Database Systems Laboratory
Indian Institute of Science

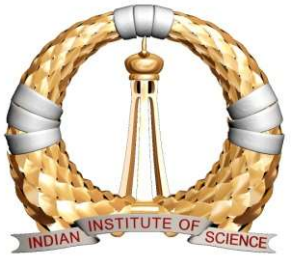

Indian Database Research: *The Journey to Sigmod 2026!*

Jayant Haritsa
Database Systems Lab
Indian Institute of Science, Bengaluru

SIGMOD countries

Country	Count
USA	39
Canada	4
China	2
Australia	1
Chile	1
France	1
Germany	1
Greece	1
Netherlands	1
India	1

1. Indian Academia

The Pioneers (1970/80s)

Nandlal Sarda, IIT Bombay

D B Phatak, IIT Bombay

A K Majumdar, IIT KGP

Veni Madhavan, IISc

D K Subramanian, IISc

Lalit Patnaik, IISc

1990s: The return of the young Turks, and the Boss

- S. Seshadri (Jeff Naughton @Wisconsin)
  - IIT Bombay 1992

- Jayant Haritsa (Miron Livny+Mike Carey @Wisconsin)
  - IISc Bangalore 1992 via UMCP
- S. Sudarshan (Raghu Ramakrishnan @Wisconsin)
  - IIT Bombay 1995 via Bell Labs

- Sunita Sarawagi (Mike Stonebraker @Berkele)
  - IIT Bombay 1999 via IBM ARC

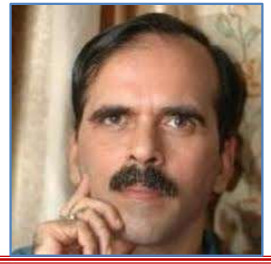

- Krithi Ramamritham
  - IIT Bombay 1998 via UMass-Amherst

SIGMOD/PODS until 2000 [7]

- SIGMOD 1985 [TIFR+IIT-B]
  - Timestamp Based Certification Schemes for Transactions in Distributed Database Systems
    Mukul Sinha, P. D. Nanadikar, S. L. Mehndiratta
- PODS 1993 [IBM ARC + IIT-B]
  - Fixed-precision estimation of join selectivity
    Peter Haas, Jeff Naughton, **S. Seshadri**, Arun Swami
- SIGMOD 1997 [IISc; IISc + UMass]
  - Secure Transaction Processing in Firm Real-Time Database Systems
    Binto George, **Jayant Haritsa**
  - Revisiting Commit Processing in Distributed Database Systems
    Ramesh Gupta, **Jayant Haritsa**, **Krithi Ramamritham**
- SIGMOD 2000 [IISc + IIT-B; IIT-B + Bell Labs; IIT-B + Bell Labs + UCB]
  - Turbo-charging Vertical Mining of Large Databases
    Pradeep Shenoy, **Jayant Haritsa**, **S. Sudarshan**, Gaurav Bhalotia, Mayank Bawa, Devavrat Shah
  - Efficient and Extensible Algorithms for Multi Query Optimization
    Prasan Roy, **S. Seshadri**, **S. Sudarshan**, Siddhesh Bhobe
  - On-line Reorganization in Object Databases
    Mohana Lakhamraju, Rajeev Rastogi**, S. Seshadri**, **S Sudarshan**

Research Institution Expansion

THEN (2000)

NOW (2026)

PPDMIIITDM (Pandit Dwarka Prasad Mishra Indian Institute of Information Technology, Design and Manufacturing) in Jabalpur

Contributions to Relational Query Processing

- Multi-query Optimization [2000+, IIT-B]
  - Cost-based heuristics (Shareable nodes, Benefit Monotonicity)
- Picasso Query Optimizer Visualizer [2005+, IISc]
  - Plan Diagrams (Product, Reduction, Applications)

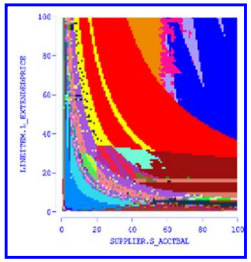

- Holistic Optimization of Database Applications [2008+, IIT-B]
  - DBridge (Batching, Asynchrony, Prefetching, Imperative → Declarative)
- Query Testing [2010+, IIT-B] [Paper in Session R17]
  - XDATA (Query Correctness Testing and Grading)
- Robust Query Processing [2014+, IISc]
  - Plan Bouquet (Quantitative Performance Guarantees wrt Ideal)

Robust Query Processing

Where Geometry beats ML!

Run-time Sub-optimality

The supposedly optimal plan-choice may actually turn out to be highly sub-optimal (e.g. a 1000 times worse!) when the query is executed with this plan. This adverse effect is due to errors in:

(a) cost model (limited)

- Reasons: Simple linear models, operator-agnostic features, fixed coefficients, system dynamics …

(b) cardinality model (large)

- Reasons: Coarse statistics, outdated statistics, attribute value independence (AVI) assumption, multiplicative error propagation, query construction, …

Proof by Authority [Guy Lohman, IBM]

Snippet from April 2014 Sigmod blog post on

"Is Query Optimization a "Solved" Problem?"

The root of all evil, the Achilles Heel of query optimization, is the estimation of the size of intermediate results, known as cardinalities. The cardinality model can easily introduce errors of many orders of magnitude! With such errors, the wonder isn't "Why did the optimizer pick a bad plan?" Rather, the wonder is "Why would the optimizer ever pick a decent plan?"

Prior DB Research (lots!)

- Sophisticated estimation techniques
  - SIGMOD 1999/2010, VLDB 2001/2009/2011, …, CIDR 2019
  - e.g. wavelet histograms, self-tuning histograms, deep learning
- Selection of stable plans
  - SIGMOD 1994/2005/2010, PODS 1999/2002, VLDB 2008, …, VLDB 2017
  - e.g. Variance-aware plan selection
- Runtime re-optimization techniques
  - SIGMOD 1998/2000/2004/2005, …, ICDE 2019 [Stonebraker et al]
  - e.g. POP (progressive optimization), RIO (re-optimizer)

Several novel ideas and formulations,
but are they robust?

New Approaches in last decade

- Learning Techniques [Rest of World]
  - Replace coarse parametrized models with fine-grained learnt models
  - Tsunami of papers over past few years!
    - RQP tutorial [ICDE 2019/VLDB 2020] lists several references
- Geometric Techniques [IISc]
  - Leverage plan cost trajectories over parameter space
  - Popular tradition in Indian academia, especially PQO
    - Sumit Ganguly@IIT Kanpur (late 1990s)
    - S Sudarshan@IIT Bombay (early 2000s)

Learning Techniques

Sample Performance

MSCN [CIDR 2019]:
Four fully-connected multi-layer neural networks

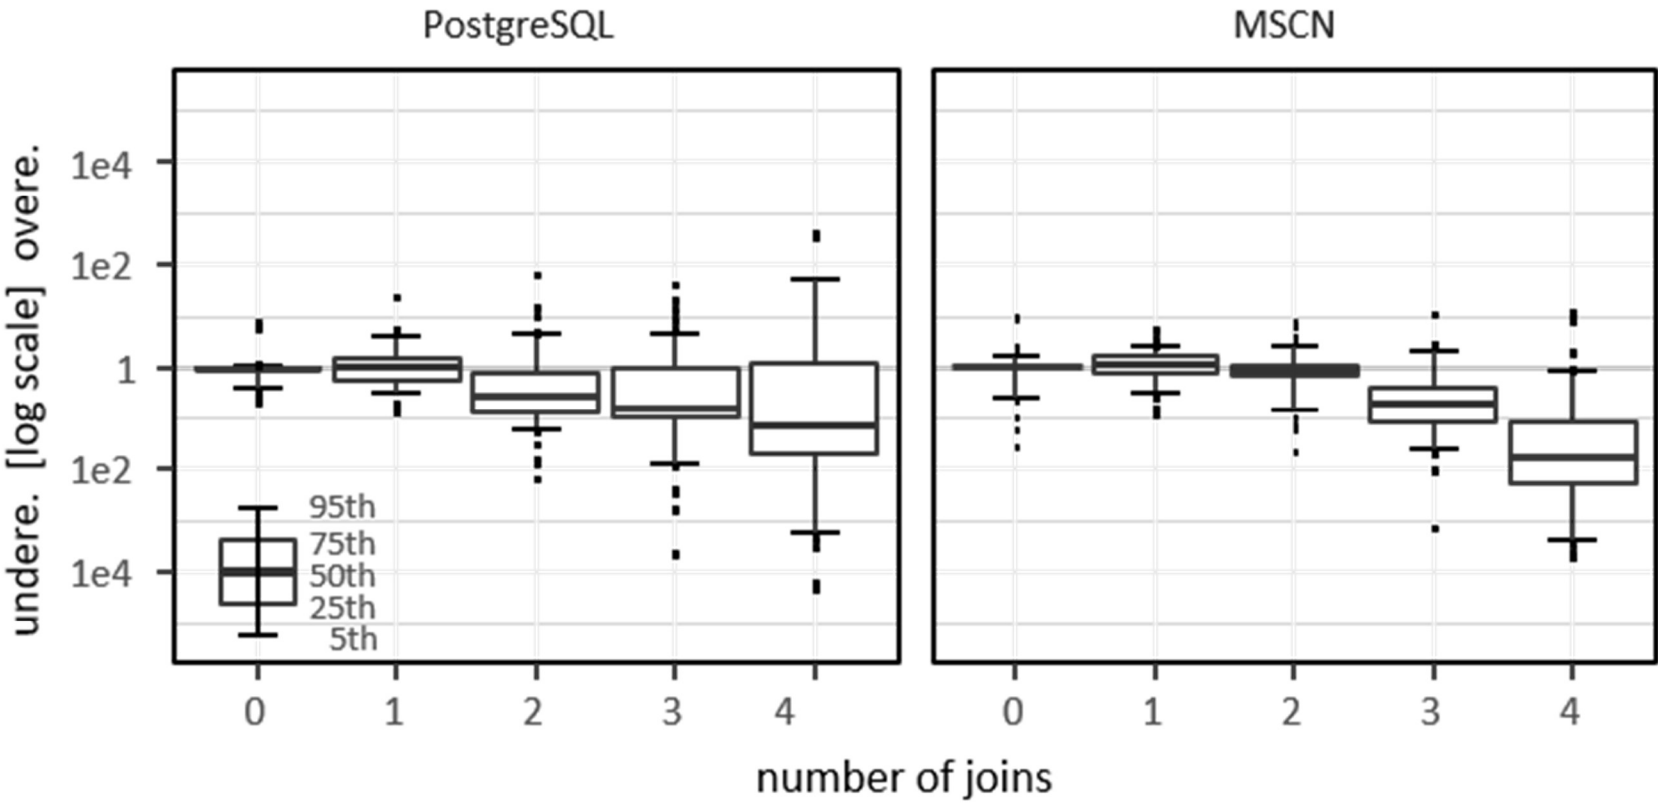

Geometric Techniques

Mechanisms

- Apply "Laws of Data" on Plan Cost Functions (PCF) over Selectivity (normalized Cardinality) Space
  - Monotonicity
    - More Data ⇒ More Cost
    - Smooth behavior (differentiable)
  - Submodularity (concavity)
    - Diminishing Costs per Additional Tuple
    - Closed under Linear Combinations
  - Bounded Growth
    - Limit on function slope (multi-linear)

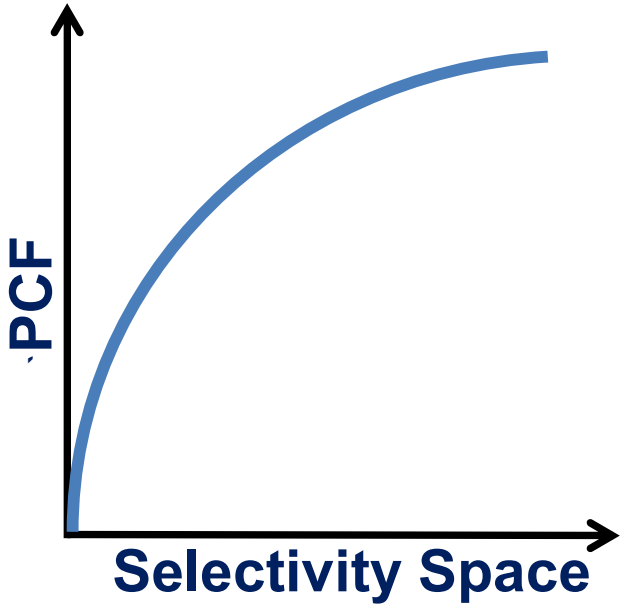

- Exploration Trajectory
  - Multi-dimensional or Manhattan (1-D)
- Leverage Query Structures
  - Chain Join

TPCH Q10

NATION — CUSTOMER — ORDERS — LINEITEM

Monotonicity: Plan Bouquet [SIGMOD 2014/TODS 2016]

- *Selectivity estimation problem cannot be solved, except by Rajanikanth!*
  - Arnold Schwarzenegger + Clint Eastwood + Elvis Presley
- Plan Bouquet is a new query processing technique that completely abandons estimating operator selectivities
- Instead, run-time selectivity discovery using compile-time selected bouquet of plans
  - provides worst case performance guarantees wrt ideal that magically knows the correct selectivities
    e.g. for single error-prone selectivity, relative guarantee of 4
  - empirical performance well within guaranteed bounds on industrial-strength environments

Parametric Optimal Set of Plans (POSP)

(Parametric Query)

```
select  S.Name, C.Title
from    STUDENT S,
        COURSE C,
        REGISTER R
where   S.RollNo = R.RollNo
and     C.CourseNo = R.CourseNo
and     C.fees < $1
```

NL: Nested Loop Join
MJ: Merge Join
HJ: Hash Join

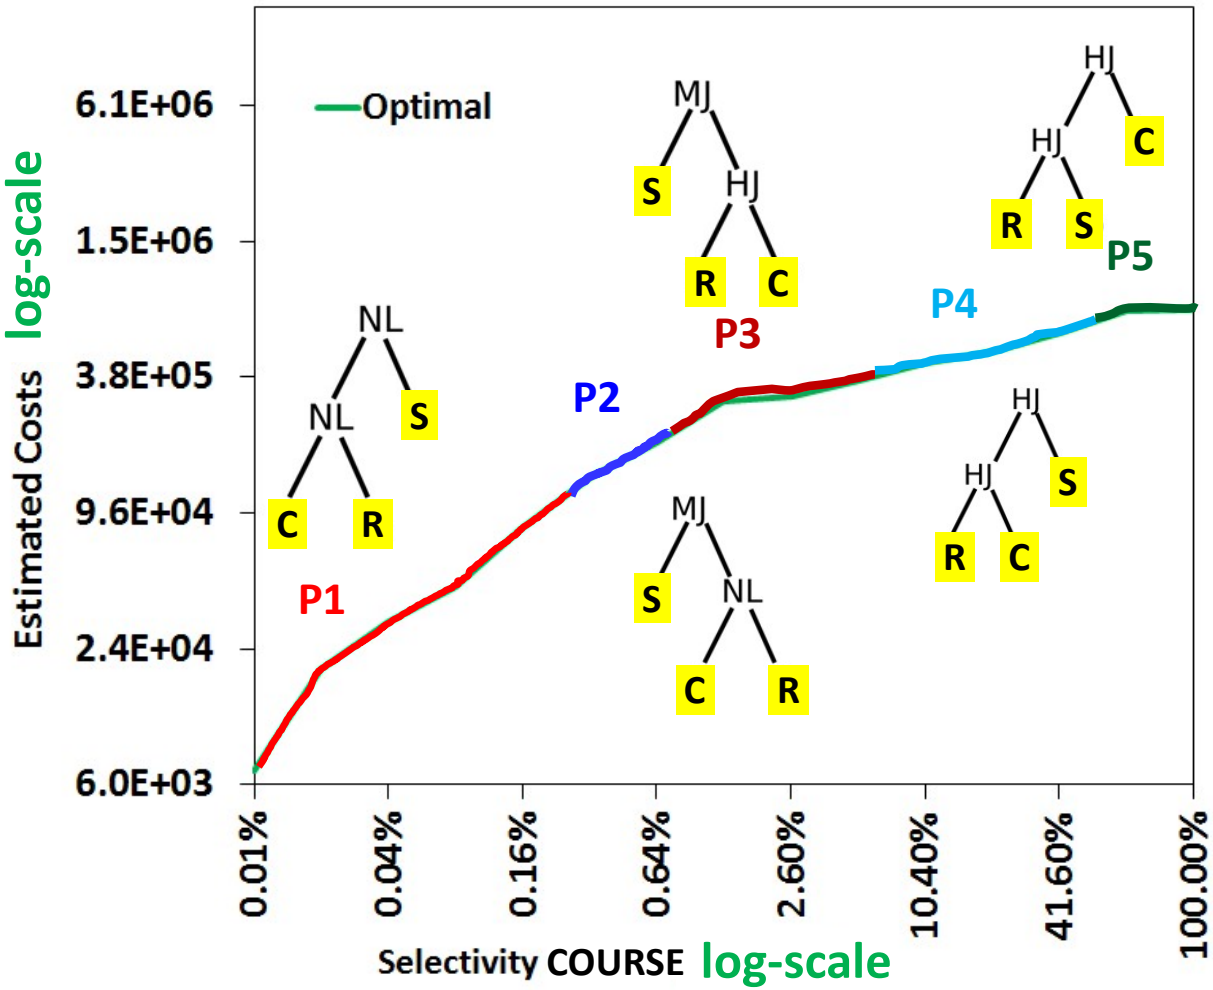

POSP Performance Profile (across ESS)

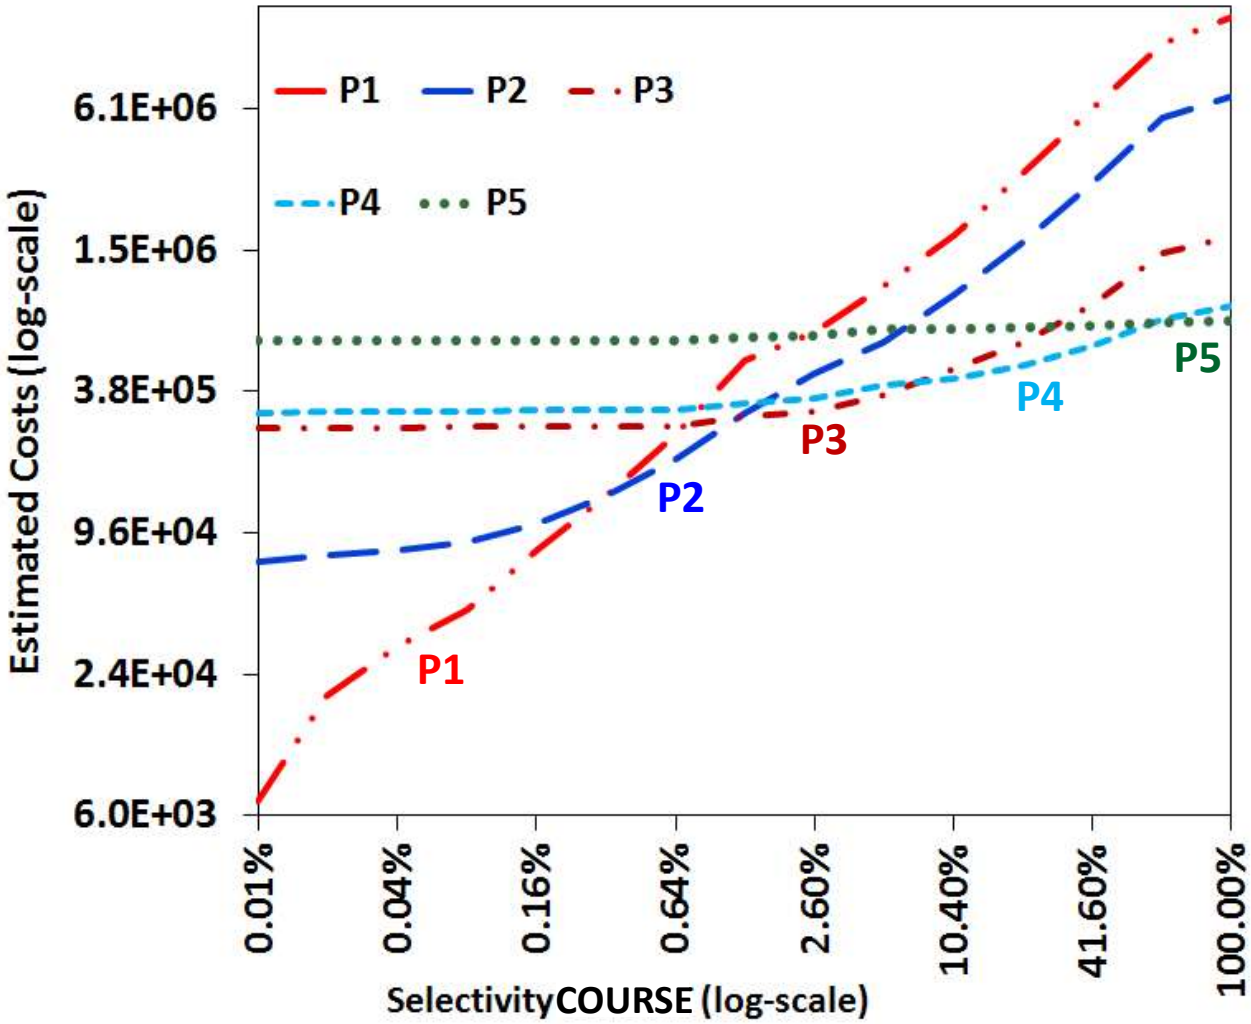

Sub-optimality Profile (across ESS)

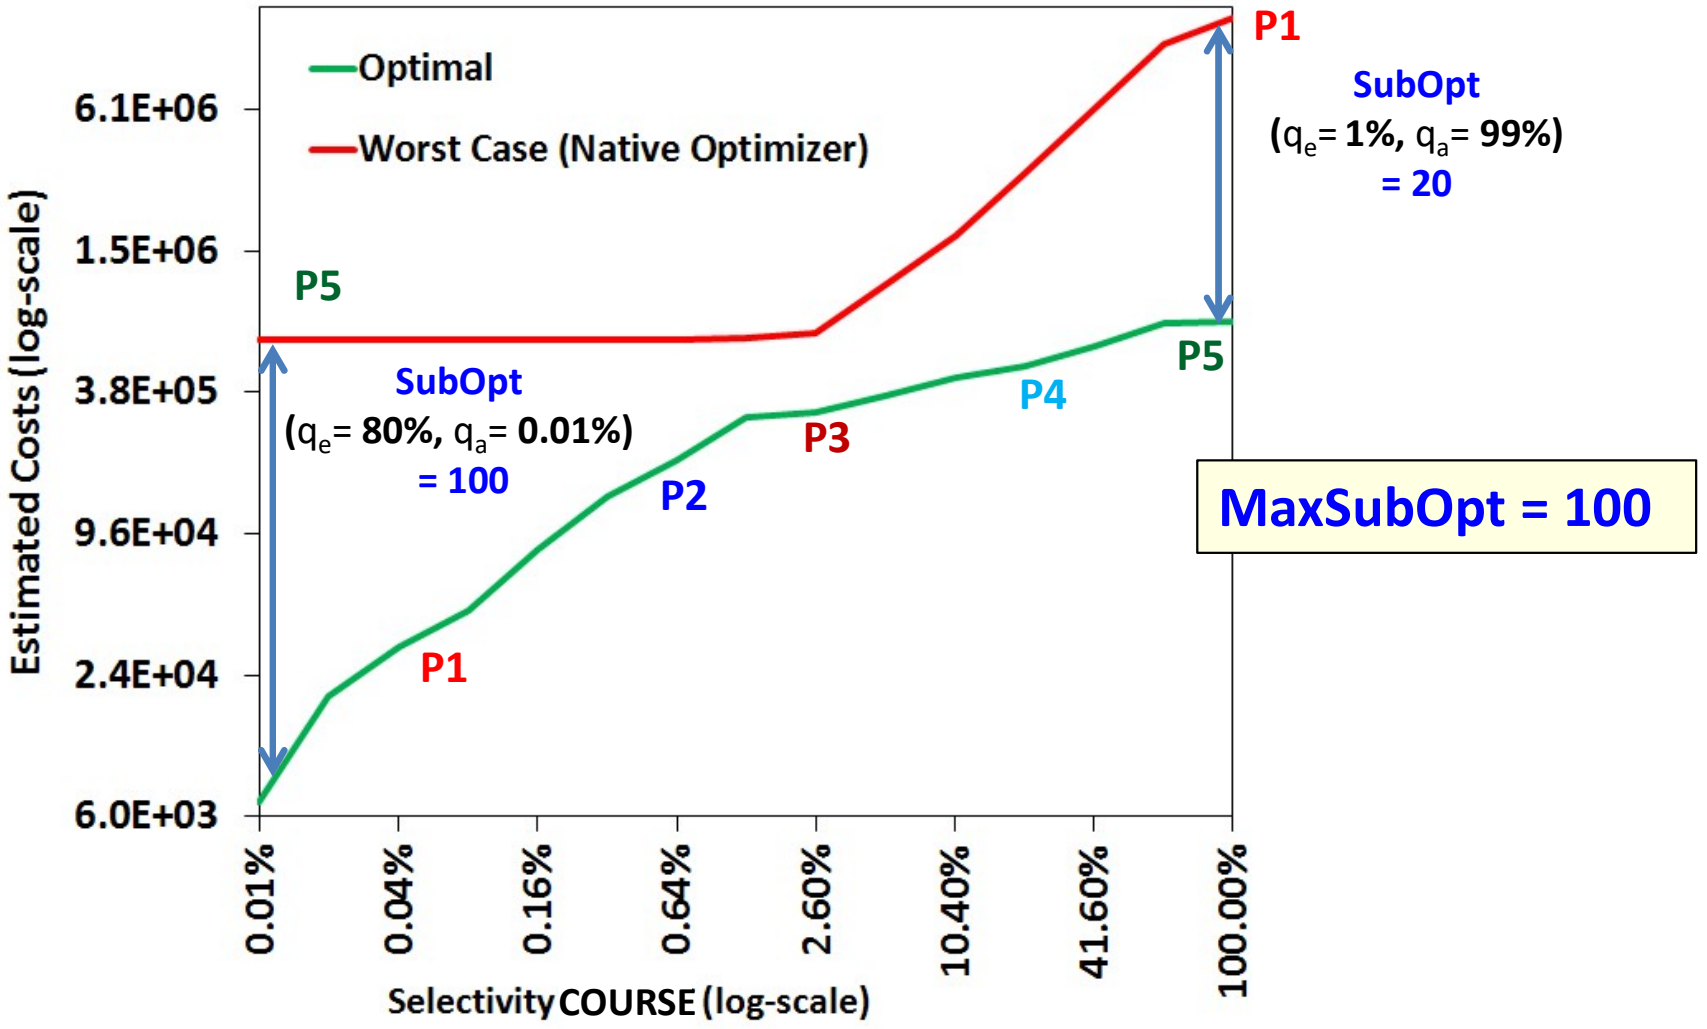

Plan Bouquet Behavior on 1D ESS

Bouquet Identification

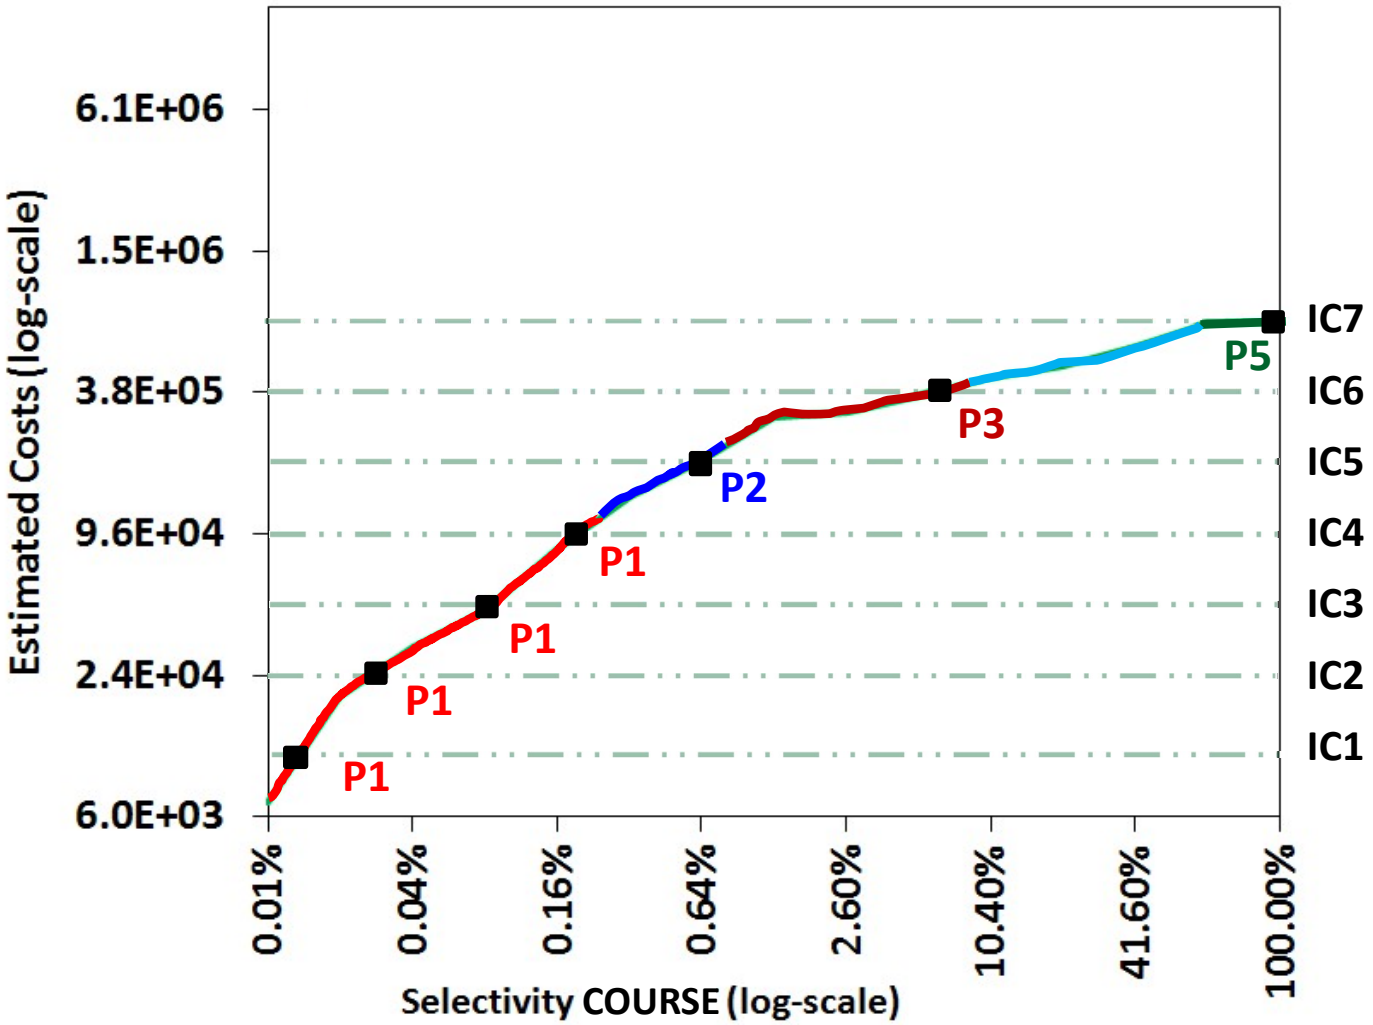

Step 1: Draw **iso-cost lines** with cost-ratio **r=2** (geometric progression).

Step 2: Find plans at intersection of optimal profile with cost steps

Bouquet = {P1, P2, P3, P5}

Bouquet Execution

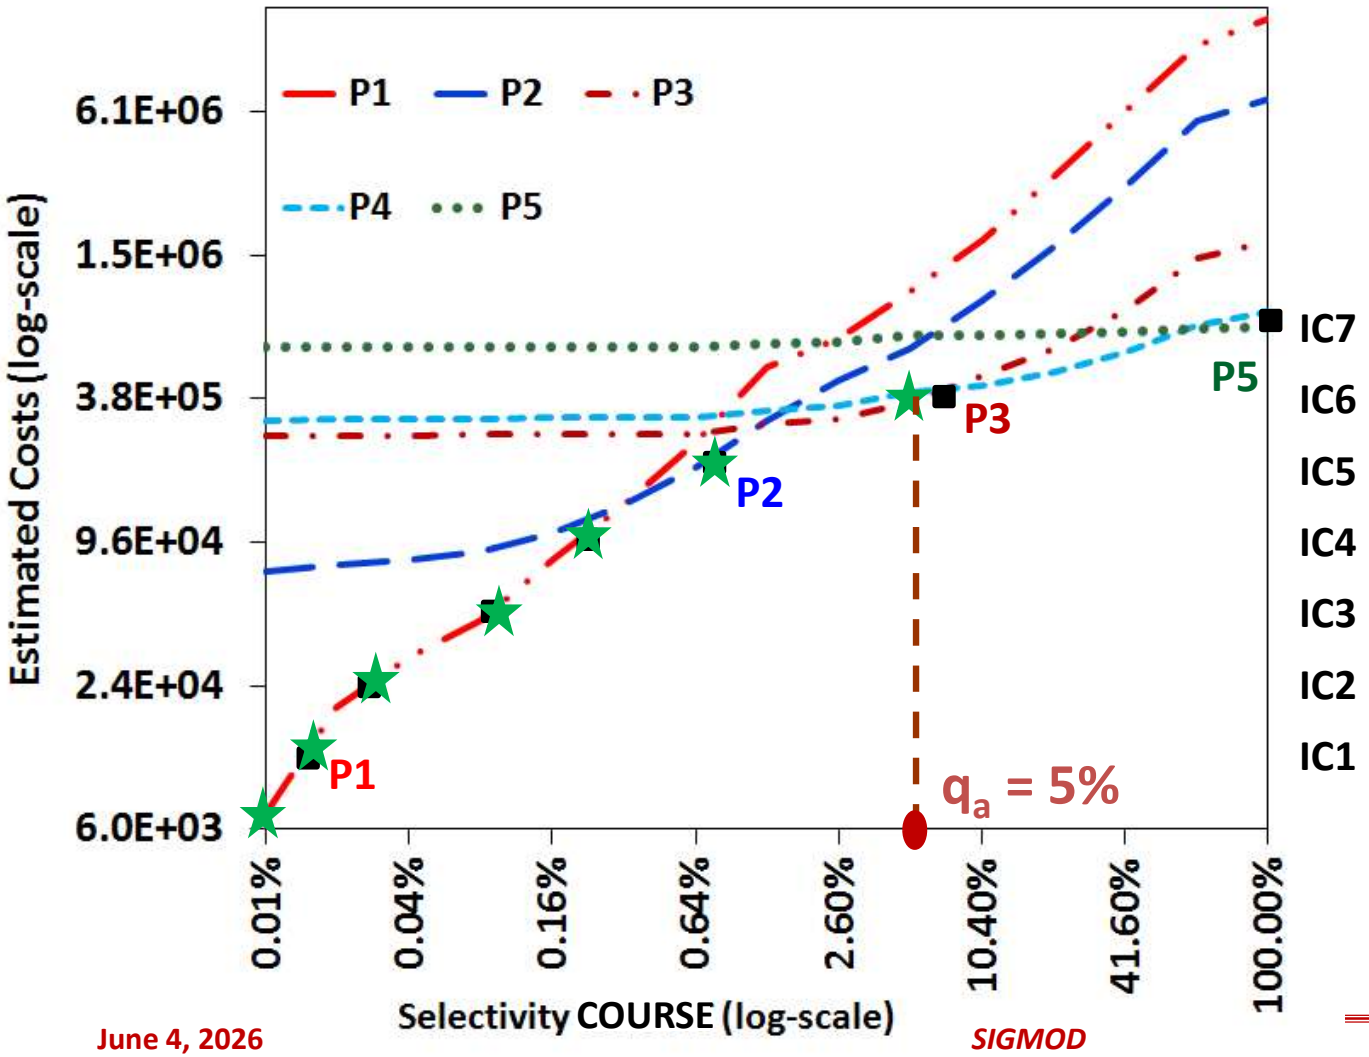

Let q_a = 5%

(1) Execute P1
with budget IC1(1.2E4)
Throw away results of P1

(2) Execute P1
with budget IC2(2.4E4)
Throw away results of P1

(3) Execute P1
with budget IC3(4.8E4)
Throw away results of P1

(4) Execute P1
with budget IC2(9.6E4)
Throw away results of P1

(5) Execute P2
with budget IC5(1.9E5)
Throw away results of P2

(6) Execute P3
with budget IC6(3.8E5)

P3 completes with cost 3.4E5

Bouquet Execution

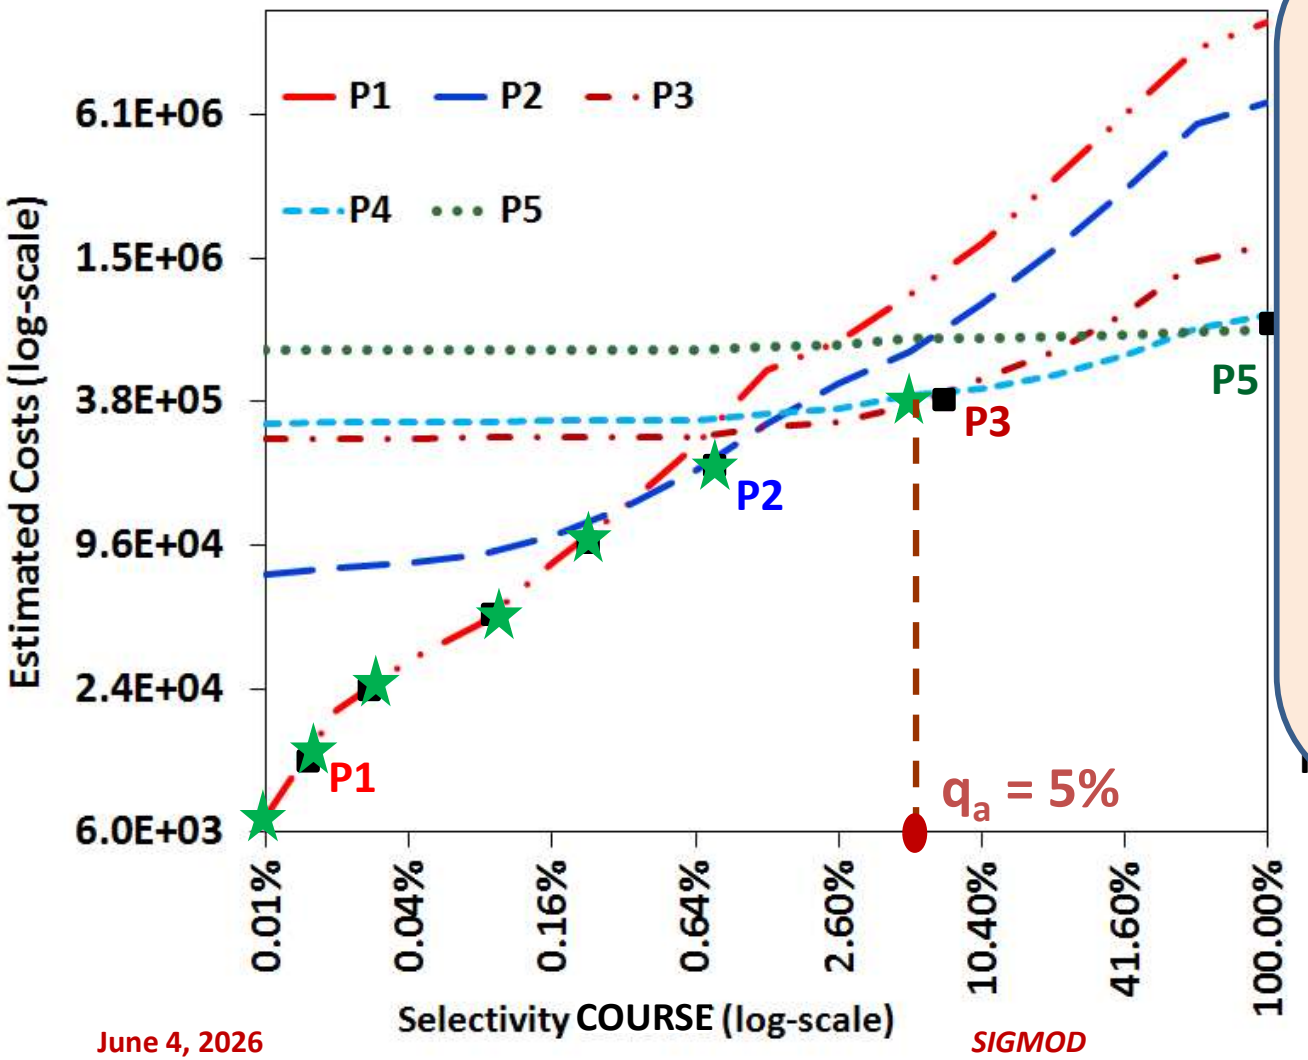

Let q_a = 5%

Bouquet Cost = 3.4 E5 (P3) +
1.92 E5 (P2) +
0.96 E5 (P1) +
0.48 E5 (P1) +
0.24 E5 (P1) +
0.12 E5 (P1)
= 7.1 E5

SubOpt (*, 5%) = 7.1/3.4 = 2.1

With obvious optimization
SubOpt(*, 5%) = 6.3/3.4 = 1.8

with budget IC6(3.8E5)
P3 completes with cost 3.4E5

Bouquet Performance

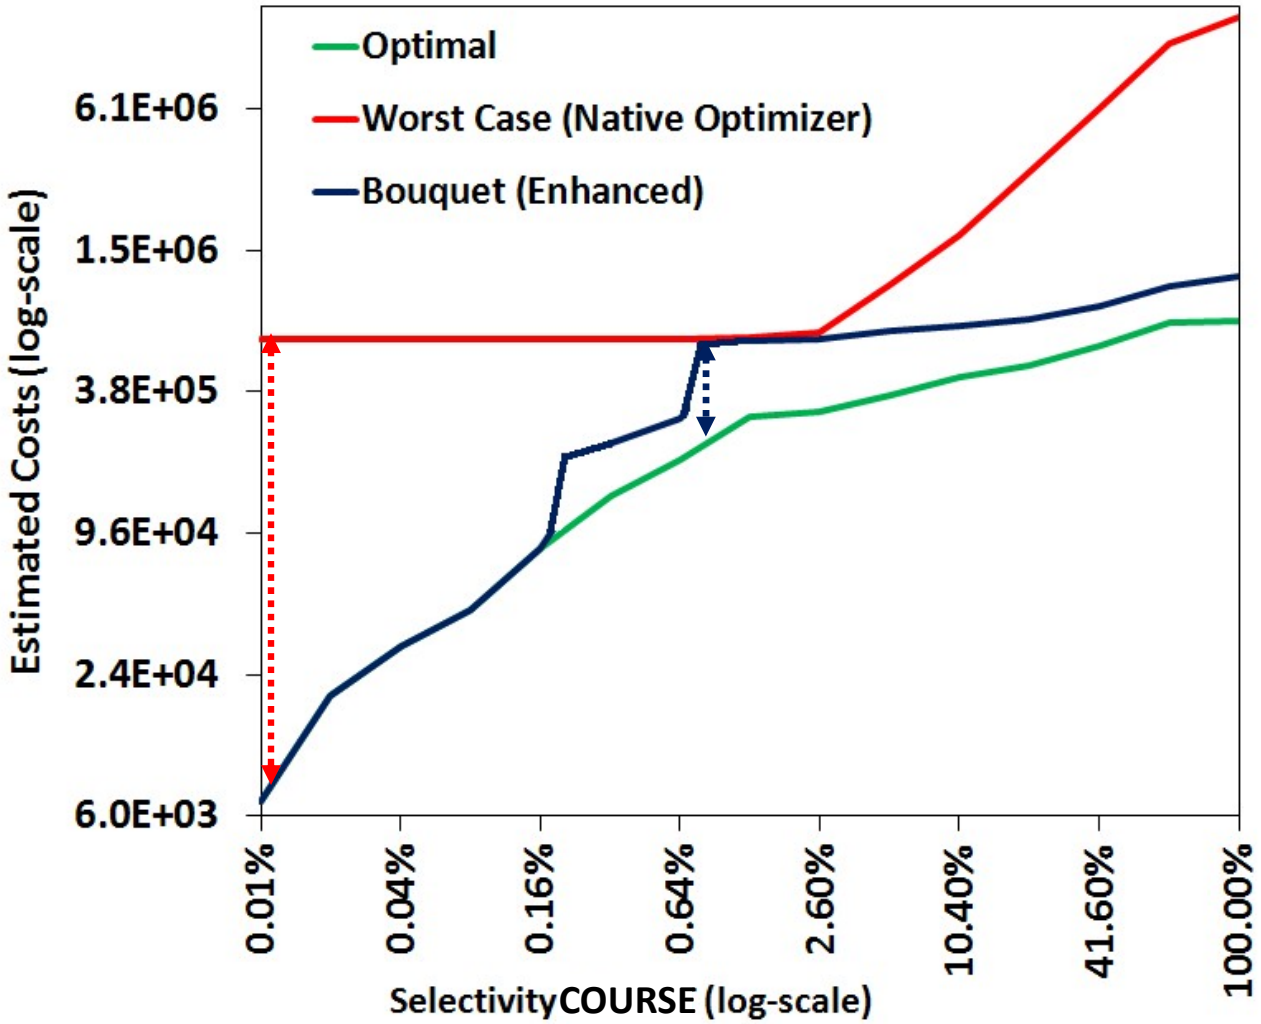

Native Optimizer

MaxSubOpt = 100

Bouquet

MaxSubOpt = 3.1

1D Performance Bound

$$\mathbf{C_{bouquet}}(\mathbf{q_{k-1}}, \mathbf{q_k}] = \mathbf{cost}(\mathbf{IC_1}) + \mathbf{cost}(\mathbf{IC_2}) + \ \ldots + \mathbf{cost}(\mathbf{IC_{k-1}}) + \mathbf{cost}(\mathbf{IC_k})$$

$$= \quad a \quad + \quad ar \quad + \ldots + \quad ar^{k-2} \quad + \quad ar^{k-1}$$

$$= \quad \frac{a(r^k - 1)}{(r-1)}$$

$\mathbf{C_{optimal}}(\mathbf{q_{k-1}}, \mathbf{q_k}] > \mathbf{ar^{k-2}}$ (Implication of PCM)

$$\mathbf{SubOpt_{bouquet}}(*, \mathbf{q_a}) \leq \frac{\mathbf{1}}{\mathbf{ar^{k-2}}} \times \frac{\mathbf{a(r^k - 1)}}{\mathbf{(r-1)}} \leq \frac{\mathbf{r^2}}{\mathbf{r-1}} \qquad \forall q_a \in (q_{k-1}, q_k]$$

- No dependency on k!
- Reaches minima at r = 2

  ➔ MSO = 4

Best performance achievable by <u>any</u> deterministic online algorithm!

PlanBouquet (SpillBound)MSO Guarantee

For D dimensions,
MSO guarantee is $D^2 + 3D$

- 1D, MSO = 18; 4D, MSO = 28; 5D, MSO = 40

Platform-independent for a query

- Geometric progression of budgets is optimal for any discovery algorithm in this class

RQP Monograph 2024

- *"Robust Query Processing: A Survey"*
  NOW Publishers, *"Foundations and Trends in Databases"* series

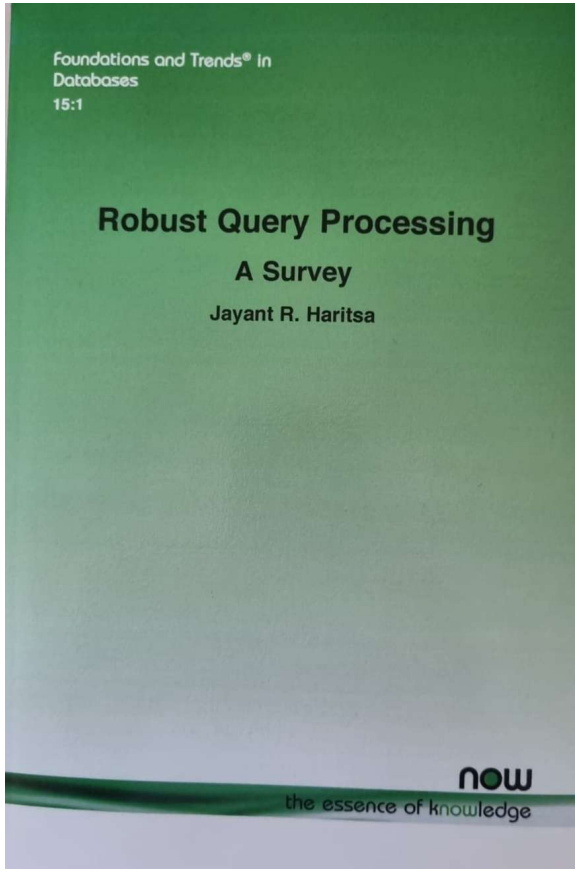

2. Indian Database Industry

The Early Years

- Limited to maintenance, documentation, support
- Knowledge of tools, but not system internals
  - drive the vehicle, but unaware of under the hood
- Sudarshan and I gave a one-week crash course for HP engineers in 2003
- Our students went to Pune, where they helped build Sybase ASE during 2004-10

Today

- Multiple teams working on query optimizers, execution engines, storage engines, transaction/logging layers, distributed SQL, database control planes, and cloud database reliability
- Google Napa, Google AlloyDB
  - core db and distributed platform engineered in Bengaluru
- Aurora PostgreSQL, Oracle Cloud, Teradata, YugabyteDB, …

- Persistent Systems: Anand Deshpande
  - Banking & Financial Services, Healthcare & Life Sciences, Software & Hi-Tech ($10B market cap)

- PGConf India conference: Started 2013, Pavan Deolasee et al
  - 2026: Over 570 attendees, 69 speakers, and 53+ technical talks. Agenda covered deep dives into database internals, performance tuning, GIS platforms, and multi-master replication.

Microsoft SQL IDC

- Spearheaded by Karthik Ramachandra (IIT-B, Jim Gray lab)

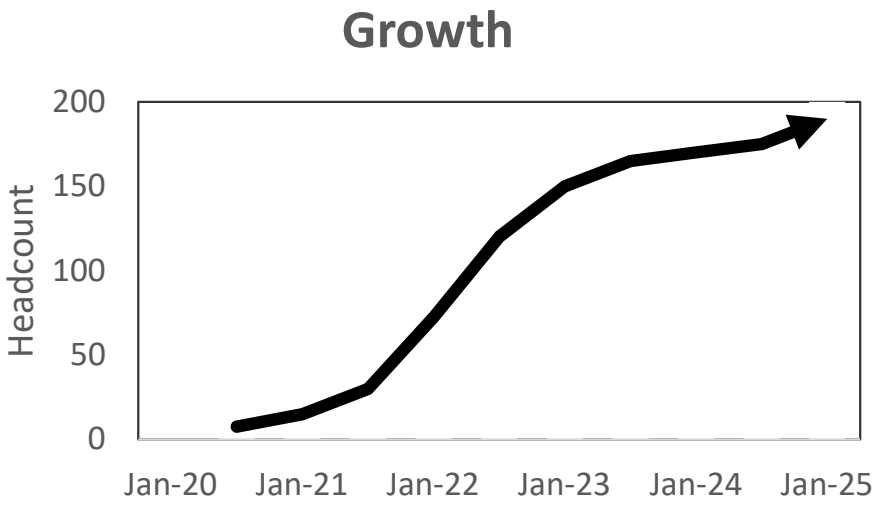

# Principals 30+	#Principal ICs 10+	D&I 26.2%
# Senior 55+	#SDE2 100+	SDE 25+

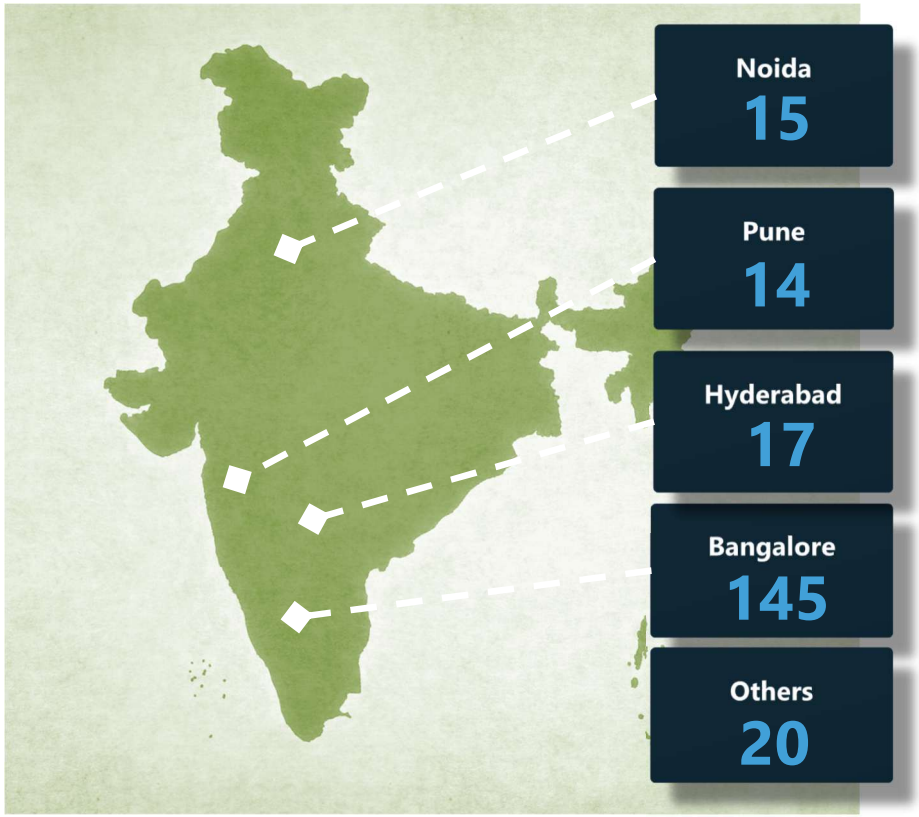

Impact

- Key contributions in features of SQL25, Hyperscale and other areas
- Implemented follow-the-sun for livesite

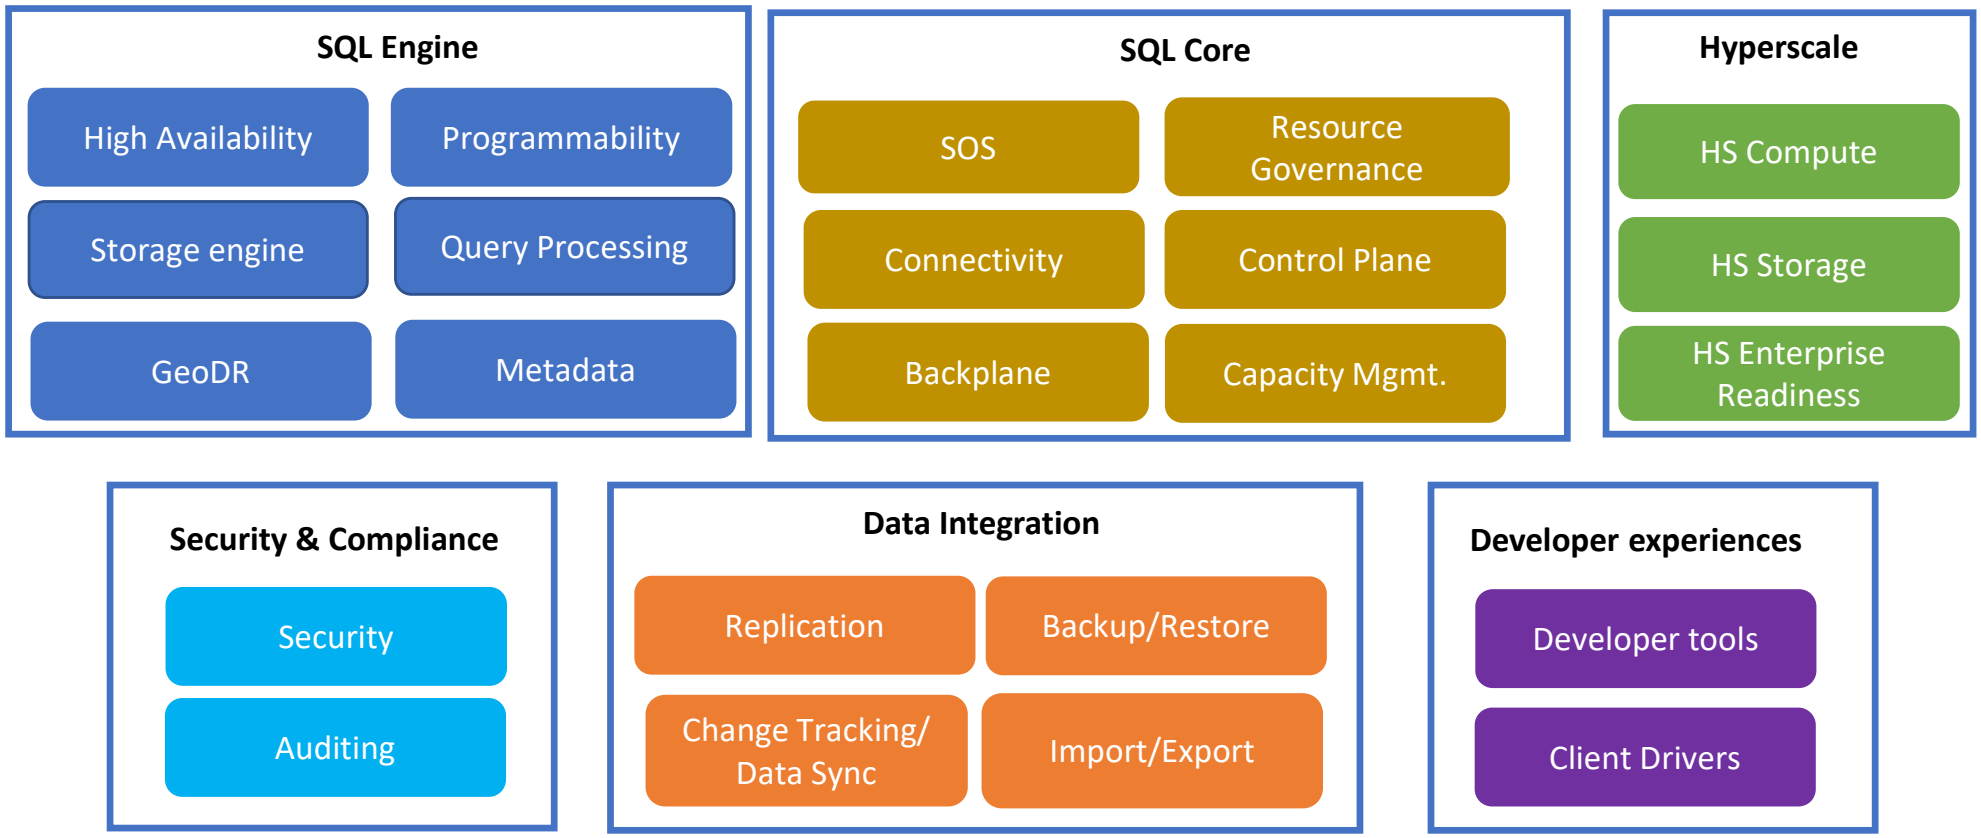

Innovation & Collaborations

- Several patents/publications from work driven by SQL IDC
- Hackathons, internship program,…
  - Socrates IDC Hackathon, LLM Hackathon
  - Internships in collaboration academia (IISc, IITD).
- Internal Collaborations ongoing with:
  - MSR India (Projects: Garuda, Blitz, SelfTune, LITHE)
  - GSL
  - MSR Redmond

3. Indian Database Deployment

The Recent Past

- In 2009, less than 20% of Indians had a bank account. Essentially, a parallel economy was running – loss of productivity, tax revenue and socio-economic development.
- 1/3 of population did not have any identity document
- Aadhaar ("foundation") – 1 billion+ in 5 years for identity using biometrics and mobile phones for 2FA

India Stack [2009+]

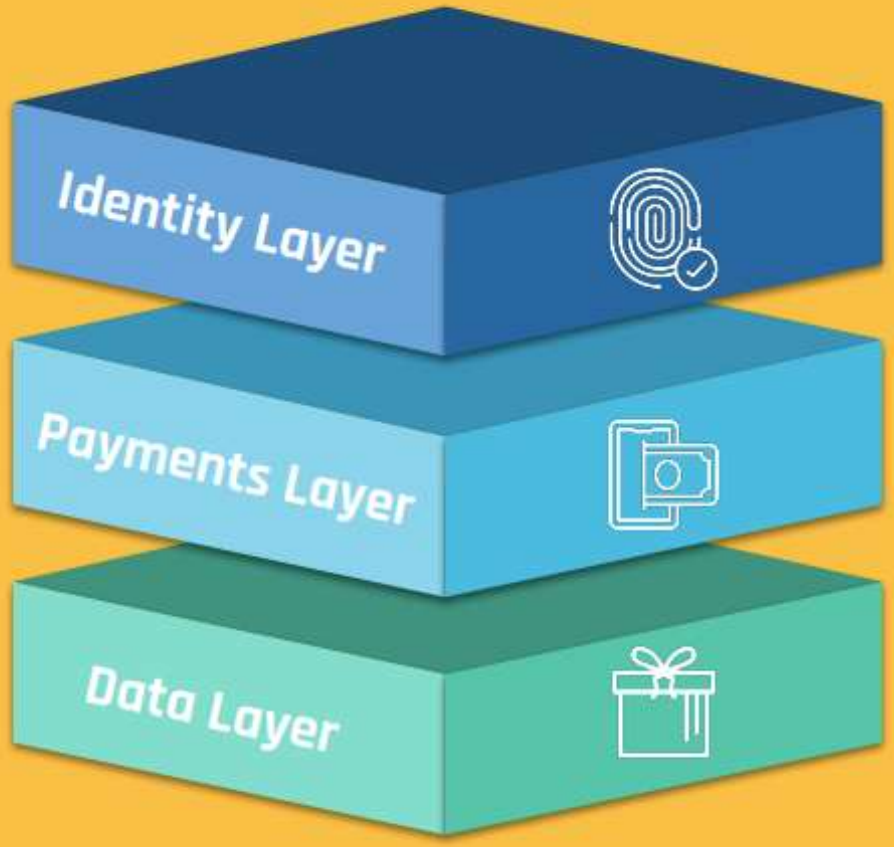

Identity Layer
Giving every resident a unique id and enabling them to prove "I am who I claim to be"

Payments Layer
Allowing anyone to pay anyone else! interoperable, fast and cheap - not just smartphones

Data Empowerment
To enable secure sharing of data

Unified Payment Interface

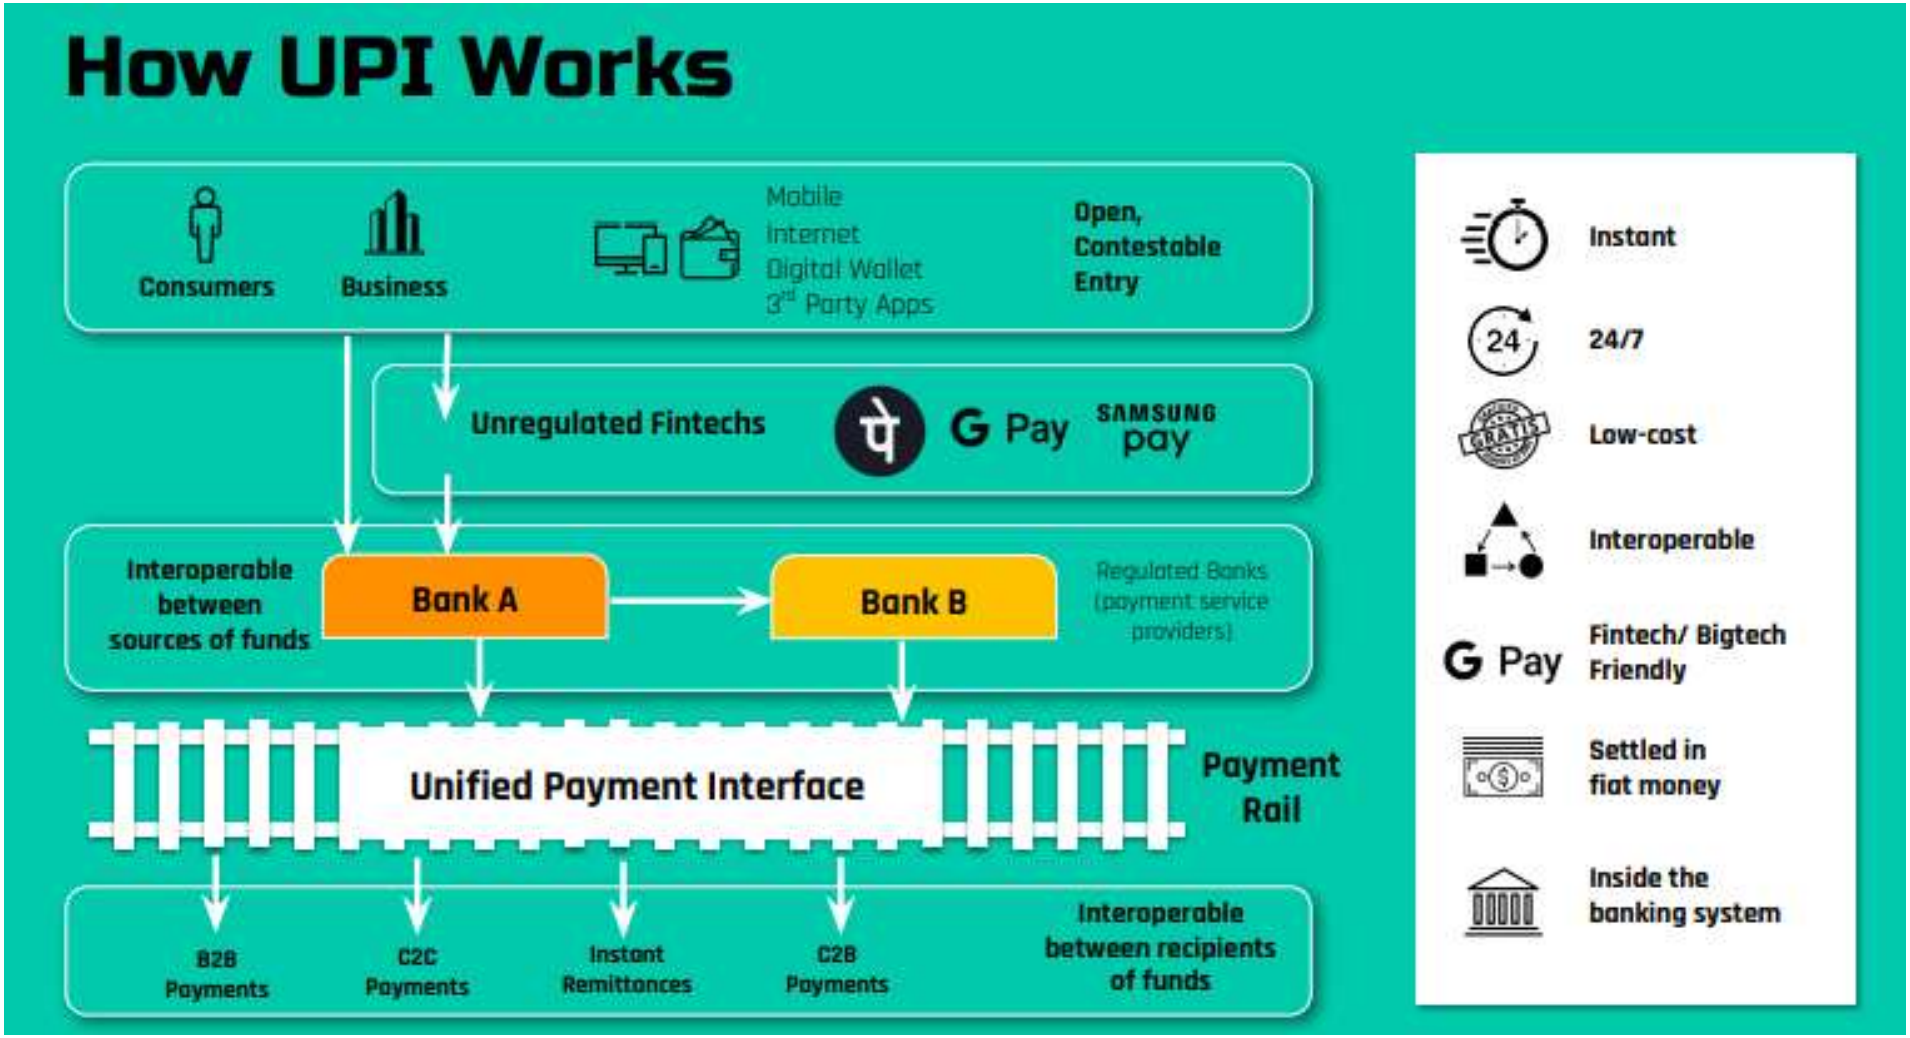

Capital Markets

- Largest volume of stock market trades per day in world
  - National Stock Exchange/Clearing in Mumbai
- ~20 billion orders and 300 million transactions on June 5, 2024
  - ~one million messages per second
- Performant, robust, fair, transparent, regulated
- T+1 settlement in 2022, USA did it only 2024, whereas Europe, Australia and Japan still in T+2!
- India rolled out T+0 since March 2024
  - China also has T+0 but restricted to shares, whereas India supports T+0 for both shares and cash.
- Disaster recovery in less than 45 minutes between Mumbai and Chennai without data loss
- Unique SAAS interoperability model with other exchanges since both PDC and DR may have same problem with faulty software

Closing Thoughts

Institutional growth, industrial capability, and nationwide deployment have combined to transform a once-limited research footprint into a multi-layered database ecosystem with global impact and societal reach.

END PRESENTATION