K-System? Try the T-System!

The name is just a play on Bob Katz' "K-System", but it doesn't have much in common with that overcomplicated mess.

Don't take the name too seriously; it's just something to refer to this process and the chosen reference headroom by, and I am not claiming to have invented any of what I am about to describe! 🙂
The "T-System" aims to establish good working practices for mixing engineers, so that mastering engineers have a perfect source to work from. It greatly simplifies the job of mixing engineers as well, by giving you a consistent reference volume and a lot of digital headroom for your transients and louder passages. This helps you balance the mix without constantly accidentally clipping things into "digital 0" or using compressors/limiters, and it gives your music incredible clarity and dynamic range -- easier than ever before!

The "T" jokingly stands for "The". So "The T-System" stands for "The The-System". But don't pronounce it that way, because people will think you're stuttering. 🙂

Adopting the system is a 3-step process, covered in detail in the article-length writeup below. Here's a brief summary of how easy it is to get started:

* If your speakers are already calibrated to output at an equal volume, you don't do Step 1 at all.

* You only do Step 2 a single time: Calibrating a song to -23 LUFS of digital headroom and then determining your preferred listening level that you want to hear -23 LUFS material at. You determine your own comfortable listening level, so you don't need an SPL meter. It's very easy. And if you want to calibrate the physical output volume to the broadcasting standard of 83 dBC SPL (or their equivalent for smaller rooms), then you're welcome to do that, so the system is fully compatible with the cinema/television/radio broadcasting industry.

* After that, you just have fun (Step 3), without needing to think *at all*. It liberates you to focus on the music without ever worrying about digital headroom or clipping anymore.

Step 1: Calibrate your speakers for even stereo balance.

To do this, you are going to need an SPL meter capable of measuring "dBC" (C-weighted decibels) with a slow response time, and you're also going to need a pink noise test signal which is bandpassed to only output 500 Hz - 2 kHz (so that it takes bass/treble reflections in your room out of the loudness equation).

1a. iPhone preparation.

If you own an iPhone, there was a recent scientific study which proved them to be as accurate as calibrated professional meters, so don't waste your money on a standalone meter. If you're on Android, you're out of luck (and will have to buy a physical meter), since every manufacturer puts different microphones in their phones and standardized Android calibrations don't exist.

For iPhones, I suggest the "SPLnFFT" app, which is the most affordable ($3.99) out of the two main test winners. When you start the app, go into the "Mic." page and click the "Mic. gain: Set" button. Turn your phone up to the loudest volume. Go into a quiet room, lay the phone down flat and press the "Restart" button. It will play a test tone and calibrate its "low gain" mode, which is what it uses if you ever want to measure really quiet sounds. Then just press "Done". Now, to set up the app for calibrating your speaker volume, you'll want to press the "Histo"(gram) tab, then tap the "FAST" label in the upper left of the white meter so that it says "SLOW" instead, and tap the "dB(C)" button at the bottom of the window to set it into C-mode.

1b. Set up the test signals.

Download these pre-made pink noise files:
* Pink Noise 500-2K -20 dBFS RMS Lt
* Pink Noise 500-2K -20 dBFS RMS Rt

They are stereo files, but only one channel plays in each (either left, or right), which makes them very easy to handle.

Import them into two separate stereo tracks in your DAW.

Set your DAW and soundcard (digitally) to 0 dB (unity gain).

Set your audio output device to 100% volume via your physical master volume knob/fader (or if you've got a "-X dB, 0 dB, +X dB" knob, then set it to the middle "0 dB" point). This is important.

Next, set your speaker's trim pot (at its back) to the lowest value.

PS: Do *not* play any regular music/sounds right now or you could possibly blow your speakers!

1c. Perform the speaker calibration.

Set your SPL meter to Slow response + dBC weighting. If your meter can only do dBA (A-weighting) then that will be okay too, since this bandpassed noise has taken out the bass/treble which would have caused a big difference between the two weighting types.

Now place the iPhone (or physical SPL meter) so that its microphone "tip" is at the normal listening sweetspot at ear level, and aim its microphone at the ceiling. Regardless of what type of SPL meter you're using (iPhone or otherwise), you must always aim at the ceiling to ensure that it doesn't accidentally point more towards any particular speaker. Try to lock the device in place (using a music stand, or someone else to hold it for you), which will simplify things.

Loop the pink noise sounds in your DAW, and play them *one speaker at a time*. Adjust the trim pot at the back of the speaker until you get ABOUT 83 dBC at the sweetspot. Different speakers will have different innate sound pressure levels for any given voltage, so yours might be really effective (loud), or really quiet, and the range on the speaker's trim pot might not be enough for that calibration. In that case, you can aim lower/higher. The point is just to pick a reference level on *one* of the speakers.

Next, play the other speaker instead and match it using its trim pot, so that it gives the exact same level +/- 0.5 dB. That's good enough to ensure that the speakers output an equally loud stereo image.

Alright, you're done with the speaker calibration, and can throw away the "pink noise" files.

Turn down your physical master volume knob to 0 now. We won't use audio again for a while, and you don't want to blow your speakers.

Step 2: Establishing your personal "reference volume" for working on music.

The "reference volume" is any volume that feels comfortable to you, and which you can imagine yourself working at day in and day out as your "optimal reference point" while mixing.

To establish the correct physical volume, we will need to output music with 23 dB of transient-headroom.

So let's get started on setting this up!

2a. Getting an "ITU-R BS.1770" loudness meter.

You will need an "ITU-R BS.1770" loudness meter, which uses the new, extremely accurate "LU" measuring unit. A "LU" (Loudness Unit) is equivalent to 1 dB of perceived volume to the human ear, and it's far superior to the older RMS measurement (which doesn't always detect the proper loudness, since RMS rises too strongly from bass and not enough from treble). In a lot of cases, the RMS and LU will be very similar, but LU is always the more accurate measurement. The term "LUFS" is loudness units relative to digital 0 (FS="Full Scale"), so "-10 LUFS" means an average loudness that's 10 dB quieter than digital zero. The "FS" suffix just clarifies that you're talking about Loudness Units relative to digital zero. The metering software you choose also needs to have True Peak (dBTP) metering, which catches inter-sample peaks.

My preferred meter is "iZotope Insight" ($499, but free if you buy Ozone Advanced). It support stereo and surround, and contains a huge amount of audio measurement tools and even allows you to print the loudness to images, files, or to an automation track in the DAW (great when working on loudness compliance during mastering). Some people use inferior products such as MeterPlugs LCAST, which is $399 for the surround version and only does basic LU metering without even having any of the other incredible measurement tools that Insight offers. If you want accurate LU metering on a really cheap budget, then get "ToneBoosters EBULoudness" (19.95 EUR).

A "Loudness Unit" meter has four important values:
* Integrated Loudness, or "I": This value builds up slowly during the playback of your whole song, and detects the highest non-momentary loudness that was discovered in your material while playing it from start to finish. It uses a gate and a slow reaction time, so that it ignores silent portions and any extremely brief (momentary) periods of extra loudness. This is the value calculated by the broadcasting industry and services like iTunes and Spotify when they need to decide how to normalize the volume of your track before putting it up on their service (more about that later).
* Short-term Loudness, or "S": Uses a 3 second sliding window for detecting the average loudness of the last 3 seconds of music; this is the most useful value for you during mixing (we'll get to that later).
* Momentary Loudness, or "M": Uses a 400 ms sliding window for the last 0.4 seconds of music; it's good for seeing the volume of quick, loud bursts, as well as being useful for quick initial volume balancing (since it reacts faster to your fader changes).
* Loudness Range: Again, this value builds up slowly during the playback of your whole song, and it detects the volume difference between the softest intro/break portions of your song, and the loudest climax. It's basically the emotional range of the piece. How much range your music will have depends on the genre, so I'm not going to talk more about this value, but be aware that it exists and can help you determine the loudness difference between loud and soft portions of your song.

2b. Setting up your meters for the correct loudness target for mixing.

We'll be working with 23 dB of transient-headroom above the average loudness of our tracks; this gives you an extreme amount of digital headroom so that you'll pretty much never, ever need any limiter on your pre-master mixdowns. The number 23 was chosen because it complies with the headroom chosen by the movie/television/radio/broadcasting industry, so that people (such as composers) working in those fields will have a very easy time integrating this system with their world; literally the only adjustment you'd make if you're working in the movie industry is that you'll also want to calibrate your output volume so that it *matches their standard of "-20 dB RMS pink noise = 83 dBC SPL at the sweetspot in a large acoustically treated room, when a single midfield/farfield speaker plays", or the equivalent loudness for your particular room size. See Appendix A (separate post) if you're interested in exact compliance with the broadcasting world, but I don't recommend it, since it's a bit complex to set up - and as a music producer it's more important that you've got a level that's comfortable for you. I'm just letting you know that the *option* of full Broadcast compliance is available to you by doing a specific physical volume calibration as well. Anyway, moving on...*

Having 23 dB of headroom in your music makes balancing tracks effortless and fun. The idea is to mix with a low digital volume, and to raise your physical volume knob to get "commercial audio levels" while working, instead of blindly slapping on limiters/compressors/ cranking things in the box. This system gives you so much headroom that you pretty much never have to watch any channel meters while working on mixes!

If you're using Izotope Insight, it has a proper preset already. Load the "Stereo Mixdown" preset. If you're using a different LU meter, you'll need to set it up properly:

* Target True-Peak: -0.2 dB
* Target Loudness: -23 LUFS

Note that we'll only treat the "Target Loudness" as the target for the *average* loudness portions of our tracks. So the louder portions of your track will be louder than that, and quieter portions will be quieter than that. So while working on louder portions of your mix, you'll be seeing "red" in your LU meters due to the target loudness. That doesn't matter *whatsoever*. The -23 LUFS is only our target for the average portions of the song.

That's different from when you use Loudness Units for mastering, where "target loudness" is instead the ceiling "Integrated Loudness" value that we never want to exceed. (And that's why the meters are displayed in red when you exceed the "target" value).

2c. Calibrating an existing song to -23 dB LUFS.

Now that you have an LU meter set to a -23 LUFS target, you'll need to calibrate an existing song to that level so that you can use it to determine your desired working volume.

I recommend that you use a well-mixed song with clear, punchy transients. Not some squashed electronic garbage with barely any dynamic range. It doesn't matter what music style you'll be producing; you just need something with clear transients and a good frequency balance when setting this up!

Import the song/project into your DAW. Find a portion of the song that represents the *average* loudness of the track. Not a quiet intro. Not a loud climax. No loud hits, either. Just the average volume where the song is chugging along and doing "its thing". Select a 4 bar loop from that section.

Now turn off your speakers/monitor controller. You do *not* want to hear anything anymore right now.

Put a Gain/Trim plugin on your master output channel, and then insert your LU meter after that. Make sure the meter is set to -23 LUFS target, as mentioned earlier.

While the song is looping, adjust the Gain/Trim plugin and look at the "Short-term Loudness" / "S" value and get it in the ballpark of ~ -23 LUFS. It can be a few decibels higher or lower. This is just to get your initial adjustment in the right ballpark.

Now, it's time to fine-tune the gain using the more stable and non-jumpy "Integrated" value. Click the LU meter's "reset" function to reset the "Integrated Loudness" measurement (in Insight, you simply click the currently displayed value to reset it), and then let it measure for 5-10 seconds until the value stabilizes. You might see something like "Integrated Loudness: -21.5 LUFS". Then you know that you're *about* 1.5 LU too loud compared to the target of -23. So adjust your gain/trim plugin by the appropriate amount, such as lowering by an additional -1.5 dB in this example case. Then click the "reset" function of the meter again to restart the "Integrated Loudness" measurement and see if you've hit -23 this time. Stop when you've finally hit -23 pretty much spot-on.

If this accurate measurement seems excessive, it's not. You'll want your reference track set up properly so that all of your music after that also has the correct transient-headroom of 23 decibels.

PS: The reason that you did this measurement with your speakers *off* is so that you don't go crazy hearing the same loop over and over. This is a scientific measurement to find a target Loudness Unit volume digitally; *not* something you need your ears for.

This is what a "-23 LUFS" calibrated track looks like in iZotope Insight:
http://i.imgur.com/fV1UP8q.png

Note the "-23.0 Integrated (LUFS)", and the yellow [] brackets around the "I" column to indicate that the volume is very stable around the chosen reference point. If we'd chosen a bad loop with both soft and loud portions, then the yellow brackets would be very wide (that'd be bad); you need to loop a section with stable volume so that the measurement becomes meaningful.

2d. It's finally time to establish your preferred working volume!

Alright, with your calibrated track ready, it's time to listen at last!

Set your *physical* master volume knob to 0%, and keep playing the loop. It is extremely important that you *only* loop the average-volume portion that has been calibrated to -23 LUFS, so that you're listening to the average working volume.

Now sit back, relax and slooowly raise the physical volume knob. It will start out super soft and quiet and then gradually rise. You *need to* rise out of the "this feels soft / I'm straining to hear the details" zone, and get into the zone of "this is a volume I can imagine working at for several hours a day; I can hear everything clearly, and it doesn't cause any ear fatigue". The louder the volume is, the clearer you'll be able to hear bass and treble details (due to the Fletcher-Munson curve of your ear, which is at its flattest around ~80-85 dB of *perceived* volume). But do *not* enter the "damn, this would quickly get annoying" area.

Remember that you can always go above your reference volume temporarily while working on projects later, if you need to hear something more clearly; you do *not* want to have such extra loudness "permanently" baked into your reference volume position! Likewise, it's very important to remember that if you choose a reference volume that's too low, then you'll tend to overcompress things to make them thicker/louder, and you'll tend to apply "smile curve" bass/treble boost EQs since your ears aren't getting the candy they want. So be sure that you settle for a volume that's "comfortable, clear/detailed, and loud enough".

When you've found a reference volume you like, then you should relax and listen to a longer portion of the average-volume section of the track. Does this sound like a volume you'd be happy to work at? Does it let you hear the details clearly? Is it non-fatiguing?

You can actually double-check your chosen reference level via your SPL meter, if you set it to "Slow" and "dBC" mode. If you're working in a home studio with midfields (or even nearfields), in a smaller room, or in a small/medium sized control room, then you'll probably end up somewhere around 72-76 dBC. The early wall reflections of small rooms cause them to psychoacoustically sound louder to your ears, which is why that number may seem low, but your brain will *perceive* it as something in the 80ish range. If you're in a big room with midfield monitors placed far away, then you'll probably end up in the 78-85 dBC range. That's because "louder" sound pressures actually don't seem "too loud" anymore when they're played in a big room, thanks to the lower amount of early reflections. Whatever you choose, it doesn't really matter. Choose what works for you.

Alright, you've found a comfortable volume that you'd be willing to use for the majority of your work? Then take a piece of tape or something and accurately mark that position on your physical master volume knob/fader. That's now your "reference volume" for how *physically* loud you want a -23 LUFS (digital) "average loudness" section of a song to be, whenever you're working and making mix decisions in the future.

Step 3: Time to enjoy the fruits of your labor!

We're almost done! You've now got a balanced stereo image, and you've established a preferred (physical) "reference volume" for -23

LUFS (digital)! You're now ready to produce music and mix songs, so let's talk about the workflow!

The workflow is the same regardless of whether you're mixing an existing song, or creating a new one.

3a. Recall your physical reference volume.

First of all, set your physical master volume knob/fader to the marked position that you determined as your reference volume earlier. You're welcome to go softer or louder at times while working on a project, but you should always be at the reference volume when making really critical mix decisions or comparing different songs. Oh, and as soon as you've finally learned your reference volume by heart, you won't even need digital meters anymore!

Here are some helpful guidelines: If you're just working on the song's arrangement or doing something else that doesn't require you to accurately hear everything that's going on, then feel free to lower the volume. If you're doing some really critical mixing work and need to be able to hear it even more clearly, then feel free to raise the volume. When you're doing general tweaks, adding new tracks to the song (and balancing their volume), doing final listening, comparing to other songs, etc, then always use the fixed "reference volume" spot.

Tip: After working with this system for a while, you should settle on preferred "soft" and "loud" positions on your master volume knob/fader, and mark those as well. That way you will always go to the exact same "soft" or "loud" volumes, and will be able to learn how they sound, and after a while you'll be able to make mix decisions (such as EQing) even at the "soft" volume. But it'll take you some time of working with the system to find your optimal "soft" and "loud" spots, so leave that decision for later!

3b. Adjust the digital volume of your project.

Insert a Gain/Trim on your Master Output channel. The gain plugin goes after any of your 2-bus' mix-compression/special processing, but *before* final limiting.

- If this is an existing project, then you're going to have to adjust the gain to whatever the project needs in order to comply with the

-23 LUFS target (during the normal-loudness portions of the track; using the same method we used for calibrating an existing song above).

- If this is a brand new project you haven't even started yet, then set the gain to -18 dB. That will ensure that all of your DAW tracks automatically get 18 dB of gain reduction before going into the master channel, thus ensuring that you can keep all of your individual channel/bus faders close to their "0" mark where their resolution is the greatest, instead of having to push them all down into the bottom ranges where their resolution is awful.

3c. Digital "target loudness" metering.

Next, insert your loudness meter, such as iZotope Insight, after the gain plugin (before any final limiting). Make sure it's in the proper mode (-23 LUFS loudness target, max -0.2 true-peak value).

Up until now, I've only showed you how to measure in LUFS (loudness relative to digital 0), but here's a tip: Good meters (such as iZotope Insight) will allow you to shift from an absolute "LUFS" scale (relative to digital 0) to a "LU" scale (shows the difference relative to your chosen target loudness). Here is a screenshot showing the "Options: Levels: Scale: Relative" setting enabled: http://i.imgur.com/6HOAq48.png (compare it to the earlier screenshot from the end of section "2c" above). It shifts the meters so that they show "LU" instead; loudness relative to your target. That way you can see that your loud portions are "+2 LU" (for instance), and that your quiet portions might be -5 LU, and so on. It's a very musical way of checking levels. I suggest making the change and saving it as a "Stereo Mixdown (Relative)" preset. There's no reason to use the harder-to-read *absolute* LUFS scale if your meter software allows you to use a relative display method instead.

- If this is an existing project that you're re-mixing, you would now adjust the gain plugin until the LU meter shows -23 LUFS for the average-volume portions of the song. If you've set your meter to relative mode, you'll of course aim at "0 LU".

- If this is a brand new project, you don't really have to do anything. Just create music that sounds physically "right" on your volume-calibrated monitors. You'll automatically be creating music where the average portions hover around the -23 LUFS mark (or if your meter is relative; "0 LU"), and where the louder climax portions probably hit about +2 LU higher than that (so around -21 LUFS). If you ever want to double-check that your project is in the correct range, simply loop an average-volume portion of it, reset the "Integrated LUFS" meter, and then loop it for 5-10 seconds and check that the Integrated LUFS meter sits at -23 LUFS. However, if you're mixing songs that are intended to be part of an album where some of the songs are intended to have a lower average volume, then you don't have to care about the LUFS value at all; you just simply mix to a volume that sounds *physically* "right" to you relative to your chosen reference volume, and you'll be alright!

3d. Enjoy that headroom!

When you're working like this, you'll find that your digital track/bus faders are suddenly a lot more useful. You've got 23 dB of transient-headroom, and plenty of room for your track to go louder or quieter, without ever battling the digital faders.

You will also discover a few things:
A) You probably won't need a limiter *at all* during your pre-mastering mixdown. If you do, it will only be to catch a few rogue, extremely short transient peaks.
B) Your mastering engineer (which may be yourself) will thank you for preserving all of the dynamics, transients and headroom, which allows them to deliver the best-sounding final master.
C) Your own mixes won't need any compression on the master bus, unless it's for a specific glue/pumping effect (such as in electronic music). Most of the time, you're better off just applying compression to individual tracks and things like the drum bus, and leaving the master channel *completely clean*.

3e. Comparing to other people's music (your reference mixes).

If you want to compare your mix-in-progress to reference tracks, then make sure that all of your reference tracks are also calibrated to -23 LUFS average loudness, so that they are directly comparable. If you do that, you'll find that all songs have equal loudness, and the only thing that differs is the transient clarity (the more headroom that reference song had, the less squashed and therefore clearer its transients will be).

3f. Gain staging in the box.

It's also helpful (as always) to apply gain-staging for individual tracks in your digital mix, but only when you use plugins that actually need it. Modern plugins use floating point maths which doesn't care about human concepts such as "volume", and therefore do not need gain staging (it has zero effect for them).

But analog-modeled plugins such as UAD (see this thread) usually operate at an internal reference level of "-18 dB RMS = the average volume", or even -12; so if you feed them a digital signal that hits something like -2 dB RMS, then their internal analog modeling will be fed a way too hot signal and you'll get way too much saturation and such out of them. This only happens for analog-modeled plugins, because those kinds of plugins have to choose a "digital volume to virtual voltage" conversion ratio, so that their virtual circuit emulation (which uses voltages) can process digital audio. So again: You only need to gain-stage when using analog emulation plugins.

To achieve the correct gain, you simply have to insert Gain/Trim plugins before the exact plugins that actually need a lower input. Do not blindly put trims on every track; because 99% of plugins out there don't need it at all. The rule of thumb is as follows: If it's a modern, digital plugin, it uses floating point math and will not need gain staging. If it's emulating analog hardware, it probably emulates the voltage characteristics too and you'll need to gain stage. Sometimes the plugin itself has an "input volume/gain" knob that's pre-analog circuit emulation and can set the correct circuit input volume for you, but other times you'll have to manually insert a gain/trim plugin before the plugin.

3g. Enjoy your incredibly clear and punchy pre-mastering mixdowns!

3h. A word about mastering...

If you're the one doing the mastering, then the process of attaining the final delivery volume is extremely simple:

- First you have to play through the track and see where the "Integrated (LUFS)" value ends up.
- Then you basically just boost the track by the amount required to reach the volume of your target genre. You'd simply insert a limiter, raise the gain and chop the transients (if any get in the way) until the overall mix volume rises from your -23 LUFS mixdown reference, up to whatever the genre requires.
- Modern genres use everything from -16 (audiophile quality) to -6 LUFS (hypercompressed pop) for the final deliverable. The higher the number, the greater the clarity; the lower the number, the louder it will be.
- If you're releasing for iTunes MFiT or Spotify, then you need to aim for a final Integrated LUFS (whole track played from start to finish) of -16 since they normalize all tracks to that volume anyway (meaning that if you go any louder than that, you'll just be silenced by their processing, and will just be losing transients you could have had if you'd kept more headroom).
- If you're doing high-quality releases on your own outside of iTunes/Spotify, then I suggest something like -10 to -14 LUFS, to maintain a great balance between loudness and good transient clarity. The final loudness choice is as always dictated by your competition in that genre.

Have fun and enjoy your awesome, newfound digital headroom! And remember to tell your friends!

(Unless you want to keep the secret to yourself.)

*- "The T-System: These go to eleven * times 2.1"!*