Sarcouncil Journal of Engineering and Computer Sciences

ISSN(Online): 2945-3585

Volume- 04| Issue- 07| 2025

Research Article **Received:** 01-06-2025| **Accepted:** 17-06-2025 | **Published:** 06-07-025

Fine-Grained Indexing Strategies for Social Feeds

Nisheedh Raveendran
Birla Institute of Technology and Science, Pilani, India

Abstract: Fine-grained indexing strategies have emerged as a transformative solution for addressing the complex challenges faced by modern social media platforms in content organization and retrieval. Traditional monolithic indexing systems struggle with the exponential growth of diverse user-generated content, including text posts, multimedia elements, and real-time interactions that demand specialized processing approaches. The evolution from simple chronological feeds to sophisticated machine learning-driven ranking systems has highlighted the inadequacy of unified indexing structures when dealing with heterogeneous content types and diverse user engagement patterns. The implementation of specialized indexing architectures enables intelligent content partitioning across multiple targeted systems, allowing for optimized processing of different content categories while maintaining system coherence through dynamic routing mechanisms. Content segregation based on activity types, engagement patterns, and social graph proximity facilitates more precise retrieval operations and enhanced personalization capabilities. The distributed nature of fine-grained systems provides superior fault tolerance and maintenance flexibility compared to traditional approaches, enabling graceful degradation during component failures and independent updates to specialized subsystems. Production deployments demonstrate substantial improvements in system responsiveness, scalability, and user satisfaction metrics. Enhanced query processing efficiency, reduced latency, and improved content relevance contribute to more engaging user experiences and better platform performance. The strategic balance between system complexity and operational benefits requires careful consideration of granularity levels, resource allocation, and maintenance overhead to achieve optimal performance outcomes.

Keywords: Fine-grained indexing, embedding systems, social media platforms, content partitioning, distributed architectures.

INTRODUCTION

Social media has changed how people communicate, but with this change came a range of unforeseen technical hurdles. The increasing availability of user-generated content took everyone by surprise on various platforms and has moved engineers back to the drawing board to seriously consider basic assumptions about how we structure and retrieve content. The simple chronological feeds have turned into complicated systems that must immediately process and rank millions of different pieces of content, while at the same time providing a unique experience for each user.

The problem extends beyond mere volume. Consider the stark differences between a short status update, a high-resolution photo album, and a live video stream. Each content type demands different processing approaches, storage strategies, and user interaction patterns. Traditional systems tried to handle everything through one-size-fits-all solutions, but this approach increasingly shows its limitations as platforms mature and user expectations rise (Ravikumar, G. 2025).

Current feed ranking systems face a perfect storm of technical demands. Users anticipate instantaneous responses, individualized experiences, and frictionless experiences across devices and contexts. At the same time, the systems underneath are analyzing histories of user behavior, mapping social networks, identifying trends, and understanding innumerable engagement signals. The mathematical complexity of representing these relationships in searchable formats has led to sophisticated embedding techniques that translate content and user preferences into multi-dimensional vectors.

Problem Statement

The standard approach involves creating massive, centralized indexes that attempt to catalog every piece of content in one comprehensive system. While this strategy appears straightforward, it creates bottlenecks that become more severe as platforms expand. The fundamental issue lies in forcing dissimilar content types into identical organizational structures.

Think about how differently people interact with various content formats. A viral meme spreads through completely different mechanisms than a thoughtful long-form post. Photo collections generate engagement patterns unlike video content. Yet traditional indexing systems treat all these formats identically, missing opportunities to optimize for specific behaviors and characteristics.

The computational burden of maintaining these enormous unified indexes creates practical problems. The time to process increases exponentially instead of incrementally, memory use rises infinitely, and routine maintenance tasks

***Corresponding Author:** Nisheedh Raveendran

DOI- https://doi.org/10.5281/zenodo.15818090

become major jobs that can affect the availability of the service. It doesn't help that in the background, the content of social media never stops, and continuously requires updates to already heavy and burdened indexes, never ceases.

Perhaps most critically, generic indexing approaches fail to recognize the nuanced ways different user segments engage with content. Power users who create and share content extensively behave differently from casual browsers who primarily consume content. Teenagers interact with platforms differently from professionals or in general adults. These behavioral patterns represent valuable optimization opportunities that unified systems simply cannot leverage effectively.

RESEARCH CONTRIBUTIONS

This investigation explores specialized indexing strategies that segment content across multiple targeted systems rather than relying on monolithic approaches. The research examines real-world deployments where these techniques have been successfully implemented at enterprise scale, providing concrete evidence of their effectiveness (Saoudi, E. M., & Jai-Andaloussi, S. 2021).

The analysis covers intelligent content separation methods that group similar content types together, enabling optimized processing for each category. Dynamic routing systems guide search queries to the most appropriate specialized indexes, maintaining efficiency while improving accuracy. Innovative synchronization techniques keep distributed indexes consistent without the performance penalties that plague traditional unified systems.

Evidence from production environments reveals substantial improvements in both technical performance metrics and user satisfaction indicators, establishing specialized indexing as a promising direction for future social media infrastructure development.

Evolution of Social Feed Ranking Systems

The early social media platforms started with a chronological ordering method as their default information organization method. Using a chronological ordering method as the default mode of organization across content appeared to be an appropriate approach in the early developmental stages of the platform, as it closely mimicked traditional consumption patterns of information, such as reading a newspaper or bulletin board systems. As these types of platforms grew and the scale of user-generated content grew, the limitations of this simplistic approach became apparent.

The chronological format of feeds caused problems with content discovery and user engagement. A post from one of your friends that was important often became irrelevant quickly as the less relevant posts from more frequently active users began to overshadow everything from your friend. This resulted in a decline in engagement metrics and user satisfaction, and platforms started considering new ranking models.

The first steps into algorithmic ranking borrowed heavily from existing algorithms in recommendation systems formed from people seeing similar movies, clothing, restaurants, and books as well as using one of the basic forms of collaborative filtering in which user engagement and interaction was looked at in terms of patterns of activity in a broader social context. Older collaborative filtering approaches eventually were really only useful within a similar prediction framework because social media content was temporal while collaborative filtering mainly relied on analysis and interactions from two dimensional lists of what people either liked or shared what lists of content depended on their relationships with that content as they basically swirled around their social network.

Value in and advancements in the sophistication of machine learning systems marked the next important phase in feed ranking technology. Modern platforms recognized that social media content exhibits distinct properties that require specialized algorithmic approaches. Unlike traditional recommendation scenarios with relatively stable user preferences, social media interests fluctuate rapidly based on current events, personal circumstances, and evolving social dynamics (Hayat, M. K. *et al.*, 2019).

Publisher: SARC Publisher

Table 1: Evolution of Social Feed Ranking Systems (Hayat, M. K. *et al.*, 2019; Ai, L. *et al.*, 2022)

System Generation	Ranking Method	Personalization Level	Scalability	Technical Complexity
First Generation	Chronological	None	Limited	Low
Second Generation	Collaborative Filtering	Basic	Moderate	Medium
Third Generation	Machine Learning	Advanced	High	High
Fourth Generation	Embedding-Based	Sophisticated	Very High	Very High

Embedding-Based Retrieval Systems

The use of embedding-based retrieval systems was a fundamental transformation in how platforms understand and organize social content. They do not use explicit categorical distinctions but mathematical representations, which model the semantic relationships or connections between users, content, and behaviors in a high-dimensional vector space that can support the bulk of an enormous social network.

For users, embeddings can encode behavioral habits, preferences, and social connection data into dense numerical representations. As a result, systems are less reliant on explicit demographic or declared interests and can discover similarity between users based on implicit behavioral signals. The mathematical nature of embeddings allows for quick calculations of similarity and allows for personalization by using the information at scale.

In contrast, everything communicated through social media in the form of content (posts) is heterogeneous and varies. Textual content requires text features to be parsed by natural language processing methods that allow the system to understand context, sentiment, and semantic meaning. Visual content requires an understanding of images, videos, and other multimedia to be resolved through computer vision methodologies. Multi-modal embedding representations combined into a single space remain a developing and research space.

Interaction embeddings introduce complications by incorporating the temporal and contextual properties of user-content relationships. The models also need to consider the dynamic nature of social relationships, the relative strength of relationships among the content being interacted with, and the contextual attributes that shape interaction routines across content types and social contexts.

Approximate Nearest Neighbor (ANN) Indexing

Large embedding retrieval invokes the need for specialized indexing structures that support fast similarity search over a high-dimensional vector space. For example, exact search techniques are possible, but can be prohibitively expensive at scale when dealing with potentially billions of items in a social media-type application. Therefore, acceptable approximations of exactness can be used to find answers that fall within accuracy/fairness and speed parameters.

Various indexing strategies demonstrate different performance characteristics depending on data properties and query patterns. Tree-based structures provide efficient search capabilities for lower-dimensional spaces but experience performance degradation as dimensionality increases. Hash-based methods offer consistent performance across large datasets while maintaining computational efficiency. Graph-based approaches have gained prominence for high-dimensional social media applications due to their superior balance of search accuracy and query response times (Ai, L. *et al.*, 2022).

Limitations of Monolithic Indexing

Single comprehensive indexing systems face increasing challenges as social platforms scale in both user base and content diversity. The computational resources required for maintaining these unified structures grow super-linearly with content volume, creating performance bottlenecks that impact system responsiveness and user experience quality.

When an indexing method supports heterogeneous content types or categories, previously discussed optimization trade-offs will create conflict when optimizing features of one type of content while degrading those of another. As platforms and applications vary, this limitation increases when also taking into account new content types beyond common text-based posts into rich media, interactive content experiences, and real-time communication features.

FINE-GRAINED INDEXING METHODOLOGY

Content Partitioning Strategies

Breaking apart massive indexing systems into smaller, specialized pieces turns out to be one of those ideas that sounds obvious in hindsight but took years to figure out properly. The traditional approach of shoving everything into one giant

Publisher: SARC Publisher

index created bottlenecks that became impossible to solve through simple optimization.

Activity-based partitioning emerged from watching how users actually behave on social platforms. Creating original posts involves completely different mental processes than commenting on someone else's content. Sharing or reacting requires minimal cognitive effort compared to crafting thoughtful responses. These behavioral differences translate directly into technical requirements that benefit from specialized handling (Wang, Y. *et al.*, 2024).

Engagement-driven separation makes sense from a resource allocation perspective. Viral content demands immediate attention and rapid processing, while niche posts can wait in slower queues without affecting user experience. This tiered approach allows systems to spend computational resources where they matter most.

Social graph proximity adds another dimension that traditional systems struggled with. Content from close friends naturally carries more weight than posts from distant connections, but determining relationship strength requires sophisticated analysis of interaction patterns and communication frequency.

Table 2: Index Architecture Component Performance Characteristics (Wang, Y. *et al.*, 2024; Gao, X., & Qiu, J. 2014).

Partitioning Strategy	Content Coverage	Query Efficiency	Maintenance Complexity	Routing Accuracy
Activity-Based	Comprehensive	High	Medium	High
Engagement-Based	Selective	Very High	Low	Medium
Graph Proximity	Targeted	Medium	High	Very High
Hybrid Approach	Complete	High	Very High	High

Index Architecture Design

A hierarchical organization solves the coordination problem associated with managing multiple specialized systems. Without proper federation, a user or application would need to know exactly which index contains the material being sought. The meta-index approach acts like an intelligent dispatcher, analyzing queries and routing them appropriately.

Content duplication across multiple indices initially seems inefficient, but redundancy serves important purposes beyond simple backup functionality. Popular content benefits from being immediately available in multiple locations, reducing access times during peak demand periods. Strategic replication also enables graceful degradation when individual components require maintenance or experience failures.

Dynamic Routing Logic

Query interpretation remains one of the most challenging aspects of distributed indexing systems. Understanding user intent from minimal input requires combining multiple signals, including search terms, user history, temporal context, and behavioral patterns. It is important to recognize that related queries raised may be interpreted very differently when we account for who poetically points to the inquisitions and at what time.

In a multi-indexed context, timing and synchronization difficulties don’t exist in a monolithic system. Various specialized indices could be slower to respond, have different levels of data freshness that may be accumulated at different times, or return even completely different data formats in their searches of the same related queries. So while simultaneous searches over a number of sources could make a more coordinated approach to data file comparisons, you do not want to combine and retrieve data from the various sources as late or delayed, and thus inconsistent.

Another terrain we have explored around the idea of multi-indexing is whether parallel processing may help to offset the slow performance of multi-index queries, as the trade-off of coordination overhead often negates any performance gains and diminishes faster time efficiency processes. Sequential approaches offer more predictable behavior at the cost of potentially longer processing times. The optimal strategy depends heavily on specific query patterns and system load characteristics.

Result aggregation from diverse sources requires understanding both the quality characteristics of each index and the preferences of individual users. Simple ranking algorithms break down when combining results with different relevance signals and confidence levels (Gao, X. & Qiu, J. *et al.*, 2014).

Publisher: SARC Publisher

IMPLEMENTATION CONSIDERATIONS AND TRADE-OFFS

Index Granularity Optimization

Index granularity decisions often make or break entire projects. The temptation to create highly specialized indices for every conceivable content type leads down a rabbit hole of complexity that few organizations can manage effectively. Real-world experience shows that the theoretical benefits of extreme specialization rarely materialize in practice.

Most engineering teams discover this lesson the hard way. Initial prototypes with dozens of micro-indices look impressive in design documents but become unmanageable once deployed at scale. The coordination overhead between systems grows exponentially, not linearly. Simple operations that once took milliseconds suddenly require complex orchestration across multiple subsystems.

The successful approach involves finding balance points where specialization provides genuine benefits without creating operational nightmares. Platforms that survive long-term typically settle on moderate granularity levels that capture meaningful content distinctions while remaining manageable by human operators.

Adaptive granularity systems represent an interesting evolution in this space. Rather than forcing engineers to predict optimal partitioning strategies upfront, these systems observe actual usage patterns and adjust their internal organization accordingly. The concept sounds sophisticated, but implementation requires considerable expertise to avoid creating systems that constantly reorganize themselves (Leung, C. H. *et al.*, 2012).

Table 3: Trade off of various index granularity techniques

Implementation Factor	Operational Complexity	Performance Benefit	Resource Requirement	Maintenance Overhead
Fine Granularity	High	High	High	High
Moderate Granularity	Medium	Medium	Medium	Medium
Coarse Granularity	Low	Low	Low	Low
Adaptive Systems	Very High	Very High	Very High	Medium

Index Refresh Mechanisms

Data freshness versus system stability presents an eternal engineering dilemma. Users expect immediate visibility of new content, but aggressive update strategies can destabilize carefully optimized index structures. The naive approach of rebuilding indices whenever content changes simply doesn't scale beyond trivial dataset sizes.

Incremental update mechanisms alleviate the immediate technical issue, yet create subtle consistency problems that are often difficult to debug. Content available in one index but not in another index presents usability challenges that are difficult to deliver to nontechnical stakeholders. Background processing helps alleviate this by pushing computationally intensive work to non-peak times, although this quickly loses effectiveness as platforms grow globally and bump into continuous traffic.

Delta indexing approaches have gained popularity as a compromise solution. Maintaining separate structures for recent content allows systems to provide immediate visibility while deferring expensive optimization operations. The periodic merging process requires careful scheduling to avoid impacting user-facing performance, particularly during unpredictable traffic spikes.

Memory and Computational Considerations

Resource allocation across multiple specialized indices resembles an elaborate balancing act that never reaches equilibrium. Popular content categories consume disproportionate resources while niche areas sit mostly idle. The asymmetric nature of social media content makes traditional capacity planning approaches inadequate.

Horizontal scaling provides the most obvious solution, but coordination costs increase dramatically as system components proliferate. Each additional server introduces new failure modes and monitoring requirements. Load balancing across heterogeneous indices requires advanced algorithms that incorporate the inherent features of different types of content.

Auto-scaling systems make it easy to worry less about resource management on behalf of their users, but they also offer end-user tuning options that should be done with the understanding of performance traffic patterns, response from the

Publisher: SARC Publisher

City, and the City's behavior generally. Aggressive scaling heuristics waste money by provisioning unnecessary resources. Conservative scaling policies leave the end-user spinning in the queue during traffic spikes. Most organizations end up with a hybrid mix of automated response along with human oversight in situations that are unusual.

Quality Assurance and Monitoring

Performance monitoring becomes exponentially more complex as systems fragment across multiple specialized components. Traditional methods of monitoring perform pretty well when we are dealing with a monolithic system. Those same methods break down however with a distributed architecture where each of the components of the distributed architecture have different behavior patterns.

Effective A/B testing is simply not possible without a purposeful experimental design that takes into account the interactions that exist between multiple components of the overall system. Simple before-and-after comparisons become insufficient when changes affect multiple indices simultaneously. Statistical significance becomes harder to achieve when effects are distributed across multiple subsystems with different baseline performance characteristics.

Long-term impact assessment reveals whether optimization efforts actually improve user experiences or just move performance bottlenecks around the system. Some improvements that look impressive in short-term testing disappear over longer time periods as users adapt their behavior to system changes (Aach, M. *et al.,* 2023).

PERFORMANCE EVALUATION AND RESULTS

Production-Scale Deployments

Running these systems in actual production environments taught some hard lessons about the gap between theory and practice. What looked elegant in architectural diagrams turned out to be messy and unpredictable when real users started hammering the system with unexpected query patterns.

The latency improvements were honestly surprising - better than anyone expected during the planning phase. Query response times dropped dramatically, especially for those complex searches that used to crawl along at painful speeds. Content ingestion became lightning fast, which made trending topic detection actually useful instead of perpetually running behind viral events.

Scalability turned out to be the real winner here. These distributed systems handle traffic spikes that would have melted traditional architectures. During major news events or viral content explosions, the system just keeps running smoothly instead of grinding to a halt (Goel, A. 2025).

Personalization Quality Metrics

User behavior changes tell the real story about whether these improvements actually matter. People started clicking on content more frequently, which sounds obvious but represents a fundamental shift in how effectively the system matches content to interests.

The diversity aspect was particularly interesting - users began exploring content types they'd never engaged with before. This wasn't just random browsing either. The system was finding genuine connections between seemingly unrelated topics that users actually found valuable. Someone interested in cooking might discover travel content focused on food culture, creating engagement patterns that traditional systems completely missed.

Fine grained indexing strategy also performs well to address the cold start problem with newer content types. Creating a separate index with fine grained indexing strategy solves the cold start problem.

Satisfaction surveys backed up what the behavioral data already suggested. People genuinely liked the content they were seeing more than before. The feedback loop became noticeably more positive, with fewer complaints about irrelevant recommendations cluttering their feeds.

System Reliability and Robustness

The fault tolerance capabilities exceeded even optimistic projections. When individual components failed, the system degraded gracefully instead of crashing spectacularly. This resilience proved invaluable during those inevitable midnight emergencies when something breaks at the worst possible moment.

Maintenance became significantly less stressful. Instead of coordinating complex system-wide updates that could break everything, teams could modify individual components without risking the entire platform. Debugging became much easier too - problems could be isolated to specific

Publisher: SARC Publisher

subsystems rather than requiring exhaustive detective work across monolithic architectures.

Trade-off Analysis

The complexity versus performance relationship revealed some uncomfortable truths about system design. More specialized indices don't always mean better performance - there's definitely a sweet spot where additional complexity starts working against you rather than for you.

Economic analysis showed favorable returns despite the higher upfront costs. The improved user engagement translated directly into better business metrics, justifying the additional engineering investment and operational complexity (Azzaz, A. *et al.*, 2025).

Comparative Analysis

Direct comparisons with alternative approaches confirmed what the theoretical analysis suggested. Fine-grained systems consistently outperformed both simple partitioning schemes and more complex hybrid architectures across most meaningful metrics. The industry adoption patterns speak for themselves - platforms that successfully implement these approaches tend to stick with them long-term, suggesting genuine practical value rather than just theoretical benefits.

Table 4: User Experience Quality Metrics and Engagement Improvements (Goel, A. 2025; Azzaz, A. *et al.*, 2025).

Performance Dimension	Traditional Systems	Fine-Grained Systems	Improvement Category	User Impact
Query Latency	High	Low	Significant	Positive
System Scalability	Limited	Extensive	Major	Very Positive
Content Relevance	Moderate	High	Substantial	Positive
Fault Tolerance	Low	High	Major	Very Positive
Maintenance Efficiency	Poor	Excellent	Significant	Neutral

CONCLUSION

The advancement of fine-grained indexing strategies marks a pivotal evolution in social media infrastructure design, fundamentally transforming how platforms manage content organization and user personalization at scale. These specialized indexing architectures address critical limitations inherent in traditional monolithic systems by enabling intelligent content segregation and optimized processing pathways tailored to specific content characteristics and user interaction patterns. The distributed nature of fine-grained systems provides enhanced resilience against failures while facilitating more efficient maintenance operations and independent component updates.

The transition from unified indexing structures to specialized distributed architectures represents a paradigm shift that acknowledges the inherent diversity of social media content and user behaviors. Through strategic partitioning based on activity types, engagement levels, and social graph relationships, platforms can deliver more relevant content recommendations while maintaining system performance under high-volume conditions. The implementation of dynamic routing mechanisms ensures optimal query distribution across specialized indices, maximizing retrieval efficiency and minimizing response times.

Production evidence consistently demonstrates the superiority of fine-grained approaches across multiple performance dimensions, including system scalability, fault tolerance, and user engagement quality. The enhanced personalization capabilities enabled by specialized indexing contribute to improved user satisfaction and platform retention, while the distributed architecture provides operational flexibility essential for managing complex social media ecosystems. The careful balance between system sophistication and operational complexity remains crucial for successful implementation, requiring thoughtful consideration of granularity levels and resource allocation strategies.

Future developments in fine-grained indexing will likely incorporate advanced automation capabilities, machine learning-driven partitioning decisions, and integration with emerging technologies to further enhance system intelligence and adaptability. The continued evolution of these strategies will play an essential role in supporting the next generation of social media platforms as

Publisher: SARC Publisher

they scale to accommodate ever-growing user bases and increasingly diverse content formats.

REFERENCES

1. Ravikumar, G. "Content Scaling: 10 Strategies for Sustainable Growth," *GravityWrite*, (2025).: https://gravitywrite.com/blog/content-scaling
2. Saoudi, E. M., & Jai-Andaloussi, S. "A distributed content-based video retrieval system for large datasets." *Journal of Big Data* 8.1 (2021): 87.
3. Hayat, M. K., Daud, A., Alshdadi, A. A., Banjar, A., Abbasi, R., Bao, Y., & Dawood, H. "Deep Learning for Social Media Analytics: a Survey." *IEEE Access* (2019).
4. Ai, L., Cheng, H., Wang, X., Chen, C., Liu, D., Zheng, X., & Wang, Y. "Approximate nearest neighbor search using enhanced accumulative quantization." *Electronics* 11.14 (2022): 2236.
5. Wang, Y., Gou, X., Xu, X., Geng, Y., Ke, X., Wu, T., ... & Wu, X. "Scalable community search over large-scale graphs based on graph transformer." *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. (2024).
6. Gao, X., & Qiu, J. "Supporting queries and analyses of large-scale social media data with customizable and scalable indexing techniques over NoSQL databases." *2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE, (2014).
7. Leung, C. H., Chan, A. W., Milani, A., Liu, J., & Li, Y. "Intelligent social media indexing and sharing using an adaptive indexing search engine." *ACM Transactions on Intelligent Systems and Technology (TIST)* 3.3 (2012): 1-27.
8. Aach, M., Inanc, E., Sarma, R., Riedel, M., & Lintermann, A. "Large scale performance analysis of distributed deep learning frameworks for convolutional neural networks." *Journal of Big Data* 10.1 (2023): 96.
9. Agrawal, P. "Optimizing Performance in Hybrid Cloud Architectures: A Deep Dive into Latency Management." *Sarcouncil Journal of Engineering and Computer Sciences* 4.6 (2025): pp 89-96
10. Goel, A. "System Design for Social Media Platforms Like Instagram," *Get SDE Ready*, (2025) https://getsdeready.com/system-design-for-social-media-platforms-like-instagram/?srsltid=AfmBOooplA7lCNn11zmoyTVKLkjJMo5PDBJbrL1r9J7xZJ8uEJkvTHcF
11. Azzaz, A., Malki, M., Slama, Z., & Dennouni, N. "Social Information Retrieval using Linked Data and Deep Learning." *Engineering, Technology & Applied Science Research* 15.3 (2025): 23360-23366

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:

Raveendran , N. " Fine-Grained Indexing Strategies for Social Feeds ." *Sarcouncil Journal of Engineering and Computer Sciences* 4.7 (2025): pp 157-164.

Publisher: SARC Publisher