
Remembering Pat O'Neil



Philip Bernstein
Microsoft Research

18th HPTS
November 5, 2019

Who was Pat?

A database innovator who had a major role in developing some of the most popular data structures and algorithms used in today's database systems

In the beginning

- Pat started his career as a mathematician
- Ph.D., 1969, in combinatorial math from Rockefeller University
- Assistant Professor of Math at MIT, 1970-1972

Pat's Erdos number is one!

- Erdős, Paul; O'Neil, Patrick E.
On a generalization of Ramsey numbers.
Discrete Math. 4 (1973), 29–35.
- <https://mathscinet.ams.org/mathscinet/search/publdoc.html?co1=AND&pg1=INDI&pg2=INDI&r=1&s1=300269&s2=189017&sort=Newest&vfpref=html>

Moved to industry

- Worked on Model 204 database system at Computer Corporation of America
- Published its bitmap algorithm developed by Bill Mann
 - Published at 2nd HPTS 1987
 - 300 citations
- Implemented B-trees (mid 1980's)

Returned to Academia

- 1988 moved to U. Mass Boston
- 1994, he published **Database: Principles Programming and Performance**, with his wife, Betty O'Neil.

Escrow Transactions

- An improvement of Gawlick's IMS FP MMDB.
 - Verify $QOH \geq \text{constant}$
Modify $QOH = QOH - \text{constant}$
- Reuter: value of QOH is a range [low, high] accounting for in-flight increments
- O'Neil: a generalized abstraction (programming model) and detailed design.
- Presented at 1st HPTS in 1985.
- Patrick E. O'Neil: **The Escrow Transactional Method.** [ACM Trans. Database Syst. 11\(4\): 405-430 \(1986\)](#)
- 300 citations

LRU-K

- Keeps track of the times of the last K references to popular database pages
- Uses this to estimate the interarrival time of references to each of those pages
- Hence, identifies frequently used pages
- The LRU-K page replacement algorithm for database disk buffering
E.J. O'Neil, P.E. O'Neil, G Weikum - ACM Sigmod Record, 1993 - dl.acm.org
- 1200 citations

The Dangers of Replication

- Jim Gray, Pat Helland, Patrick E. O'Neil, Dennis E. Shasha:
The Dangers of Replication and a Solution. SIGMOD Conference 1996: 173-182
- If every server broadcasts all updates to every other server, the explosion in message count is intolerable
- 1600 citations

A Critique of ANSI SQL isolation levels

- At Microsoft, Hal Berenson wrote a memo to clear up confusion about isolation levels
- Pat was on sabbatical at Microsoft, read Hal's memo, and explained where the confusion came from: the ANSI SQL Standard.

A Critique of ANSI SQL isolation levels (cont'd)

- ANSI SQL-92 [MS, ANSI] defines Isolation Levels in terms of phenomena: Dirty Reads, Non-Repeatable Reads, and Phantoms. This paper shows that these phenomena and the ANSI SQL definitions fail to characterize several popular isolation levels, including the standard locking implementations of the levels. Investigating the ambiguities of the phenomena leads to clearer definitions; in addition new phenomena that better characterize isolation types are introduced. **An important multiversion isolation type, Snapshot Isolation, is defined.**
- Hal Berenson, Philip A. Bernstein, Jim Gray, Jim Melton, Elizabeth J. O'Neil, Patrick E. O'Neil: **A Critique of ANSI SQL Isolation Levels.** SIGMOD Conference 1995: 1-10
- 1150 citations

Log-Structured Merge Tree

- **The log-structured merge-tree (LSM-tree) 1996 ACTA INFORMATICA** Patrick O'Neil, Edward Cheng, Dieter Gawlick, Elizabeth O'Neil
- Developed while on sabbatical at DEC DB group, Mountain View
- Paper completed while at Microsoft
- Used in Bigtable, HBase, LevelDB, SQLite4, RocksDB, WiredTiger, Apache Cassandra, and many others
- 800 citations

Isolation Level Definitions

- ANSI isolation levels were defined with locking in mind
- Pat collaborated on developing implementation-independent definitions
- Atul Adya, Barbara Liskov, Patrick E. O'Neil:
Generalized Isolation Level Definitions. ICDE 2000: 67-78
- 200 citations

ORDPATH

- A decimal numbering of XML nodes that enable efficient storage and retrieval
- Patrick E. O'Neil, [Elizabeth J. O'Neil](#), [Shankar Pal](#), [Istvan Cseri](#), [Gideon Schaller](#), [Nigel Westbury](#):
ORDPATHs: Insert-Friendly XML Node Labels. [SIGMOD Conference 2004](#): 903-908
- Shipped in Microsoft SQL Server
- 550 citations

Making Snapshot Isolation Serializable

- Characterizes when nonserializable executions of applications can occur under SI.
- Apply this theory to demonstrate that the TPC-C benchmark application has no serialization anomalies under SI
- Alan Fekete, Dimitrios Liarokapis, Elizabeth J. O'Neil, Patrick E. O'Neil, Dennis E. Shasha: Making snapshot isolation serializable. ACM Trans. Database Syst. 30(2): 492-528 (2005)
- 350 citations

Web Search

Microsoft Search results for **patrick e. o'neil**

16,600,000 Results Any time ▾

Including results for **patrick e o'neill**.
Do you want results only for **patrick e. o'neil**?


Patrick O'Neil - Wikipedia
https://en.wikipedia.org/wiki/Patrick_E_O'Neil ▾
Patrick Eugene O'Neil is an American computer scientist, an expert on databases, and a professor of computer science at the University of Massachusetts Boston. O'Neil did his undergraduate studies at the Massachusetts Institute of Technology, receiving a B.S. in mathematics in 1963.
Born: 1942, Mineola, New York Died: September 20, 2019 (aged 77), ...
Fields: Computer Science Known for: Distributed Database, SQL Isolation, ...



Patrick O'Neal - IMDb
<https://www.imdb.com/name/nm0641929> ▾
Patrick O'Neal, Actor: Under Siege. Dark, dashing and coldly handsome with intense, penetrating eyes, Patrick O'Neal was known for walking that fine line between elegant heroics and elegant villainy during his five-decade career.

Patrick O'Neil - Wikipedia
https://en.wikipedia.org/wiki/Patrick_O'Neil ▾
Patrick Eugene O'Neil is an American computer scientist, an expert on databases, and a professor of computer science at the University of Massachusetts Boston. O'Neil did his undergraduate studies at the Massachusetts Institute of Technology, receiving a B.S. in mathematics in 1963.
Born: 1942, Mineola, New York Died: September 20, 2019 (aged 77), ...
Fields: Computer Science Known for: Distributed Database, SQL Isolation, ...

Obituary | Patrick E. O'Neill | Osheim & Schmidt Funeral Home
<https://www.osheimschmidt.com/obituary/5715446> ▾
Patrick E. O'Neill, 95, passed away peacefully on June 30, 2015, surrounded by loved ones at Rapid


Patrick O'Neal
Film Actor



Patrick Wisdom O'Neal was an American television, stage and film actor and New York restaurateur.

 
Wikipedia IMDb

Lived: Sep 26, 1927 - Sep 09, 1994 (age 66)
Height: 6 feet 1 inches
Spouse: Cynthia O'Neal (m. 1956 - 1994)
Children: Fitzjohn O'Neal (Son) - Maximilian O'Neal (Son)
Siblings: Michael O'Neal (Brother)
Education: University of Florida - Riverside Military Academy

Interesting stories


 Time Out for the Ginger Man Pub Is Only Part of Actor Patrick O'Neal's Movable Feast
People · 1 year ago

 A Tribute to Patrick O'Neal - HOTCHKA
HOTCHKA · 4 years ago

[Feedback](#)

Ichtheology

- **Fishes of Alabama and the Mobile Basin, 2004**
by Maurice F. Mettee, Patrick E. O'Neil,
J. Malcolm Pierson
- “Fishes of Alabama and the Mobile basin proved indispensable when I worked on a watershed assessment that included parts of the Coosa and Tallapoosa River drainages. We used this reference almost exclusively - the dichotomous keys are superb!”, review by A. Earnest
- 300 citations
- A different Pat O’Neil, of course.



A Uniformly Funny Joke Book



- A man enters a bank holding a gun and says nobody should move, then demands money from the tellers. Once he has the money, he turns to a customer and asks, "Did you see me rob this bank?"
- "Yes sir, I did," the man replies. And the robber shoots him in the temple, killing him instantly.
- He then turns to a couple and asks the man, "Did you see me rob this bank?"
- "No sir, I didn't," the man answers. "But my wife did!"

Final Remarks

➤ Weiwei Gong

➤ Betty O'Neil

