The long, interesting tail of Indie TV

Daniel Cutting, Björn Landfeldt[1]*, and Aaron Quigley[2]

[1] School of Information Technologies, University of Sydney, Australia,
{dcutting,bjornl}@it.usyd.edu.au
[2] School of Computer Science & Informatics, University College Dublin, Ireland
aquigley@ieee.org

Abstract. We describe a scenario of independent TV programme dissemination where viewers receive their own tailored channel created from content in which they're interested. We focus primarily on the problem of distributing the content from the producers to all interested viewers. Our approach is based on a multicast scheme called *implicit group messaging*, realised as a distributed structured peer-to-peer network.

1 Introduction

Anderson's "Long Tail" [1], is a term used to describe the shape of the curve typically formed when products, such as books or CDs, are plotted in order of popularity. This shape, essentially a power-law distribution, has analogues in many fields, notably the Zipf, Pareto, Benford and Bradford distributions, amongst others. Anderson notes that physical stores can only afford to reserve shelf space for products that are guaranteed to be popular, whereas online stores can typically afford to sell rarer products, since delivery of the product can be arranged on an ad hoc, on-demand basis and there are no physical storage constraints.

A similar situation arises in the realm of television and films. It is generally very expensive to create, programme and broadcast a television channel, which means the content must be targeted at a broad range of viewers in order to recoup the costs and make a profit. However, if the cost for creation and dissemination were substantially reduced, it would be possible for many more in-depth channels to exist catering to niche interests and demographics.

The cost of creation has already dropped substantially. Consumer-level equipment is now of excellent quality, and the tools necessary for production have become affordable and capable of running on a standard desktop PC or laptop. Distribution too is becoming cheaper. With the rapid growth of the Internet and mass adoption of broadband connectivity, dozens of specialised TV channels such as Sail.tv and YUKS TV[3] have started to stream video to small audiences of highly interested viewers. Platforms like Democracy[4] also allow people to make their own channels, based on RSS feeds, to which anybody with a reader/viewer can subscribe.

* This work was supported by the Smart Internet Technology CRC.
[3] http://sail.tv, http://yuks.tv
[4] http://www.getdemocracy.com

2 Indie TV

We suggest a model of independent ("Indie") programming taken to its logical conclusion — a tailored TV channel for each viewer based specifically on their distinct and variable interests.

Such a personalised channel could adapt to a viewer's context by interfacing with other components such as calendars and "to do" lists in order to know what sorts of interests will be most applicable for any given point in time. This could allow the channel to act as a news/stock ticker while the viewer is working, a dance music video channel when they are hosting a party or a movie channel when relaxing after a day of work. A calendar event for a tramping trip to New Zealand could allow the channel to tailor the content over the preceding week to include imagery and information about the Great Walks, florae and faunae. Synchronising and interfacing with portable devices such as iPods would also allow the channel to keep track of what has already been seen (when driving or on the bus for example) and continue playback from where the viewer left off.

This concept requires several components to work in concert. Firstly, content needs to be **created**. Secondly, it needs to be **disseminated** to interested viewers. Thirdly, it needs to be **blended** into a coherent channel for viewing on a TV or computer according to the viewer's context.

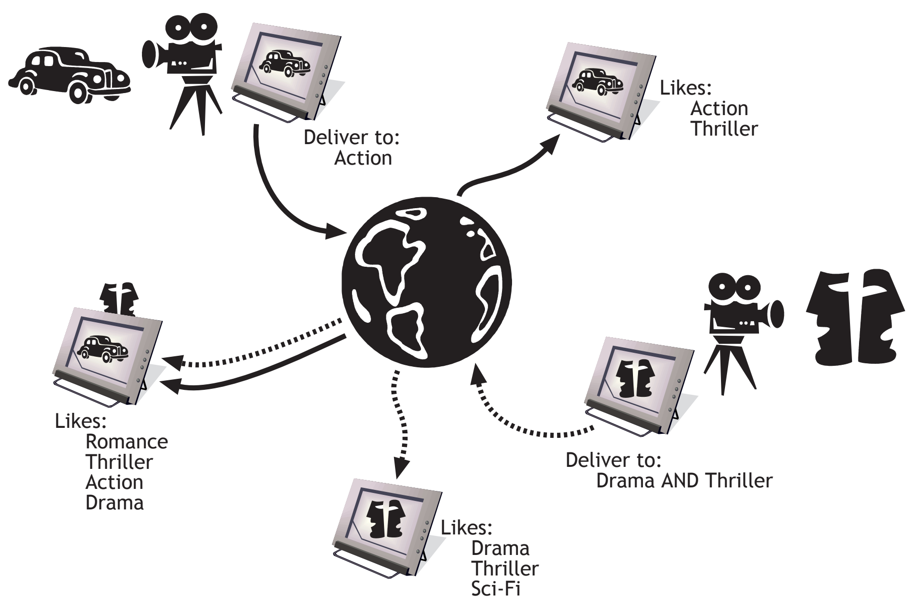

Fig. 1. The Indie TV architecture. Peers express their interests and producers target their content to groups of peers that are most appropriate. Movies are queued by recipients for blending and playback.

The first requirement is outside the scope of this article, but we assume that as the cost of creating content decreases, more varied niche content will be created. The third requirement could be tackled in several ways. One straightforward approach is to simply queue content as it becomes available and display it in sequence, although the precise technique employed could vary from viewer to viewer and is also outside the scope of this article. We focus on the problem of distributing content from producers to interested viewers.

Our approach to this problem is based on a form of many-to-many multicasting called *implicit group messaging* that allows any participant in a peer-to-peer (P2P) network to send a message to all members of a dynamically described group of peers, simply by specifying a combination of their interests. This fundamental mechanism (described in more detail in section 3) is used by producers in the Indie TV model to disseminate their content to all interested consumers as it is created. An agent on each consumer's device then takes this content and packages it as a continuously playing, tailored television channel for playback.

Figure 1 demonstrates the general functionality of this approach. The peers connect to the P2P network and express their interest in certain genres. One producer has made a short action film and thinks it will be appreciated by a wide range of action movie fans so it is sent to the implicit group of peers in the system that have indicated they enjoy action movies. Another producer has created a combination dramatic thriller, and delivers it to only those viewers that have expressed interest in **both** of these genres. Note that peers that have not expressed all of the required interest do not receive the film. One of the viewers has received both films and has queued the dramatic thriller for showing subsequent to the action movie.

3 Dissemination

Our P2P system is grounded in a theoretical model that provides certain guarantees for message delivery. In particular, it guarantees that all connected peers matching an implicit group specification will receive a message within a bounded number of hops. The number of non-members involved is minimised by the clustering of matching peers and amortisation of delivery paths.

The model is a distributed, structured overlay network comprising homogenous peers residing at locations on the surface of a d-dimensional torus, in the style of CAN [2]. Peers own patches of the surface (called *extents*) and are able to communicate with their neighbours. Messages can be routed across the surface to specific coordinates simply by passing them from neighbour to neighbour. Figure 2 shows an example of a surface consisting of several hundred peers.

A peer's location is calculated by hashing its interests into a Bloom Filter [3] and mapping the resultant bit string to a position on the surface (the black dots in the figure) using a PR quadtree representation [4]; thus a peer's location inherently encodes its interests. A new peer can join the network by requesting the peer currently managing the region covering its address to divide its extents and accommodate it as a neighbour [5].

When a peer wishes to publish content to a group of viewers (called *CAST*ing), it calculates the set of all possible extents where members of that group may reside using the same address encoding scheme, and then geometrically routes the message from neighbour to neighbour over the surface, branching as necessary to reach all matching peers. At each step, the targeted extents that are yet to receive the CAST are clustered together according to their angular difference from the peer processing the message, thus allowing amortisation of messages into single paths that branch only when the angle between extents exceeds a threshold (measured in radians) called the *branch factor*.

Intuitively, if multiple unexplored extents lie in the same direction away from the current peer (and hence have minimal angular differences), then only one message needs to be sent in that direction. As the message gets closer to the extents, the angular differences will become more marked, and once they exceed a threshold the message can be cloned and the route can branch as necessary.

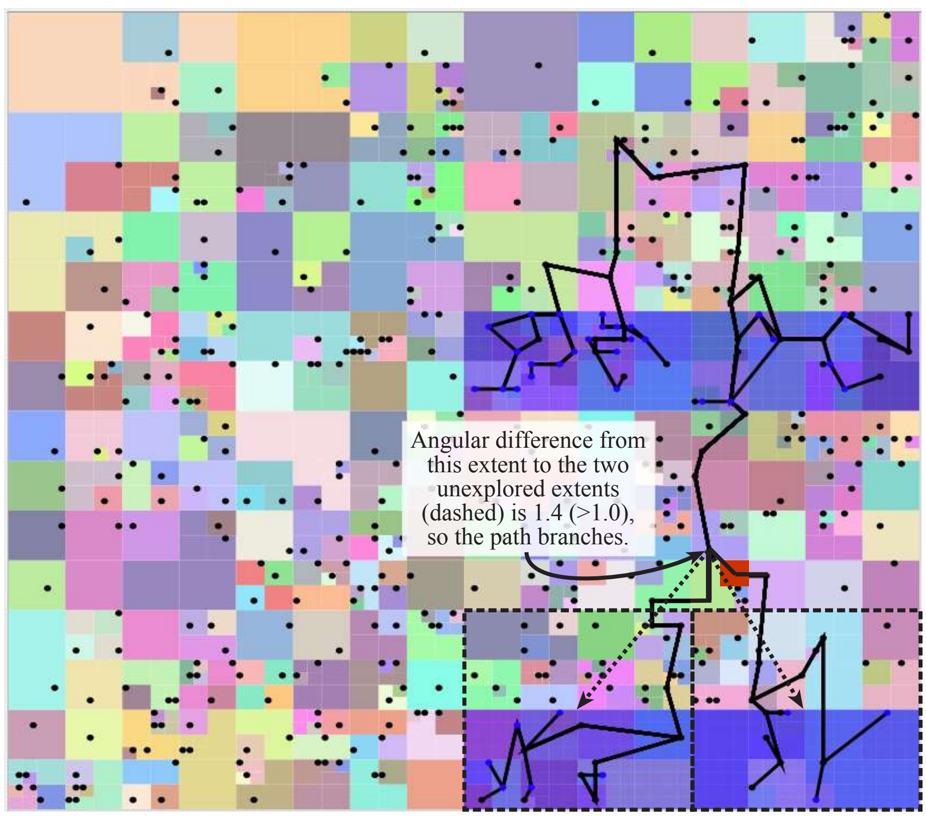

Fig. 2. A 2D surface containing several hundred peers. A CAST has traversed from the top-right of the surface to all peers in the shaded region which, due to their location, are calculated to be interested in dramatic thrillers.

Figure 2 shows a CAST with a branch factor of 1.0 radians. A producer at the top-right of the surface has CAST a new film to the implicit group of peers with an interest in drama and thrillers. The shaded region represents the extents of the surface where peers with both of these interests must reside, and the message is propagated towards them. At the branch point indicated in the figure, the CAST forks towards the two remaining regions since at this point the angular difference of the two clusters exceeds 1.0 radians.

4 Current and future work

We have investigated the performance of the model in terms of the number of peers that participate in delivering CAST messages and have found that as the specificity of the implicit groups increases (i.e. as more interests are combined), proportionately fewer peers partake in their dissemination. We have also found that the amount of storage required of each peer to maintain the system is unrelated to the size of the network, implying high scalability. We note that although implicit group messaging could deliver the actual video content to viewers, a more efficient approach may be to use IGM for notifying interested peers of content availability, and have them automatically download the content using a traditional content delivery network such as Akamai[5].

Current work includes constructing the Indie TV scenario, using our P2P model as the dissemination component, and strengthening the theoretical analysis of the model. We are also investigating adjustments to the model to improve the usage of the underlying physical network, and comparing our approach to other possibilities such as centralised server-based systems and unstructured P2P networks.

Future work will include using the implicit group messaging concept in other scenarios such as social collaboration software, and possibly constructing complete software for actual usage by the general public.

References

1. Anderson, C.: The Long Tail. http://www.wired.com/wired/archive/12.10/tail.html (2004)
2. Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S.: A scalable content addressable network. In: SIGCOMM'01, San Diego, California, USA, Berkeley, CA, ACM (2001)
3. Bloom, B.H.: Space/time trade-offs in hash coding with allowable errors. Commun. ACM **13**(7) (1970) 422–426
4. Orenstein, J.A.: Multidimensional tries used for associative searching. Inf. Process. Lett. **14**(4) (1982) 150–157
5. Cutting, D., Landfeldt, B., Quigley, A.: Implicit group messaging in peer-to-peer networks. Technical Report 583, University of Sydney, School of Information Technologies, University of Sydney 2006 (2006) ISBN: 1 86487 804 5.

[5] http://www.akamai.com