# Content-Based Networking


Antonio Carzaniga and Alexander L. Wolf

Department of Computer Science
University of Colorado
Boulder, Colorado USA
{carzanig,alw}@cs.colorado.edu


A prominent class of applications is emerging, one that is putting increasing pressure on the communication service models of traditional unicast and multicast networks. These applications are characterized by various combinations of the following properties:

- large and dynamically varying numbers of autonomous clients engaged in multi-party communication;
- wide geographical distribution;
- dynamic and unpredictable interaction patterns; and
- high rates of message traffic.

Examples include auctioning (e.g., eBay), information sharing (e.g., Gnutella), information fusion and dissemination (e.g., the U.S. Air Force Joint Battlespace Infosphere), sensor grids, personalized news distribution, service discovery, and multi-player games. What these applications have in common is a communication style in which the flow of messages from senders to receivers is determined implicitly by characteristics of the receivers, rather than explicitly through knowledge of destinations by senders.

## *A New Network Infrastructure*

We have developed a new kind of advanced communication infrastructure to support this class of applications. We refer to the infrastructure as a *content-based network*. A content-based network is a communication network based on a novel connectionless service model. Clients declare their characteristics to the network by means of *predicates*. A predicate, sometimes also referred to as a “profile”, describes the interests of a client. A particular message is delivered to a particular client if the client's predicate applied to the information *content* of the message yields true. In this way, the predicate determines which messages will flow to the client via the network. Together, the predicates of all the clients form the basis of the routing and forwarding functions in the network. A content-based network therefore avoids traditional destination address configuration in favor of a more abstract and powerful means of identifying intended recipients.

The service is provided as a network in the fullest sense. It is supported by an interconnection of store-and-forward message routers to achieve scalability, reliability, security, and administrative autonomy superior to that of any centralized service. The routers are robust and efficient, and are deployable along side traditional unicast and multicast services in the context of the existing network infrastructure.

Although we are introducing a new kind of network, it is conceived of in terms of well-established networking principles.

### Network Configuration

The address of a receiver in a content-based network is its predicate, indicating what messages are of interest. Similarly, the address of a sender in a content-based network is a predicate that describes the kinds of messages it can generate. Taken together, the service can perform a variety of optimizations in the creation, maintenance, and use of routing and forwarding information based on the predicates. The network is configured (i.e., client addresses are made know) through the presentation of predicates to the network.

### Subnetting

Subnetting simplifies network configuration (and reconfiguration), as well as reduces the computational costs of forwarding. Subnets in traditional networks are built into the structure of statically assigned addresses. Subnets in content-based networks, on the other hand, arise dynamically from similarities among the active predicates. For example, if a new client is added to the network whose predicate is the same or more specific than a predicate already advertised by a router from that side of the network, then there is no need to advertise the new predicate, since a message matching the more specific predicate would also match the more general and so reach the new client anyway. Thus, the more general predicate acts as a subnet address that subsumes the more specific predicate.

### Routing

Given our conception of addresses and subnets in a content-based network, the next challenge is to formulate an efficient and reliable means to determine routes, propagate the routing information, and maintain the routing information in the face of a highly dynamic constellation of predicates. We leverage existing protocols for the discovery and maintenance of basic network topology information (e.g., distance-vector and link-state protocols). We then map the subnet information, which arises from the predicates, onto the basic network topology. This mapping will evolve as the topology changes (e.g., because a link fails) and as the predicates change (e.g., because a client revises its predicate

### Forwarding

Routing decisions are distilled into forwarding tables located in each content-based router. The space occupied by these tables and the speed at which they can be used to make forwarding decisions are key to the efficiency of the content-based network. We are targeting the design of routers that can handle predicates numbered in the millions, and throughput rates of tens to hundreds of thousands of messages per second.

### Quality of Service

As in traditional networks, the basic content-based network is a best-effort service. Nevertheless, it is important to also provide differentiated services and levels of reliability. In particular, we envision priority and bounded-time message delivery mechanisms, as well as various reliability guarantees, such as “at most once” and “at least once”.

## *Comparison*

The term “content-based routing” is used in a number of areas with varying connotations. One example is the mechanism used in clusters of Web servers for mapping HTTP requests to individual servers within a cluster. The mechanism is called content-based routing because each HTTP request is directed by a gateway of some kind (e.g., a master server, a scheduler, or a switch) to a particular server based on the content of the request. Another example is that of multimedia information systems, in which the term content-based routing describes the processing, querying, and navigation of various documents based on the nature of their content. In both cases the “routing” is being performed, not to establish a communication path, but simply to demultiplex an incoming stream of content for specialized local processing. Moreover, in both cases something other than the network infrastructure is implementing the particular notion of content-based routing.

Application-level publish/subscribe middleware has a service model similar to that of a content-based network. Under this model, clients subscribe for messages either by attaching themselves to an identified stream of messages (“channel-based” subscription), by posing a predicate over a required field of all messages (“subject-based” subscription), or by posing a predicate over all fields of a message (“content-based” subscription). Clients send messages by publishing them to the common service, which is responsible for delivering the messages to the appropriate receivers based on their subscriptions. The CORBA Notification Service and the Java Message Service are both based on a publish/subscribe model, as are commercial products such as SoftWired's iBus, TIBCO's TIB/Rendezvous, Talarian's SmartSockets, Hewlett-Packard's E-speak, and the messaging system in Vitria's BusinessWare. Existing publish/subscribe middleware is built around a centralized message dispatcher that registers all subscriptions and processes every message. This solution is conceptually simple and allows for various optimizations in the matching function, however it cannot scale to large, wide-area, high-demand applications. A few of these systems have begun to offer a “federated” service, by which they mean that centralized dispatchers can be statically configured to share information. This is a stopgap approach that cannot adequately support the demands of highly dynamic and large-scale applications.

A multi-source multicast service (such as IP multicast) can be viewed as a special case of a content-based networking service. In fact, joining a group is equivalent to defining a receiver predicate that selects messages (i.e., datagrams) addressed to the group. Indeed, we believe that some ideas developed for multicast routing may be applicable to content-based routing. For example, content-based routing may use propagation trees built using the IP multicast infrastructure (or similar protocols). It should be clear, however, that the service model of a content-based network is substantially different from a network-level multicast, both in the expressiveness of the service interface and in the ultimate goals of the design. The primary use of IP multicast is as a vehicle for the transmission of streaming media channels. This is reflected in the simplicity of the interface, whereby clients can simply tune into one or more channels (i.e., join multicast groups). By contrast, the service model of a content-based network is oriented towards message-based applications that require more precise selection capabilities.