A refinement to this protocol, BGP-2, was described in RFC 1163 in June 1990,[7] and a further refinement, BGP-3, was described in RFC 1267 in October 1991.[8] The current version, BGP-4, was first deployed within the Internet in 1993. The RFC describing this protocol, RFC 1771,[9] was published in March 1995, and was subsequently refined with the publication of RFC 4271 in January 2006.[1] The core protocol has been stable for some years now, although further refinement has been undertaken through the use of negotiated capabilities undertaken at BGP session startup.

BGP is an instance of what we commonly refer to today as a *Bellman-Ford Distance Vector* routing algorithm.[10,11] This algorithm allows a collection of connected devices (BGP speakers) to each learn the relative topology of the connecting network. Its basic approach is very simple: each BGP speaker tells all its other neighbours about what it has learned if the new learned information alters the local view of the network. This scenario is a lot like a social rumour network, where everyone who hears a new rumour immediately informs all their friends. BGP works in a very similar fashion: each time a neighbour informs a BGP speaker about reachability to an IP address prefix, the BGP speaker compares this new reachability information against its stored knowledge that it gained from previous announcements from other neighbours. If this new information provides a "better" path to the prefix, then the local speaker moves this prefix and associated next-hop forwarding decision to the local forwarding table and informs all its immediate neighbours of a new path to a prefix, implicitly citing itself as the next hop. BGP keeps track of the propagation of route advertisements across the inter-domain space by recording the sequence of network *Autonomous Systems* (ASs) that propagate the route in a route attribute called the *AS Path*. A "better" route is one with a shorter AS path, and a loop is detected when a BGP speaker sees its own AS in the received AS Path (Figure 1).

Figure 1: The Propagation of a route in BGP

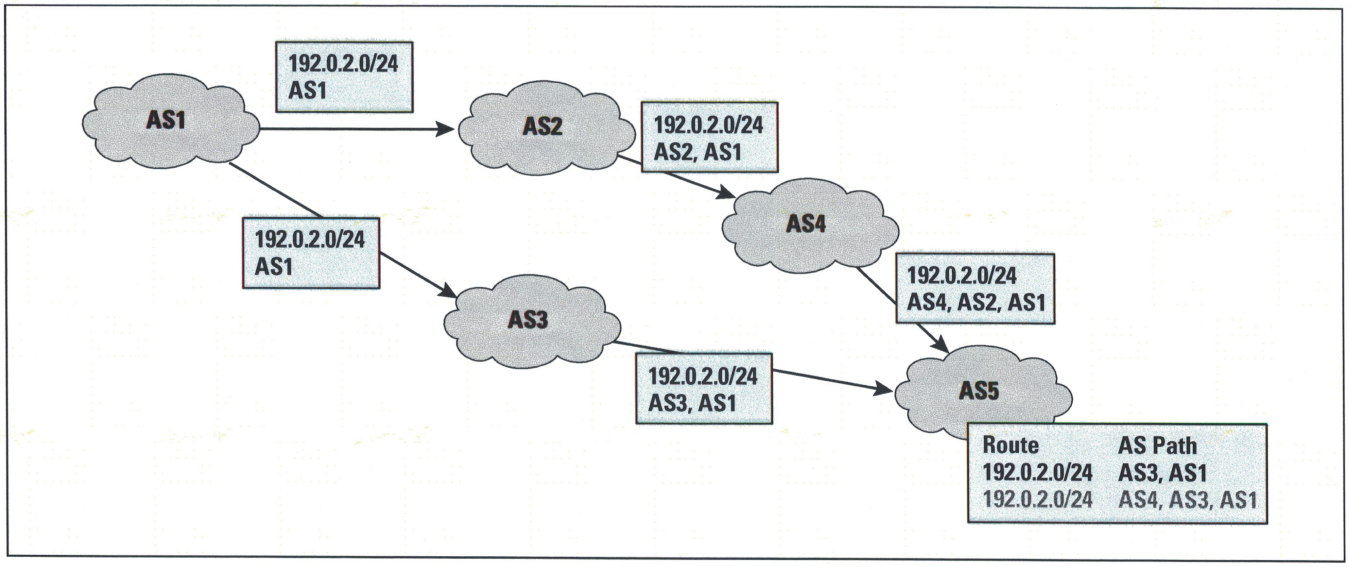

A refinement to this protocol, BGP-2, was described in RFC 1163 in June 1990,[7] and a further refinement, BGP-3, was described in RFC 1267 in October 1991.[8] The current version, BGP-4, was first deployed within the Internet in 1993. The RFC describing this protocol, RFC 1771,[9] was published in March 1995, and was subsequently refined with the publication of RFC 4271 in January 2006.[1] The core protocol has been stable for some years now, although further refinement has been undertaken through the use of negotiated capabilities undertaken at BGP session startup.

BGP is an instance of what we commonly refer to today as a *Bellman-Ford Distance Vector* routing algorithm.[10,11] This algorithm allows a collection of connected devices (BGP speakers) to each learn the relative topology of the connecting network. Its basic approach is very simple: each BGP speaker tells all its other neighbours about what it has learned if the new learned information alters the local view of the network. This scenario is a lot like a social rumour network, where everyone who hears a new rumour immediately informs all their friends. BGP works in a very similar fashion: each time a neighbour informs a BGP speaker about reachability to an IP address prefix, the BGP speaker compares this new reachability information against its stored knowledge that it gained from previous announcements from other neighbours. If this new information provides a "better" path to the prefix, then the local speaker moves this prefix and associated next-hop forwarding decision to the local forwarding table and informs all its immediate neighbours of a new path to a prefix, implicitly citing itself as the next hop. BGP keeps track of the propagation of route advertisements across the inter-domain space by recording the sequence of network *Autonomous Systems* (ASs) that propagate the route in a route attribute called the *AS Path*. A "better" route is one with a shorter AS path, and a loop is detected when a BGP speaker sees its own AS in the received AS Path (Figure 1).

Figure 1: The Propagation of a route in BGP

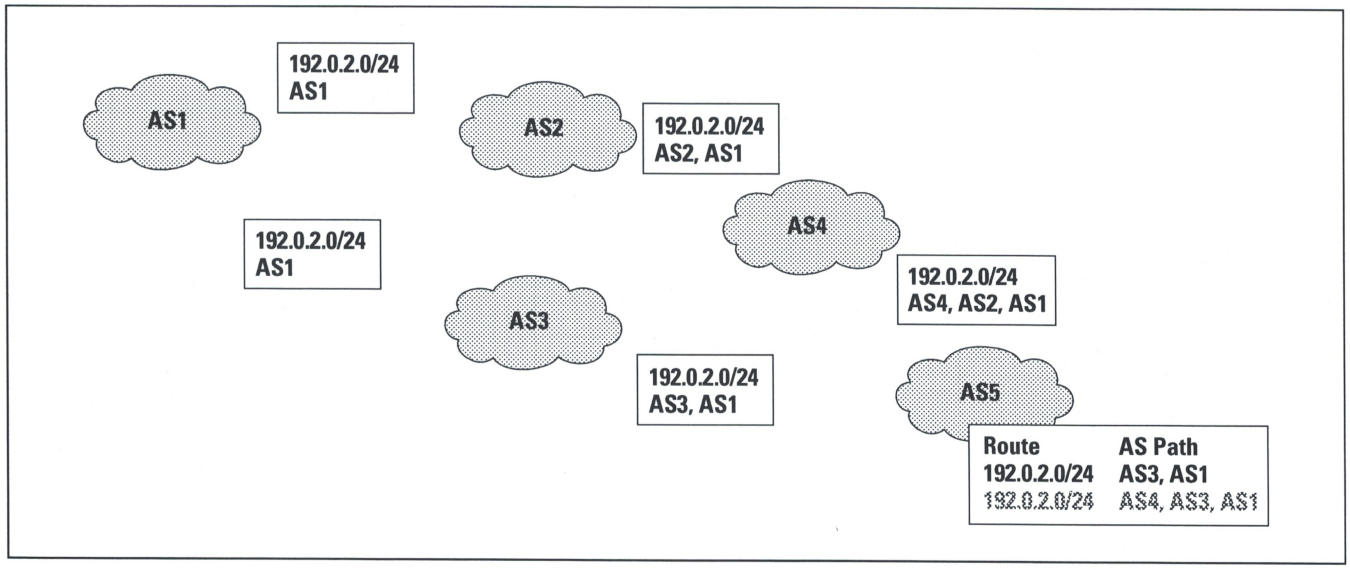

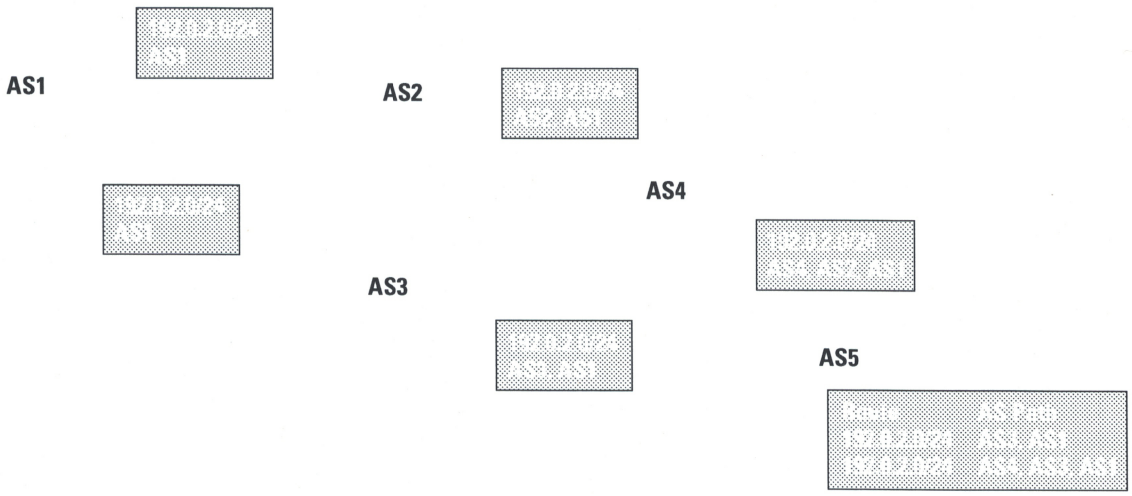
AS1
AS2
AS4
AS3
AS5

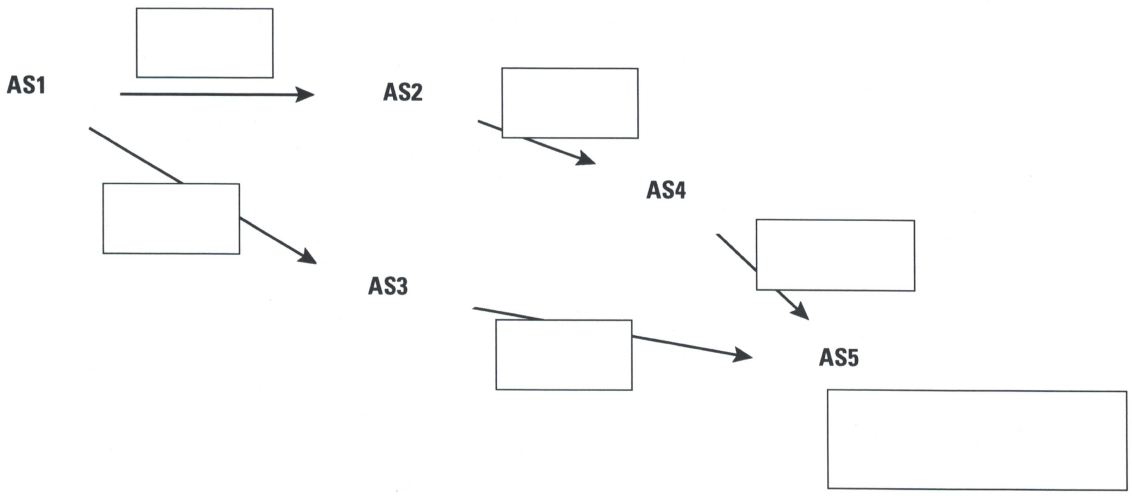
AS1
AS2
AS4
AS3
AS5