Chapter 6

Demonstration of network design methodology

6.1 Scope

The scope of the design methodology demonstration in this chapter is defined as follows: The developed integrated design methodology, as presented in section 5.1.3, has as its most novel component the clustering approach to finding hub nodes from economic statistics, as presented in section 5.2. The use of this clustering approach in finding the number and positions of hub nodes on the backbone level of the multi-level network model, as presented in section 4.1, will be demonstrated.

In contrast to the simulation experiment presented in section 5.3.4, this demonstration will utilise actual economic activity statistics and geographical coordinates. The network to be designed will cover the whole of South Africa and is designed with an aggregate network capacity of 1 Tbps. This capacity refers to the sum of all logical connections that the network would be able to carry simultaneously.

As indicated in figure 5.4, a demand matrix reducer function is performed due to the memory constraints discussed in section 5.3. Results for the network under design will be presented for the following implementation configurations:

1. Demand matrix reducer limited to 1 MB of memory and clustering function limited to 1 MB of memory.
2. Demand matrix reducer limited to 1 MB of memory and clustering function limited to 8 MB of memory.
3. Demand matrix reducer limited to 8 MB of memory and clustering function limited to 1 MB of memory.
4. Demand matrix reducer limited to 8 MB of memory and clustering function limited to 8 MB of memory.

These scenarios were chosen due to the coverage that they afford to the problem space of how much memory to allow. The vast amount of time required to run these scenarios, several hours in some cases, made the use of scenarios requiring greater amounts of memory impractical. In the discussion section at the end of the chapter it would also be interesting to note that the obtained results exhibit asymptotic behaviour, which seems to indicate that consideration of greater memory scenarios would not provide significantly different results.

6.2 Input statistics

The economic activity metric used to weigh the network nodes, as described in section 3.1.3, was the remuneration of employees and turnover according to the levies received by district councils, metropolitan councils and regional councils by magisterial district published by Statistics South Africa [60]. Due to the sheer volume of statistics, table 6.1 only shows the 20 most economically active magisterial districts. Figure 6.1 shows a 29 bin logarithmic histogram of the economic activity of the 349 magisterial districts.

One global demand matrix was initially generated with nodal add/drop traffic values determined by a modified gravity model, as described in section 3.1.3, with 1 Tbps as the chosen aggregate network capacity requirement of the network under design.

Figure 6.2 shows how a 39 bin logarithmic histogram of the nodal add/drop traffic from the generated demand matrix follows the distribution of economic activity, shown in figure 6.1, on which it is based. The greatest nodal add/drop traffic was calculated for Johannesburg at 242.5 Gbps and the smallest nodal add/drop traffic was calculated for Mutale at 9.6875 kbps. The generated demand matrix has the constraint of being symmetrical, due to the modified gravity model used in its creation.

Geographical coordinates for the main town or city of each of the 349 magisterial districts in South Africa, shown in figure 6.3, were obtained from *www.gpswaypoints.co.za*.

A demand matrix reduction function is utilised to circumvent memory limitations, as discussed in section 5.3, when implementing the methodology on a computer with finite memory. For the network under design it has been decided that 50 super network nodes will be created from the original 349 network nodes representing the magisterial districts of South Africa. These super nodes are allowed to contribute intra-nodal traffic to the demand matrix, since they represent several other network nodes.

When reducing the original matrix a clustering process is used to aggregate network nodes to the new super network nodes. This clustering process has been designed to utilise limited amounts of memory by managing the node multiplication process, which is discussed in section 3.1.3. Results will be presented for where the memory for the demand matrix reduction process was limited to 1 MB and 8 MB respectively. The greatest and smallest add/drop traffic values of the reduced demand matrix's super nodes are presented in table 6.2.

6.3 Results

6.3.1 Evaluation of intra/inter-cluster traffic ratio

The results obtained for the 4 implementation configurations, as defined in section 6.1, are presented in figures 6.4 to 6.7. The mean intra/inter-cluster traffic ratio value for each number of hub nodes was calculated by using equation 5.18. The standard deviation value for the intra/inter-cluster traffic ratio for each number of hub nodes was calculated by using equation 5.19.

6.3.2 Discussion

In order to determine the number of hub nodes required on the backbone level of the multi-level network model for each of the four implementation configuration, as defined in section 6.1, the results obtained in section 6.3.1 have to be evaluated against a target

District	**Employees (R1000m/yr)**	**Institutions (R1000m/yr)**	**Total income (R1000m/yr)**
Johannesburg	48.5	262.1	310.7
Pretoria	31.2	148.7	179.9
Durban	20.7	119.7	140.4
Randburg	16.9	105.8	122.7
Cape	14.8	81.2	96.0
Germiston	7.6	58.6	66.3
Port Elizabeth	9.2	51.8	61.0
Kempton Park	6.8	48.8	55.7
Lower Umfolozi	5.0	33.0	38.1
Rustenburg	5.2	31.9	37.2
Bellville	3.8	28.9	32.7
Highveld Ridge	3.4	26.3	29.8
Pinetown	3.9	24.0	28.0
Wynberg	4.3	23.0	27.4
Goodwood	2.6	23.9	26.5
East London	4.7	19.8	24.5
Pietermaritzburg	6.4	18.1	24.5
Witbank	3.5	20.3	23.8
Middelburg MP	3.6	20.0	23.6

Table 6.1: The 20 most economically active magisterial districts in South Africa [60].

Memory used	Greatest add/drop traffic		Smallest add/drop traffic	
(MB)	**Node**	**Gbps**	**Node**	**Mbps**
1	Johannesburg	473.4	Molteno	18.5
8	Johannesburg	242.5	Edenburg	27.6

Table 6.2: The add/drop traffic values for the super nodes created by the demand matrix reducer, utilising either 1 MB or 8 MB of memory.

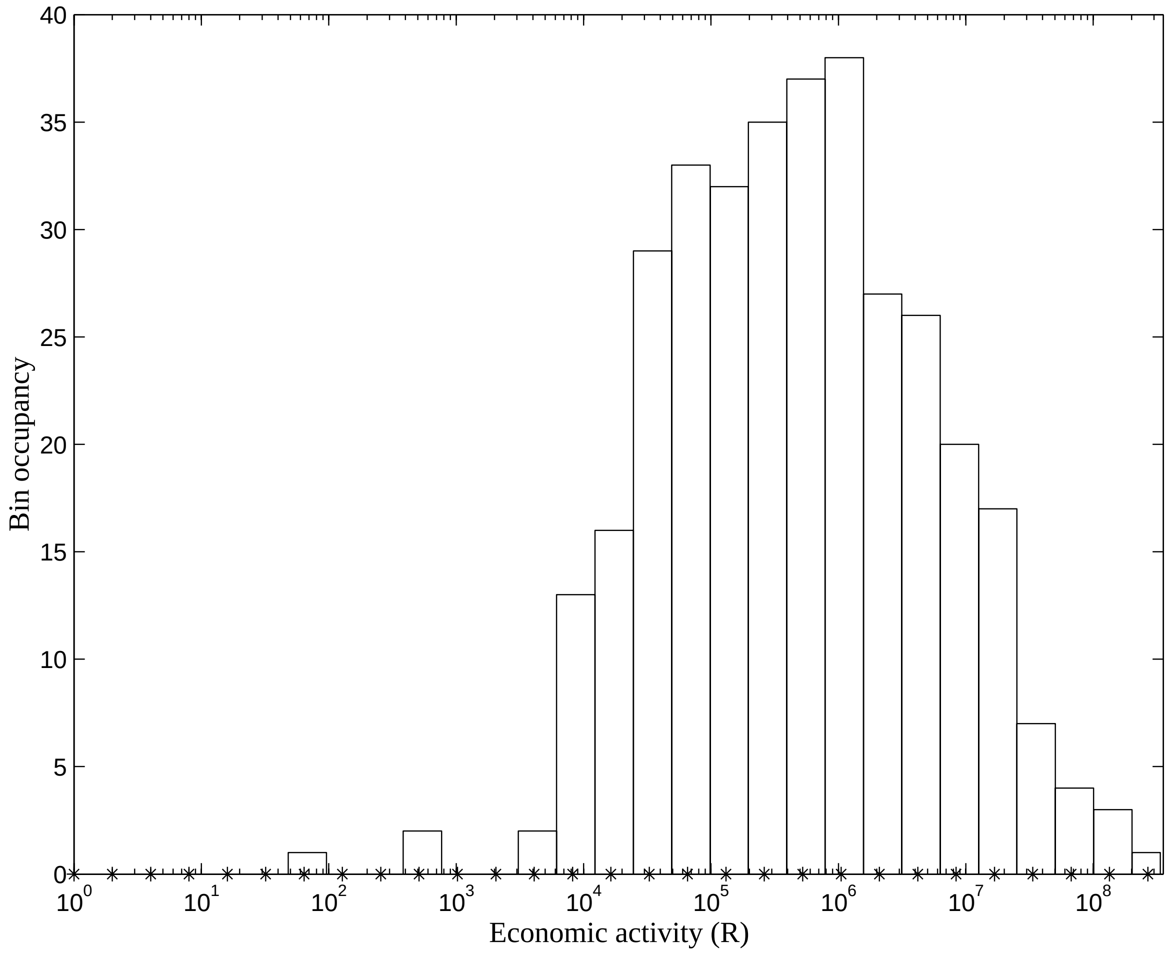

Figure 6.1: Histogram with 29 bins on a logarithmic x-axis, showing the distribution of economic activity over the magisterial districts of South Africa.

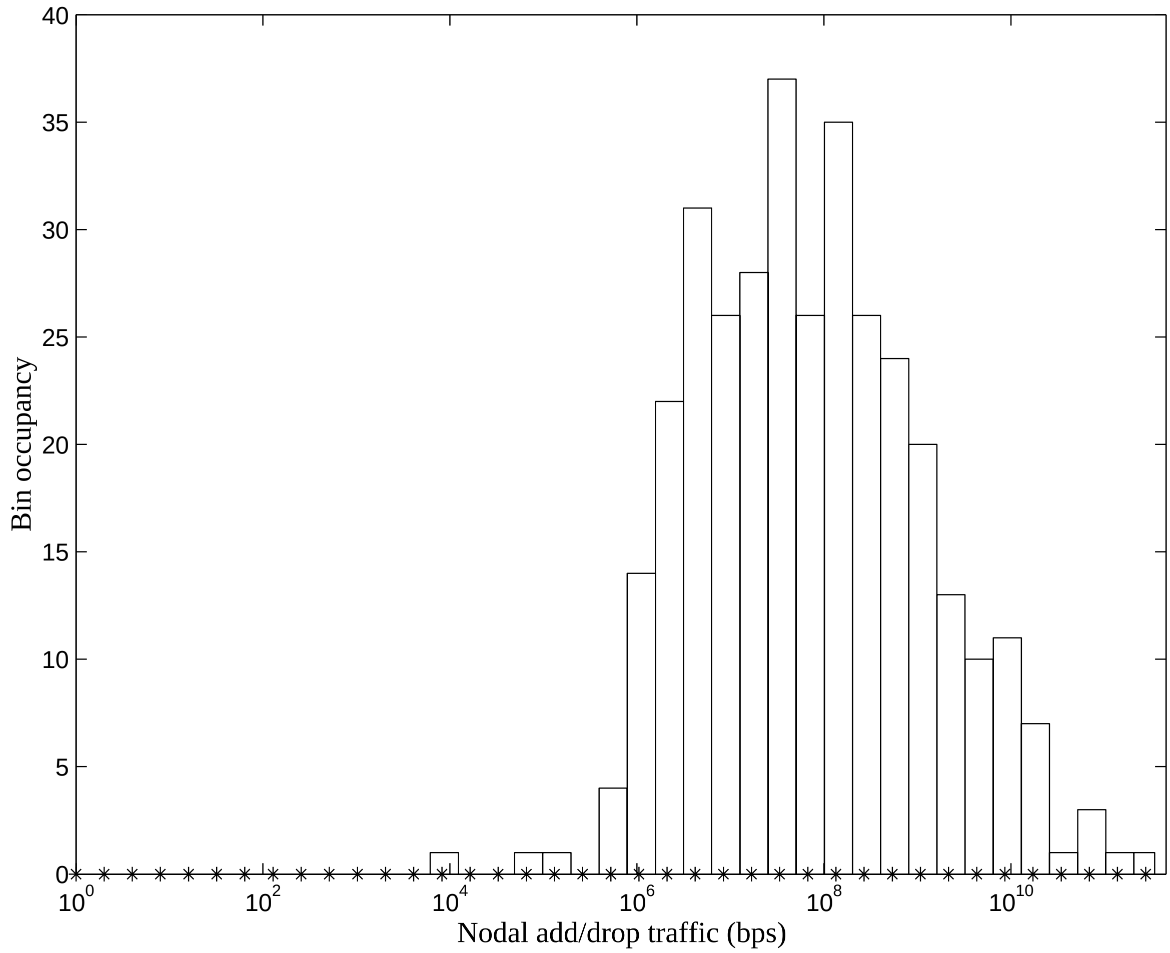

Figure 6.2: Histogram with 39 bins on a logarithmic x-axis, showing estimated nodal add/drop traffic over the magisterial districts of South Africa.

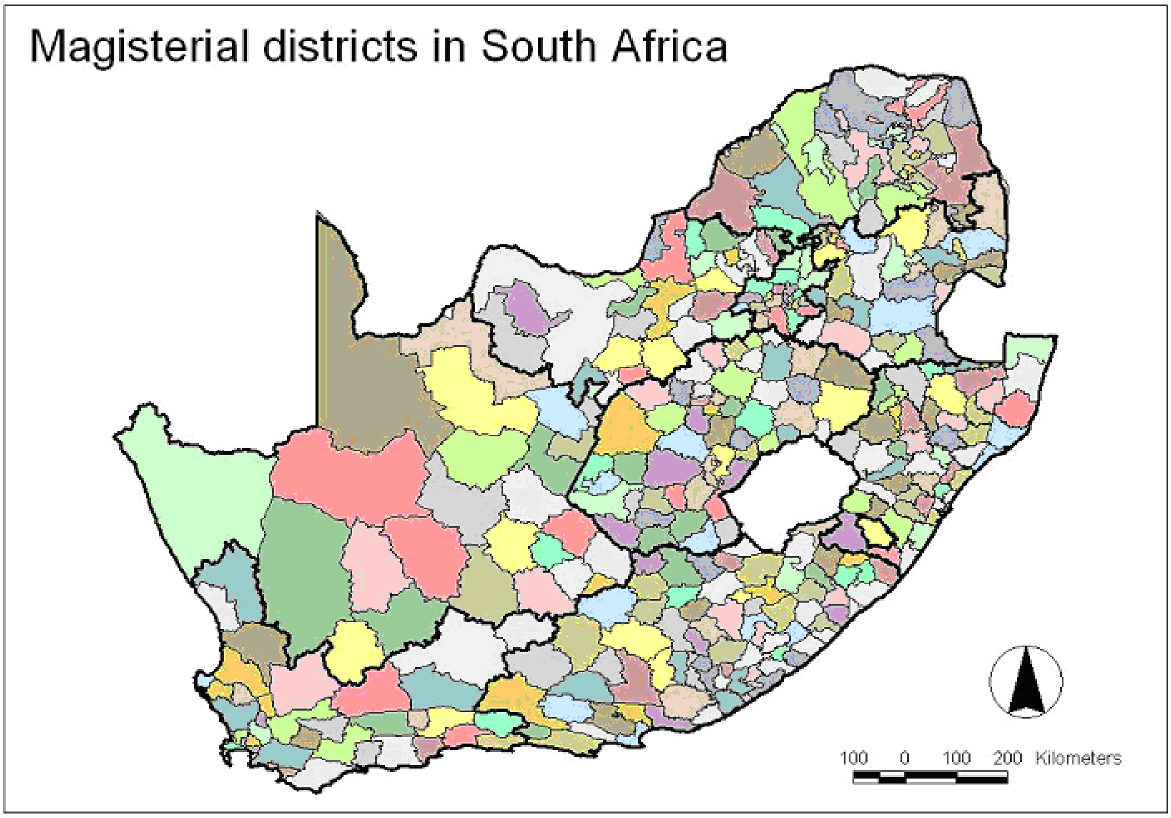

Figure 6.3: Map of South Africa indicating its 349 magisterial districts [61].

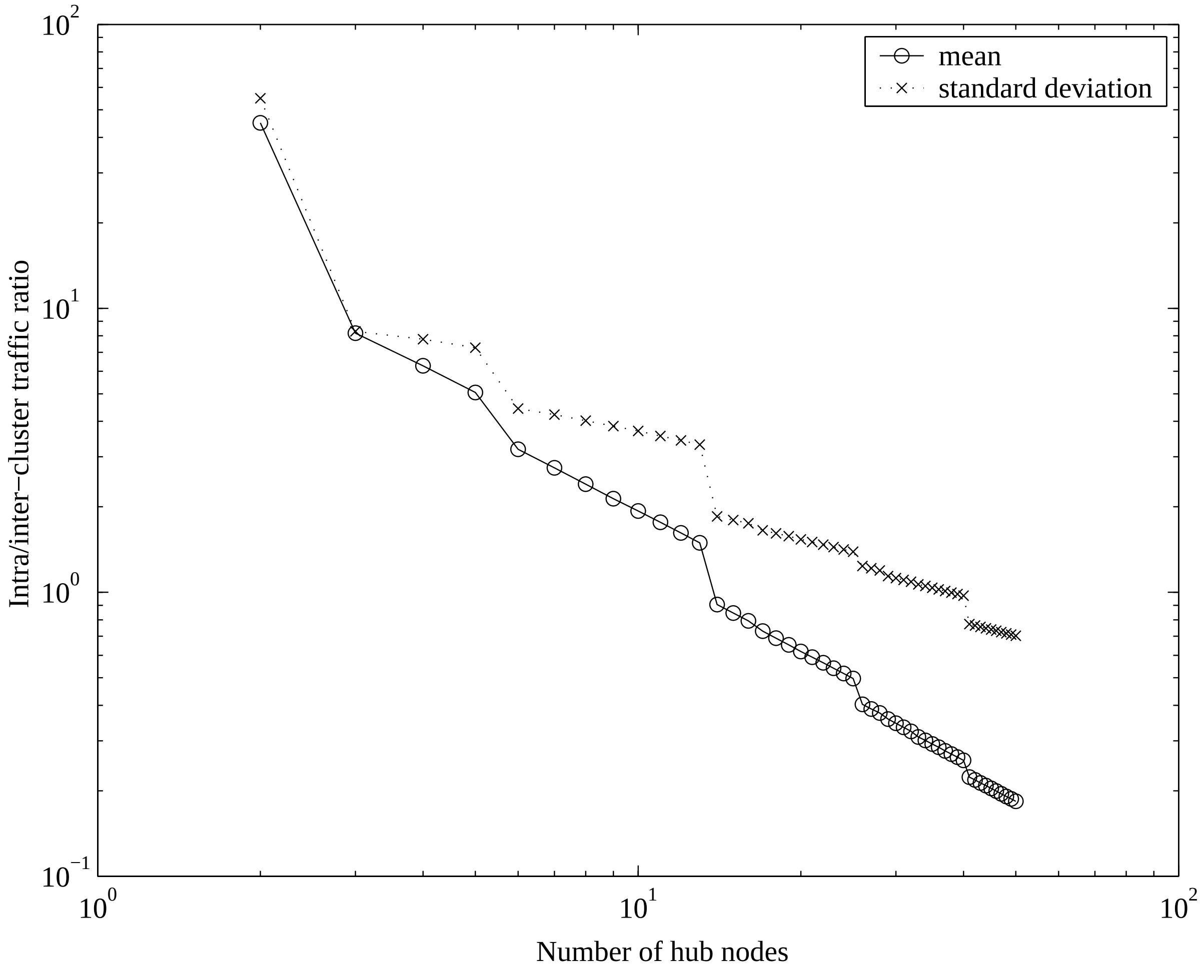

Figure 6.4: Means and standard deviations of the intra/inter-cluster traffic ratio as a function of the number of hub nodes, for implementation configuration 1 as defined in section 6.1.

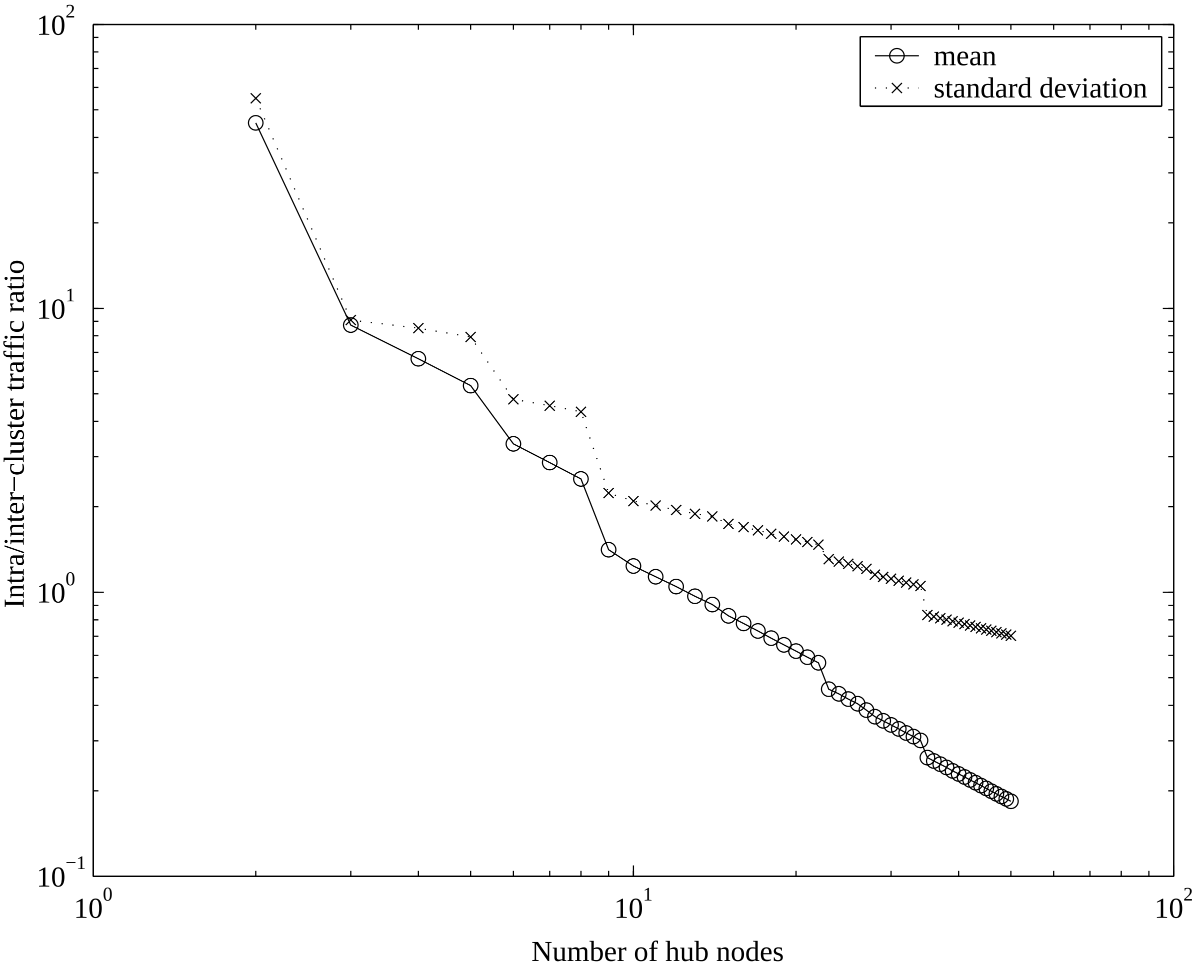

Figure 6.5: Means and standard deviations of the intra/inter-cluster traffic ratio as a function of the number of hub nodes, for implementation configuration 2 as defined in section 6.1.

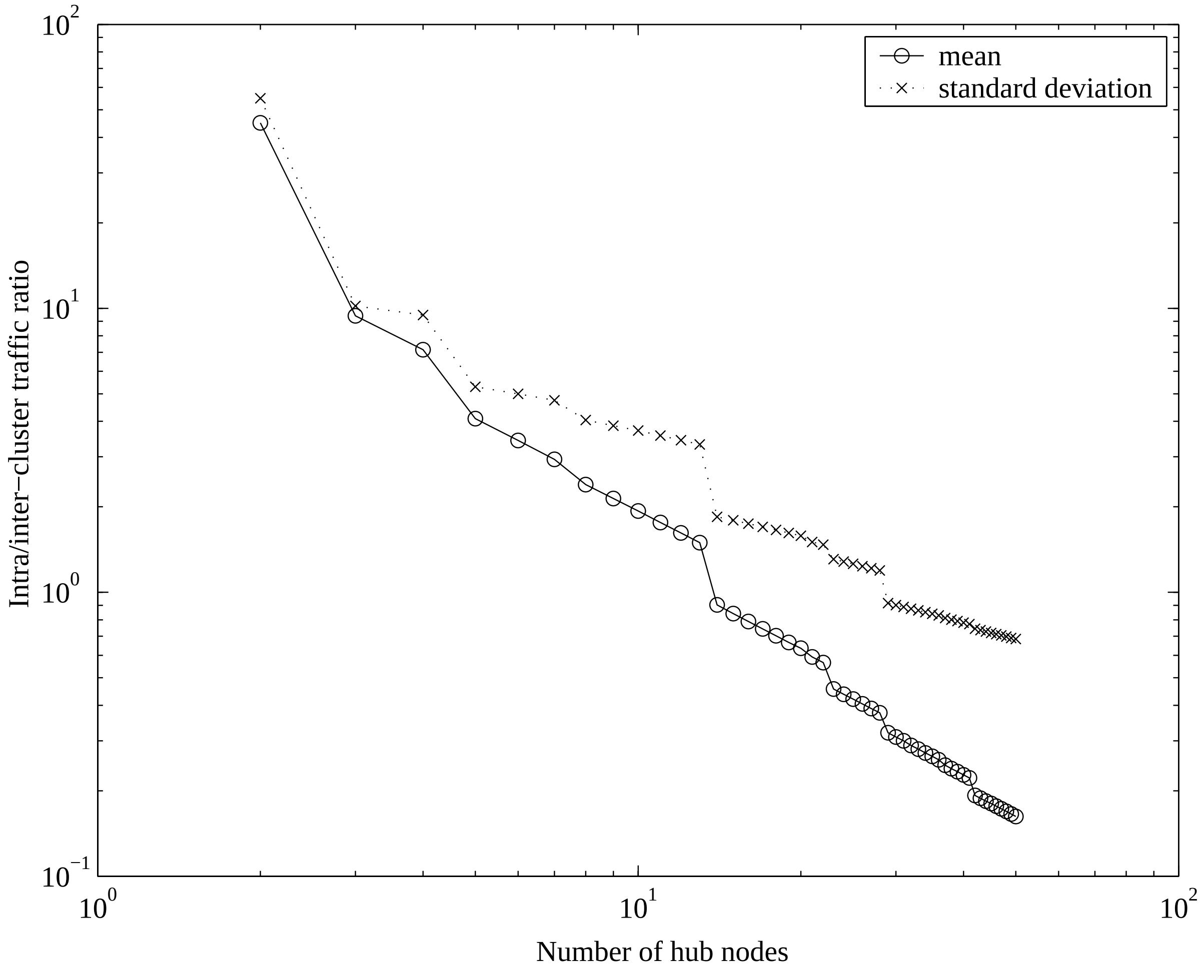

Figure 6.6: Means and standard deviations of the intra/inter-cluster traffic ratio as a function of the number of hub nodes, for implementation configuration 3 as defined in section 6.1.

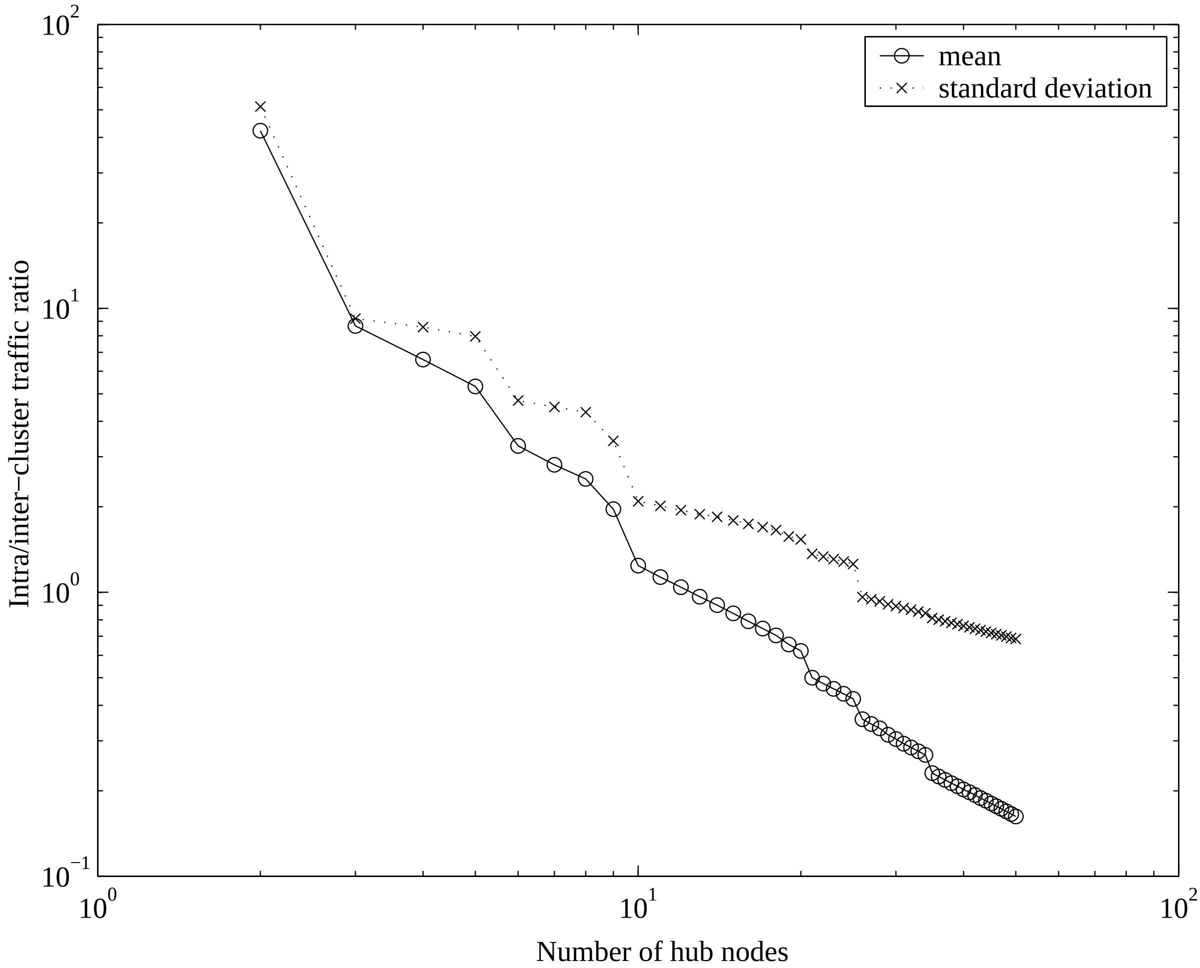

Figure 6.7: Means and standard deviations of the intra/inter-cluster traffic ratio as a function of the number of hub nodes, for implementation configuration 4 as defined in section 6.1.

Implementation configuration	$R_t = 1$	$R_t = 5$	$R_t = 9$
1	14	6	3
2	13	6	3
3	14	5	4
4	13	6	3

Table 6.3: The number of backbone level hub nodes for each of the four implementation configurations, defined in section 6.1, according to various target intra/inter-cluster traffic ratios indicated by R_t.

or preferable intra/inter-cluster traffic ratio. The guidelines provided in the discussion of section 5.3.3 were considered in selecting the target intra/inter-cluster traffic ratio to be 5. This target value is lower than the value of 9 recommended in section 5.3.3 for the selection of hub nodes on the backbone level of the multi-level network model.

Table 6.3 shows the number of hub nodes required on the backbone level of the multi-level network model for each of the four implementation configurations outlined in section 6.1, for various target intra/inter-cluster traffic ratios, including the target of 5. As expected it can be seen that the number of hub nodes required decreases as the target intra/inter-cluster traffic ratio increases.

The implementation configurations are quite evenly matched when it comes to the number of backbone hub nodes required at different intra/inter-cluster traffic ratio targets. It is thus not possible to identify any trends about the sensitivity of the number of hub nodes to the amount of memory available for the demand matrix reduction and clustering processes.

6.3.3 Clustering results

The backbone hub nodes and their corresponding clusters are shown in figures 6.8 to 6.11 and tables 6.4 to 6.7 for the 4 implementation configurations defined in section 6.1. The network nodes with the large circles around them are hub nodes of the backbone network, and the similarly shaded nodes in their vicinity belong to the same cluster. Only the networks designed for the target intra/inter-cluster traffic ratio of 5 are shown.

Three of the four figures show a network comprising of 6 backbone hub nodes. It is interesting to note that non of the implementation configurations resulted in identical network designs. The two implementation configuration pairs that utilised the same amounts of memory in the demand matrix reduction function had the same 50 network nodes. The difference in the memory limit for the clustering process also influenced the selection of hub nodes and creation of clusters. Of these four implementation configurations, the one that utilised the most memory in both the demand matrix reduction and clustering processes is regarded as the best, since it could have solved the problem with more resolution than any of the others.

It is however satisfying to observe that the network designed by the implementation configuration that was most severely limited by memory constraints, performed very good and managed to produce a solution very similar to that found by using a lot more memory. The effective implementation and successful operation of the demand matrix reduction function have been identified as the reasons for the success of the network design methodology under limited memory conditions.

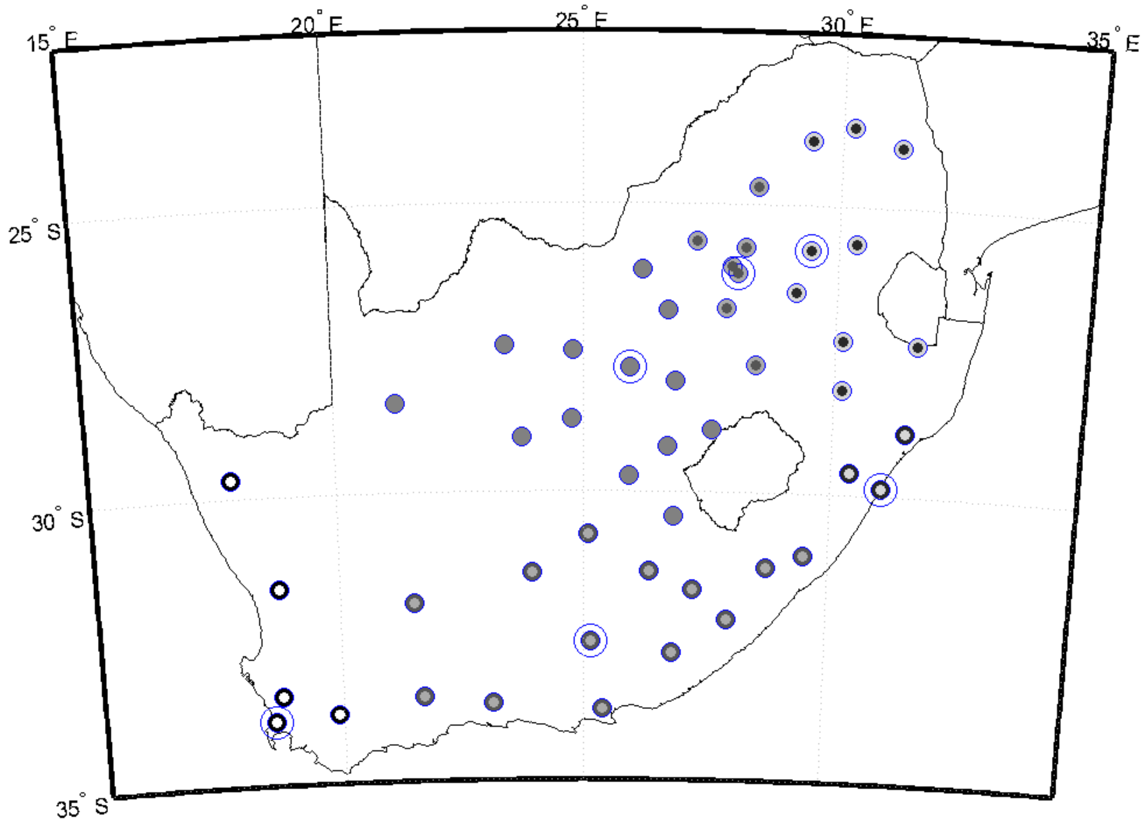

Figure 6.8: A map of South Africa showing the backbone hub nodes and clusters for the network designed according to implementation configuration 1, as defined in section 6.1.

	Hub node	**1**	**2**	**3**	**4**	**5**	**6**
1	Durban	X	1.571	1.753	15.192	3.449	1.842
2	Goodwood		X	0.961	5.712	0.917	1.713
3	Hoopstad			X	13.253	1.443	0.805
4	Johannesburg				X	27.446	4.723
5	Middelburg MP					X	0.797
6	Pearston						X

Table 6.4: The hub nodes and resultant symmetrical backbone demand matrix, in Gbps, for the network designed according to implementation configuration 1, as defined in section 6.1.

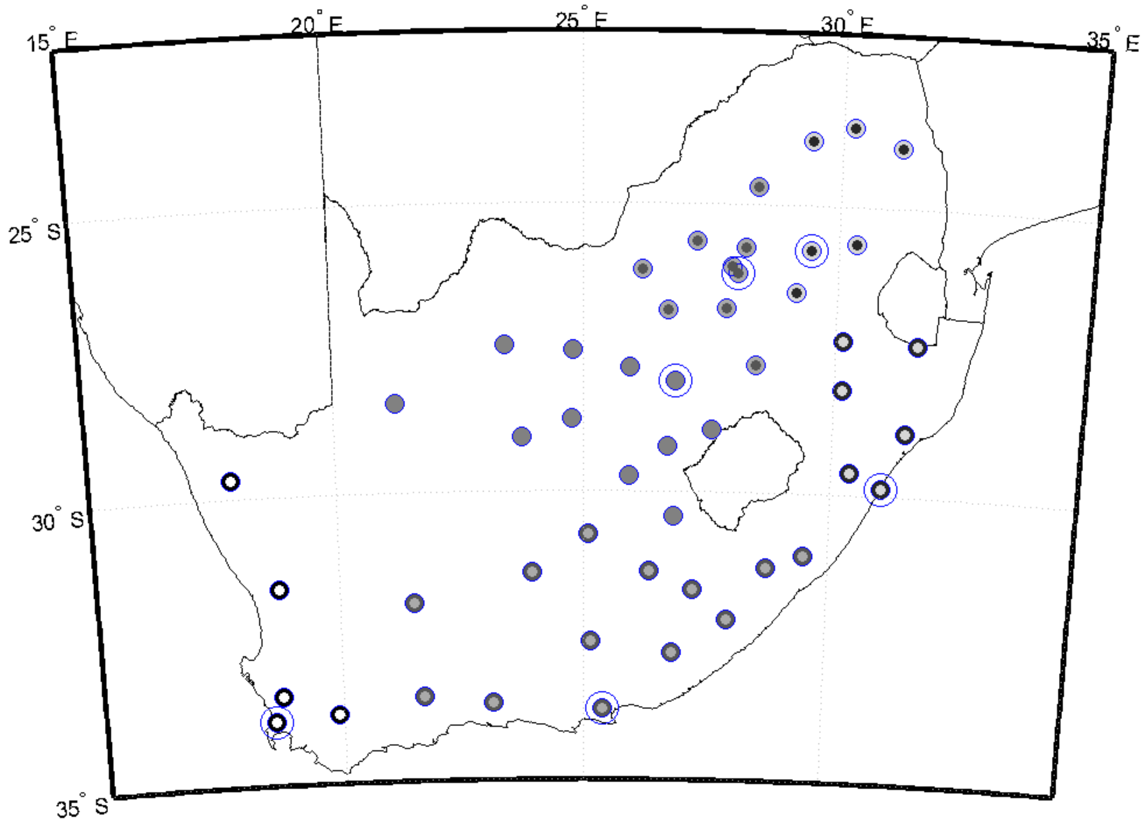

Figure 6.9: A map of South Africa showing the backbone hub nodes and clusters for the network designed according to implementation configuration 2, as defined in section 6.1.

	Hub node	**1**	**2**	**3**	**4**	**5**	**6**
1	Durban	X	1.735	18.610	2.996	2.005	1.426
2	Goodwood		X	5.942	0.753	1.713	0.731
3	Johannesburg			X	25.051	4.910	7.556
4	Middelburg MP				X	0.634	0.747
5	Uitenhage					X	0.619
6	Virginia						X

Table 6.5: The hub nodes and resultant symmetrical backbone demand matrix, in Gbps, for the network designed according to implementation configuration 2, as defined in section 6.1.

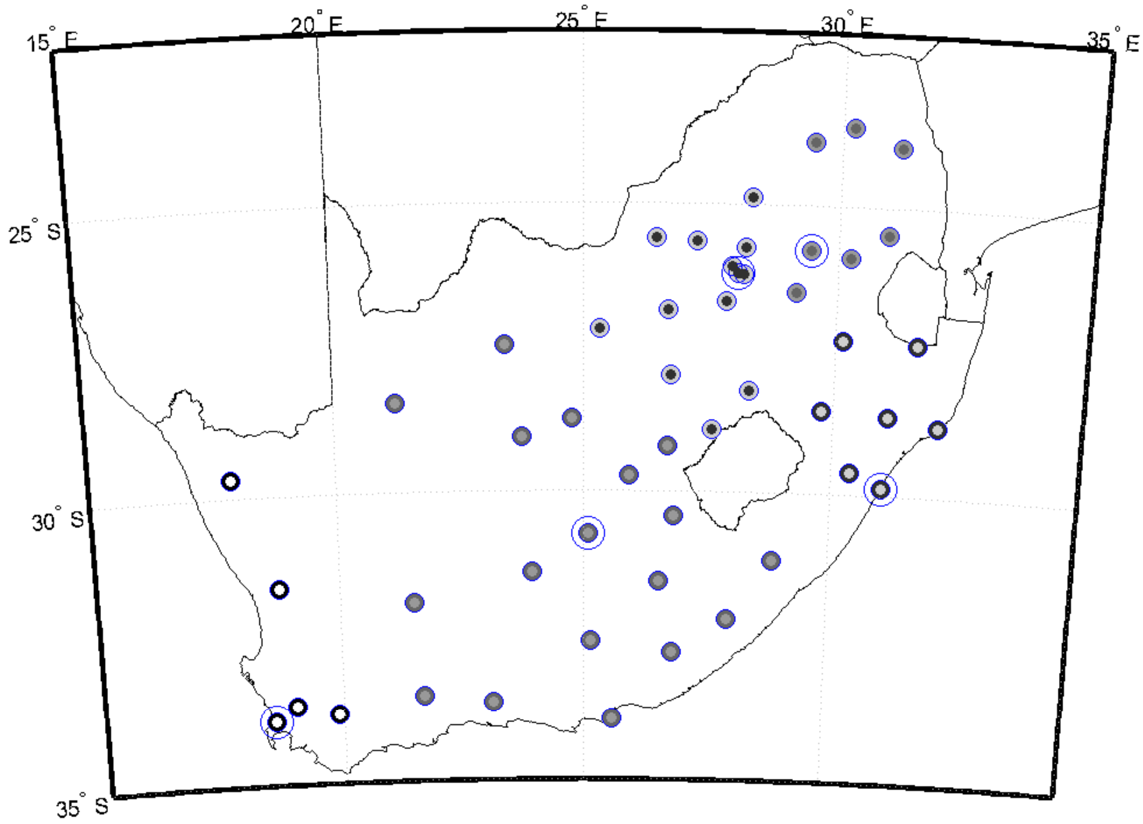

Figure 6.10: A map of South Africa showing the backbone hub nodes and clusters for the network designed according to implementation configuration 3, as defined in section 6.1.

	Hub node	**1**	**2**	**3**	**4**	**5**
1	Colesberg	X	2.781	2.201	8.964	1.047
2	Durban		X	1.714	18.687	2.941
3	Goodwood			X	6.184	0.773
4	Johannesburg				X	25.957
5	Middelburg MP					X

Table 6.6: The hub nodes and resultant symmetrical backbone demand matrix, in Gbps, for the network designed according to implementation configuration 3, as defined in section 6.1.

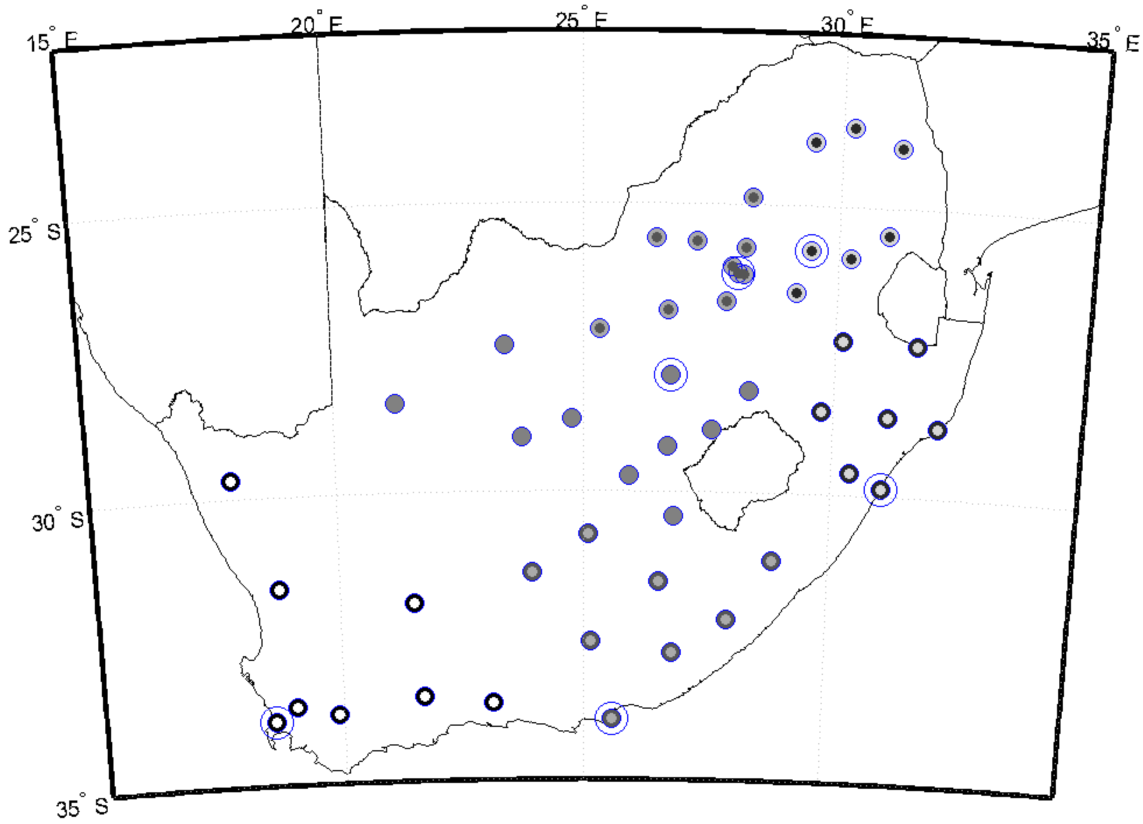

Figure 6.11: A map of South Africa showing the backbone hub nodes and clusters for the network designed according to implementation configuration 4, as defined in section 6.1.

	Hub node	**1**	**2**	**3**	**4**	**5**	**6**
1	Durban	X	1.901	17.914	2.941	1.800	1.567
2	Goodwood		X	6.534	0.850	1.484	0.808
3	Johannesburg			X	25.513	4.305	7.954
4	Middelburg MP				X	0.575	0.839
5	Port Elizabeth					X	0.554
6	Welkom						X

Table 6.7: The hub nodes and resultant symmetrical backbone demand matrix, in Gbps, for the network designed according to implementation configuration 4, as defined in section 6.1.