



TV INSIGHTS – APPLICATION OF BIG DATA TO TELEVISION

Bhavan Gandhi, Alfonso Martinez-Smith, & Doug Kuhlman

ARRIS, USA

ABSTRACT

Television continues to evolve from pure consumptive linear viewing to web-like interactive experiences. Emerging new applications, such as multi-screen discovery, consumption, and network DVR, all allow and support user interaction. This is possible due to the evolution of TV application architectures from dedicated hardware-centric functionality to a combination of hardware appliances controlled by modular software services. Big Data technologies can support these rich, interactive TV experiences by collecting, storing, and analysing federated events and creating usable information for end-consumers, operators, and programmers. This paper provides an application's view into how Big Data enables insight from user behaviour, network / appliance operations, and content analysis.

INTRODUCTION – INTERSECTION OF TV & BIG DATA

Overview of Big Data

The advent of the Internet and the resulting volume of data that needs to be managed, stored, and analysed led to the emergence of modern day Big Data technologies. Yahoo and Google developed breakthrough technologies on processing Internet scale data and storing the processed results on distributed commodity servers to enable high availability and scalability. Google File System (GFS) [4] and Big Table [5] showcase early pioneering work that was developed for storing data. This is the precursor to some of the mainstream, open source big data related technologies that have been developed and evolved by the open source community like Apache. Apache HBase [6] and Apache Hadoop Distributed File System (HDFS) [7] are open source equivalents of Big Table and GFS respectively. A typical Big Data technology stack is comprised of an event intake layer facilitating the input of multi-formatted information from a variety of sources, and a core Hadoop eco-system for storing and processing collected information. Analysis of the collected data provides meaningful understanding and insight through dashboards, reports or application-facing interfaces.

TV – Sources for Big Data

The end-to-end TV ecosystem is a complex audio and video delivery system that carries source multimedia content that is encoded and distributed to a head-end. Within the head-end, the content is further manipulated or formatted for delivery to users on their customer premise equipment (CPE) device via the core network. This end-to-end functionality can be logically thought of as a vertical decomposition into application, control, video, and network layers.




Figure 1 – The TV Ecosystem

Instrumenting these various layers allows access to rich information that can be harnessed by Big Data. Client-facing applications resident on a CPE device or a web-application can be instrumented to collect user interaction and content consumption information. This provides insight into user behaviour. An instrumented control plane, used for managing subscribers and creating network centric applications such as nDVR and targeted ad insertion, can help drive operational understanding and efficiencies. Instrumented video and network layers provide insight into sub-program information, appliance operation (encoders, network recorders, packagers, and delivery) and network QoS.

DERIVING INSIGHT FOR EMERGING TV SERVICES

Quantifying User Interaction

Instrumented client applications, whether they are on traditional CPE or multi-screen devices, help quantify user interaction behaviour. They enable measuring traditional events like user navigation, content discovery, session start, tuner change, and direct interaction during content play. It is also possible to analyse the collected events to ascertain more abstract and aggregated consumption behaviour such as program popularity by geographic region.

Event monitoring and tracking systems facilitate user interaction logging in an unprecedented way. Keeping in mind user privacy concerns, known hashing and security techniques are employed to protect subscriber anonymity and provide data security. Even in this form, meaningful operational and content utilisation information can be determined for prescriptive and predictive purposes. Some examples of the types of analyses include number of active users at any given date and time, program/channel popularity, channel tuning persistence rates, and recorded content peak playback statistics. This information itself is operationally useful and is typically presented via dashboards and reports.

Collecting and analysing subscribers' consumption and usage is an essential step in creating targeted experiences. Consumer-facing applications and services may also access actionable information for individual subscribers through defined interfaces, allowing customised user experiences and applications. Subscriber consumption patterns can be supplemented with appropriate subscriber profile information accessed through APIs to filter and construct customised application views for each user. An example of this

would be dynamically creating a “My Favourite Channels” view by combining all the above with a user’s channel usage patterns. This particular example would be possible through the use of recommender systems using machine-learning techniques, applied to the mined data and usage patterns.

Operational Analytics for quantifying network DVR (nDVR)

As TV continues to move from linear consumptive to one that offers interactivity, nDVR is an emerging application that offers end-consumers flexibility to record in the cloud and subsequently playback content on a variety of devices. This drives efficiencies for the operator in providing thinner end-clients, while optimising content storage and delivery. Conceptually, an nDVR system is comprised of consumer facing application(s) that interfaces with Backoffice services, which control physical recording and delivery sub-systems.

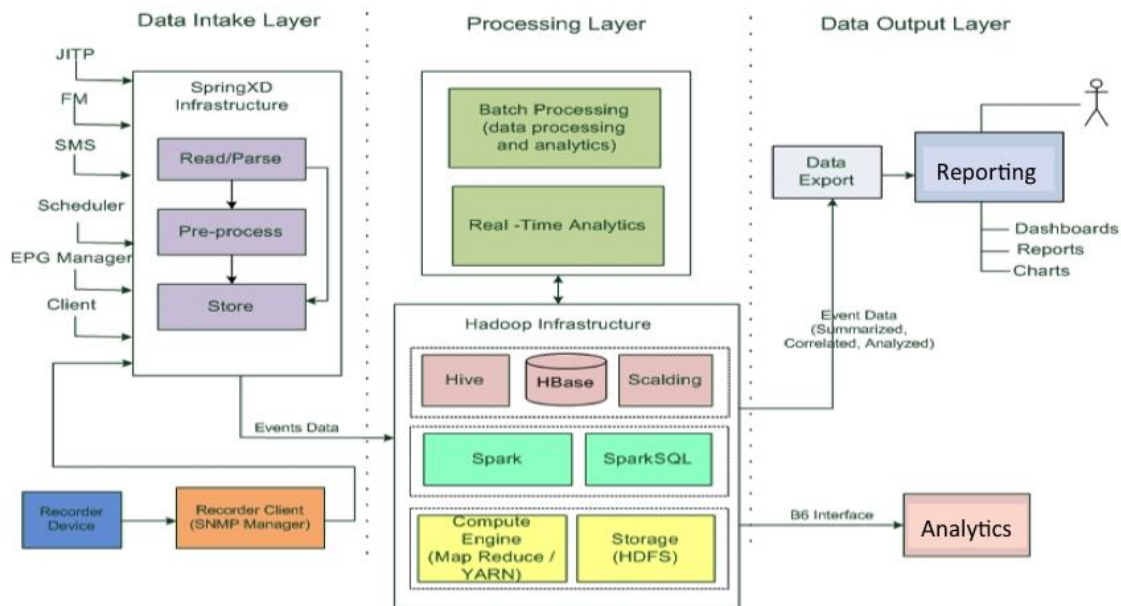


Figure 2 – nDVR High-level Analytics Architecture

Figure 2 shows a high-level architecture of an Analytics system for nDVR. The nDVR sub-components are instrumented to push events into the Data Intake Layer in a federated manner. The instrumented sub-components or services are:

- Client Application – The client or user facing application captures user interactions with the application / content including navigation and trick play behaviour
- Backoffice Services – These are comprised of subscriber (SMS), content guide / electronic program guide (EPG) information, content recording (Scheduler), and content playback / fulfilment (FM) transactions
- Data Plane Appliances – The recorders are instrumented to provide physical recording statistics and the just-in-time packaging (JITP) system provides physical delivery information

Events are collected from the various sub-components and services in a variety of formats. These transactions are stored in HDFS and processed real-time or batched to create actionable insight. The specific analysis jobs result in creating meaningful dashboards and reports providing useful operational and usage information.

The objective of any analytics system is to drive meaningful understanding into how the application, services, and hardware are being used. Operationally, such a system could be used to predict system scaling to meet demand.

Media Analysis – Deriving Sub-program Content Structure

Beyond usage and operational transactions, usable information can be derived from the content itself by employing media analysis techniques on the various components of the content [8].

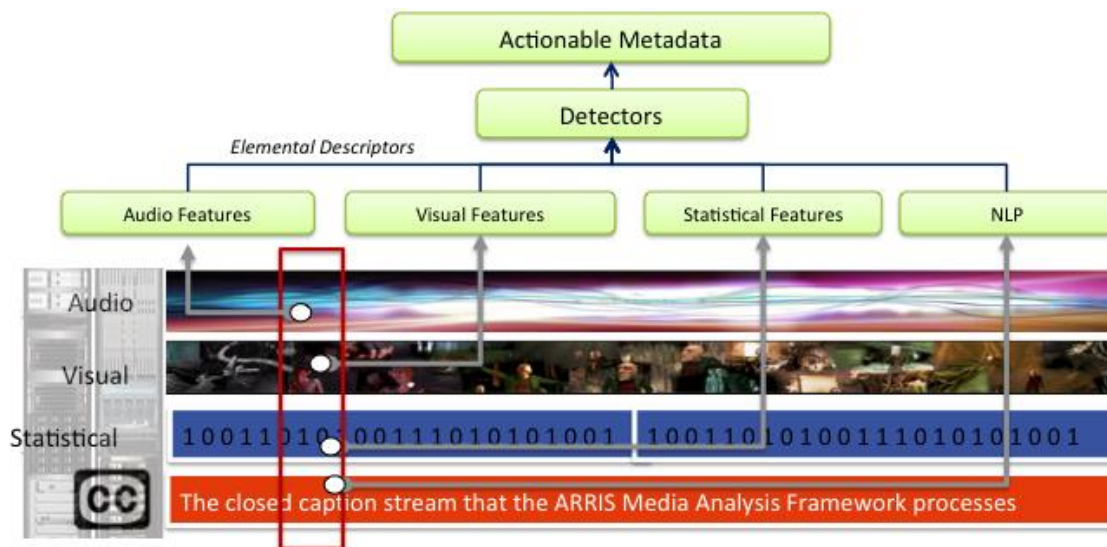


Figure 3 - Media Analysis – High Level Process

Figure provides a high-level processing framework for analysing content. The Media Analysis process is a suite of event detectors (audio, video, statistical, natural language) that feed into a fusion engine with the goal of semantically labelling the content with actionable metadata. Combining different types of detectors provides a simple yet powerful mechanism to align events from different modalities to yield interesting results (e.g., detecting black video frames and quick silence audio intervals to infer ad boundaries). Other detector combinations can be used, often in hierarchical structures, to obtain semantically rich events that are present in the content and would be otherwise hard to extract. Examples include creating sports highlight reels, harvesting news, or classifying weather stories by topic. Combining logos, camera angles, telestrations, and speech patterns to infer a particular channel’s creative features in TV productions of different genres does this.

While different TV program features are interesting in and of themselves, combining the usage behaviour and operational information with semantically labelled content can provide meaningful insight across the various vertical layers of the TV ecosystem.

EXAMPLE INSIGHTS

Information that is collected and processed must provide actionable value to the user, operator, or programmer. This section shows example dashboards derived from actual data providing examples of meaningful insights.

User Recording & Playback Behaviour

Understanding how users interact with content on their CPE device / application is a fundamental starting point to gain a better understanding of user behaviour. Figure 4 shows an aggregate view of typical user behaviour over a time span. It shows DVR Recordings, DVR Playback, and VOD Plays.




Figure 4 - User Activity Dashboard

This dashboard shows DVR usage is predominant in comparison to traditional VOD consumption; ten times as many programs are viewed via a user's DVR system than are viewed via VOD. Also, the number of recordings outnumber playback consistently indicating that not all recordings are played back. Lastly, there is a cyclic pattern with recordings whereas playbacks are uniformly distributed; this corresponds to more recordings being instantiated during the week than on weekends.

Detailed, empirical evidence such as this provides insight into how DVR systems should be sized, bandwidth efficiencies that could be garnered if an operator moves to an nDVR system, and understanding detailed recording and playback indirectly provided insight into content affinity.

Operational Insight - nDVR

Operationally, it is important to understand how the data plane is being utilised to help with capacity planning and management. Peak usage information is more insightful than average activity or total activity.

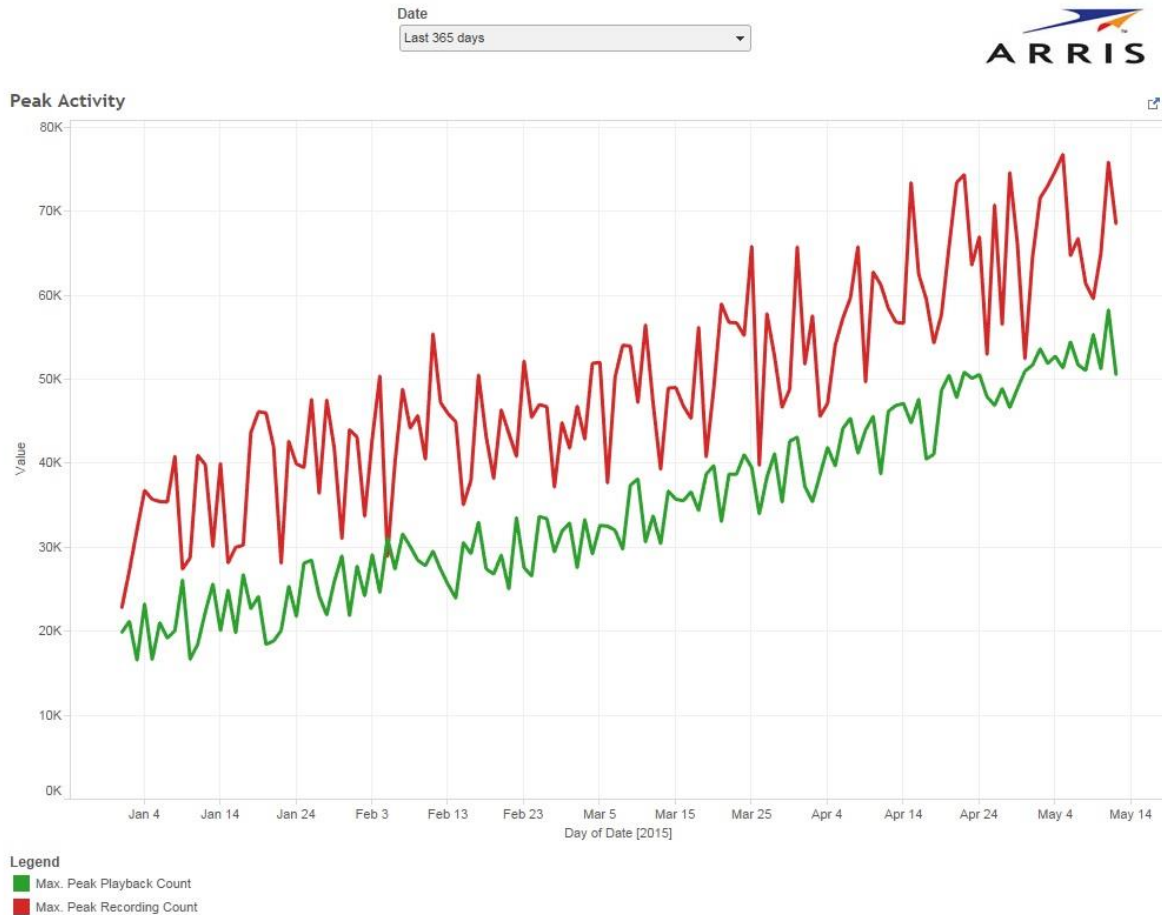


Figure 5 – Peak User Recording & Playback Activity

Figure 5 shows the trend of peak user recording and playback rates over time. Analysis of this data indicates that both playback and recording rates are increasing. Unsurprisingly, recording peaks surpass playback peaks and show a similar cyclical behaviour as in Figure 4.

Analysis of this sort is invaluable for capacity management and planning of additional resource needs. In this example, if the system can support a peak for both recordings and playbacks of 100K concurrent sessions, the danger of hitting that maximum for recordings is imminent. Peaks tend to present spikes and a single major content program event could surpass the system's limitations. Playback rate is safer, currently, though it is growing faster (additional analysis shows rate of growth) and will probably surpass 200K concurrent sessions before recordings will.

Bridging User Behaviour with Content Structure

Big Data enables analysing information from disparate sources to help glean insight across different verticals in a TV ecosystem.

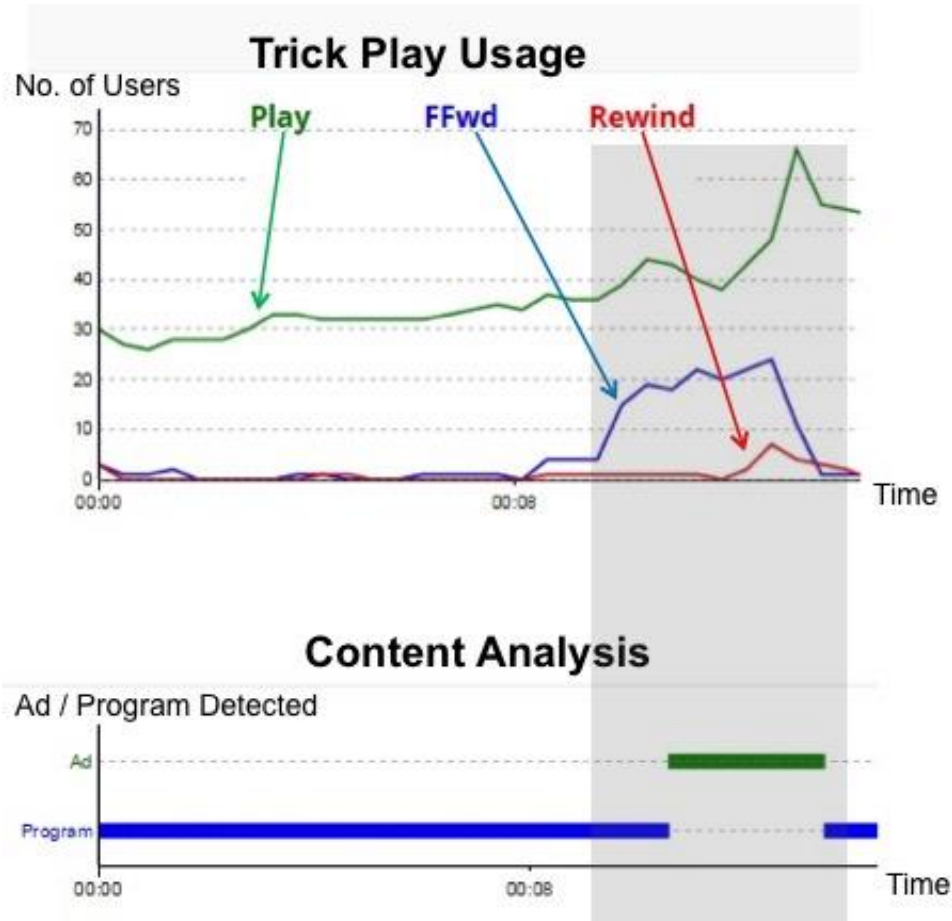


Figure 6 - User Behaviour Correlated with Content Structure

Figure 6 shows the insight afforded by combining content structure information with usage behaviour. Content analysis, as performed by the Media Analysis, was used to automatically characterise and annotate content segments as Ad or Program segments. Usage analytics for the same programming provides insights into the way consumers interact with the content. Examining information in a superimposed manner from the two sets of processed events provides content owners, producers, and operators with actionable insight into the actual value placed by consumers on specific content segments; through this dashboard, it can be determined what percentage of the viewers skipped through the ads. The data presented here is for a specific asset from a two-week window of operational data in 2015.

Beyond simply characterising ad-skipping behaviour, information about content structure and ad segments allows operators and programmers to surmise the relative value of sub-program segments, by tabulating the number of viewers who watched, who fast-forwarded through it, or who performed a rewind operation to watch it again.

CONCLUSIONS

Big Data technologies were developed primarily to target specific needs of the Internet. However, as our TV delivery systems grow more complex and as the services evolve to be more web-like, there is direct applicability and flexibility in using these technologies for television systems. Actionable insight derived from collecting, processing, and analysing federated sets of events provides several levels of valuable information. User behaviour prediction, operational planning, and program/sub-program affinity are just some of the areas that can be forecast with this approach.

In this paper, we provided an overview of the Big Data technology stack, discussed the sources of events in the TV system, and provided examples of meaningful conclusions that can be ascertained from analysed events across the various vertical layers of application, control, network, and content planes. The combination of these insights, together with state-of-the-art recommender systems and machine learning strategies, will be essential in driving operational efficiencies.

REFERENCES

1. Arthur, L. (2013, August 15). What is big data? Retrieved from <http://www.forbes.com/sites/lisaarthur/2013/08/15/what-is-big-data/>
2. Coutts, A. (2012, September 25). Breaking down 'big data' and Internet in the age of variety, volume, and velocity. Retrieved from <http://www.digitaltrends.com/web/state-of-the-web-big-data/>
3. Laney, D. (2001). 3D data management: Controlling data volume, velocity, and variety. In *Application Delivery Strategies*. Stamford, CT: META Group, Inc.
4. Ghemawat, S., Gobioff, H., & Leung, S. (2003, October). The Google file system. ACM - SOSP '03, Bolton Landing, NY.
5. Chang, F., Ghemawat, S., Hsieh, W., Wallach, D. A., Burrows, M., Chandra, T., Fikes, A., & Gruber, R. E. (2006, November). Bigtable: A distributed storage system for structured data. OSDI'06: Seventh symposium on operating system design and implementation, Seattle, WA.
6. Apache hbase. (2014, March 17). Retrieved from <http://hbase.apache.org>
7. Apache hadoop. (2014, March 14). Retrieved from <http://hadoop.apache.org>
8. Liu, Z., and Wang, Y., Audio Feature Extraction and Analysis for Scene Segmentation and Classification, *Journal of VLSI Signal Processing Systems for Signal, Image, and Video Technology*, Kluwer Academic Publishers, Vol. 20, Issue 1-2, pp. 61-79, 1998.

ACKNOWLEDGEMENTS

The authors wish to thank their colleagues in the Applied Research Center at ARRIS for their contributions to this work.