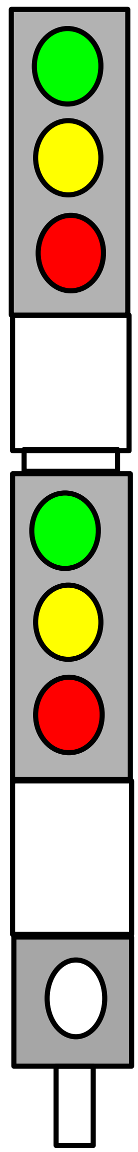

SIGNALS

Questions and Answers about the Signals in use by the Toronto Transit Commission

1.01 Control Length -- the Purpose of Signals

Q. How do signals know when to change and what is meant by control length?

A. The charter of rapid transit is evident in its name: to move people safely as rapidly as possible. While trains, motors, and propulsion power are responsible for the rapidity, signals are responsible for the safety.
A rapid transit signal's job is to keep trains at a safe distance from dangers, including each other, via three closely-related functions,

* To indicate to train operators whether the track ahead is clear,
* To instruct the train operator to proceed (green), proceed prepared to stop at the next signal (yellow), or to stop (red), accordingly, and
* To forcibly stop a train via its train stop should that train fail to comply with an indication of "stop."

An extent of track being "clear" means that it is proven free of dangers such as other trains, trailing-point switches set the wrong way, conflicting routes from other signals, and so on. The extent of track ahead of a given signal for which it makes this check, which is different for different signals, is called its control length, as it controls the signal's indication. If, at a given time, any of these dangers appear within any portion of a signal's control length, that signal is required and designed to display an indication of "stop". A "clear" indication confirms the proven absence of dangers within a signal's control length.

Signals are placed regularly along a track, at insulated joints, subject to considerations to be discussed. Each signal's indication instructs the train operator what to do up to the next signal. An indication of "proceed", or "proceed prepared to stop at next signal" is only valid up until the next signal, when that signal's indication assumes validity.
Consider the following "single line" signaling diagram. As per standard, the lines under the track starting in little circles at signals and ending in arrows represent those signals' control lengths. The "lazy house" shaped object represents a stopped train. The big arrow indicates the direction of traffic. The signals here are 200' apart.

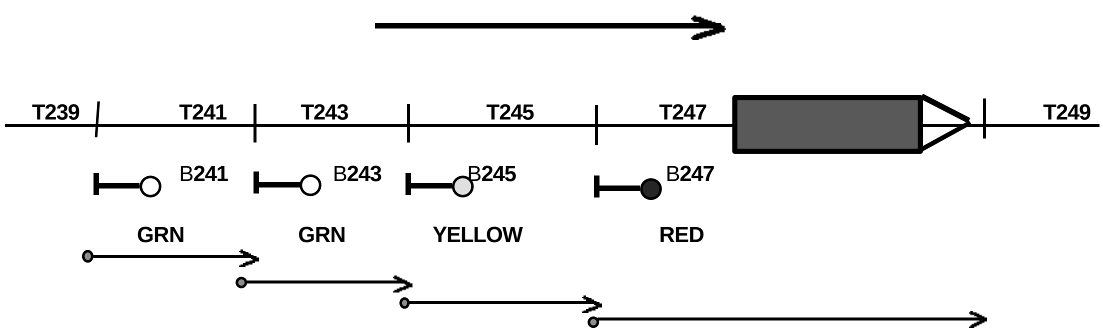

Here, each signal is controlled solely by the occupancy of the "block" in front of it, extending to the next signal (that is, whether it is clear or not. 245 is yellow because 247 is red). Assuming that 245 is visible early enough for the operator of an approaching train to observe and obey it and slow down, and posted speed restrictions are being obeyed, that train will stop before arriving at the red signal 247, and the rear of the stopped train, and this scenario is adequate.

But in real life, this will not do: an oncoming train might not be operating within the posted speed limit, and the train operator might be temporarily or permanently physically, mentally, or morally incapacitated, in which case the signal system, via its train stops, must forcibly stop the train.

In these cases, the yellow aspect (indicating "prepare to stop at next signal") might be ignored, and a train going at the fastest possible speed will arrive at signal 247 at speed.

Assuming signal 247 had a train stop that acted as expected (assumed "tripping" position when the end of a train passed it, although this not quite the case in reality, see train stops), and the oncoming train's brakes were fully applied automatically by it, a collision would still result, because trains moving at 40 or 50 miles per hour cannot "stop on a dime:" typically, hundreds of feet are required for a train moving at full speed to come to a halt via a brake application. On this account, a train must be commanded to stop hundreds of feet before an obstruction at which it must stop. This distance is called the worst-case (or "safe") braking distance, and is a function of the maximum speed of trains, the terrain, the weather conditions, and so on. It is called worst-case because it must take into account brakes in the worst condition (short of total failure), trains going at the maximum achievable, impermissible, speed, tracks at maximum slipperiness, and so on.

It can be seen that simply increasing the block length to the worst-case braking distance, or many times it, does not improve the situation -- a train approaching another train at the near end of a block at speed will not have enough space to stop, even if "tripped", no matter how long the block. In track-circuit based signaling, train position can be reckoned only in terms of which track circuits the train occupies. The solution to this problem is *overlap* of control lengths, that is, having each track circuit (section) control not one but several signals in advance of it -- were there only one red signal behind a train, the situation is always possible in which it is too close behind the train for safe braking. Consider the following far more realistic scenario:

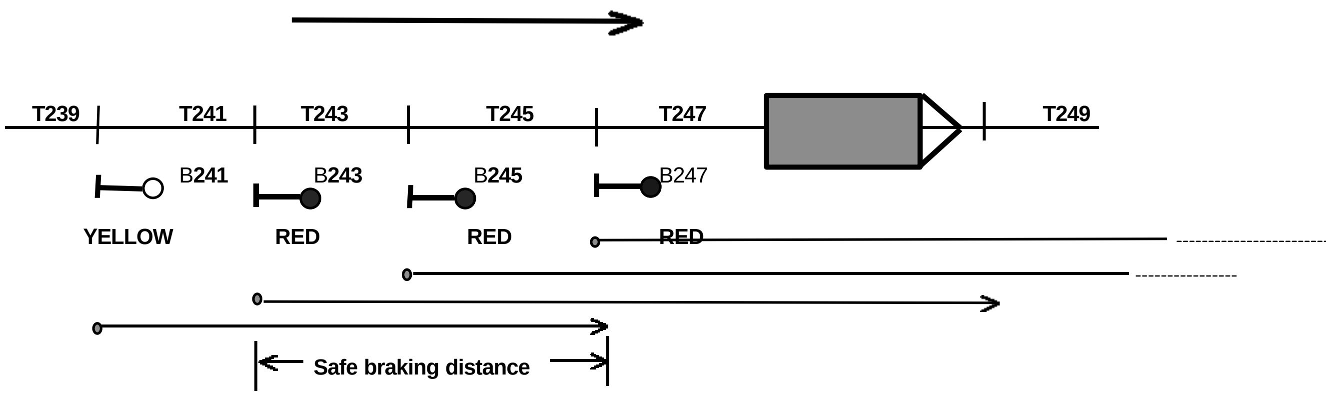

Overlapped signal control lengths

Here, the control lengths of all signals are overlapped. 243, 245, and 247 indicate "stop", and 241 indicates "prepare to stop at next signal" (yellow aspect). The worst-case braking distance here is somewhere between 200 and 400 feet. An oncoming train barreling ahead at full speed ignoring signal 241 will be tripped by 243, and come to a halt within the safe braking distance indicated, that is, short of block 247 and the train it contains. From this it can be seen that the key to safety is a control length overlap at least equal to the worst-case braking distance.

The key to understanding this is to think (in this case) about signal 241, not 243. By giving a "clear" (albeit "prepare to stop at next") indication, signal 241 has asserted that should a train accept its indication, the track which is encompassed in its control length is safe, in particular, such that if the *next signal* (i.e., 243) within that control length indicates "stop", and that indication is obeyed or enforced via a train stop, that train will stop within that control length (i.e., before IJ 247) and clear of any danger. In other words, the 241 asserts by its clear indication that its control length is safe, even if you are forcibly stopped within it.

From the preceding illustration it can be seen that a signal must have a control length of at least one safe braking distance beyond the *next signal*, i.e., its overlap. Because of this design, and the use of train stops, the signal system ensures that trains are kept a safe distance from each other and from other, transient, dangers. (Note that this has very definite implications for interlockings: an interlocking cannot allow a dangerous condition to be thrown in the path of a train which has already accepted a signal which has spoken for the absence of such conditions in its control length. See approach locking.).

In the preceding scenario, it is assumed that the operator of an oncoming train will see the yellow aspect of 241 and begin to brake sufficiently soon to be able to stop the train before reaching 243, lest it train be tripped. Given that the distance from 241 to 243 is markedly less than the worst-case braking distance, we can assume that the distance from the place where 241 can first be seen to 243 must be the expectable braking distance at the posted operating speed. A train obeying 241 under these conditions will not be tripped at 243. If 241 is not clearly visible far back enough, either inter-signal spacing should be longer, or more than one signal back should display yellow ("prepare to stop at next signal") if 243 is red (a feature known as overlapped distant control.) Another common technique used at interlockings and places where closely-spaced signals are needed is the employment of signals that have only yellow and red aspects, i.e., always, when clear, indicate "prepare to stop at next signal."

Note that if one can "prove" that a train is going at reduced speed, required braking distance will be less, and the enforced separation between trains, and hence signal control lengths, would be less. This is the idea underlying station time signals, described in the referenced section.
In much "routine" automatic signal territory, inter-signal spacing is equal to or greater than one safe braking distance, and control lengths two track circuits (two safe braking distances). Closer signals allow for finer control, especially when combined with the station time feature.

The situation is slightly different for home signal call-on. The control lengths of such signals and moves do not check for trains, but do check for switches, conflicting routes, etc. (See Types of Signals.). In these cases, what trainmen (like pilots) call "visual rules" apply: trains must operate at very low speed, prepared to stop within vision. Since a signal being clear indicates that there are no dangers in its control length, it must also be so that no two signals not in the same direction on the same track may clear movement into the same track section. This is one of the fundamental principles of interlocking, and is enforced by disallowing the calling of a signal (via its PB, in NX/UR interlockings an abstraction implemented in relays rather than a physical lever) when the PB's of other signals are now so set that clearing this signal would create such a situation. Even below the level of the PB's, the relay circuitry of the signals prohibits opposing and conflicting signals from being clear simultaneously.

1.02 Types of Signals

Q. What are the different types of signals and their purpose?

A. There are four basic categories of signals used in the Toronto subway, whose distinctions must be understood if one is to be a motorperson, let alone a Tower Controller! While some attributes, such as being a time signal, are applicable to several of them, these distinctions between these categories underlie the notion of interlocking.

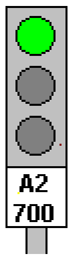

Automatic Signal

The most fundamental kind of signal is an automatic signal. It has one head with two or three lenses, at least one of which is not red. An automatic signal is not controlled by an interlocking, for, by definition, its state does not affect the interlocking. Therefore, an automatic signal is controlled only by track occupancy (i.e., the presence, absence, and movement of trains in its control length), including timers dependent on track occupancy, as with time signals. Automatic signals are the only kind present in the vast majority of the subway system, that is, between interlockings. When represented on an interlocking panel, automatic signals do not have GK lights. Automatic signals and approach signals implement automatic key-by (AK).

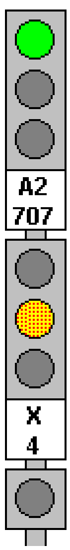

Home Signal

The second most fundamental kind of signal is a home signal. A home signal is always controlled by a PB of an interlocking, in addition to track occupancy and possible timers. The signal's PB number appears on its lower head. A home signal always has more than one multi-lens head. The upper head has the same meaning as with an automatic signal (proceed, proceed prepared to stop at next signal, stop). The lower head tells whether a normal or diverging route is set, yellow being the latter. Red always appears together in both heads. A home signal is used whenever a switch is protected, facing or trailing, including all cases of a choice of routes, or potential conflicting movements are involved. When a home signal indicates double-red, the indication is "stop and stay:" there is never any Automatic Key-by to clear the train stop of a home signal. There is, however, call-on, which is indicated by the single yellow light at the bottom of the signal, which allows the train stop to be cleared by cooperation of the motorperson and the Tower Controller in highly restricted circumstances (see the call-on description for details). The auto-call on for the X-18 signal approaching end terminals is controlled by station timing (ST). The X-18 is known as an outer home signal and is really an *approach signal*.

The third most common kind of signal is an approach signal (X-20, X-22, X-18). An approach signal looks just like an automatic signal; there is no easy way for a motorperson to tell the difference , but that is not a problem, as the meaning of either to him or her, and the behavior of their stops, is the same. Approach signals are controlled by interlocking levers (actually push buttons, PB, on an NX panel), and can be forced red by being cancelled by the Tower Controller at any time. Approach signals never actually protect switches or govern conflicting routes, but their control lengths frequently encompass trailing-point switches or interlocking exits. Calling an approach signal can move a switch; canceling an approach signal may be necessary before certain routes can be set up. Although to a motorperson an approach signal appears the same as an automatic signal, to the Tower Controller it is represented in the same way as a home signal, that is as a GK light with a PB number. Approach signals play an important role in approach locking. Approach signals are sometimes combined with Station Time (ST) timing.

The fourth type of signal that maybe encountered, especially when turning back from an interlock area is the Backup signal.

Backup signals, such BA-851 at Warden Station, are dummies that are always red, and are there to forbid movement unconditionally and serve as placeholders in the interlocking. They are extremely uninteresting because they never change state. Backup signals have train stops.

1.03 Basic NX Interlocking Operation

Q. What is an *Interlocking* and how does it work?

A. The basic idea of an NX ("ENtrance-EXit") interlocking is to route trains over the interlocking by specifying the starting point (entrance) and ending point (exit) of each such path. On a NX interlocking panel, buttons are pressed at the points representing the entrance and the exit; signal GK lights for the entrances, and the exit lights that subsequently light up for the exits.

When you initiate a route at a home signal by pushing the appropriate button such as signal X-16 at Finch, the interlocking will respond by lighting the GK light red. It will send out seeking-tendrils over all routes in that direction emanating from that point, and offer a choice of exits by displaying exit lights at each possible exit. Every possible consideration of switches that are locked and conflicting routes, i.e., the current situation, as described by other sections in this chapter, will be factored into the choice of which exits are offered. If no exits are possible, the interlocking will (usually) not let you initiate at all (the GK light will not turn red, and no exit lights will light up).

At that point you can push the exit-button: this is called route completion, and causes electrical tendrils to spread out from the exit to seek the tendrils of the initiating signal: when they meet, they will enwrap each other, cancel all other exit lights, and turn on the entrance light (which looks just like an exit light) at the point of initiation, so you know that this has happened. At this point, the interlocking will cause all switches to move and all encompassed train stops to be lowered such that that route is actually set up and the initiating signal to be called and probably cleared. This is called "route selection," and is the basic operation of an NX/UR interlocking.

At that point, when the route is all set up, white lines of light will indicate the track sections which form the route in the proper order. However, you might find, when initiating a route, that not all the exits you expected are offered: you will then have to figure out why and cancel conflicting routes and/or move trains out.

1.04 Signal Calling (Requesting)

Q. How does the Tower Controller control the signals in an Interlock?

A. Signals that can potentially affect each other and be affected by switches must be controlled by the interlocking such that their interactions be safe. Each signal is conceptually associated with a PB, as is with the Hillcrest Tower NX control panel, which in earlier interlockings was an actual lever, controlled by the Towerman. In an "all relay" interlocking such as Toronto's, the lever *is a push-button*, and its effect is simulated by logic. If the PB is in its normal position ("cancelled"), the associated signal is guaranteed to be red. Only when the button is pushed to its other, called, position, is the signal able to clear in response to train movement.

Therefore, the PB's controlling signals were interlocked with each other -- they couldn't be moved out of their normal (signal may not clear) position unless other PB's controlling other, conflicting signals were in the normal position. Of course, exactly which signals conflict is highly dependent upon switch position, which is in turn constrained by signals and so on.
In an NX interlocking implementation such as Toronto's, signals are not operated directly, but are controlled automatically by the route selection mechanism (see Basic NX Operation). It is possible to "cancel" any signal, though, that is, pull up its button to the "normal" position simply by unfleeting and pulling up the appropriate button of the signal you want to cancel.
Note that the GK light for a home signal will come on (red) when a route is initiated there; you can tell that it is not really called yet, i.e., the route not yet set up, by the possible presence of exit lights and the absence of white lines of light, indicating routes successfully set up.

A called signal will not clear until certain conditions are met. These conditions are the lack of trains (except for call-on) and conflict along the signal's control length.

1.05 Auto-cancel and Signal Fleeting

Q. What does fleeted mean?

A. At active interlockings, it is quite common to set up a different route for every train, that is, very few successive trains take the same route. In these cases, it would be convenient if the train could "cancel its own route," that is, clean up behind itself to facilitate subsequent setup for the next train. In other cases, especially at inactive interlockings, it would be convenient to establish routes and leave them set for many successive trains.

NX/UR interlockings provide both modes of operation. The default is auto-cancel: when a train passes a home or approach signal, the signal will be cancelled and the GK light on the panel will go dark. Although the first track section beyond the signal will show red (occupied), subsequent sections will still show white and remain part of the route, lit up in white, held by route locking until the train passes through.

To cause a signal to remain called, it must be fleeted. On a NX panel, this is done by turning the signal button (which is normally pushed) in the direction of traffic. When the signal is fleeted, the call for the signal, will not be cancelled by the track occupancy resulting from train motion. Automatic and backup signals may not be fleeted, approach and home signals may.
The fleeted status on Toronto's NX Panel is shown by a little white lights underneath the button for that signal. Turning the button on a fleeted signal un-fleets it, but does not cancel it. Turning and lifting up on the button on a fleeted signal, however, both un-fleets and cancels it simultaneously.

1.06 Switch Locking

Q. What do is meant by “the switches are locked”?

A. Moving the points of a switch while a train is on that switch, or about to hit it, is, obviously, about as unsafe as blowing a hole in an airliner in flight. Switch movement must be inhibited when the track sections of a switch are occupied, or train movement is cleared into it.
Interlockings provide a facility called "switch locking" which permits the movement of the switch only when the proper set of conditions prevails. When the switch may not be moved, it is said to be "locked;" it may only be moved when "unlocked."

If you see a line of red lights on the switch or any part of its track sections, that is a train, and that is why it is locked. The switch will be unlocked only when that train clears the switch track section limits, if all else is OK. If you see a line of white lights, that means that movement is either routed by some signal further back along the white lights, or already in progress (see Route Locking), in which case the switch will not be unlocked until the train is gone.
If the switch still will not move, it may be because it is within the far extent (overlap) of the control length of some signal further back, and the proposed movement would set a trailing-point switch against that signal. In this case, you must cancel that signal, which locks the switch, before attempting to move it. The interlocking will usually flash the GK light of that signal and the points of the switch in the position they are locked when you try to initiate: see Locking Conflict panel indications for more on this.

Again, in an NX interlocking, you do not move switches explicitly, but via route selection (see Basic NX Operation). Nevertheless, the route selection mechanism might refuse to offer certain exits because the switches that would have to be moved to route to that exit are locked, and may not be moved to that position.

1.07 Approach and Time Locking

Q. Why do I have to sometimes wait to at a double red (Yonge W/B X-58) after a train or work car has gone into Lower Bay and is no longer in sight?

A. Approach locking and time locking are consequences of the fact that trains are big and heavy and, when moving fast, take a long distance and time to come to a halt, no matter how hard the brakes are applied. Consider the following schematic scenario:

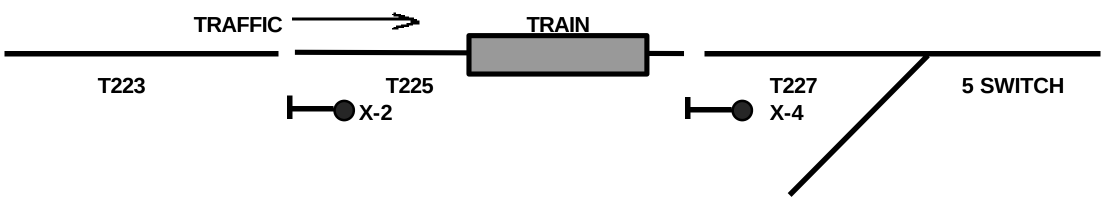

Approach Locking

Imagine that switch 5 is set to a diverging route. By all that should now be understood, this means that signals 2 and 4 must be red (of course, 2 must be red because of the train), and cannot even be routed.

Now suppose that switch 5 was set to the normal route, with the train on it, and both signals were clear. Suppose the train were moving rapidly, and as the train approached signal 4 and the switch, as shown, the Tower Controller cancelled signal 4, making it red, and then caused the switch (5) to move, while the train was still barreling towards 4 and 5. Nothing so far would prevent this -- the train would indeed be tripped by the stop and signal at 4, but would take a long distance to slow down, and would doubtless run the *trailing point* switch and damage it and/or be derailed.

Cancelling a signal at *any time* cannot be made impermissible or prevented -- it is a safety feature that a train may be stopped at any time. But telling the interlocking that "all is OK, 4 is no longer called" when there is a train coming at it is not satisfactory.
The solution to this problem is *approach locking.* When a train is approaching a *home signal*, and it is cancelled by the tower, the signal will de-route and go red, but the rest of the interlocking

will not "believe" it, and the route will lock out other routes and hold switches locked via *route locking.*

Of course, it might be so that the train has been sitting there for a while, and is not moving rapidly. As just explained, approach locking would make it impossible for the Tower Controller to change his or her mind with a train facing a signal. The solution to this is *time locking*. Approach locking will release after a sufficiently long period of time. Approach locking will also "quick release" as soon as a train gets into the route ("accepts" it), so that route locking can take over.
The approach locking of *approach signals* is one of the chief reasons for their existence. Consider again this illustration:

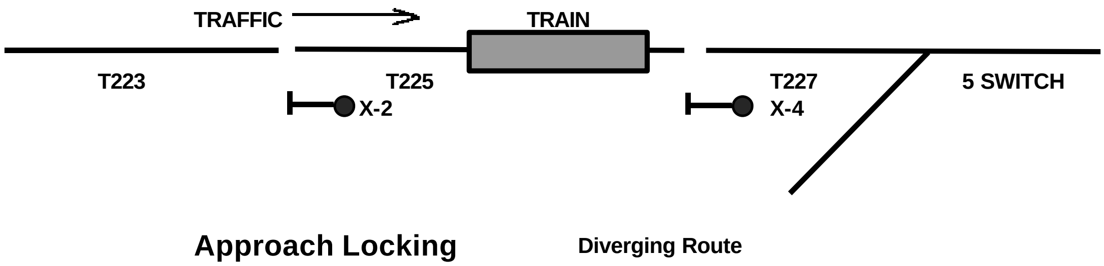

Approach Locking

Imagine that the train is not there, that signal 2 has been clear for a long time, and 4 has been de-routed for a long time, and the switch is set for the main route. The interlocking has "believed" that 4 has been de-routed (approach locking has been satisfied) for a long time. Again, 2 is routed. Along comes the train at high speed. As soon as the train passes 2, the Tower Controller cancels 2 and moves the switch. Nothing now prevents the train from tripping at 4 and derailing at high speed.

The solution to this is *approach locking* for the approach signal 2. The interlocking will not "believe" that 2 is cancelled if there is a train past it, i.e., in section 225, until time-locking for 2 expires. Time locking timing starts when a signal is cancelled. Thus, it would be impossible for the switch to be moved or a conflicting route to be set up unless 2 was either no longer called for "a long time" or there was no train in the "approach section." The problem stops there-- the train cannot pass 2 without coming to a full stop, and thus, its speed at 4 would be limited.

1.08 Route Locking

Q. What do they mean by “the switches are timing down”?

A. Route locking is a interlocking feature that allows all the side-effects of a routing to continue to hold once a train has entered ("accepted") the route, even if the signals are, as is common, cancelled shortly after that. That is, a train having accepted a route is "just as bad," as far as the locking of switches (see Switch Locking) in the route and potential conflicting movements (see Signal Control Length) are concerned.
Route locking is one of the criteria used in Switch Locking.

1.09 Timed Signals

Q. What are timed signals?

A. Sometimes, besides the functions discussed under Control Lengths, signals equipped with timers are used to restrict the speed of trains. This occurs in two forms, Grade Time (GT), which restricts the speed of trains unconditionally, and Station Time (ST), which allows a signal to shorten its control length if a train approaches it at restricted speed. GT is easier to understand, and will be explained first. ST will be explained subsequently.

1.10 Grade Time (GT) Signals

Q. What is Grade Timing?

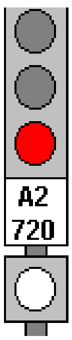

Automatic signal displaying Lunar White (GT) aspect

A. GT (Grade Time) signals are so called because they are used on grades (slopes) and around curves. They operate by not clearing until the train has spent a sufficient number of seconds in the track sections in front of the signal to be evidence of low enough speed.

The simplest kind of GT signal is the "one shot" GT, which, in the T.T.C. signal system, has a single bright white ("lunar white") lens, which, when it is displayed with a yellow or red signal indicates, that if the next signal is approached sufficiently slowly, that it will clear. This is called a "one-shot" GT signal.

Q. Approaching a terminal with trains present there are no lunars, why?

1.11 Station Time (ST) Signals

Although Station Time (ST) signals are extremely common in the T.T.C, they are not identified as such. Station Time facilitates trains "closing in" on stations to keep them moving, although slower, when trains are stopped ahead in a station. Here is the canonical situation:

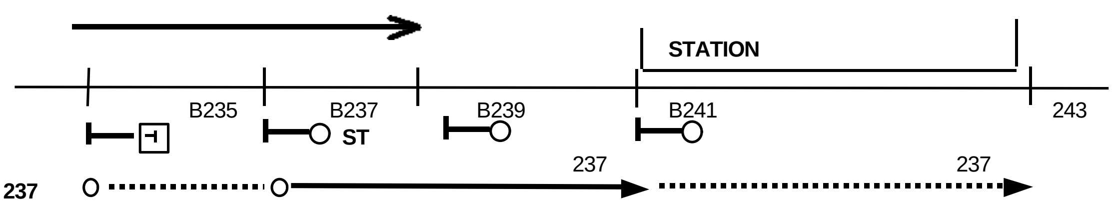

Illustration of ST (Station Time) signal timing

The control length of signal B-237 is shown as extending from the signal itself down track to IJ 243. That is, if any portion of a train appears between insulated joints 237 and 243, 237 will display a red, "stop" indication to oncoming traffic.

However, if there is a train, or portion of a train occupying the "dotted portion of the control length," that is, between 241 and 243, but no portion of any train in the "solid portion of the control length" between 237 and 241 (that is, the end of a train is in the station), a new condition obtains: a second train approaching 237 sufficiently slowly to "time out" the section between 235 and 237 will cause signal 237 to "cut back" its control length to the solid portion, effectively ignoring the train in the dotted portion. Of course, this is not unsafe, because signals 239 and 241 still enforce a separation.

An ST signal allows trains to "close in" on each other by acting like a GT signal if only the dotted portion of its control length is occupied. Some signals are GT and ST simultaneously; the ST timing must obviously be slower than the GT timing if this is to be meaningful.

1.13 (Automatic Train) Stops

Q. How do train stops work and what is the purpose of auto key bys and call ons?

A. Unlike street traffic lights, subway signals can force disobedient traffic to stop. Next to each signal (except certain yard signals) at track level lies a T- or hammer-shaped apparatus, controlled by motors, which rises (via a counterweight or springs, for safety reasons, when not actively held down) to an upright position more or less when the associated signal is "red," and wants trains to stop. This trip, or train stop, will engage a brake valve on the underside of a passing train and trip it, i.e., cause the train to go into air emergency and stop.

The "more or less" is due to the fact that a signal becomes red as soon as the front of the first car of the train passes it, and care must be taken not to trip the rest of that very train. Also, trips for signals that are being passed backwards, such as at interlockings, must be carefully cleared (put "down", i.e., not "tripping") when routes are cleared over them, and restored when the train is passed or the route cancelled.
On the subway, stops take a second or two to "operate", i.e., come up or down. All signals (except a few yard signals) have stops. *Automatic* and *approach* signals implement a feature called *automatic key-by*, ("AK"). This feature allows you to pass a red automatic or approach signal if you creep up to it extremely slowly: the relative placements of the train wheel, trip valve, IJ and stop are sufficiently carefully worked out that if a train crawls past the IJ of such a signal, the stop will go down. You may not, and will not be allowed to, "key by" a home signal.

Because the standard implementation of AK optimizes the use of the same track relay contact to support reverse-direction motion, you may observe that the stops for automatic signals do not come up immediately behind a train, but at a distance of one track section behind the train. For this reason, overlap of control length is critical if stops for automatic signals are to have any efficacy.
There is also the *manual call-on* feature, which allows a *"manual key-by"* past home signals in certain very special circumstances. The motorperson and the Tower Controller can cooperate to lower the stop of a home signal when it is possible and necessary.

1.14 (Home Signal) Call-On

A manual Call-on is a feature which allows trains to pass red home signals under very special circumstances by means of close cooperation between the Tower Controller and the motorperson.
The idea of a manual call-on is to allow trains to close in on each other past home signals, and or be able to pass a section of track that is showing occupancy due to a trackdown, just like with automatic key-by on automatic and approach signals, except that both the Tower Controller and motorperson must cooperate via taking special explicit action: if the circumstances are acceptable for a manual call-on, the Tower Controller displays a call-on indication on a home signal, and the motorperson creeps up to the signal and accepts the call-on by pressing a special button, plunger or lever that causes the train stop to be lowered. He or she may then pass the

signal prepared to stop within vision. The manual call-on "aspect" (the way it looks) of a home signal is red over red over yellow; there is special yellow light at the bottom of each home signal reserved for this purpose.

If a manual call-on is to be cleared on a home signal, there can be no conflict of switches, signals, or routes in the control length of the signal -- if there are trains in the control length, a call-on is permitted, but if there is any other problem, it is not. What is more, the track section in front of the signal, the "approach section," must be occupied -- a train must be at the signal. To "clear the call-on," the Tower Controller sets up the route as usual, but at some time after initiating, presses the call-on button for that signal. When the call-on is cleared, the GK Light will blink yellow, and the call-on indication will be displayed. At that time, the "call on is being offered," as a full signal display will reveal. The signal, as always, can be cancelled at any time by canceling it.

1.15 Locking Conflict panel indications

Q. How does the interlocking notify the controller of potential conflict of routes?

A. An NX/UR interlocking has the ability to tell you why it refuses to set up a route that you ask it to. Normally, this is because either an exit you want is looking into a cleared (or approach-locked) route in the opposite direction, or a switch that must be moved to establish that route is locked in the wrong position. The former case can usually be discerned easily by inspection. The latter case is equally easy to discern when the track section containing the switch is occupied or routed -- lit up in red or white, the reason why the switch will not move is shown plainly.

1.16 Relays

Q. What do relays have to do with signals?

A. A relay is an electromechanical switch, a set of switches operated together by an electromagnet. Electricity flowing through the switches can be used to control any electrical apparatus, or, more relays. Relays are neurons, and form "logic networks" implementing complex functionality. Although relays as logic elements have largely been obsolete by computer technology, the earliest electrical computers, including phone switches, were built of relays, and railroad signaling and interlocking is still implemented in relays today. This is because their modes of failure, unlike those of software, can be fully enumerated and designed around.

Relays sport two types of switches, or "contacts," those that allow electricity to flow through them when the "coil" (electromagnet) is energized and not when not, and those that allow current to flow when the coil is not energized, and not when it is. In railroad terminology, the former are called front contacts, and the latter back contacts. A contact is said to be closed when electricity can flow through it, and open when not. A typical railroad signaling relay might have a dozen contacts total, some front, some back. Railroad signaling relays are of extremely high quality, rugged, and quite expensive.

A relay is said to have a state at any time: it is said to be picked or up if the coil is energized (and front contacts closed and back contacts open), or dropped or down if the coil is not energized, front contacts open, and back contacts closed.

1.17 Further Information

If anyone requires any further information or questions about the signals and how they work you can contact Russ Hilder at pax 2403. Any questions the Tower Controllers can't answer will be forwarded on to the Signal Department and a reply will be sent. All operators and Supervisory Staff are invited to come up to Hillcrest Tower at any time to visit and learn more about our system.

1.18 Credits and Acknowledgments

Most of the above text and illustrations were obtained from the documentation of the NXSYS Signal Simulator program, and used by permission of its author, Bernard S. Greenberg, of Boston, Mass., USA. NXSYS is a highly interactive application for Microsoft Windows(tm)that simulates virtually all signalling and interlocking features as used in Toronto and New York City, including everything described in this manual, and is geared toward being a learning tool for exploring and experimenting with these concepts. It is for free (but copyrighted by its author), and comes with full documentation. 16 and 32-bit versions are available.

The NXSYS (public domain) program is available at TTC from Russ Hilder. The program is under active development, and current versions, status information, and the like are available at its World Wide Web page:

http://www.basistech.com/bsg/nxsys.htm.

Mr. Greenberg's email address is bsg@basistech.com.