

#### The Well-Tempered Elephant

Gianni Ciolli Global VP, HA Practice Lead PGConf.De 2025 – Berlin, 8-9 May Who are you?



#### PostgreSQL 16 Administration Cookbook

olve real-world Database Administration hallenges with 180+ practical recipes and best practices



#### Gianni Ciolli

• Global VP, Practice Lead HA

**EDB** Postgres for the AI Generation

- Co-founder of **2ndQuadrant**
- Author of the **PostgreSQL** Administration Cookbook
- Contributor to PostgreSQL (Hot Standby) and related tools such as PGD, Barman, TPA, repmgr

# What will you present?



- The Well-Tempered Clavier
- What is a "fugue"?
- Why PostgreSQL?
- What and how?
- Really?



### The Well-Tempered Clavier

- Published by Johann Sebastian Bach
- The **most important** piano music
- Two books: in 1722 and in 1742
- In total, 48 Prelude and Fugue pieces, numbered 846-893 in the BWV catalog
- 48 = 2 (books) x 12 (keys) x 2 (modes)



### The fugues in the Well-Tempered Clavier

- Fugues follow a precise format and can be analyzed with a software
- Kyle Rother (University of Cape Town) transcribed the 48 fugues in **digital format** (Lilypond)
- We import these fugues into PostgreSQL
- Lilypond parser moved to a separate ly2pg extension



### What is a Fugue?

### What is a "fugue"?

- A piece of music composed by 2+ concurrent **voices** 
  - Each voice is a sequence of **notes**
  - Standard voice names, related to the type of singer

2		Terminal - gciolli@l	aptop680:	~/git/pgwtc				× I I X
gciolli=#	SELECT	DISTINCT(vox)	FROM	pgwtc.notes	ORDER	BY	vox	DESC;
vox								
soprano								
alto								
mezzo								
tenor								
bass								
(5 rows)								
	_							
gciolli=#								



### Music Quickstart #1

- There are 7 notes: C D E F G A B (Italian: Do Re Mi Fa Sol La Si)
- They repeat endlessly above and below
- Notes are drawn as circles on a set of lines (the "score")
- Symbols for silence ("rest")
- Stems indicate duration





### What is a "fugue"?

- A piece of music composed by 2+ concurrent **voices** 
  - Each voice is a sequence of notes
- The fugue begins with one voice alone, playing the "subject"
- The other voices join later, one at the time
  - They also start with the same subject
- The subject occurs in all the voices several times



### We want an example!

### Music Example #1

• This is from BWV 846





### Music Example #1

• This is from BWV 846, with two subject occurrences





### Music Quickstart #2

• Notes are placed on lines and on the spaces between



• In a sequence of note, **intervals** are important





### Music Example #2

• This is from BWV 849





### Music Example #2

• This is from BWV 849, with three subject occurrences



### Music Quickstart #2

• The start of the fugue



• The subject





### Music Quickstart #2

• The start of the fugue



- The subject
  - Intervals: -1 +3 -1 -1





### Where is PostgreSQL?

### Well-Tempered Elephant, a.k.a. pgwtc

- Use **PostgreSQL** to analyse fugues
- The WTC is the obvious first ambition, but pgwtc should be usable for other fugues too
- Load notes in a Postgres table
  - Fugues have between 2 and 5 voices
  - Slightly more than 1k notes per fugue



### Example: notes (pretty view)

2_			Ter	minal - gciolli@laptor	680: ~/git/pgwtc Screenshot_2		ted)-5.0 AGE_coOurX	
gcio	gciolli=# SELECT * from pgwtc.notes_pretty							
WHER	E src = 'B'	WV846'	AND vo	x = 'alto'				
ORDER BY ord LIMIT 10;								
id	src	vox	ord	lilypond	lilypond_full	initio	durations	
	++	+	· +	+	•	++		
20	BWV846	alto	1	r	r8	1:1.000	{8}	
22	BWV846	alto	2	c'	c'8	1:1.500	{8}	
24	BWV846	alto	3	d '	d ' 8	1:2.000	{8}	
26	BWV846	alto	4	e'	e'8	1:2.500	{8}	
28	BWV846	alto	5	f'	f'8.	1:3.000	{8.}	
30	BWV846	alto	6	g'	g'32	1:3.750	{32}	
32	BWV846	alto	7	f'	f'32	1:3.875	{32}	
34	BWV846	alto	8	e'	e'8	1:4.000	{8}	
36	BWV846	alto	9	a'	a'8	1:4.500	{8}	
40	BWV846	alto	10	d '	d ' 8	2:1.000	{8}	
(10 )	rows)							
gcio	lli=#							



### Example: notes (the actual table)

	Terminal - gciolli@la	ptop680: ~/git/pg	wtc	[term1] (exported)-3.0 (RGB colour & it gir 🖬 🗙 r			
gciolli=# SELECT * from pgwtc.notes							
WHERE src = 'BWV846'	AND vox = 'alto	o '					
ORDER BY ord LIMIT 1	0;						
id   src   vox	ord   nota	start	ticks	durations			
++	+++	+ +	+				
20   BWV846   alto	1   (,0)	0	48	{8}			
22   BWV846   alto	2   (35,0)	48	48	{8}			
24   BWV846   alto	3   (36,0)	96	48	{8}			
26   BWV846   alto	4   (37,0)	144	48	{8}			
28   BWV846   alto	5   (38,0)	192	72	{8.}			
30   BWV846   alto	6   (39,0)	264	12	{32}			
32   BWV846   alto	7   (38,0)	276	12	{32}			
34   BWV846   alto	8   (37,0)	288	48	{8}			
36   BWV846   alto	9   (40,0)	336	48	{8}			
40   BWV846   alto	10   (36,0)	384	48	{8}			
(10 rows)							
				5			
gciolli=#							



## Why PostgreSQL?

### Why PostgreSQL?

- Database **tables** to **store** music easily
- **Views** to display music in a friendly way
- Window functions to aggregate sequences of notes
- **Custom operators** to simplify the interface
- Extensions to facilitate reuse of the above tools for the analysis of other fugues



### Example: custom operators

2	No. No.	Same 13	Termin	al - gciolli@lap	top680: ~/git/pg	wtc (led) 8.6	(RGB colour 8-bit	gamma mteger,	GTMP BUILT IN SI	RGB, TlayerA 1128/D3 X 1
gciolli=#	ciolli=# SELECT # nota									
, nota ##	nota ## durations									
, start @	start @ tempo('4/4')									
, *										
FROM pgwtc	.notes									
WHERE src	= 'BWV846'	AND vox = 'al	lto'							
ORDER BY O	rd LIMIT 10	;								
?column?	?column?	?column?	id	src	vox	ord	nota	start	ticks	durations
	+	+	+ +		++	·+		+	+	
r	r8	1:1.000	20	BWV846	alto	1	(,0)	0	48	{8}
c '	c'8	1:1.500	22	BWV846	alto	2	(35,0)	48	48	{8}
d '	d'8	1:2.000	24	BWV846	alto	3	(36,0)	96	48	{8}
e'	e'8	1:2.500	26	BWV846	alto	4	(37,0)	144	48	{8}
f'	f'8.	1:3.000	28	BWV846	alto	5	(38,0)	192	72	{8.}
g'	g'32	1:3.750	30	BWV846	alto	6	(39,0)	264	12	{32}
f'	f'32	1:3.875	32	BWV846	alto	7	(38,0)	276	12	{32}
e'	e'8	1:4.000	34	BWV846	alto	8	(37,0)	288	48	{8}
a'	a'8	1:4.500	36	BWV846	alto	9	(40,0)	336	48	{8}
d '	d'8	2:1.000	40	BWV846	alto	10	(36,0)	384	48	{8}
(10 rows)										2

200

JU I DRIVUTO I ULLU

gciolli=#

1 1 1 2 1 2 1

### Listing Subjects

- A fugue begins with one voice alone, playing the "subject"
- How long is the subject?
   We encode that information in the metadata
- Recursive query starting from the first note



### Finding Subjects

	mind - geoditoria subjects.sql - GNU Emacs at laptop680 [ed]-11.	0 (RGB colour 8-bit gamma integel GIMP built-in sRGB, 1 layer) 834×323 -🖬 II 🗙
File Edit Options Buffers Tools SQL Help		
📄 🖹 🗶 🏝 Save 🤙 Undo		
CREATE VIEW subjects AS	UNION ALL	, aggregated_subject_ids AS (
WITH RECURSIVE first_notes AS ( SELECT DISTINCT ON (src) min(start) AS start , src , vox FROM pgwtc.notes WHERE (nota).tono IS NOT NULL GROUP BY src, vox ORDER BY src, min(start)	SELECT n.src , n.vox , n.id , n.ord , s.depth + 1 AS depth , n.nota , n.ticks FROM pgwtc.notes n JOIN pgwtc.metadata m USING (src) DOIN gwbict ide c	<pre>SELECT     src , vox , array_agg(id ORDER BY ord) AS ids , array_agg(nota ORDER BY ord) AS note , array_agg(ticks ORDER BY ord) AS rhythm FROM subject_ids GROUP BY src, vox ORDER BY src, vox )</pre>
<pre>), subject_ids AS (    SELECT     src    , vox    , id    , ord    , 1 AS depth    , n.nota    , n.ticks    FROM first_notes f    JOIN pgwtc.notes n    USING (src, vox, start) ]</pre>	<pre>OUN SUBJECT_105 S ON n.src = s.src AND n.vox = s.vox AND n.ord = s.ord + 1 WHERE s.depth &lt; m.subject_length ) </pre>	SELECT src , vox , start , tempo , clavis , ids , note , rhythm FROM aggregated_subject_ids s JOIN pgwtc.metadata USING (src) JOIN first_notes USING (src, vox) ORDER BY src, vox;
U: subjects.sql Top L27	( U: subjects.sql 36% L54 (SQL	U: subjects.sql Bot L81 (SQL[ANSI])
wrote /nome/gciolli/git/pgwtc/sub	jects.sql	1



### Example: subjects

2	Terminal - gciolli@laptop680: ~/git/pgwtc 👘 "Untitled -130 (RGB colour & bit gamma integer, GINA built 🖬 🗙 🕯
gciolli=# SELECT	<pre>FROM pgwtc.subjects_pretty WHERE src = 'BWV846' \gx</pre>
-[ RECORD 1 ]+	·
src	BWV846
vox	alto
initio	1:1.500
clavis	C
lilypond_voice	c'8 d'8 e'8 f'8. g'32 f'32 e'8 a'8 d'8 g'8 ~ g'16 a'16 g'16 f'16 e'16
gciolli=# SELECT -[ RECORD 1 ]	<pre>r * FROM pgwtc.subjects WHERE src = 'BWV846' \gx</pre>
STC   BWV846	
vox   allo	
tempo $  (4, 4)$	
clavis   C	
ids   {22,24.	.26.28.30.32.34.36.40.46.48.50.52.54}
note   {"(35.0	))","(36,0)","(37,0)","(38,0)","(39,0)","(38,0)","(37,0)","(40,0)","(36,
0)","(39,0)","(4	40,0)","(39,0)","(38,0)","(37,0)"}
rhythm   {48,48,	48,72,12,12,48,48,48,72,24,24,24,24
gciolli=#	



### Finding Subject Occurrences

- We build a list of **patterns** = for each subject, all possible forms in which that subject can appear
- Then we look for patterns in notes with a **recursive query** joining notes and patterns
- Relational power!



#### Subject entries in BWV 846







gciolli=#	SELECT s	rc,vox,initic	)				
,pattern_id,lilypond_voice							
ROM pgwtc.subject_occurrences_pretty							
WHERE src	= 'BWV84	6';					
src	vox	initio	pattern_id				
+		++	•••••••				
BWV846	alto	1:1.500	RØ				
BWV846	soprano	2:3.500	R4				
BWV846	tenor	4:1.500	R - 3				
BWV846	bass	5:3.500	R-7				
BWV846	soprano	7:1.500	R7				
BWV846	tenor	7:2.500	R-3				
BWV846	alto	9:1.500	R4				
BWV846	bass	10:3.500	R-10				
BWV846	alto	10:4.500	R1				
BWV846	tenor	12:1.500	R-5				
BWV846	alto	14:1.500	RØ				
BWV846	tenor	14:2.500	R-3				
BWV846	bass	15:1.500	R-10				
BWV846	soprano	16:2.500	R7				
BWV846	alto	16:3.500	R4				
BWV846	tenor	17:1.500	R-2				
BWV846	bass	17:3.000	R-6				
BWV846	tenor	19:1.500	R-2				
BWV846	alto	19:2.500	R2				
BWV846	soprano	20:4.500	R4				
BWV846	tenor	21:3.500	R-1				
BWV846	tenor	24:1.500	R-7				
BWV846	alto	24:3.500	R3				
(23 rows)							

Time: 652.507 ms gciolli=#



### Status and Limitations of version 0.1

- We import only 24 fugues (25499 notes)
  - Lilypond syntax is complex
  - Future ly2pg versions will parse the other 24
- We recognise some, but not all, the patterns in the WTC
  - We do not find all the occurrences yet







### Thank you!

# gianni.ciolli@enterprisedb.com https://2025.pgconf.de/feedback



©EDB 2025 – ALL RIGHTS RESERVED.