Creating Aggregate Objects from Normalised Data We'll get some nulls. (An additional note)

David J. Greaves

(with grateful thanks to Timothy G. Griffin)

Computer Laboratory University of Cambridge, UK

Michaelmas Term, 2022-23

djg11 (cl.cam.ac.uk)

Creating Aggregate Objects from Normalised

1A DB 2022 1/8

TGG's Thought Experiment

- Let's do a thought experiment.

- In the next few slides imagine that we intend to use a relational database to store read-optimised tables generated from a set of write-optimised tables (that is, having little redundancy).
- We will encounter some problems that can be solved by representing our data as aggregate/semi-structured objects.

Start with a simple relationship ...



S is related to T by relation R:

S	R	Т
	<u>A</u> <u>B</u> Y	B 7
<u>A</u> X	a1 b1 y1	\underline{D} \underline{Z}
a1 x1	a1 b2 y2	DI ZI
a2 x2	a1 b3 y3	02 22 b2 72
a3 x3	a2 b1 y4	03 23 b4 74
	a2 b3 y5	U4 Z4

Imagine that our read-oriented applications can't afford to do joins!

djg11 (cl.cam.ac.uk)

Creating Aggregate Objects from Normalised

Implement the relationship as one big table?

BigTableOne: An out	er joi	n of	<i>S</i> , <i>R</i>	l, and	<i>T</i> b	
	<u>A</u>	Х	B	Ζ	Y	
	a1	x1	b1	z1	y1	
	a1	x1	b2	z2	y2	
	a1	x1	b3	z3	уЗ	
	a2	x2	b1	z1	y4	
	a2	x2	b3	z3	y5	
	a3	xЗ				
			b4	z4		

- Since we don't update this data, there will be no problems associated with redundancy.
- o Notice the null values starting to pop up.

However, there may be many further relationships ...



S is further related	further related to T by relation Q: S Q T \underline{A} X \underline{A} \underline{B} W $\underline{a1}$ $x1$ $a1$ $b4$ $w1$ $a2$ $x2$ $a3$ $b2$ $w2$ $a3$ $x3$ $a3$ $b3$ $w3$				
S A X a1 x1 a2 x2 a3 x3	Q A B W a1 b4 w1 a3 b2 w2 a3 b3 w3	T B Z b1 z1 b2 z2 b3 z3 b4 z4			

1A DB 2022 5/8

< 6 b

Implement with another big table?

BigTableTwo: An out	ər joi	n of	<i>S</i> , <i>G</i>	, and	d <i>T</i>
	<u>A</u>	Х	<u>B</u>	Ζ	W
	a1	x1	b4	z4	w1
	a3	xЗ	b2	z2	w2
	a3	xЗ	b3	z3	w3
	a2	x2			
			b1	z1	

Having two tables makes reading a bit more difficult!

A

- B

Combine into one big table?

BigTable: Derived	from	<i>S</i> , <i>F</i>	R, Q,	and	Τ		
	A	Х	B	Ζ	Y	W	
	a1	x1	b1	z1	y1		
	a1	x1	b2	z2	y2		
	a1	x1	b3	z3	уЗ		
	a1	x1	b4	z4		w1	
	a2	x2	b1	z1	y4		
	a2	x2	b3	z3	y5		
	a3	xЗ	b2	z2		w2	
	a3	xЗ	b3	z3		w3	

Note that our big tables use a **composite key**.

Problems with BigTable

- We could store BigTable and speed up some queries.
- Note: if we'd also stored further attributes of *A* in *S*, that data would be replicated |*B*| times.
- But suppose that our applications typically access data using either *S*'s key or *T*'s key.
- Creating indices on the A and B columns could speed things up, but our applications may still be forced to gather information from many rows in order to collect all information related to a given key of S or a given key of T.
- It would be better to access all data associated with a given key of S or a given key of T using only a single database lookup.

Potential Solution

- Represent the data using aggregate forms.
- This example was highly structured, but the aggregate forms also support semi-structured objects.