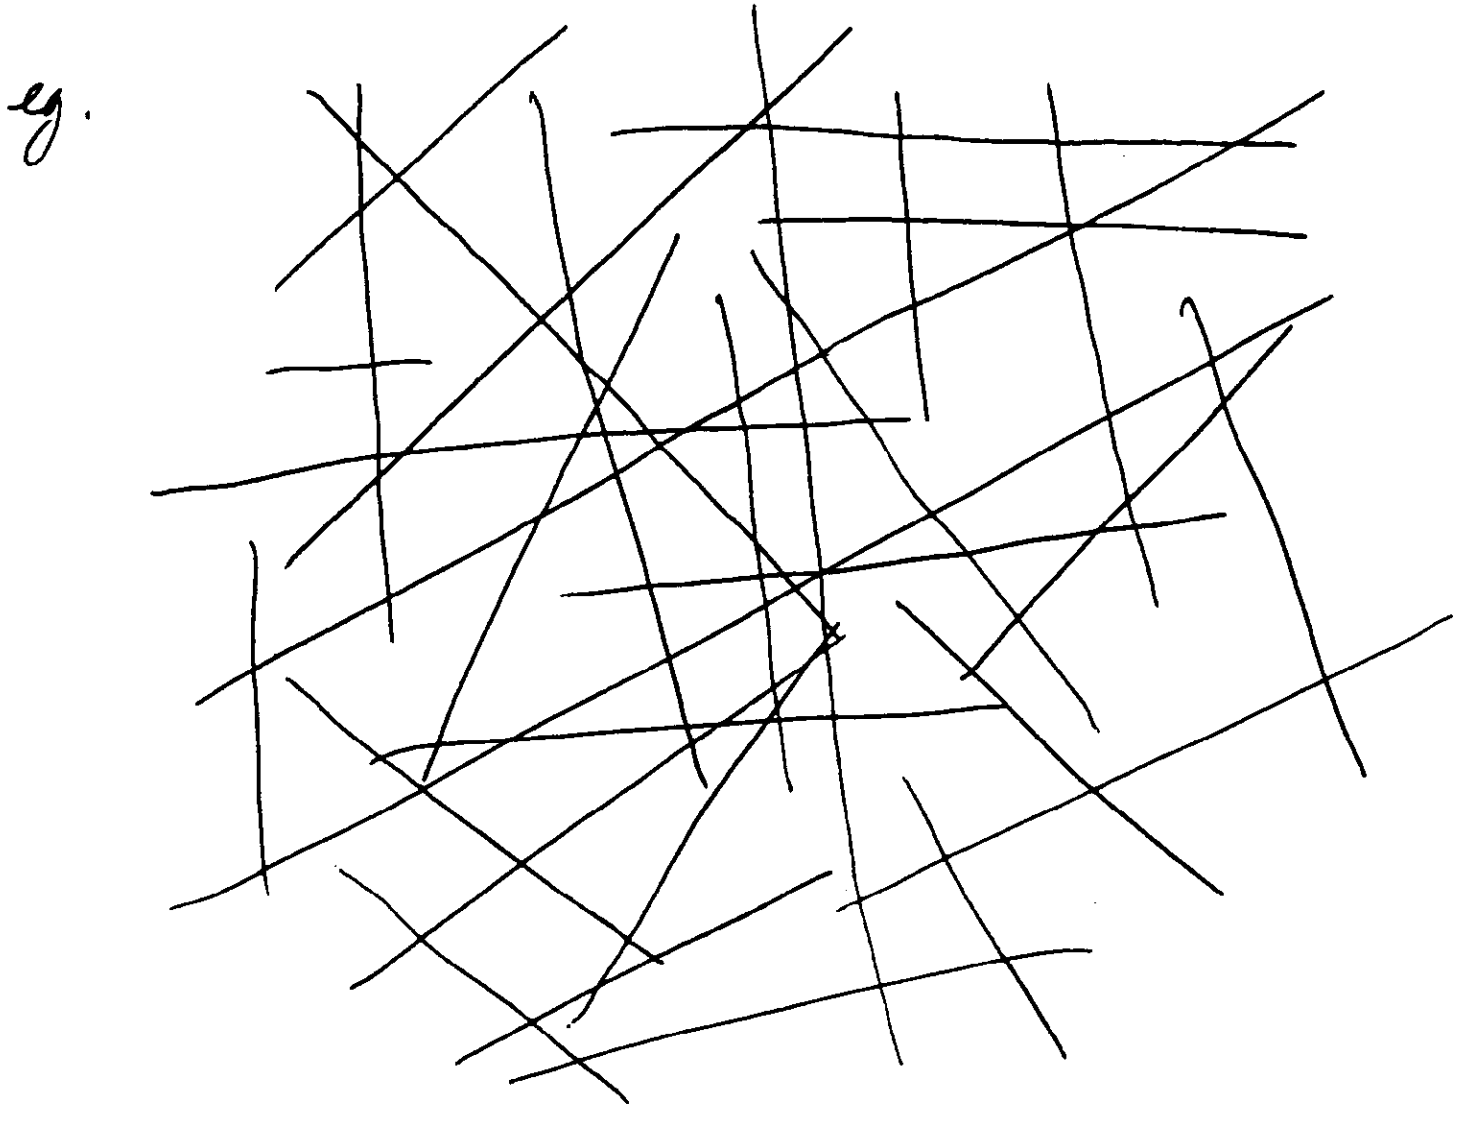
eg.



Real Genome Graphs

eg. 25,000 nodes
 1,000,000 edges

Structure: Like a plate of spaghetti.



Research Problems

- (1) Transforming genome-mapping data into a graph.
- (2) Identifying contiguous paths (contigs) within the graph.
- (3) Generating a genome map from these paths.

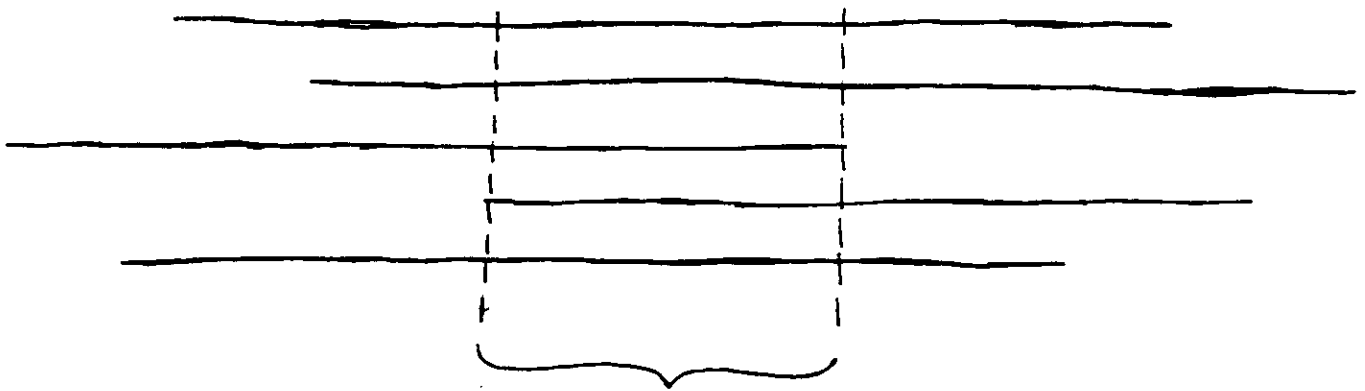
Our approach uses numerous graph algorithms and graph visualizations.

Graph Generation (Data Integration)

Many forms of physical mapping data determine whether two clones overlap.

Definition: A cluster is a maximal set of mutually overlapping clones.

eg.

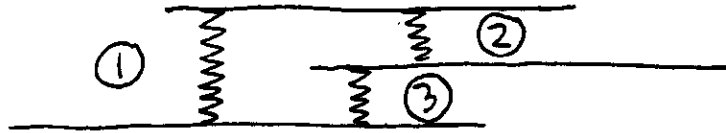


Common
Genomic
Region

Observation 1

Clusters Filter out False overlaps.

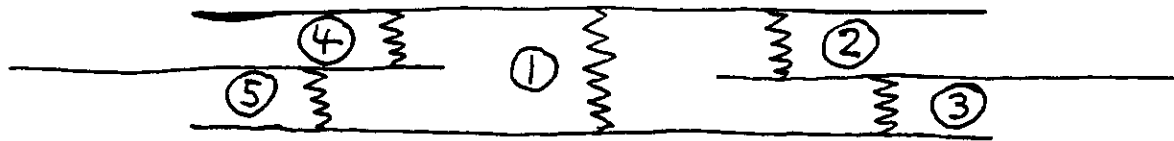
eg.



Overlap ① is corroborated by overlaps ② and ③.

Larger clusters have more corroboration.

eg.

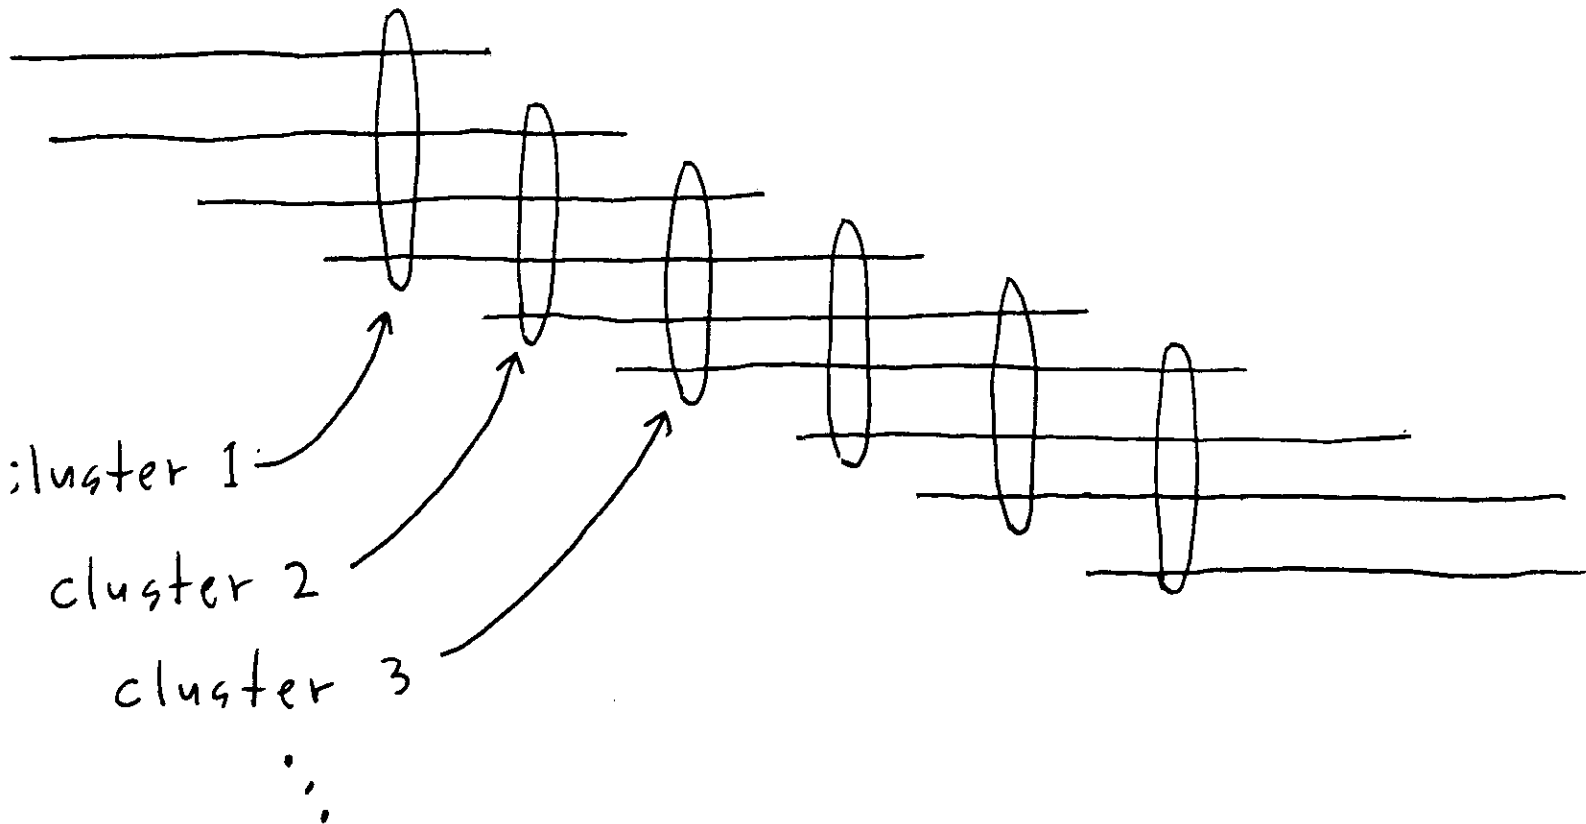


Overlap ① is corroborated by overlaps ② and ③, and by overlaps ④ and ⑤

Observation 2:

Average cluster size
= Average depth of clone coverage.

eg.

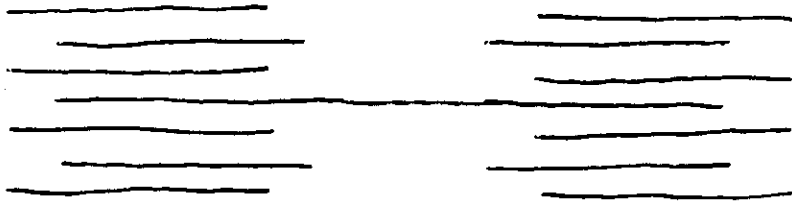


Cluster Proximity

Single Linkage:

Two clusters are close if they have a clone in common.

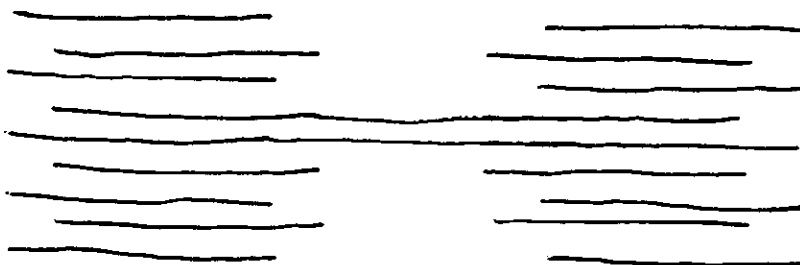
eg.



error
prone

Double Linkage:

Two clusters are close if they have two clones in common.

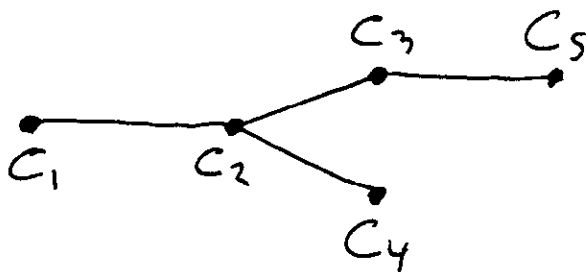
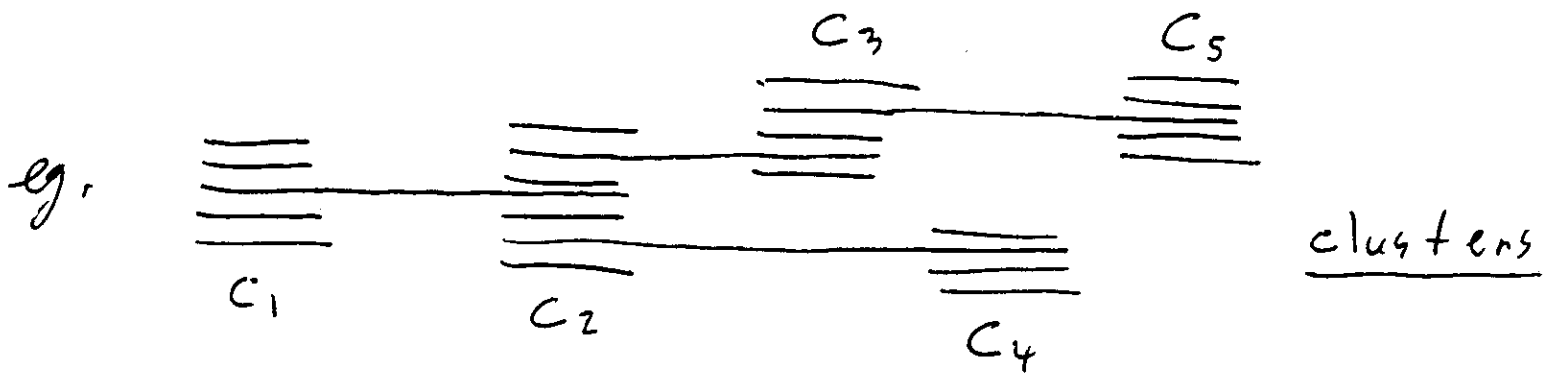


much less
error
prone

Many kinds of linkage are possible.

Cluster Graphs

- Each node is a cluster.
- An edge between two nodes means the two clusters are close.



single-linkage
cluster graph.

Map - Assembly Phases

- (1) Overlap (cluster formation)
- (2) Linkage (graph manipulation)
- (3) Ordering (cluster ordering)

Example 1

- Using all sts content data from

- Whitehead / MIT
- Ceph / Genethon

with the double-linkage strategy,
the proximity graph has

- 1,463 nodes
- 11,017 edges.

- Overall graph structure:

- Many connections between sts probes on different chromosomes.
- 1 large graph instead of 23 smaller ones.
- \hat{u} , a big mess.

... graph of the ...



Example 1

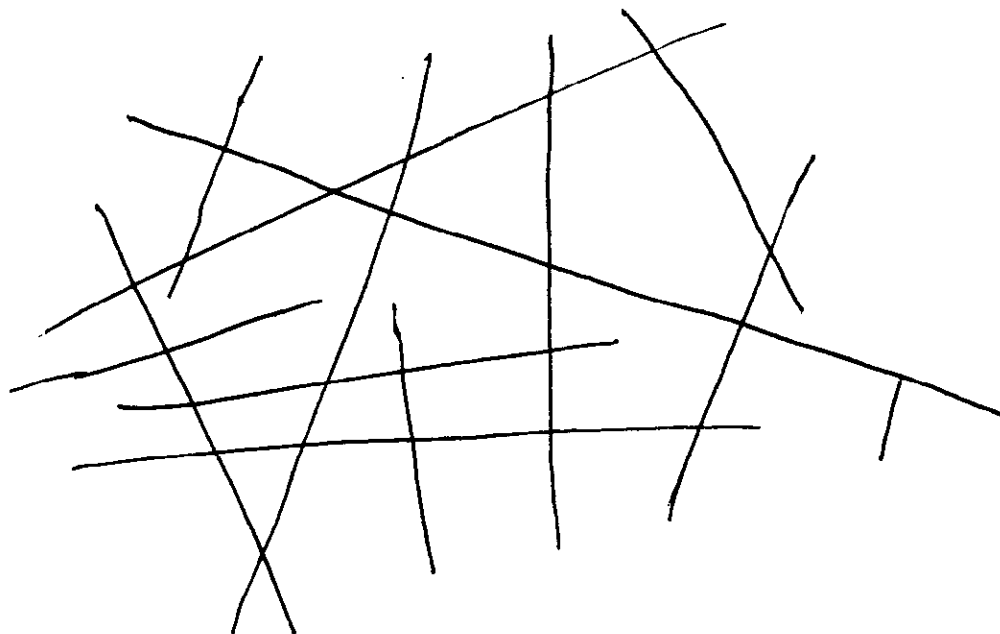
1,463 nodes

11,017 edges

Graph Abstraction (Simplification)

Although the graph is messy, it has a simple "piecewise linear" structure.

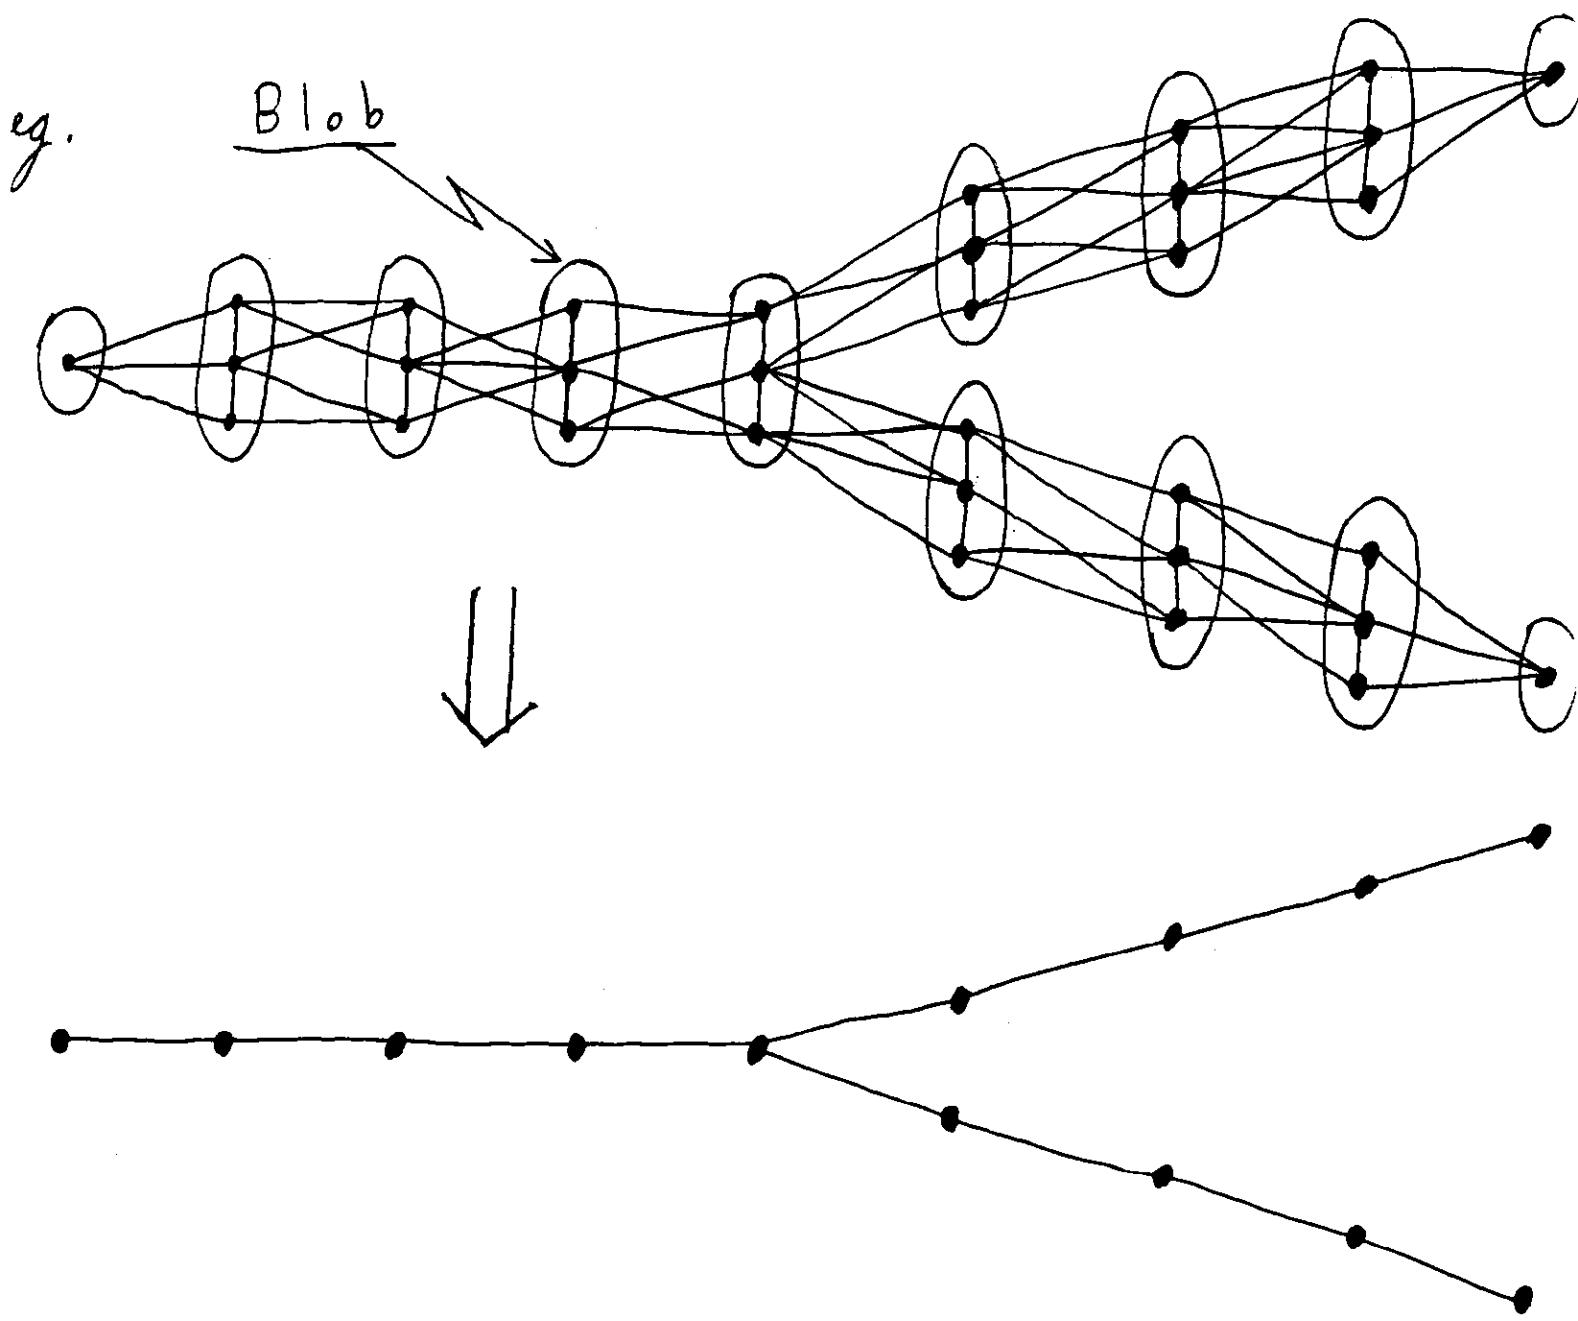
eg.

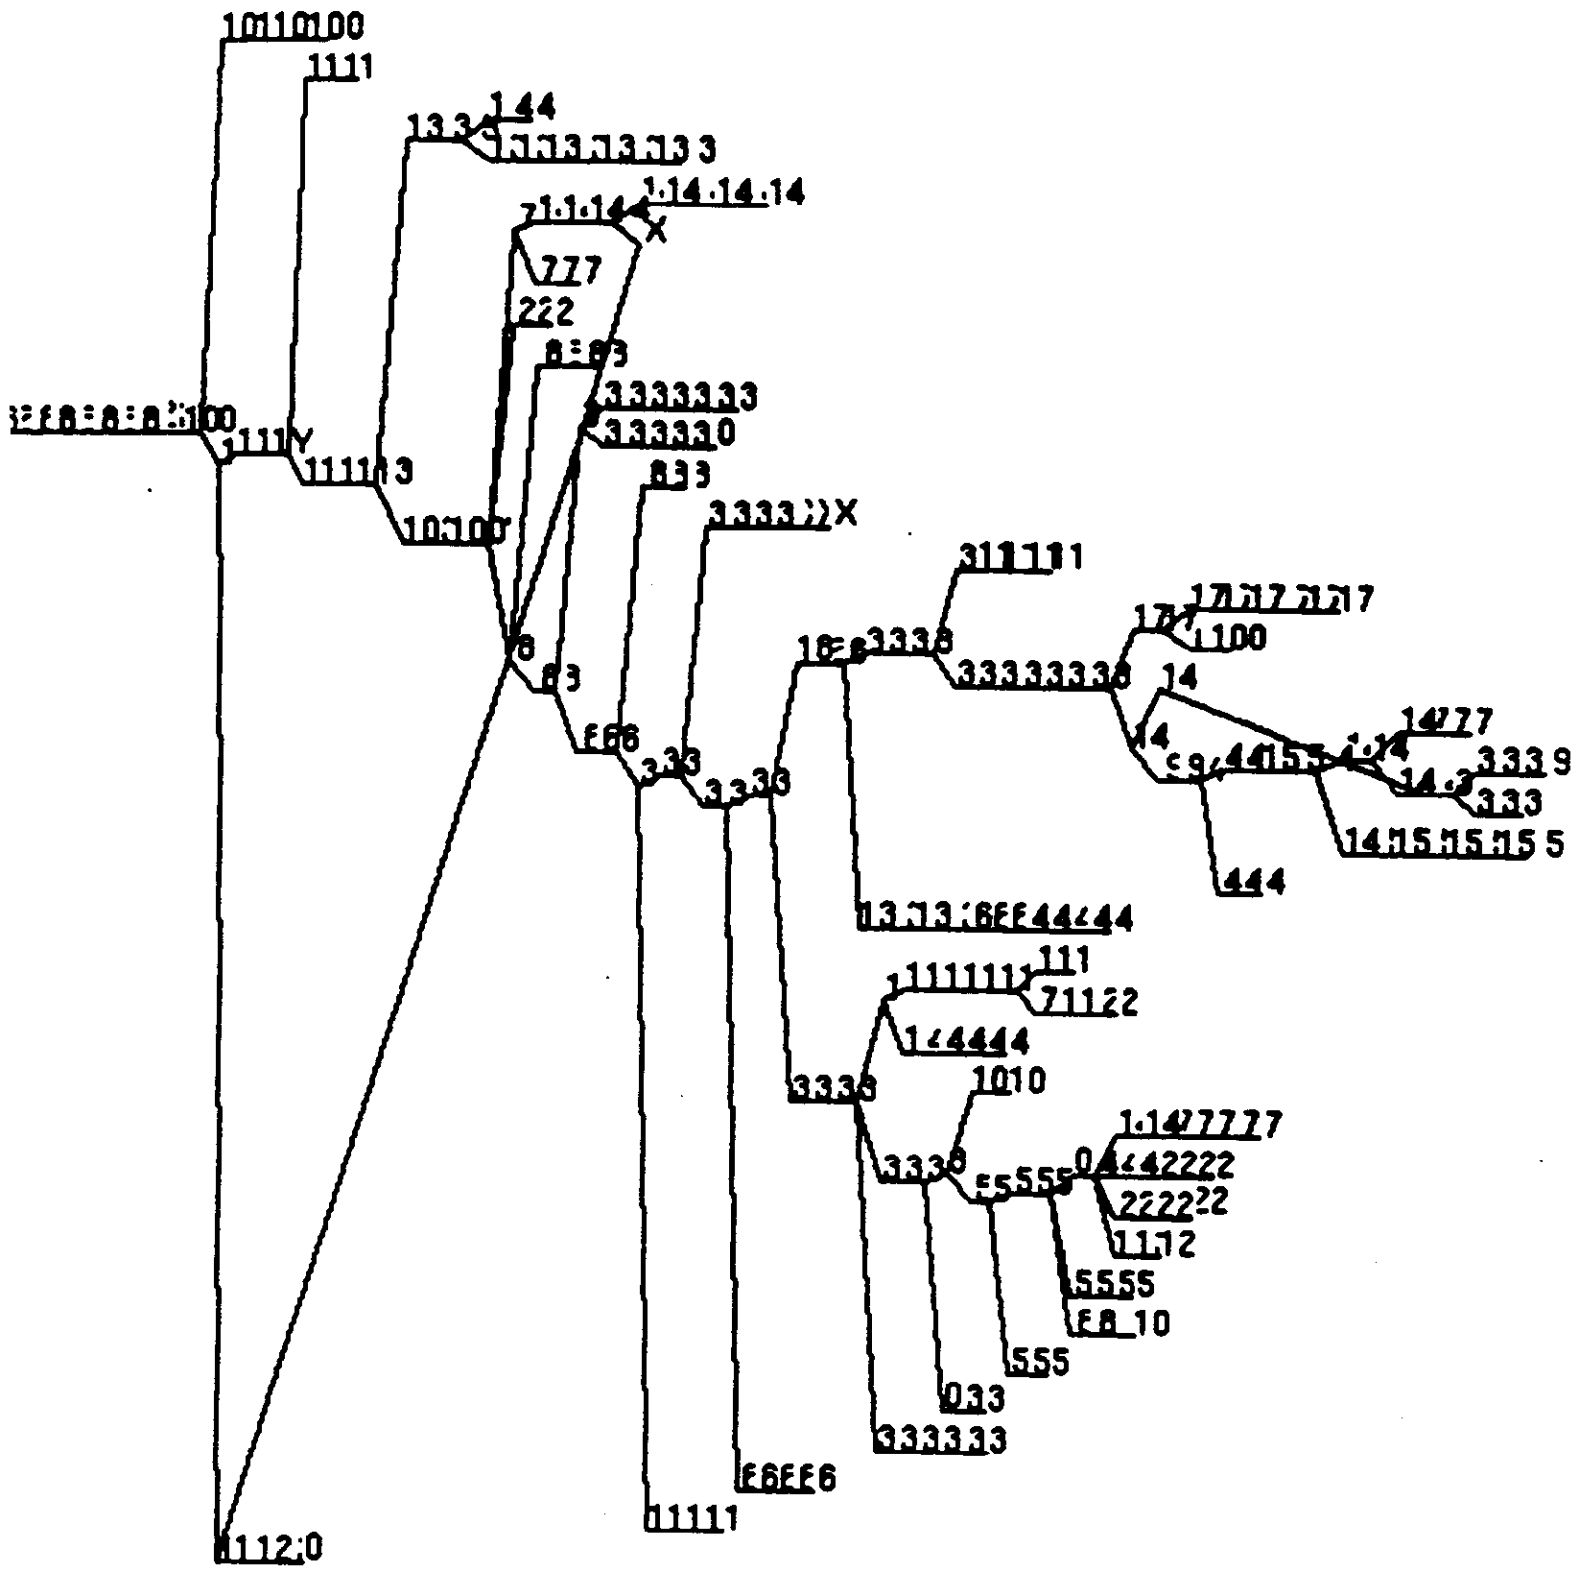


Next Step: Extract and display this structure.

Graph Abstraction

Idea: Coalesce groups of neighboring nodes into single nodes.





Abstracted sts graph of the human genome.
 (301 nodes, 302 edges)

Example 2

Using all human mapping data from

- Whitehead/MIT
- Ceph/Genethon

including

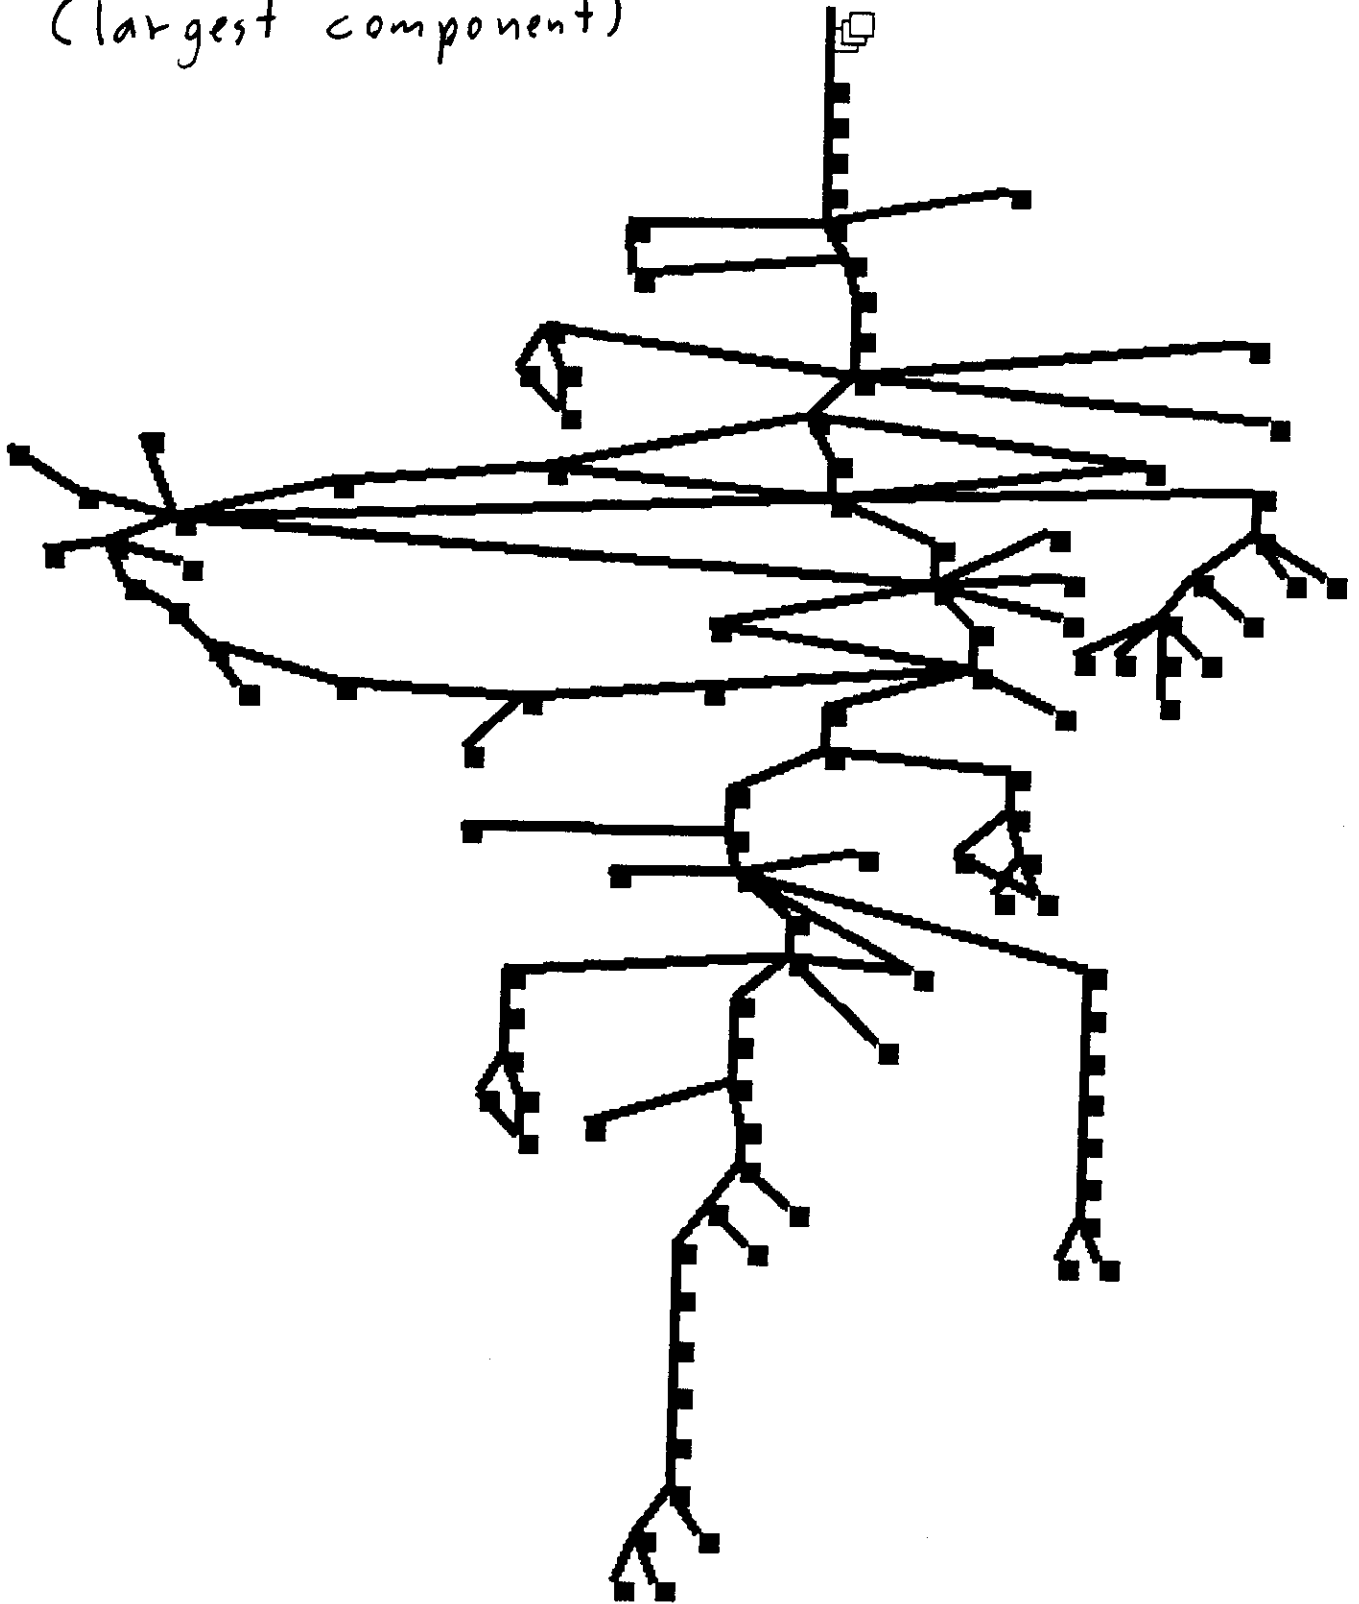
- sts content data,
- Fingerprint data,
- ALU-PCR data,

the cluster graph has

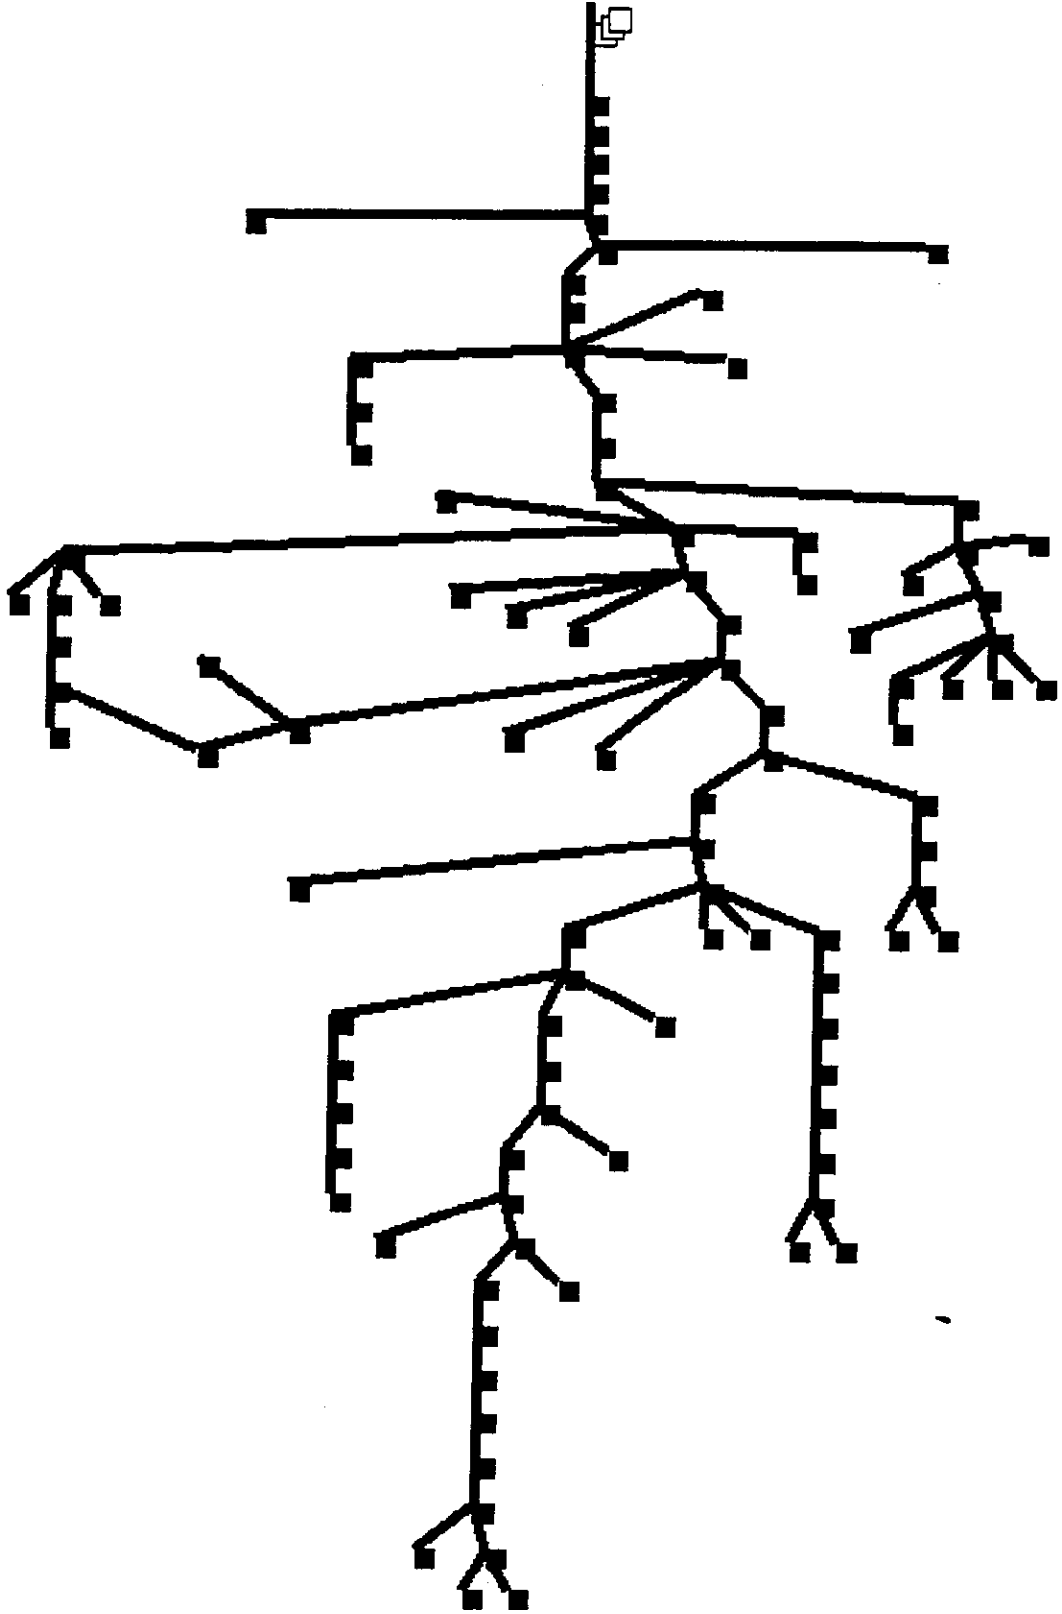
- 450,342 edges,
- 27,379 nodes.

Most edges (432,888) are concentrated in a single connected component.

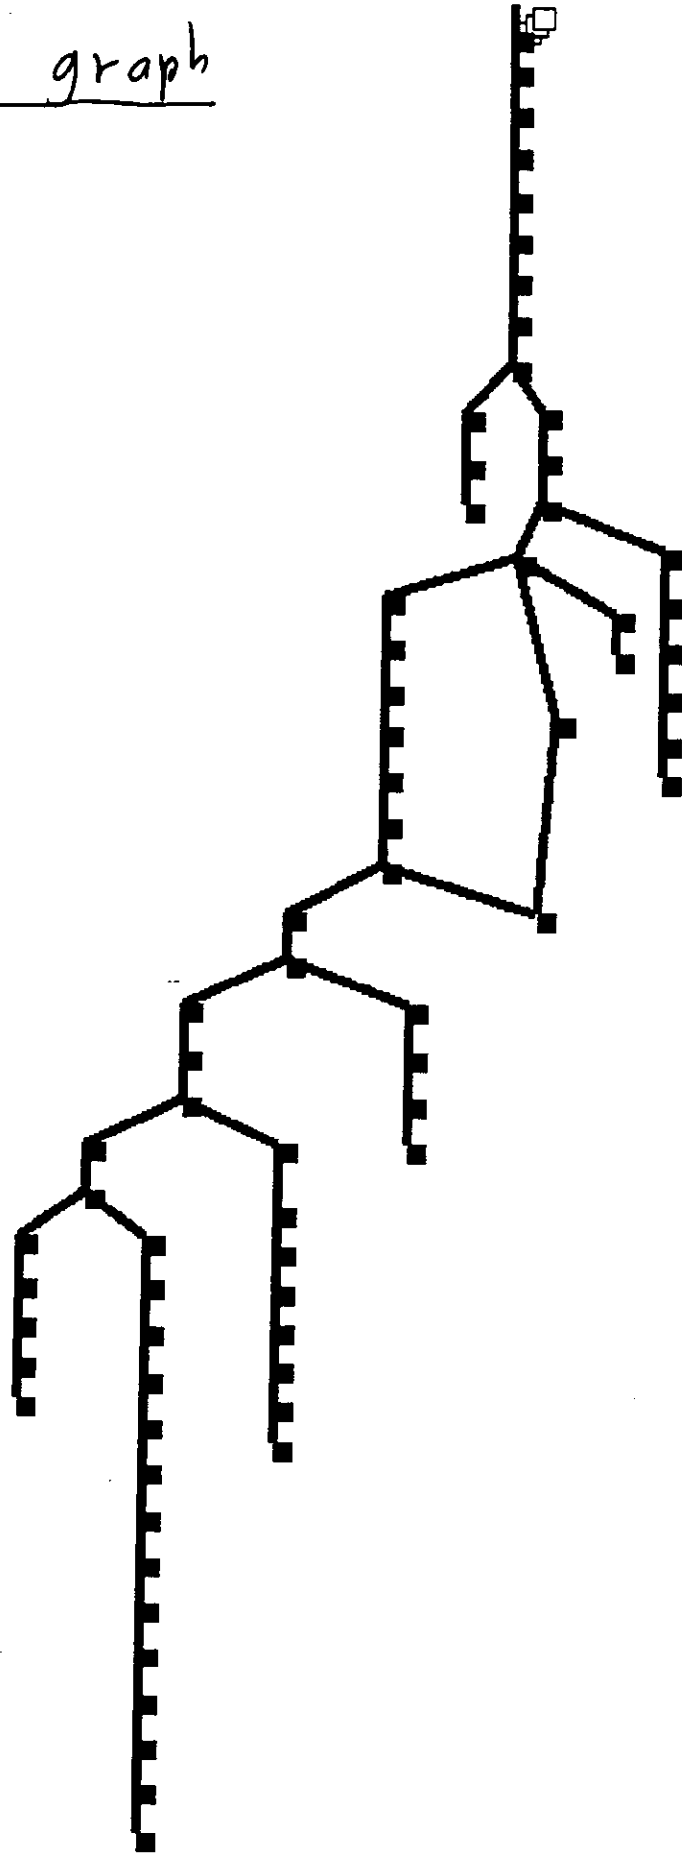
Abstracted cluster graph
of Human Chromosome 7
(largest component)

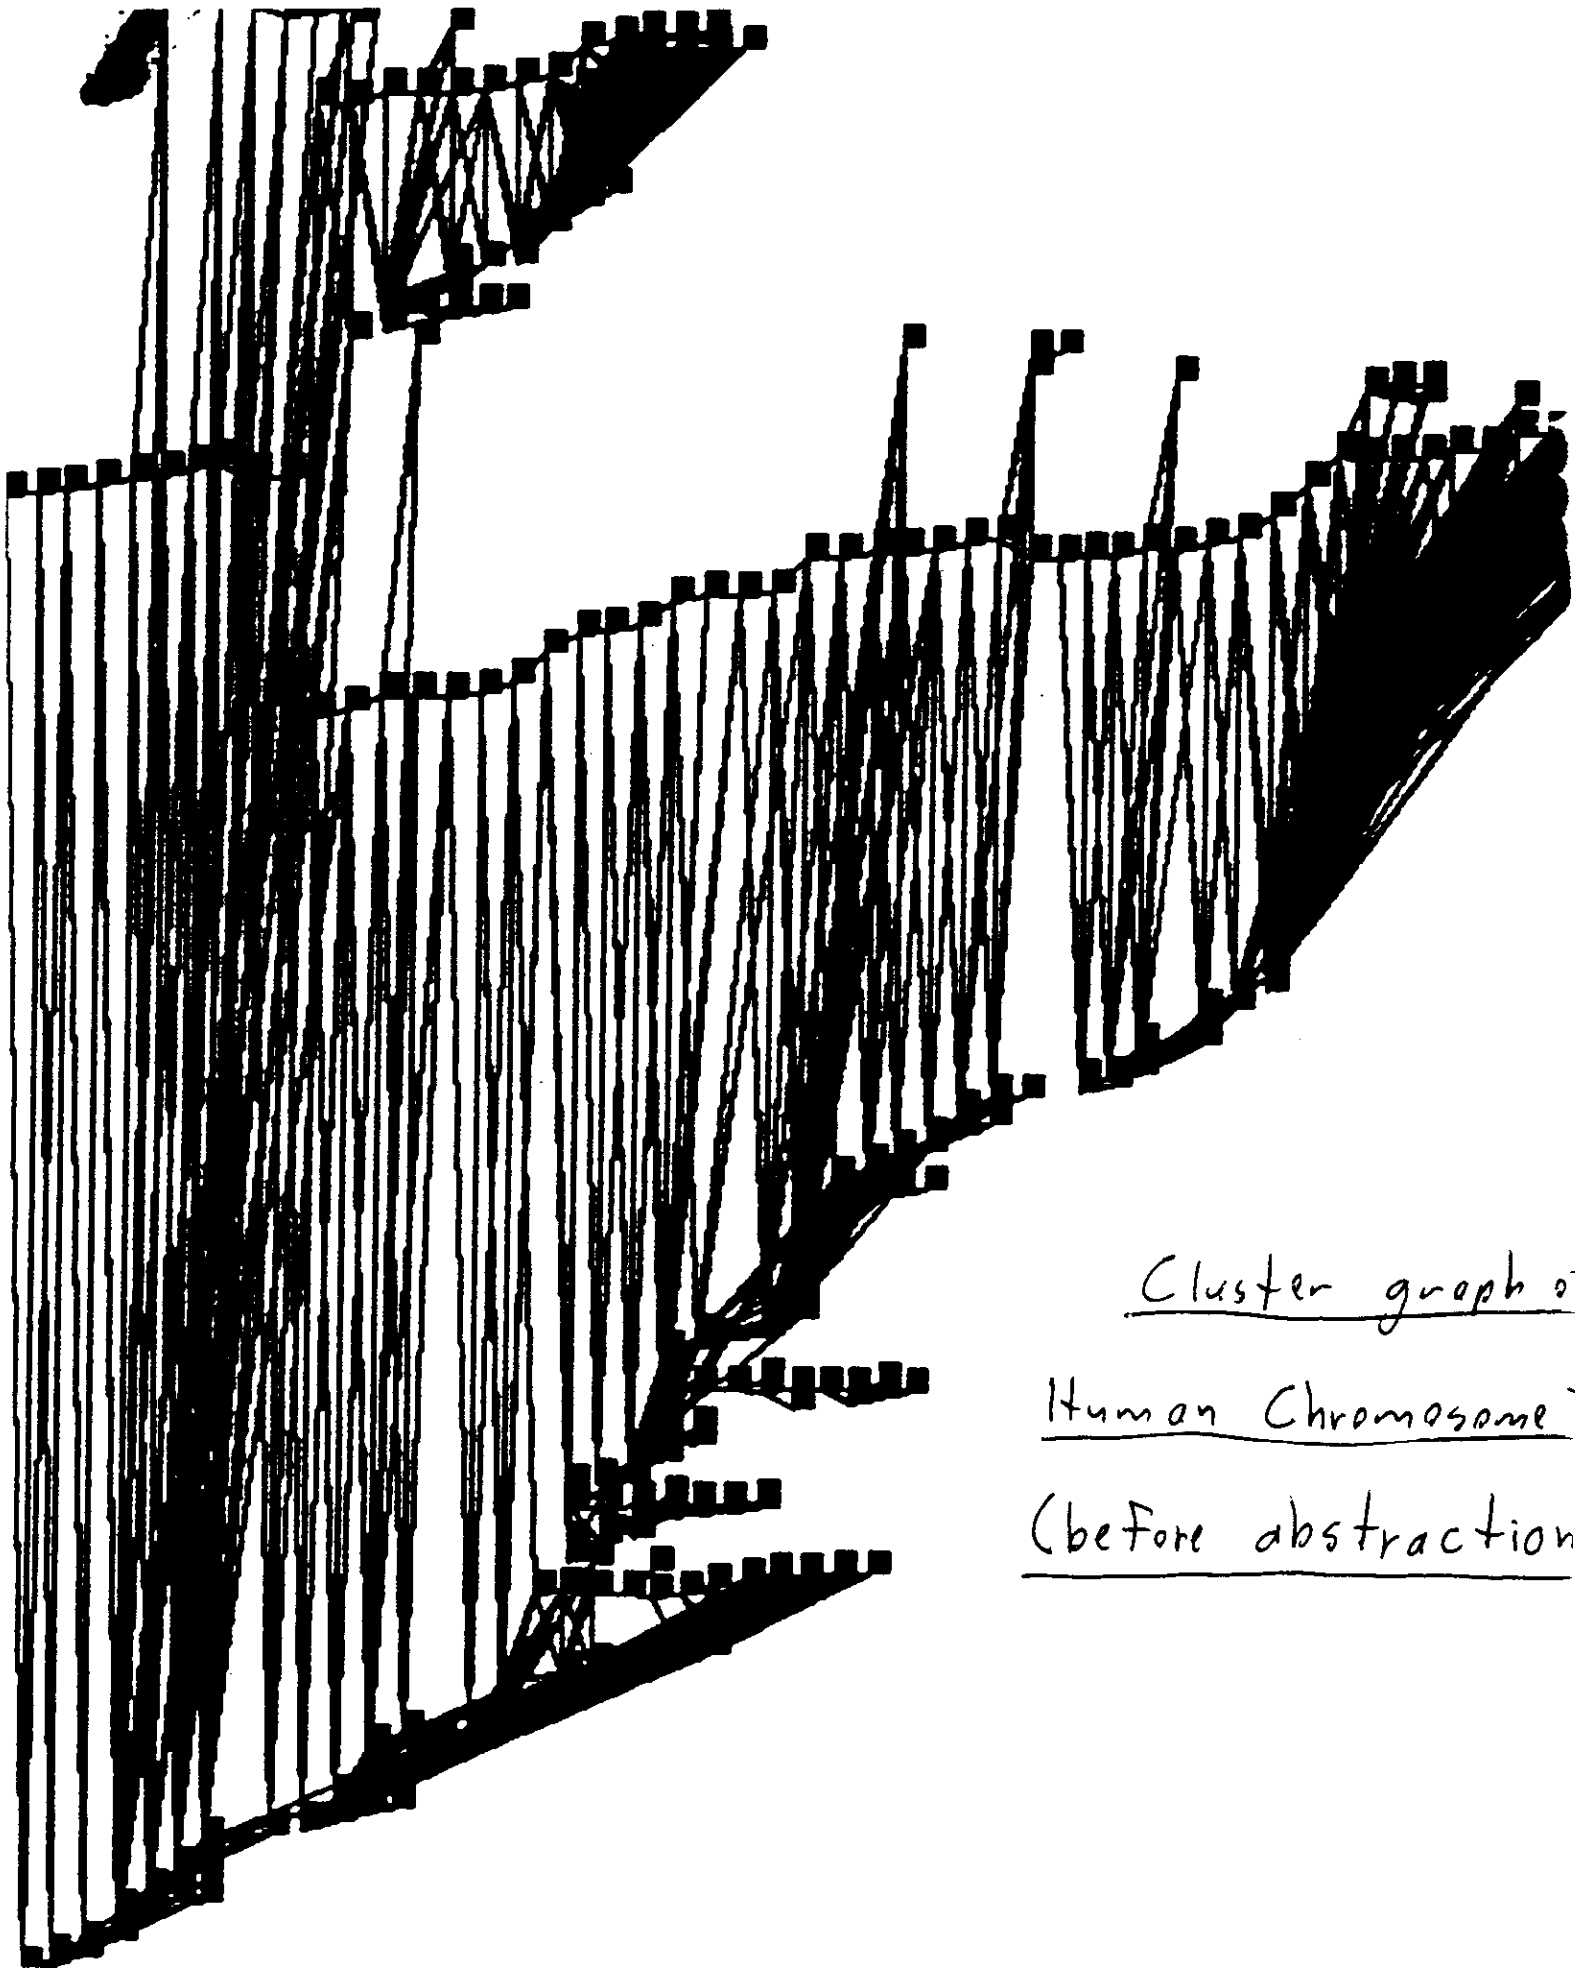


Partially trimmed graph



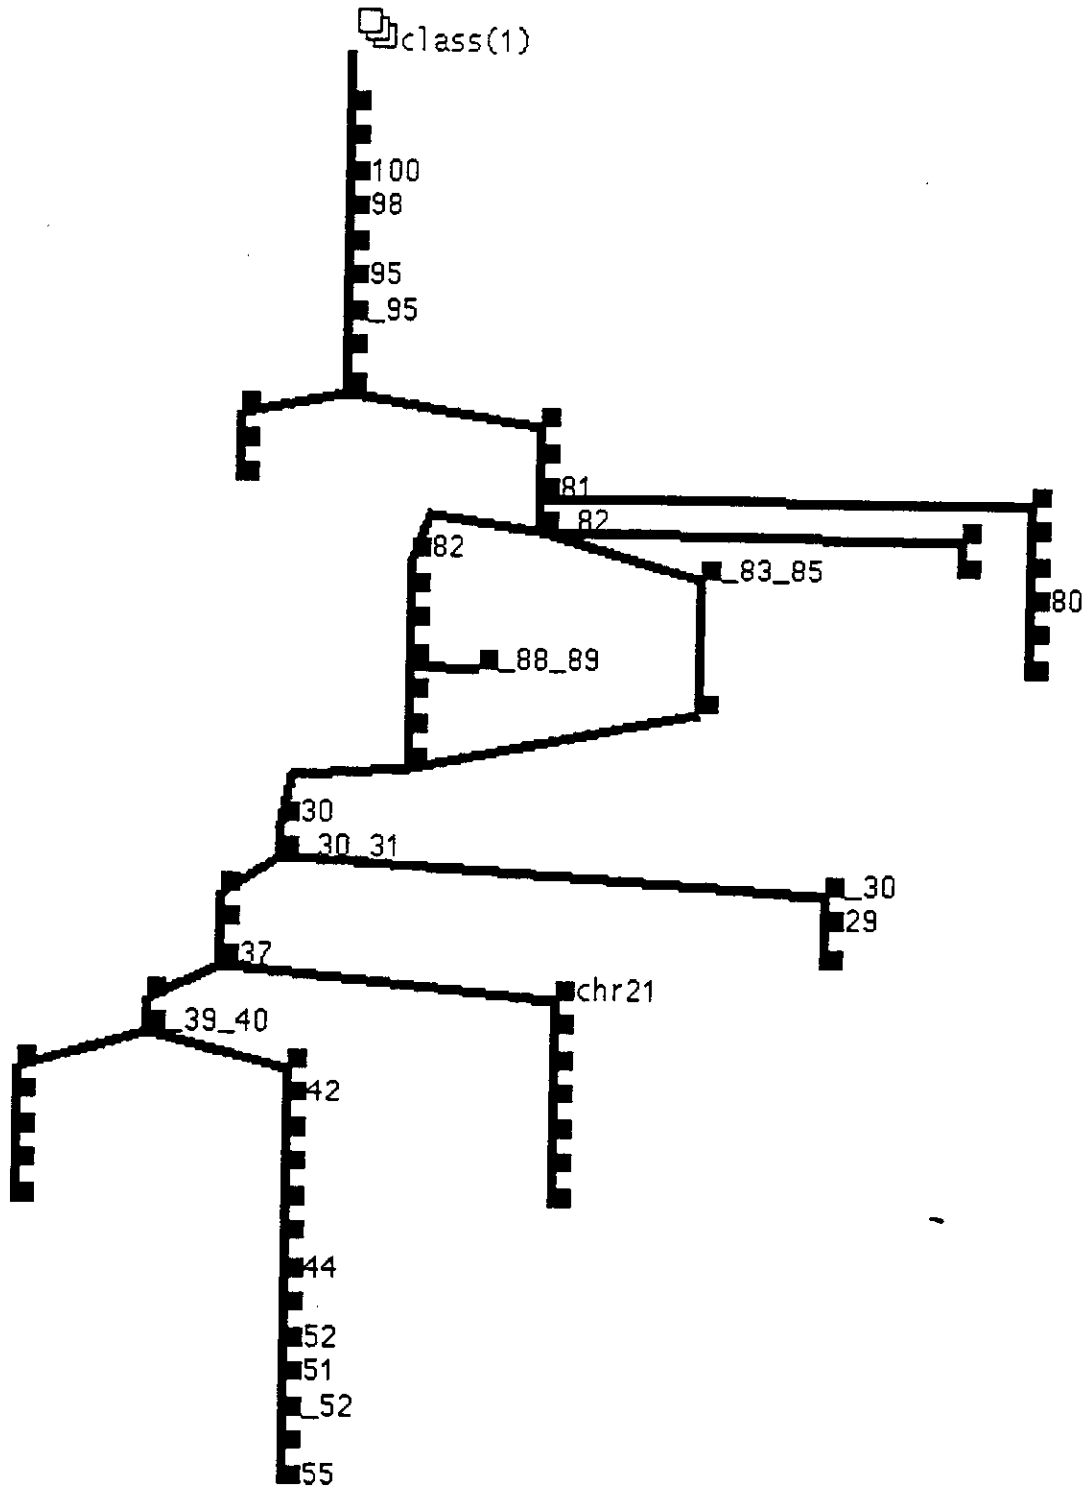
Fully trimmed graph



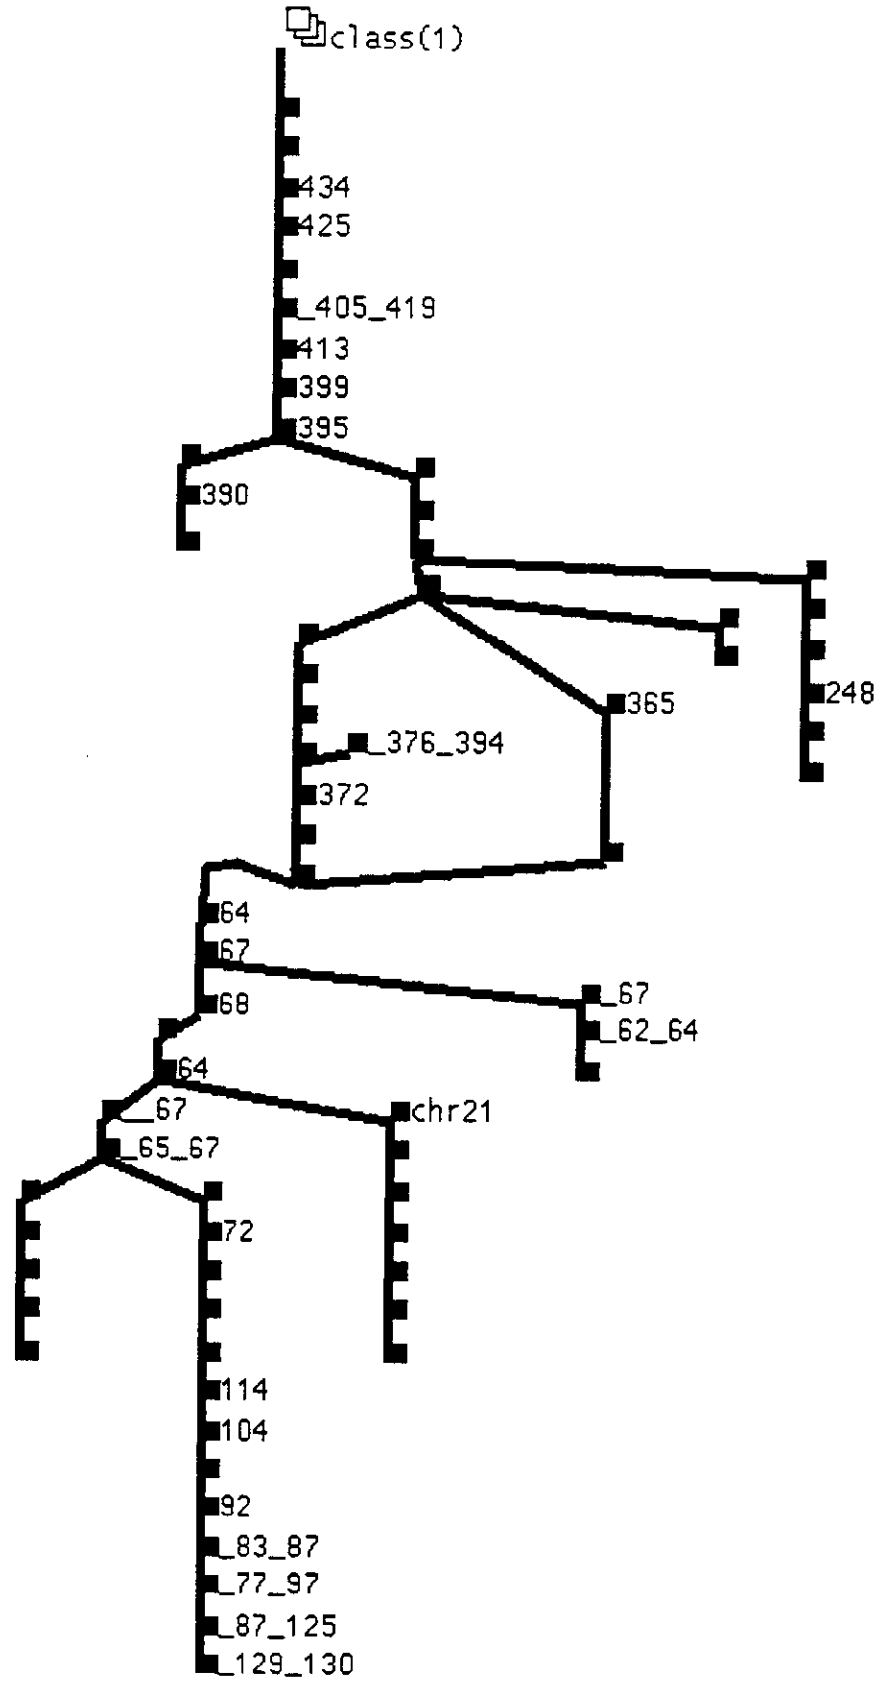


Cluster graph of
Human Chromosome
(before abstraction)

Graph of Chromosome 7 with genetically-mapped positions

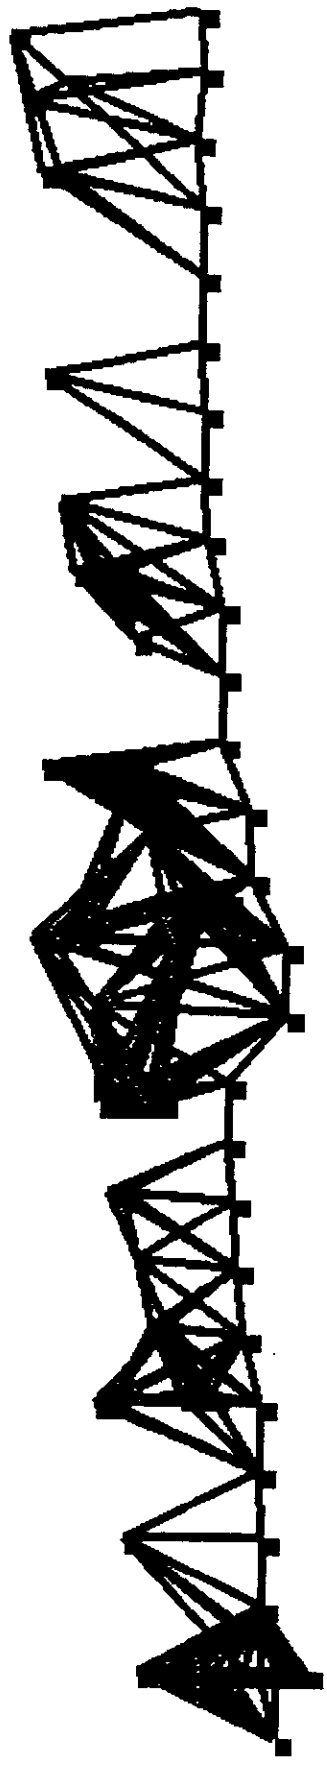
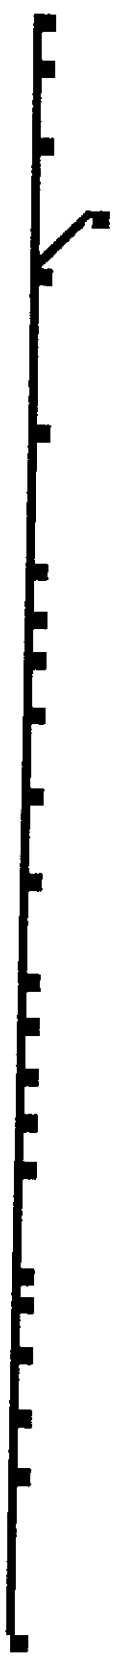


Graph of Chromosome 7 with radiation-hybrid positions



Contig of Human Chromosome 7

Abstracted
cluster
graph



Detailed
cluster
graph

Cluster Graphs:

- Many forms of physical-mapping data determine whether two clones overlap.
- A cluster is a maximal set of mutually-overlapping clones.
- A cluster represents a point on the genome.
- clusters filter out many false overlaps.
- clusters reduce map-assembly to three phases:

overlap

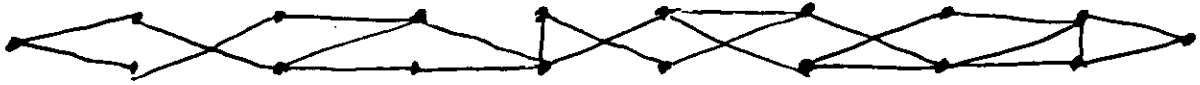
linkage

ordering

Errors and Anomalies

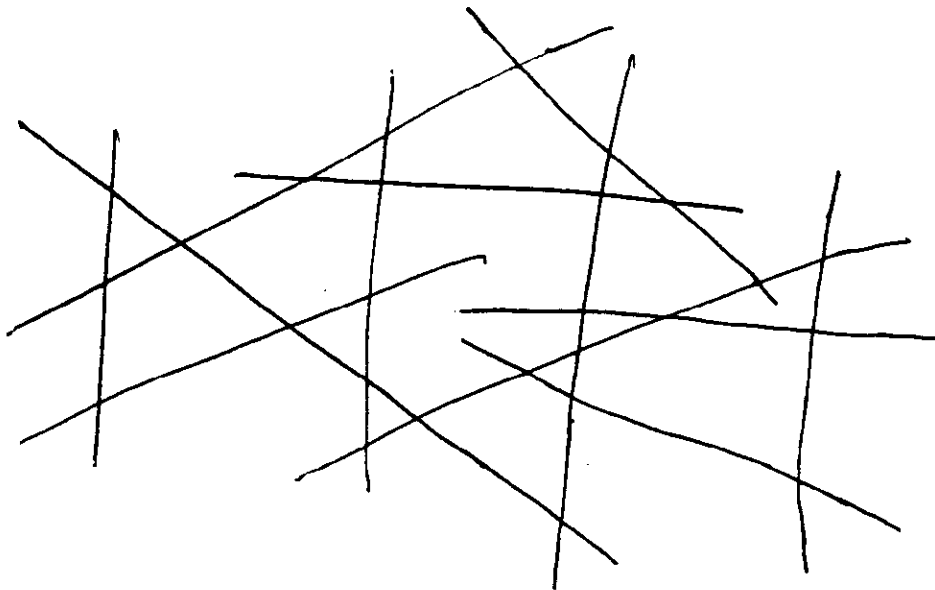
- Given ideal data, a cluster graph would be "nearly linear," i.e., long & thin (in fact, an interval graph).

eg.

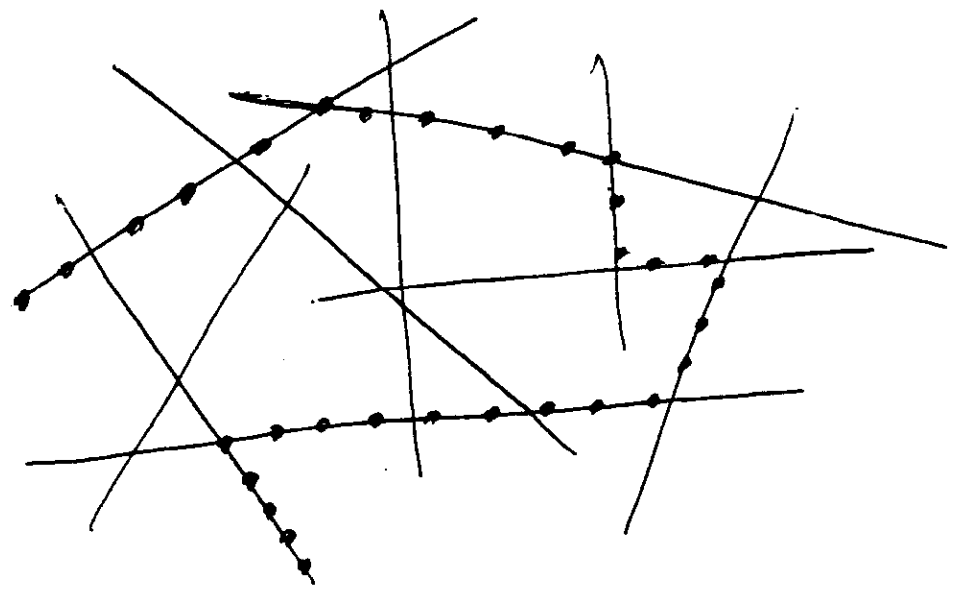


- However, because of experimental error (mainly false positives), the graphs are "piecewise linear".

eg.



- This piecewise linear structure can be automatically extracted and displayed.
- Using high-level mapping data (eg, chr. assignments, genetic maps, rh maps, long linear subgraphs can be extracted

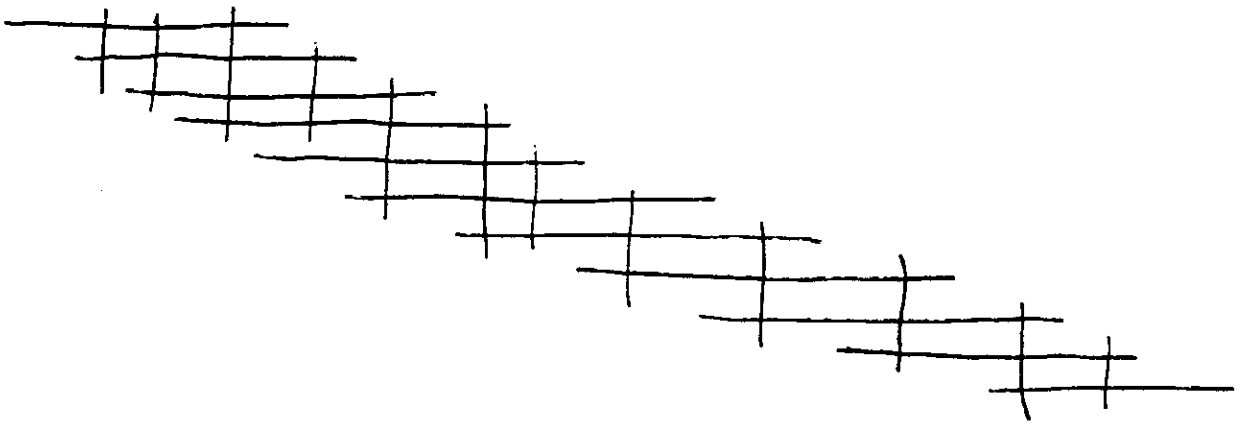


- These subgraphs represent genomic contigs.

Map Assembly

- These "nearly linear" subgraphs are the input to an algorithm that generates a physical genome map, i.e., an ordering of clones & sts's.

eg,



i.e., Cluster graphs help to clean up data and remove anomalies.

Our Publications

Available on the Web, at

www.cs.toronto.edu/~bonner

under three categories:

- Genome Mapping
- Sequence Databases
- Laboratory Workflow