Routing in the traditional DDD telephone network

Douglas A. Kerr

Issue 1
September 20, 2025

Introduction

The long distance (toll) telephone network in the form that was used from about 1950 through perhaps the 1980s was initially put in place to support *operator toll dialing* (where the "originating" long distance operator sets up the entire connection for a call without the aid of any other operators), but was soon utilized for *direct distance dialing* (DDD), in which the users could dial long distance calls themselves, without any operator being involved.

An important aspect of the plan was the algorithm that guided the route of a connection as it unfolded from the calling party to the called party. This exploited the concept of "alternate routing", in which the extension of the connection from a given switching office forward could in many cases take any of several routes to different offices (with an order of preference). This supported a network operation that was very advantageous economically, and in addition provided redundancy that could help to overcome such things as the substantial failure of transmission facilities.

This article describes the underlying premises of this routing strategy, in its original form, and how it was actually typically executed, based on a somewhat simplified picture. Some background in important principles of teletraffic engineering is included.

1 A SIMPLIFIED PICTURE

The DDD network structure and routing scheme described here are somewhat simplified from the full reality in order to best illustrate the concepts and principles involved. The network structure illustrated is a simplified form of the structure as defined through perhaps 1975. Later changes through perhaps the 1980s do not in any significant way affect the concepts and principles described here.

And there are of course a zillion "ifs, ands, or buts" in this area which, similarly, do not in any significant way affect the concepts and principles described here.

2 TELETRAFFIC CONCEPTS

2.1 Introduction

The routing plan I will describe is highly influenced by some principles of teletraffic engineering (known in the telephone business as just "traffic engineering") Here is a brief review.

Suppose between central offices A an B there is a single group of interoffice trunks. Suppose we can predict the statistics of the calls that would hopefully travel over those trunks (lets say on the average over the iconic "busy hour of the average busy day"). I say "hopefully", because it might be that some of these "wannabe" calls will arise when there is not an idle trunk in that trunk group, and they are thus "thwarted" ("blocked" is the technical term).

2.2 An important performance metric

An important metric is the probability that (during this iconic measurement window) a call that is "offered" will not find an idle trunk to carry it, and thus is said to be "blocked".

Usually, an acceptable limit for this probability is established to use in guiding such things as "how many trunks should be provided in this group". A common such "bogey" is 1%; that is, during the iconic measurement window, of 1000 calls that would like to advance over these trunks, not over 10 would be "blocked" for lack of an idle trunk.

We may see this level of blocking written as "P.01". We might think that "P" means "probability", but in fact it is shorthand for a particular statistical model of the random arrival of call "wannabes" that is assumed.[1]

Not surprisingly, if we start with some number of trunks in the group, and estimate the blocking probability, and then (the "offered" traffic being unchanged) we increase the number of trunks, the blocking probability will decrease. Broadly speaking, we typically choose to implement the number of trunks that would make that probability equal to our "bogey"–no more trunks, as that would be "overkill".

The probability of blocking, either was predicted for some specific number of trunks or as a "bogey", is called the "blocking grade of service" (or, more commonly, just the "grade of service") over that trunk group.

[1] It indicates the "Poisson" model of call arrival, based on the Poisson Distribution, characterized by French mathematician Siméon Denis Poisson, a major contributor to the development of statistical mathematics (among many other things).

Of course, for any given number of trunks, on the average they are not carrying calls "full time". The average fraction of time that a trunk is actually carrying a call (over the iconic time window) is said to be the *occupancy* of the trunk. We can think of that as a metric of the "bang for the buck" that the trunks in that group provides.

Of course, for the same offered traffic, getting a lower probability of blocked calls requires a greater number of trunks. Since the overall load is the same (if we put aside the matter of a variable small number of the offered calls not actually being carried), it is easy to see that a better "grade of service" (that is, a lower probability of blocking) leads to "less bang for the buck". Not surprising. "Insurance costs".

3 TELEPHONE NUMBERS

During the era of interest to us here (and at present), North American telephone numbers had a format as seen in this example:

311-555-2368

where "311" is the area code (formally, the *numbering plan area code,* (NPA), 555 is the *office code (*formally the *central office code)* (COC[2]), and 2368 is the *station number*.

For many years, the standard long distance dialing plan was that the NPA was only dialed if it was different from the NPA of the caller. If no NPA was dialed, the switching systems took it to be that of the caller.

During the area of interest, all numbers with a certain NPA and a certain COC were served by a certain central office, and the routing algorithm described here is predicated on that.

4 NETWORK HIERARCHICAL ORGANIZATION

4.1 Concept

In the toll network plan being described, the various central offices (switching points) are assigned hierarchical "ranks" (actually called "classes"), as seen in Figure 1.

---

[2] That is not the official abbreviation for the central office code part of the telephone number. There really isn't any (that is itself a long story). I use "COC" here for convenience.

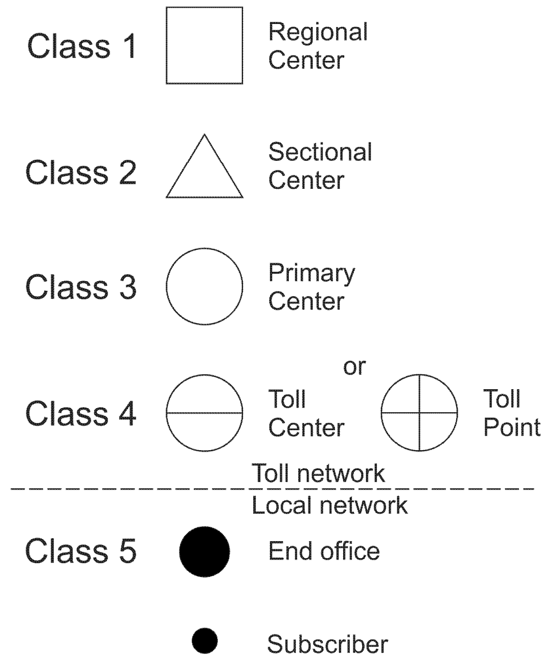

Figure 1. Central office classes

The *end office* is a local switching office that actually serves the subscribers' lines. Note that it is not a part of the toll network, yet for completeness it is assigned a class in the toll network hierarchical system (Class 5), just as if it were.

> Note that in actuality, at the Class 4 level, we can have two kinds of office, the Toll Center (specifically identified as "Class 4C"), and the Toll Point (specifically identified as "Class 4P'). The distinction has to do with what sort of operator services are provided at that office. This has no impact on the description in this article, and for conciseness I will show all Call 4 offices as Toll Centers.
>
> I also note that some time subsequent to 1975, a new "level" was defined, lower than Class 4, called "Class 4X', the offices at which were called "Intermediate Points". Again, this has no impact on the concepts and principles discussed here, and so I will ignore it from here on. It did effect the details of execution of the routing concept.

This scheme in fact implies a certain hierarchical structure (called the "homing" structure), which we see in Figure 2.

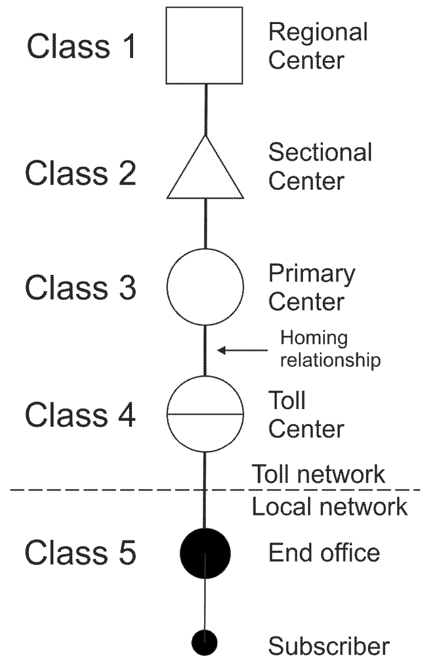

Figure 2. Central office homing structure

Each office, regardless of class (except for Class 1 offices) is (in the ideal model) assigned an office of the next higher class on which it is said to “home”.

Note that this does not suggest that all traffic from one office to somewhere else would go next to the office on which it homes (although the homing arrangements do confer an important property on the routing plan).

A bigger picture of the overall arrangement is given in Figure 3. we see one Regional Center (Class 1 office) and 2 Sectional Centers (Class 2) with several of their “subordinate” offices at each descending “class” level.

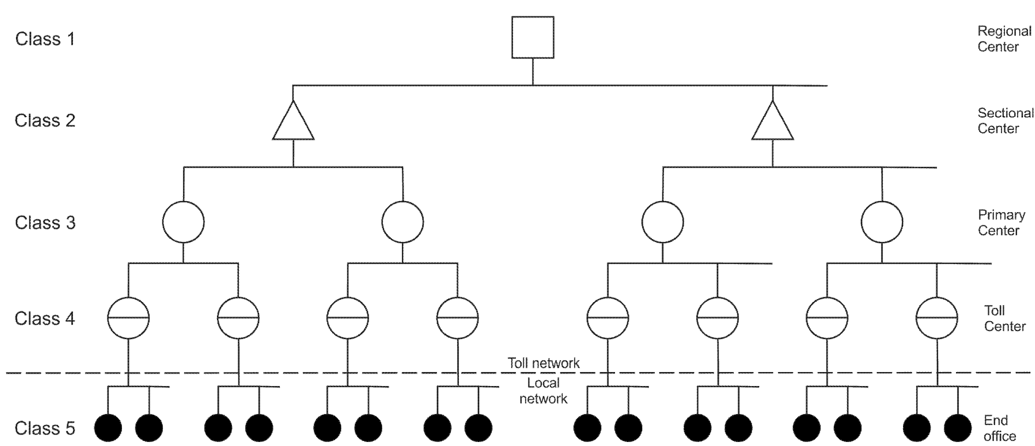

Figure 3. Central office homing structure—a larger picture

Note that the lines here show the hierarchical homing structure, and do not necessarily suggest preferred routes through the network, nor do they represent all trunk routes. (But as we will see, there will be trunk routes along all those "homing" lines.)

The open ended lines remind us that, in general, there are "more of the same" homing on each office in the toll network

4.2 Provision of trunk groups

In fact, trunk groups may be provided between offices of whatever class based on estimations of traffic and in light of the routing algorithm I will discuss. An example might be as seen in Figure 4.

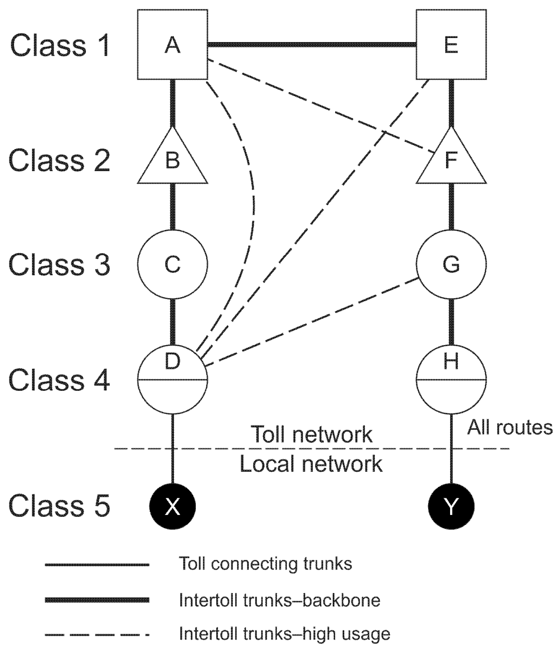

Figure 4. Illustrative trunk groups

This might not be a realistic portrayal of all the trunk groups that might exist (which could make the drawing way too cluttered). I have chosen here to just show those trunk groups that will shortly figure into my discussion of the working of the routing algorithm that has been described.

The trunks that connect a Call 5 office (part of the local network) to the (typically) Class 4 office on which it homes are called *toll connecting trunks*. The trunks that run between offices in the toll network proper are called *intertoll trunks*.

In the toll network proper, there will always be a trunk group between an office and the office on which it homes, and as well trunk groups between each pair of Class 1 offices. These are called *backbone* trunks. And in fact, the "route of last resort" between the caller's end office and the called party's end office will be entirely along this backbone chain.

In addition, it is essentially "obligatory" that there be a trunk group from every Class 4 office (Toll Center) to the superior Regional Center.

But there will also be other trunk groups between various offices, as justified by the expected traffic incidence. These trunk groups will not be "sized" to attain the grade of service "bogey" as to blocked calls that is established for the network. Rather, they have fewer trunks than would create that blocking probability.[3] The major ramifications of that are:

- The fraction of calls whose route prospectively would utilize those trunks that are blocked because all those trunks are busy will be greater than contemplated by the network blocking bogey. But that does not degrade the service to the callers because there are other routes (perhaps even along the backbone) that can come into play if all trunks in such a group are busy, likely preventing the call itself from actually being blocked.

- These smaller trunk groups are of course less costly than the larger groups that would be required to of themselves attain the network blocking "bogey", thus there is a cost saving in such groups.

- These trunks have a greater average busy hour occupancy than would the trunks of a larger group for that same traffic. Thus they are described generically as "high usage" trunk groups.

- These trunks give a greater "bang for the buck" each than if enough were provided to meet, as to traffic over these trunks, the "bogey" for call blocking established for the network overall.

In contrast, the backbone trunk groups are "sized" so that the fraction of traffic "offered" to them that is blocked because all the trunks in the group are busy is essentially the call blocking "bogey" for the network.

The backbone trunk groups are sometimes called, in contact to the "high usage" moniker for the other trunk groups, "liberally engineered" groups, but that term is not always used. They are often called (in contrast with the "high usage" trunk group) just "backbone" trunk groups. But because they are the "last choice" at that office for calls toward certain destination numbers, they are also often called "final" trunk groups.

---

[3] They might be sized to give an expected blocking probability of 3%.

That low blocking probability criterion leads carries the corollary that those backbone trunks each give a lesser "bang for the buck" than we get for the "high usage" trunk groups. However, only a small fraction of the overall calls through the network use these backbone trunks, so the overall economic picture for this plan is favorable. But for any given situation, the lowest overall cost for handling the projected traffic occurs for a certain number of trunks in the high usage group.

5 THE ROUTING ALGORITHM

5.1 Basic principle

The algorithm for controlling routing through the toll network was based on this principle:

> At the office to which the connection has progressed, a ongoing route should be chosen to the office that is the one farthest along the "backbone chain" toward the ultimate destination office to which there is an available trunk.

Figure 5 shows what this might mean in an example situation.

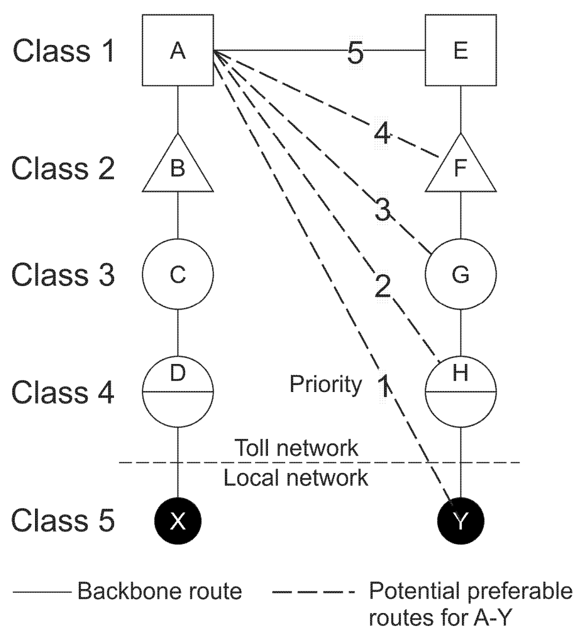

Figure 5. The underlying principle of the routing algorithm

We assume here that the connection for a call from a subscriber served by office X, and destined for a number served by end office Y, has progressed to office A. The most desirable route from office A toward office Y (labeled as "priority 1") would be directly to office Y (the office "farthest along the backbone to the destination office"—all the way there, in this case).

If that can't actually be followed (there is no trunk group from office A to office Y, or if there is, there is no idle trunk in that group at the moment), the next most desirable route would be to office H (priority 2), the "next farthest along the backbone route". And so forth.

The onward "route of lest resort" here (and always) is in fact next over a backbone trunk, if one exists, in this case to office E.

5.2 An important property

Note that this algorithm does not require the office that is currently "holding the baby" to have any knowledge of how the call got to it, and important consideration given the signaling systems and protocols in use at the time.

5.3 About "idle" trunks

In what follows, I will speak of trunks that are "available" at the moment to be part of an ongoing connection as "idle" trunks, since that is the term most often used in this area.

The term "idle" is not ideally descriptive here. A trunk that is not handling any calls, but which has temporarily been taken "out of the batting order", perhaps for maintenance, is "idle" in the common use of the word, but is of course not "available" to handle a prospective call. But I will use "idle" rather than "available" for consistency with the most common usage elsewhere (not to mention that it is shorter).

5.4 The algorithm fully at work in place

In figure 6 we can see an example of this algorithm actually fully at work for the network portion seen in Figure 4.

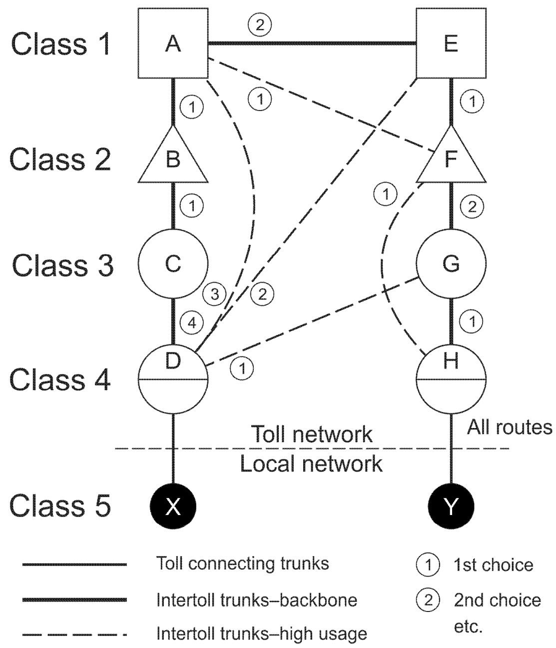

Figure 6. The routing algorithm at work

The lines (solid and dotted) show all the trunk groups we assume are in place between these offices. This arrangement might not be realistic; it is contrived to be a handy vehicle for discussing the working of the routing algorithm. Similarly, the routing patterns themselves may not be realistic. They are also chosen to best illustrate the working of the routing algorithm.

The circled numbers (①,②) indicate the various alternate routes toward the destination number from that office (in order of precedence). If there is only one route (no "alternates"), it is nevertheless marked ① for consistency.

The call whose connection we follow is from a subscriber served by end office X to a number served by end office Y. We assume that office X will route such a call to Class 4 office D (it likely would route all toll calls there). This is in its effect its "gateway" to the toll network, and might well be where the billing data for the call is collected. As to routing, all that office D knows that is of importance here is the called telephone number.

Based on what we can think of as a "lookup table" operating from the numbering plan area code (NPA) and the central office code (COC) parts of that number, office D learns that its "1st choice" onward route would be to office G (via a "high usage" trunk group).[4]

---

[4] This assignment in the "lookup table" was made in accordance with the "farthest along" concept.

Would not a trunk to office H be preferable under that algorithm? Yes, but we did presume that there is a no trunk group from office D to office H. So that possibility is moot.

Suppose there is in fact an idle trunk in that “first choice” trunk group (①) from office D to office G. Office D routes the call there. Office G is now “holding the baby” It looks up the NPA[5] and COC and learns that its first choice (in fact its only choice) for a call to such a number is to office H, over a backbone trunk group.

If no trunks in th G-H trunk group are idle, the call is blocked, and the caller receives a tone signal advising of that.

Office H unavoidably routes the call to office Y over a direct trunk. If all trunks in that group are busy, the call is blocked, and the caller receives a tone signal advising of that.

Now, assume that office D learned that its 1st choice for this call is over a trunk to office G but it finds that no trunk in that group is idle. The “alternate route” return from the “lookup table” (which again was based on the “farthest along” concept) tells that its 2nd choice is a trunk to office E, and it advances the connection over a trunk to that office.

Office E is now “holding the baby”. Looking up the NPA+COC of the called number in its lookup table, it learns that its 1st (and only) choice of route toward that number is a trunk to office F.

At office F, its 1st choice for a route to that number is to office H. That office finds that its 1st choice route toward the target number is via office H. The rest of the route is preordained.

If there is no idle trunk in the group from office F to office H the second choice is to office G. We already know what would happen there and onward.

Now we go back to where the connection had proceeded to office D and office D learned that its 1st choice route was to office G. There was no idle trunk in its 1st choice route (to office G), and it learns that its 2nd choice is to office E.

But this time there are no idle trunks in that group. Office D learns that its 3rd choice route toward the target number would be via office A (over a backbone trunk). If there is no idle trunk in that group, the 4th (and final) choice is over a trunk to office C (over a “backbone” trunk

[5] If that is not dialed it is in effect inferred to be the NPA of the calling party.

At office C, the 1st (and only) choice for a call to this number is to office B, and at office B, the 1st (and only) choice is to office A.

I hope the reader can by now grasp how the routing algorithm is typically executed.

Note that this process might progress so that the route actually followed is the "route of last resort", wholly over backbone trunks (that is, X-D-C-B-A-E-F-G-H-Y).

6 REDUNDANCY

I did not call attention to this earlier so as not to derail the flow of the main story, but it is important to note that the alternate routing capability discussed above as (a) leading to an economically advantageous network structure also (b) provides a way to circumvent the failure of a significant number (maybe all) of the trunks between two offices.

It is for that reason that always there are at least two trunk routes toward a certain destination at every office in the toll network (even though, for simplicity's sake, I did not show that in my illustrations).

-#-