ECE-345 Design Exercise 3

Brian Frost-LaPlante and Nikola Janjušević

May 9, 2018

Introduction

In this exercise, a Phase Locked Loop (PLL) was designed for a center frequency of 100MHz and rail voltage of 1.8V. At this relatively low frequency, special attention was paid to the PLL's characteristics of spectral purity and phase jitter by a heavy focus on the symmetry of its operation. The PLL implementation consists of the classic four functional blocks: the Phase Frequency Detector, Charge Pump, Loop Filter, and Voltage Controlled Oscillator. This report first goes through the design considerations made for each block in the order they were designed. Next, the layout of our design, problems incurred by our layout, and techniques we used to mitigate them are discussed both in the setting of the PLL's functional blocks and in the layout as a whole. Finally, our PLL's performance is measured in both its transient response and spectral output under 9 conditions of varying process, temperature and supply voltages. Schematics, layout images, transient responses, and spectral plots are provided in the appendices at the end of this report.

Phase Frequency Detector

The first major design decision in the construction of the phase frequency detector (PFD) was the choice of topography. The circuit we chose, as seen in Appendix I, is a fairly large circuit; it contains 60 transistors arranged in NAND and NOT gate functional blocks. There exist smaller implementations of a PFD, but they either use some analog circuitry (our implementation is purely digital) or use gates other than NAND gates. The former of these is undesirable as analog is simply less resilient than digital, as it is far more affected by layout effects, temperature, process and voltage changes. It is also less compact generally, requiring larger transistors, whereas minimum-length transistors may be used without error for digital logic. The use of gates other than NAND and NOT is not preferred here, as NAND gates are simply faster than most other digital logic gates, and switching speed is quite crucial in a PFD. This circuit, although it contains many transistors, is easily broken up into functional blocks, works quickly and reliably (due to being entirely digital), and is uniform in transistor size.

The transistors are picked so that both NFETs and PFETs are minimum length. The NFETs are set to minimum width as well, and the PFETs are then matched in k by having longer widths. The result is very small logic gates, leading to nicely compact circuitry despite the large number of transistors. This is also beneficial as the switching time of transistors is faster as they become smaller.

The two-NOT-gate buffers along the top and bottom paths are to ensure certain outcomes in race condition states, and other race conditions are handled by conscious decisions in order of inputs to NAND gates. The input on a NAND gate corresponding to an NFET which is closer to the

output of the gate actually switches the output considerably faster, and this is taken into account for each input to each NAND gate in this design. The decisions therein are spared in their entirety for the sake of brevity.

Some important features of the PFD include its symmetry of output and its deadband in either pulse. At a logic level, the UP and DOWN pieces of the circuit are identical, and thus there should be some symmetry in their output and deadband, however due to the effect noted earlier with the order of inputs to the gates, this is not entirely the case. There is an inherent asymmetry which one may choose to move in the favor of one pulse at the detriment of the other. The detrimented pulse was decided to be the UP pulse, and this was taken into account when designing the charge pump. The effect of this detriment, however, was seen to be minor; assuming that a 1V or higher pulse can drive the charge pump, a deadband of about 25ps for the DOWN pulse and 30ps for the UP pulse was seen. These values are fairly small in the first place, and are not incredibly different from one another, as the signals merely differ in signal path by 4-input NAND gate input positions.

Voltage-Controlled Oscillator

The voltage-controlled oscillator (VCO) is comprised of a ring oscillator, control circuitry, and inverters at the output. The inverters at the output are simply to turn the output of the ring oscillator into a clean square wave, and are the same as the inverters used in the PFD; they are not of particular interest to the design of the VCO.

The design of the VCO was started with the consideration of the ring oscillator; this is simply an odd number of current-starved inverter blocks, each with the same intrinsic biasing, in feedback. A ring oscillator uses the natural switching times of each NOT gate to produce an oscillatory signal, whose frequency can be determined via observing the effective RC networks introduced by transistors MS2 and MS3 in each inverter and considering the effects of propagation. As the gate-drain capacitances of transistors of different sizes can be easily obtained through Cadence, the determination of the frequency of a ring oscillator for a given current was simply an exercise in arithmetic. We fix the number of current-starved inverters to three for compactness, and pick a fairly high current, 86μA, as we calculated that average-sized transistors will cause oscillation of approximately the correct frequency at such a current. The design of the ring oscillator past this point comes down to one major design decision: at what bias voltage should the VCO output 100 MHz? As we've fixed the current, and answering this question gives us the overdrive voltage, this will decide the sizes of all transistors in the ring oscillator. Answering this question outside of the context of the control circuit is a fairly poor idea, as one wants to assure this bias corresponds to the center of the total VCO's linear region.

The range of linearity for the VCO is determined most specifically by transistor MV1, which is affected by R2 so as to remain in the linear region for a given range of control voltages. A natural decision for the center of the control voltage range is 0.9V – mid-rail – but a bit of arithmetic, along with the fact that V_T is around 0.5V for all transistors, shows that this is on the low end of viable linear-region gate voltages. Given the center current of 86μA and fixing a resistor size of 1kΩ, we find that MV1 is in the linear region for voltages between 0.85V and 1.35V, the center of which is 1.1V. This is deemed reasonable, as it provides a large overdrive to the current starved inverter, allowing it to achieve 86μA without massive transistors. This decided the size of our current-starved

inverter transistors, as described above, and we simply matched the sizes of all NFETs and PFETs in the control block with those in the current-starved inverters.

With these sizes fixed, the schematic produced a signal of 100MHz with an input voltage of 1.106V, near exactly as desired, and was linear between 0.85V and 1.35V as calculated. The VCO gain was seen to be about 50MHz/V, a respectable number giving about 25MHz of bandwidth in the linear range.

Charge Pump

In designing the charge pump, we first decided upon a reasonable pump current. Available approximations for the feedback properties of the PLL, such as settling time and damping ratio, show that making the pump current large has a positive affect on the circuits transient response. However it is not the only such factor, and it is important not to sacrifice the use of reasonably sized transistors so as to simply maximize the current. For example, one can make transistors very wide with minimum length, but doing so introduces short channel effects, and generally makes the circuit less resilient to process, temperature and voltage changes. In analog designs in the past, we had seen that currents in the tens of microamperes could be implemented in this process without using incredibly large transistors, so long as the overdrive voltages of the source and sink transistors (MC1 and MC6, in this case) were relatively large. We chose a pump current of about 40 microamperes, as we could reasonably achieve this with $|V_{GS}|$ of 0.8 V for each of MC1 and MC6. We matched the k values of these two transistors, and biased them with independent 2-transistor biasing networks, providing 1V at the gate of MC1 and 0.8V at the gate of MC6.

An important aspect in the functionality of a charge pump beyond the pump current is that it is able to respond quickly to UP and DOWN signals from the PFD. These signals may be very short in duration, and may not reach rail voltage. As such, transistors MC2,MC3, MC4 and MC5 must act as fast digital switches. Smaller transistors are generally much better for these purposes, as they have fast switching times. The sizes of these four transistors were decided by using the minimum width and varying the lengths so as to match the k values between the four transistors. For the chosen pump current, we could not use true minimum size transistors, but we could get rather close, so the switching times of these transistors are nearly as fast as they can be for the process.

To evaluate the performance of our charge pump design in schematic, we tested first that its response to UP and DOWN pulses was relatively symmetric. When used in tandem with a loop filter (described in the following section), this means that UP and DOWN pulses of equal length should have the same effect on the charge stored on a capacitor, and thus on the control voltage to the VCO. Transistor sizes were tweaked minorly so that the magnitude of the current response to UP and DOWN pulses were approximately equal. Unsurprisingly, as the digital NMOS transistors are a bit larger than the digital PMOS transistors, their switching time was seen to be slower, yielding a slightly worse response to small DOWN pulses than to small UP pulses. As the NMOS and PMOS transistors could not be made to be the same size, and this asymmetry acts to oppose the asymmetry introduced by the PFD, this effect was merely noted for later analysis. More specific performance evaluations of the charge pump were saved until the loop filter had been designed, as described in the following section.

Loop Filter

Although the loop filter consists of only three components, its design is absolutely critical to the control system properties of the PLL. This portion of the design process was saved until last as its properties depend on both the gain of the VCO and the pump current. Our choices of component values was initially guided by Gardner's *Phaselock Techniques* and *Charge-Pump Phaselock Loops*, which provides a large amount of algebraic expressions for the approximate second-order characteristics of the PLL. In the second-order approximation, it is assumed that the smaller capacitor (C2) has no effect on the frequency response of the system. The value of C2 was set always to be a tenth of the value of C1, and this approximation was evaluated in testing, as explained in the section titled Performance.

We started with the aspects of the circuit that had already been decided on; as explained in the sections above, we knew that our VCO had a gain of 50MHz/V with a bandwidth of 25MHz, and that our pump current was 40μA. We then decided on a desired value of the damping ratio, ζ. We chose $\zeta \approx 0.9$, as this number is between the threshold for no overshoot ($\zeta = 1/\sqrt{2}$) but lower than 1, at and above which the response is no longer underdamped. The choice of this middling value was so that, should unexpected layout effects minorly affect our gain, pump current or passive component values, the response should still remain underdamped with $\zeta > 1/\sqrt{2}$. This should also, theoretically, make the transient response more resilient to changes in temperature and process. We pick a number larger than the middle of this range, as our assumption is that these unwanted changes will likely lower our pump current and VCO gain rather than raise them.

We then decided on the value of R by use of Gardner's upper bound: $R < 2\pi f_{BW}/K_0 I_P$. This equation gave an upper bound of about 78kΩ, and the value of 50kΩ was decided upon. This value was the same order of magnitude but substantially smaller than the upper bound, so that changes in the process and temperature could not push the component value above the upper bound. From here, we simply used Gardner's $C_1 = \zeta^2 8\pi/K_0 I_p R^2$ and overcompensated slightly for a capacitance of 4pF. We chose C2 to be 400fF – a tenth of C1 – as mentioned earlier. These values give a natural frequency of approximately 1.4MHz and thereby a settling time of approximately 498ns.

With the loop filter designed as such, the charge pump can better be tested. We tested how the control voltage changed with response to equal duration UP and DOWN pulses, and noted only minor asymmetries. More specifically, as expected, the control voltage was less affected by the DOWN pulse than it was by the UP pulse, for reasons explained in the Charge Pump section above. The effect of this is more easily quantified now, as the voltage difference is seen to be in the single-digit millivolts. Given the gain of the VCO and the desired frequency range of the PLL, along with the fact that this does not represent a hard deadband but rather a slower reaction, this level of asymmetry ought to be fairly negligible. The visible effects of this are further described in the section on Performance, as are the visible transient response effects of the chosen loop filter component values.

Layout

The layout procedure was first carried out in functional blocks, and then pieced together rather consciously. There were several effects seen in the layout of each functional block, as well as the

total circuit, that require some discussion.

The first and most apparent problem is to do with the available passive components, as is defined by the process being used. Unsurprisingly, the majority of the circuit area is taken up by passive components, with the largest such component being the 4pF loop filter capacitor. One unforeseen problem was that the TSMC18 process did not allow for single metal-insulator-metal capacitors (MIMCAPs) above a certain capacitance. This turned out to be fairly unimportant, as we were able to simply connect four 1pF MIMCAPs in parallel for an effective capacitance of 4pF. In laying out the charge pump and loop filter as a functional block, it was made obvious that the footprint of our entire circuit would be entirely decided by these passive components. It was decided that the transistor elements would fit in a rectangle defined as follows: the bottom of the circuit would be bounded by the 4pF capacitor block, the left would be bounded by the 50kΩ resistor, and the right would be bounded by the 400fF capacitor. We wound up being able to meet this goal, yielding one of the most compact possible theoretical circuit areas given these passive values.

Given our past experience with analog circuit layouts, we were quite wary of well-proximity effects when laying out this circuit. These are effects that change the threshold voltages of NFET transistors based on their orientations and physical proximity to NWELLs, which can have massive negative effect on transistors with carefully chosen or necessarily matched bias voltages. While in an operational amplifier, threshold voltage swings in the range of 50mV can have large effects on gain stages, the PLL is much more resilient to such changes. For one, much of the PLL is digital, meaning that the gate voltage of many of the transistors should always be at one of ground or rail. The behavior of such digital transistors is not majorly affected by threshold voltage changes, and as such, digital blocks are laid out for compactness rather than well-proximity effect mitigation. Care was taken so that the VCO and charge pump's analog-behaving transistors were symmetric with respect to NWELL. Some cases arised in which doing so without compromising compactness was challenging, and the use of dummy transistors was considered. However, in such cases, testing was done on the schematic level to determine sensitivity to threshold voltage changes of $\pm$50mV, and in all cases the operation of the circuit was left unchanged. Much of this is due to the natural stability induced by the feedback loop, but it also can also simply be due to the fact that many of these functional blocks are resilient in their own right; the design was such that, for example, a small change in pump current would not compromise the PLL's overall behavior, etc.

The behavior of the PFD and charge pump functional blocks were seen to be near entirely unchanged by the layout process on their own, however the VCO did present some rather large shifts in performance. More specifically, the gain curve was shifted upward in voltage by approximately 100mV, leaving the gain fixed but changing the control voltage corresponding to the center frequency. As described in the section on the VCO, this block was designed so that a control voltage of 1.1 V would correspond to an output frequency of 100MHz. In layout, this value shifts up to 1.2V, but this has no critical affect on the operation of the circuit under standard conditions. That is to say that the entire linear region of the gain curve still lies in the range of 0V to 1.8V. Unsurprisingly, this does have negative effects with respect to both temperature and voltage, as described in the section on Performance. The origin of this fairly large control voltage change was isolated through several tests on the controlling currents, and it was determined to be some combination of NWELL effects and parasitic effects within the ring oscillator. Parasitics were minimized via use of only the metal 1 layer where possible.

In layout we also decided to implement two power rails so as to decrease noise at VDD. The circuit's spectral content was observed in two states: one layout in which there was a single VDD and a second in which the power rail of the charge pump was separated from that of the PFD and VCO. A 20-30 dB drop in the noise floor was seen when the rails were separated, so we decided upon the design with separate power rails.

As a final, general note, the MIMCAPs use metal layers 5 and 6. As such, all 6 metal layers must be used in the final PLL layout. Despite this, for the minimization of parasitics, as well as good practice, only metal layers 1 and 2 were used in the circuit outside of connections to the capacitors.

Performance

The majority of simulations were performed at 25°C with a supply voltage of 1.8V, using the TT process. The input signal is a square wave generated by the Cadence analoglib "pulse" vsource, with rise and fall times of 1fs. The signal outputs from many of these simulations can be seen in Appendix II. While the performance of individual blocks of circuitry have been explained in the previous sections, we will now discuss their effects on the performance of the PLL as a whole. The most notable of these effects is that of the shifted range of the VCO. While the loop still locks with excellent spectral purity at 100MHz under these standard conditions, its upper bandwidth has been reduced in layout relative to the expected schematic behavior. More quantitatively, the loop locks still at 110MHz, but it begins to exhibit nonlinear behavior at frequencies higher than this. This also causes the PLL to not lock at much lower supply voltages, despite its ability to do so in schematic; it can, however, handle supply voltage drops of at least 100 mV.

The step response of our circuit is evaluated based on its settling time and overshoot. In the section on the Loop Filter, we described that we designed this subcircuit to have no overshoot and a settling time of approximately 500 ns. We also described that this design process did not directly take into account the second, smaller capacitor. The settling time of a step response from 100MHz to 111.111MHz was seen to be 516 ns; a value extremely close to that given by our calculations. The overshoot was seen to be minimal, although non-zero, as is to be expected given that the damping ratio is merely an application of the second-order approximation.

Looking at both the UP and DOWN pulses in lock, as well as the spectrum of the PLL output, we are able to evaluate the overall asymmetries within our design. As noted, the PFD itself has a larger deadband for UP pulses than for DOWN pulses, which should contribute to the jitter of the PLL. Specifically, the noise floor at lower frequencies should be higher than that at higher frequencies due to this effect alone. However, the charge pump introduces a second asymmetry, which ought to affect the jitter in the opposite way (as described in the Charge Pump section). In observing the UP and DOWN pulses in lock, the DOWN pulses were seen to be considerably larger in amplitude. This corresponds to the smaller relative response time for DOWN pulses at the input of the charge pump, effectively increasing the DOWN pulse deadband. This, along with the small current asymmetry in the charge pump, ought to counteract the spectral asymmetry introduced by the PFD. In observing the spectrum of the output in the range 90MHz to 110MHz, the noise floors were seen to be nearly symmetric, with that on the lower-frequency end being slightly larger

with that on the higher frequency end. As the deadband was small to begin with (merely 35 ps at maximum), the noise floor is still below -110 dB, however this slight asymmetry shows that the PFD deadband effects are stronger than the asymmetries introduced by the Charge Pump.

One other feasible source of error could be leakage at the output of the charge pump. Should this have proven to be an issue, it would have informed a change in the transistor sizes in this block. However the control voltage was seen to be incredibly stable in lock, and this error was not taken into consideration. The phase error in our PLL was seen to manifest in rising edge differences of about 30ps, when measured from midrail. This value corresponds near directly to the PFD deadband, meaning that the PFD error dominates the error of the entire PLL.

In addition, the extracted PLL was tested under 9 conditions of voltage, temperature, and process. Performance was measured in terms of frequency-stepped response settling time and the Spurious-Free Dynamic Range (SFDR) as seen in the spectra of the feedback/output signal. The settling time listed is the time after which the control voltage remains within a 2% window of its final value. Two sets of conditions at the low supply voltage give loops that do not lock at 100 MHz. Table 1 summarizes the results of these simulations.

Condition	Voltage: V	Temperature (°C)	Process	Settling Time: μs	SFDR (dB)
1	1.80	25	TT	.516	114.64
2	1.80	0	SS	.260	58.82
3	1.80	70	FF	2.30	60.14
4	1.62	0	SS	>2.5	98.00
5	1.62	25	TT	n/a	n/a
6	1.62	70	FF	n/a	n/a
7	1.98	0	SS	.277	62.95
8	1.98	25	TT	.208	60.65
9	1.98	70	FF	.230	96.01

Table 1: PLL Characteristics under Voltage, Temperature, and Process variations

The circuit's behavior with respect to changes in supply voltage, temperature and process are unsurprising. With respect to supply voltage, recall that layout effects push the VCO 100MHz control voltage significantly higher. This leads to downward shifts in supply voltage pushing 100MHz outside of the bandwidth of the VCO under some temperature conditions. As downward shifts in temperature have the opposite affect on control voltage, it is to be expected that the 0° simulation still locked at 100 MHz, even with the lower supply voltage. Note that the circuit under these conditions still locked at lower frequencies, as is to be expected. Upward shifts in supply voltage push the control voltage corresponding to 100MHz down, and as noted earlier, layout effects on the VCO give the control voltage plenty of downward room to move. Thus, at this voltage, 100MHz lies closer to the center of the control voltage range. This makes the circuit more resilient to temperature and process changes.

The best performance in terms of SFDR is seen at the standard testing conditions. As the circuit was designed with these conditions in mind, and with spectral purity as the main consideration, this is expected. However, a noise floor below 50 dB is seen for all conditions under which the

loop locks, attesting to the spectral resilience of our design. The spectral width of the main lobe in all locking conditions was also seen to be rather small. The settling time, on the other hand, is seen to improve in many of the alternative simulation conditions. It is challenging to cite a single reason which can justify this improvement, but changes in process can alter passive component values, having fairly substantial effects on the control system parameters of the circuit. Similarly, temperature can have the same effect, while also altering the mobility of carriers, and thus the current magnitudes within the circuit. Higher voltages also clearly increase current magnitudes, increasing the pump current, which ought to directly decrease the settling time (as seen). In most cases, a settling time below 1μs was achieved.

Conclusion

The designed PLL under standard conditions boasts excellent spectral purity in lock, as well as a settling time lower than a single microsecond. Jitter is seen to be determined mostly by the PFD deadband. The most notable flaw in the final PLL is the VCO gain curve, which has been moved from its desired schematic position by layout effects. However, this flaw does not detriment to PLL performance at 100MHz under standard conditions. The PLL's resilience to process, temperature, and supply voltages are very heavily affected by this shift in the VCO gain curve. This noted, the PLL still functions with reasonable spectral purity and settling time at several alternate combinations of process, temperature and supply voltage. The overall design of the PLL, outside of this VCO anomaly, matched well in theory to the final extracted circuit.

Appendix I – Schematics

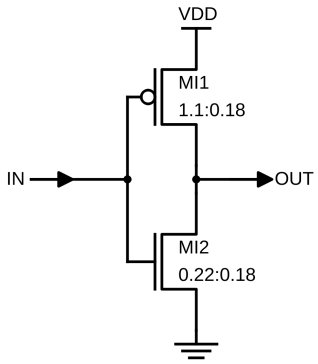

Figure 1: NOT Gate

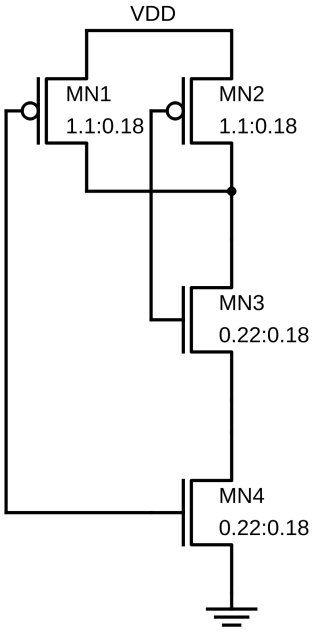

Figure 2: 2-Input NAND Gate

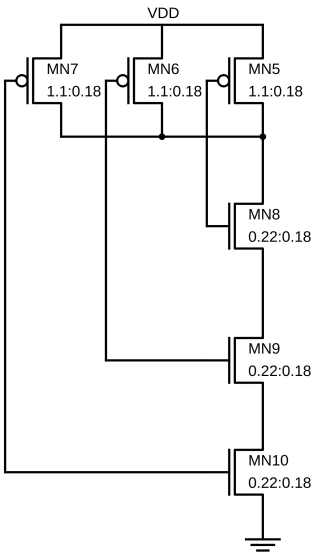

Figure 3: 3-Input NAND Gate

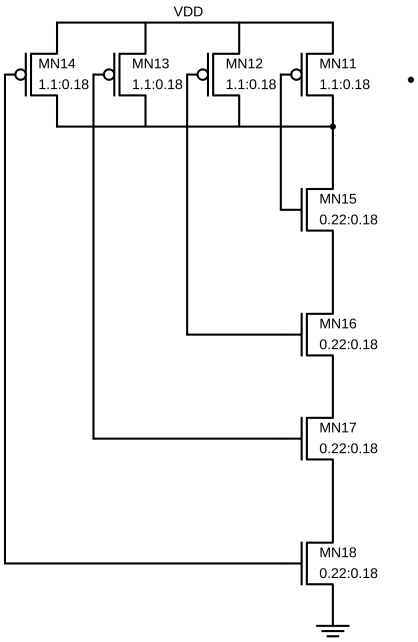

Figure 4: 4-Input NAND Gate

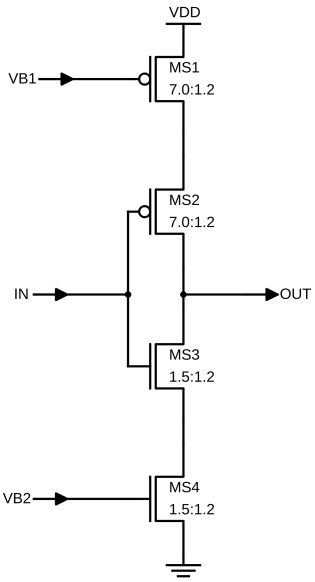

Figure 5: Current-Starved Inverter

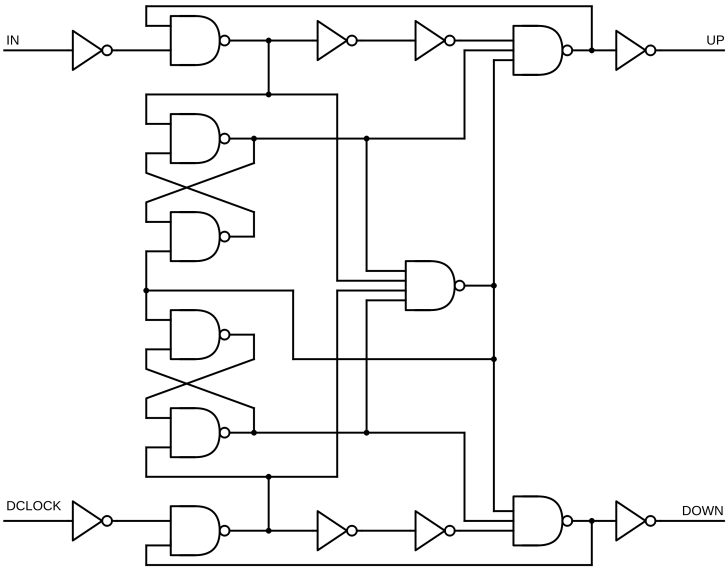

Figure 6: PFD Schematic

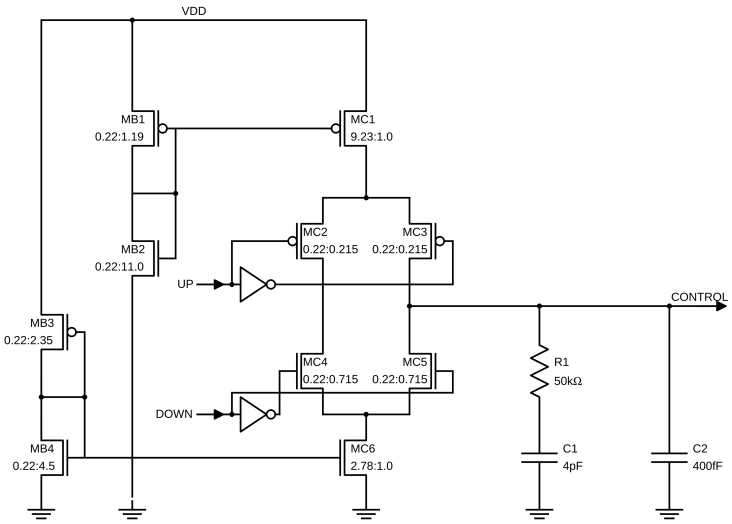

Figure 7: Charge Pump and Loop Filter Schematic

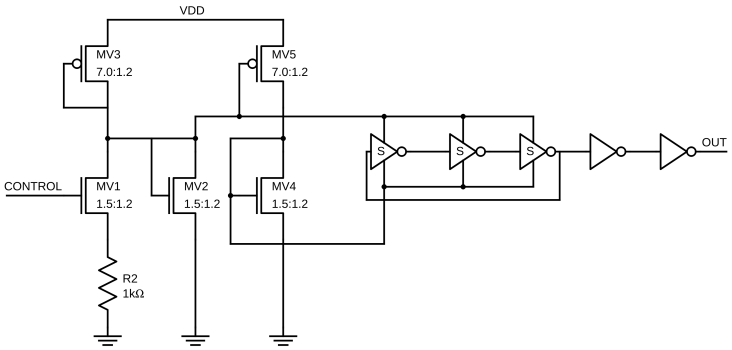

Figure 8: VCO Schematic

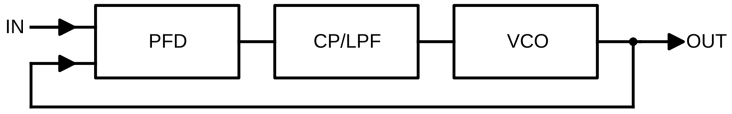

Figure 9: Total PLL

Appendix II – Plots

The first 5 pages of plots in this section are all created from simulation of the extracted PLL at 25°, 1.8V supply voltage and using the TT process. Beyond these plots, several spectral outputs were plotted in the range of 97MHz to 103 MHz under varying conditions. These plots were made using the Fourier object from Cadence's analoglib, with fundamental frequency 50kHz and 2060 harmonics. The following step responses are due to input jumps, where the simulation is allowed 2.5μs to lock at 100Mhz, after which point the input frequency is instantaneously changed to 111.111MHz. The step responses and spectra for all but one of the 1.62V supply voltage tests are omitted, as the loop generally did not lock at this voltage. In the 0° case at 1.62V, the loop locks at 100MHz but not at 111.111MHz, and so the step response plot is omitted here.

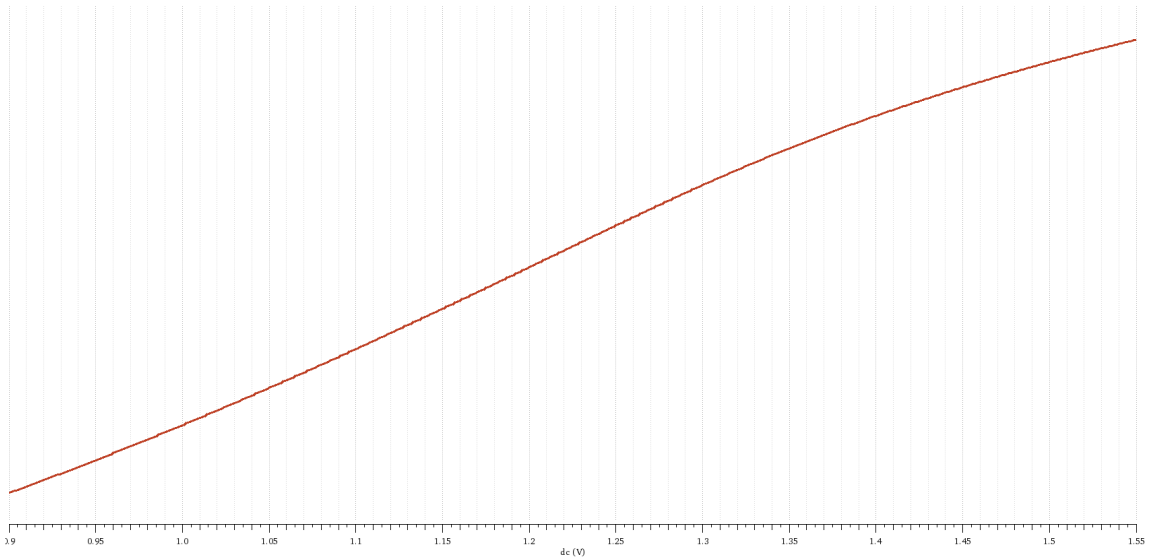

Figure 10: VCO Gain Curve in Layout from 80MHz to 120MHz

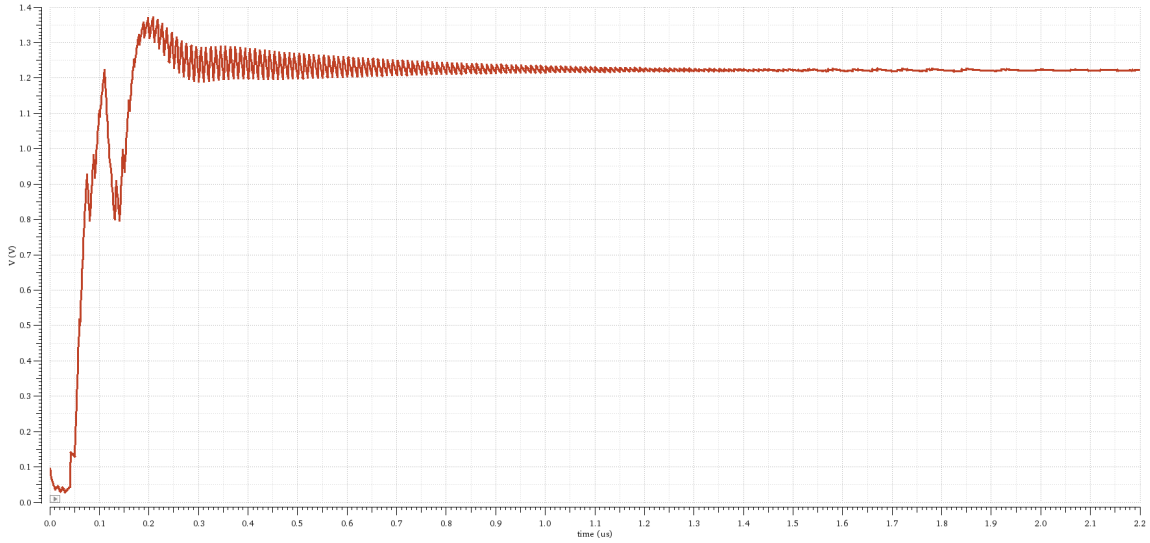

Figure 11: Control Voltage from Startup to Lock

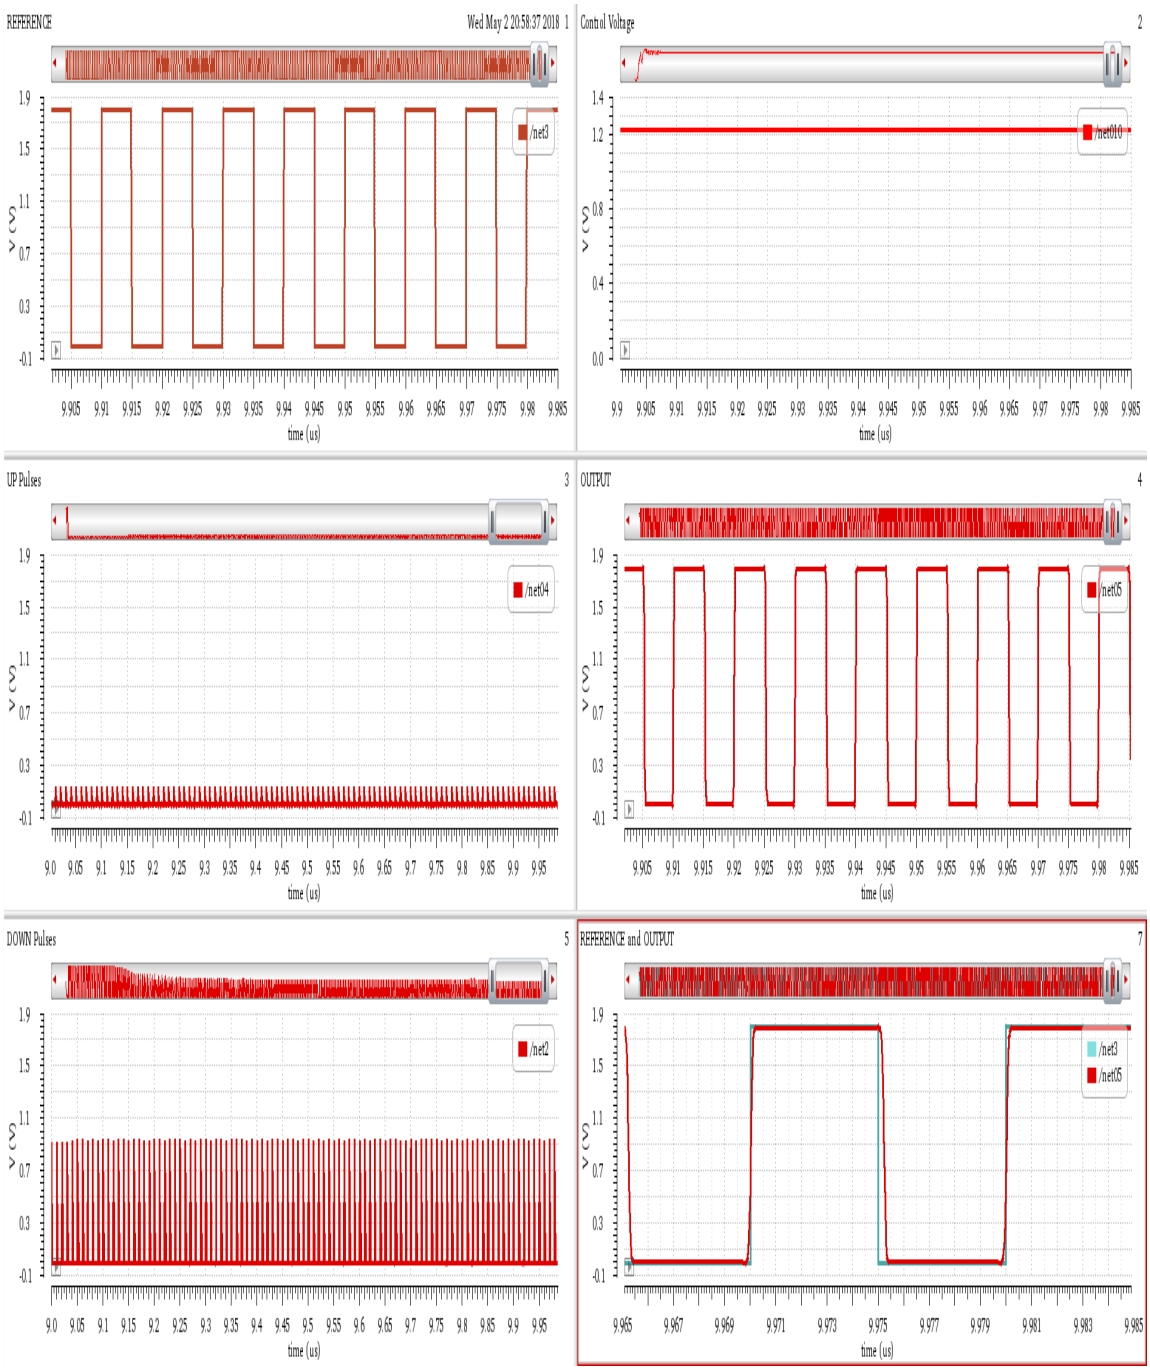

Figure 12: Salient Signals while Locked; From Top-Left to Bottom-Right: Reference, Control, UP, Output, DOWN, Reference and Output Overlayed

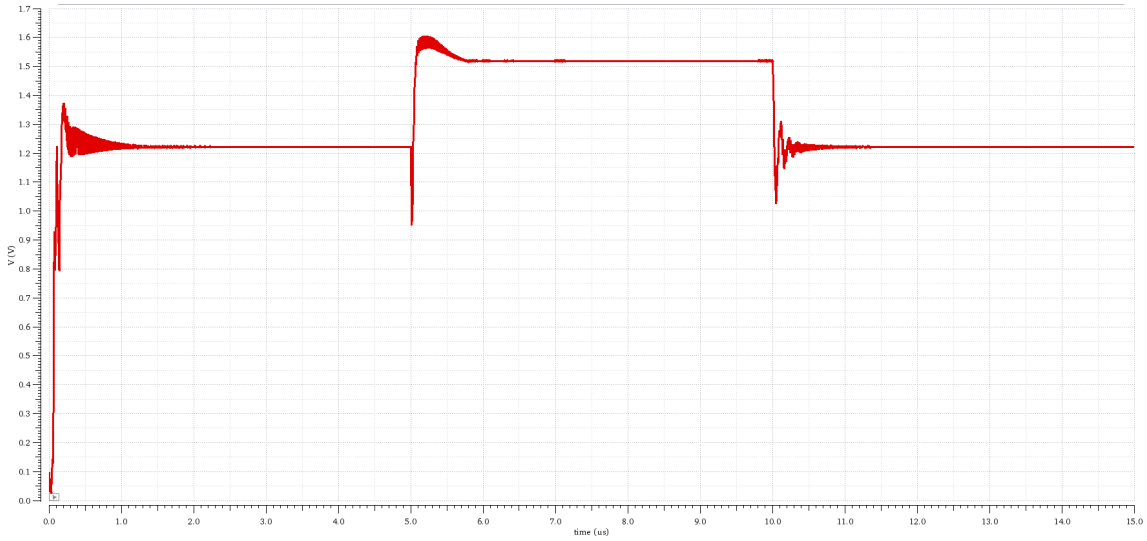

Figure 13: Frequency Jump from 100MHz to 111.111MHz and Back

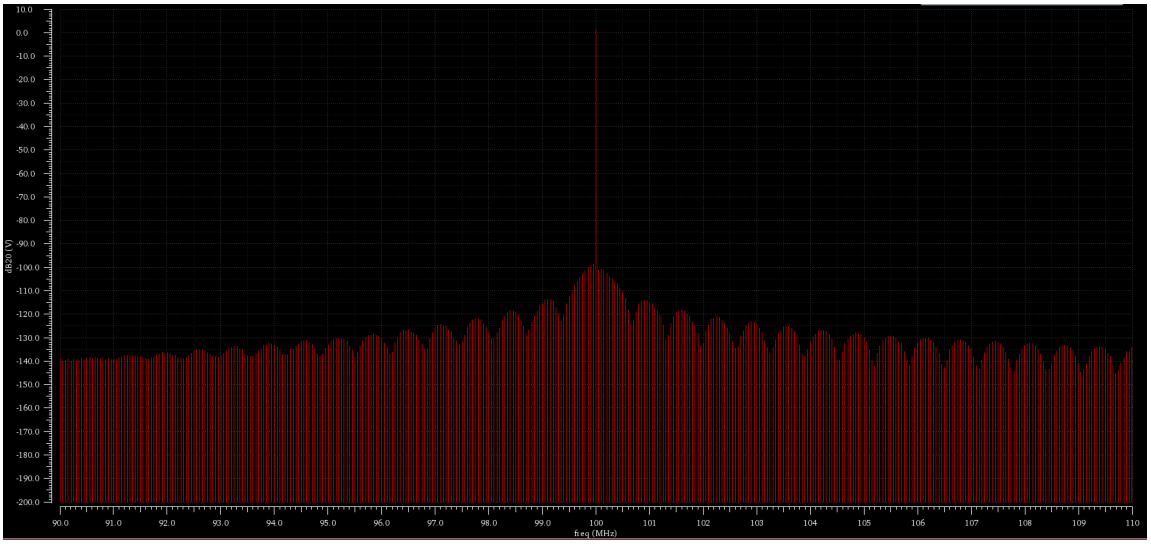

Figure 14: Spectral Content of Output in Lock

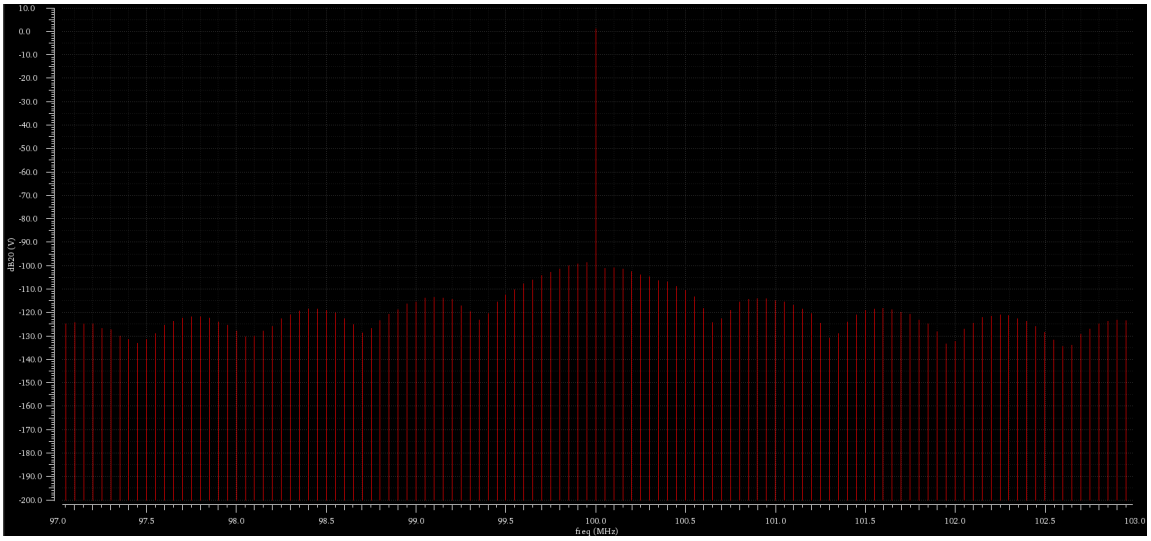

Figure 15: Spectrum at 1.80V, 25°, TT Process

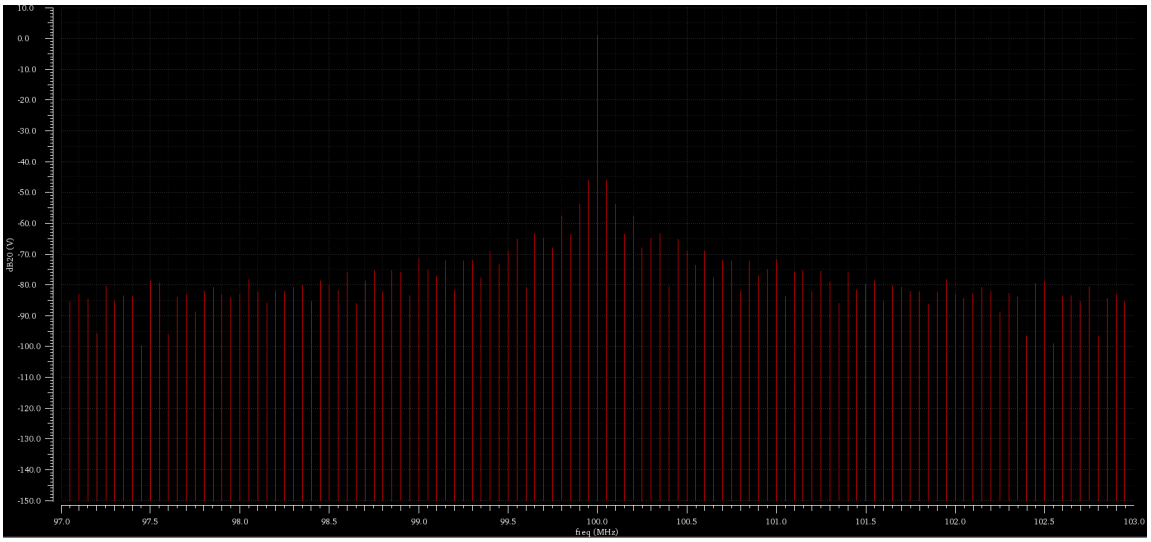

Figure 16: Spectrum at 1.80V, 0°, SS Process

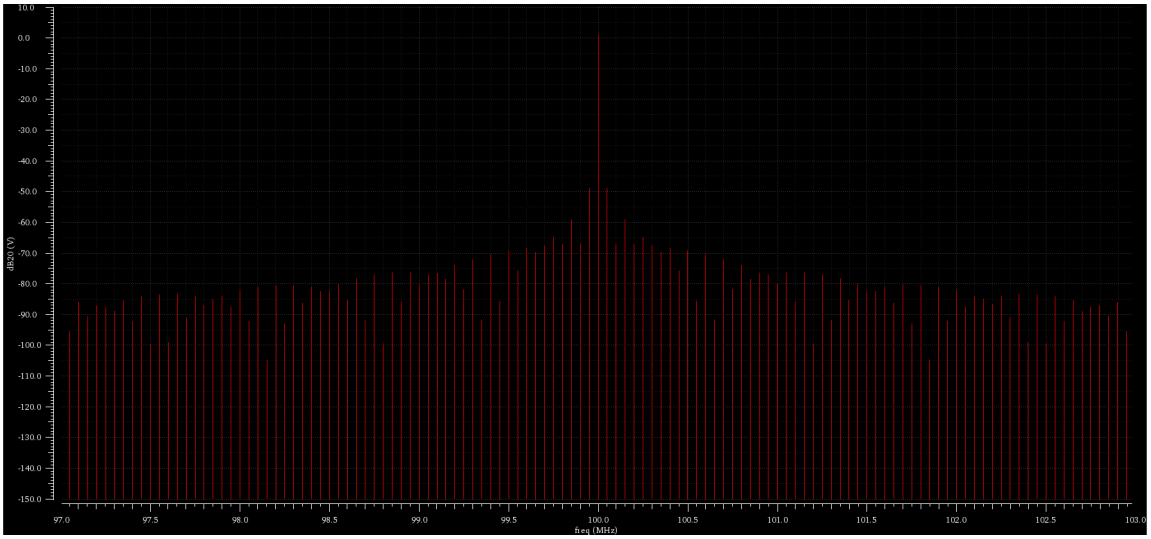

Figure 17: Spectrum at 1.80V, 70°, FF Process

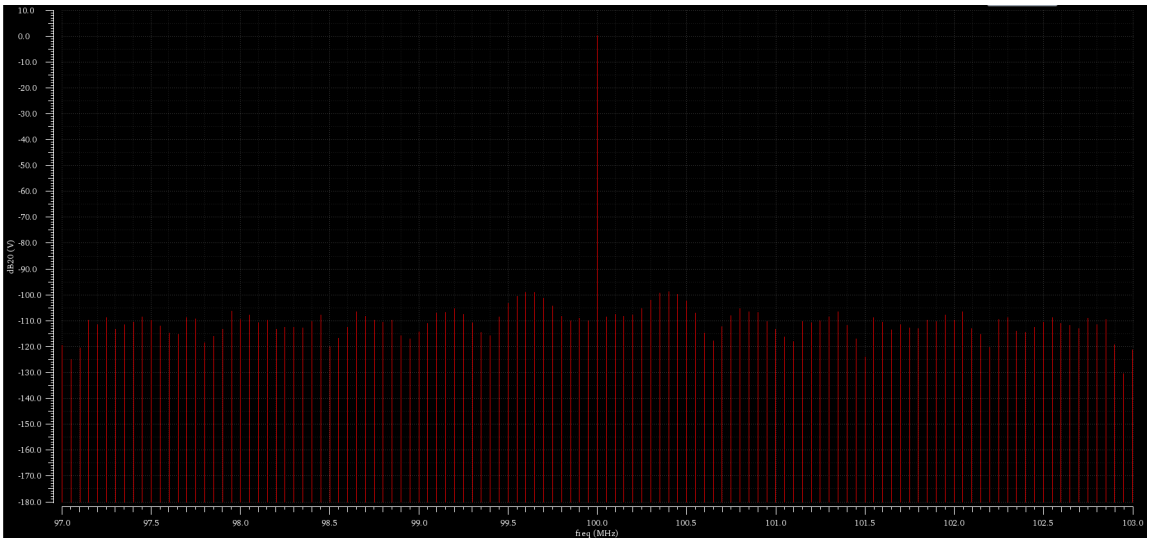

Figure 18: Spectrum at 1.62V, 0°, SS Process

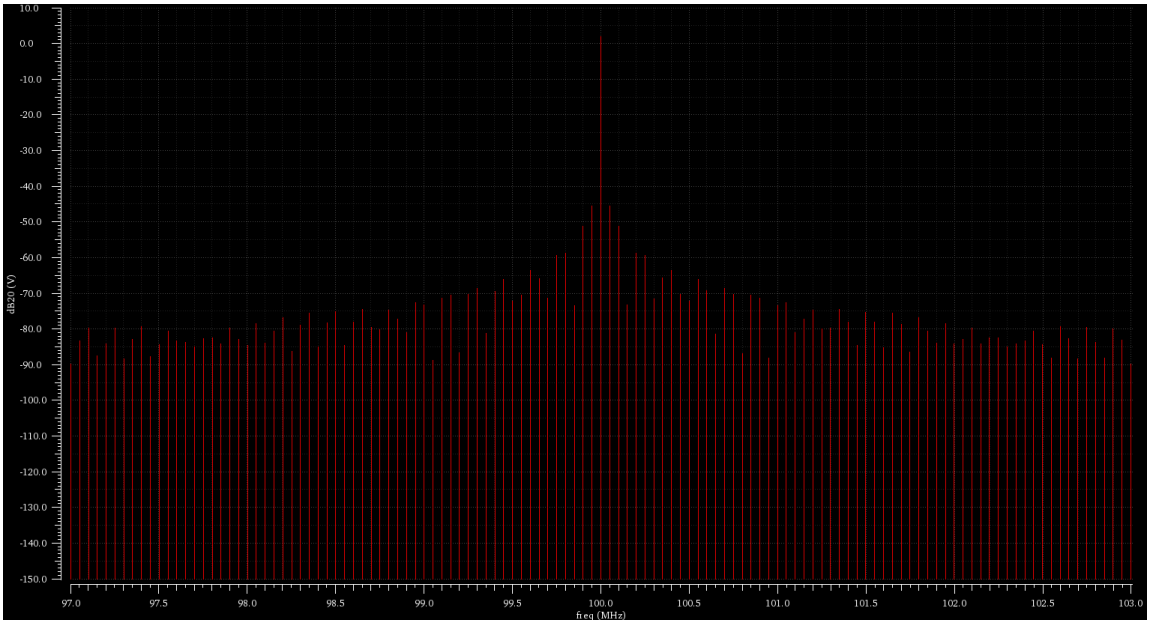

Figure 19: Spectrum at 1.98V, 25°, TT Process

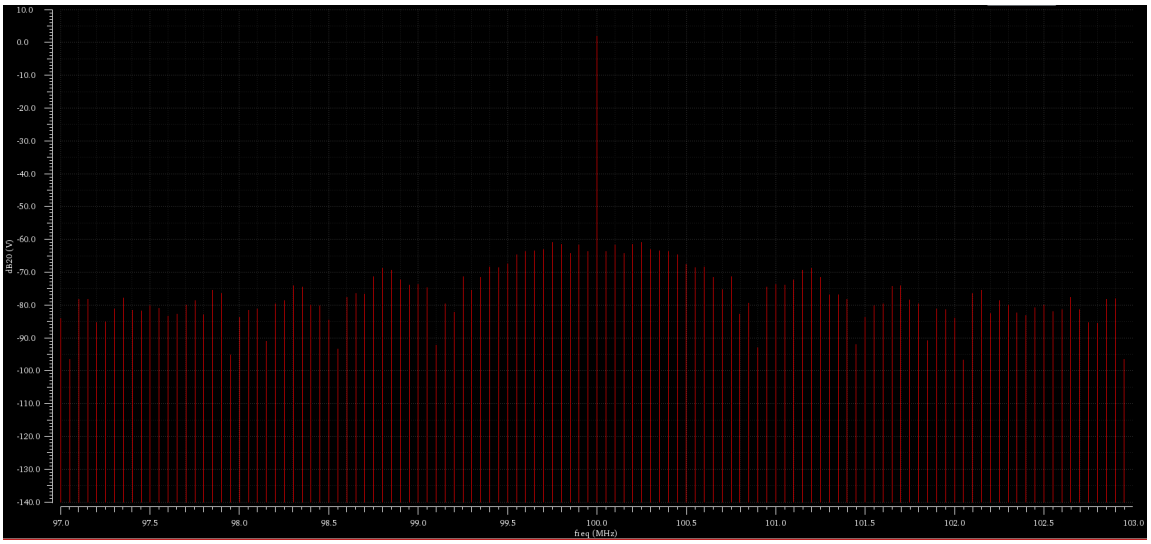

Figure 20: Spectrum at 1.98V, 0°, SS Process

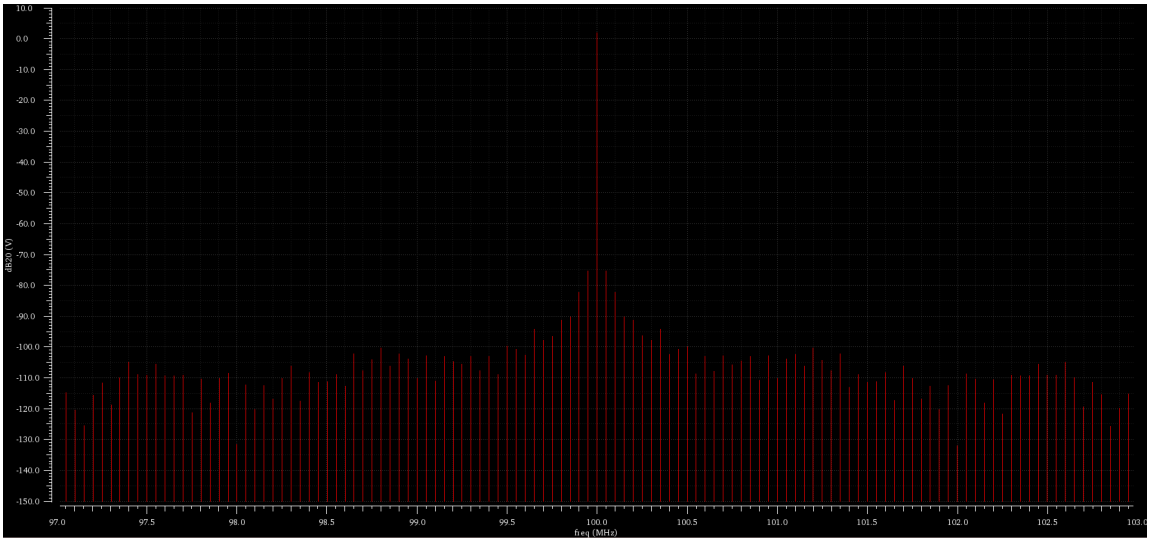

Figure 21: Spectrum at 1.98V, 70°, FF Process

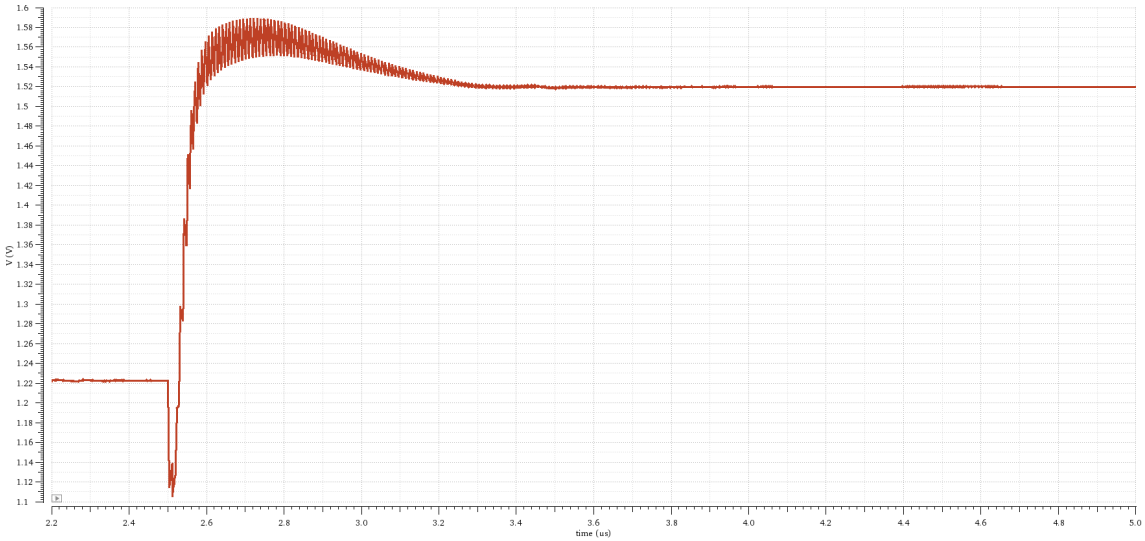

Figure 22: Step Response at 1.80V, 25°, TT Process

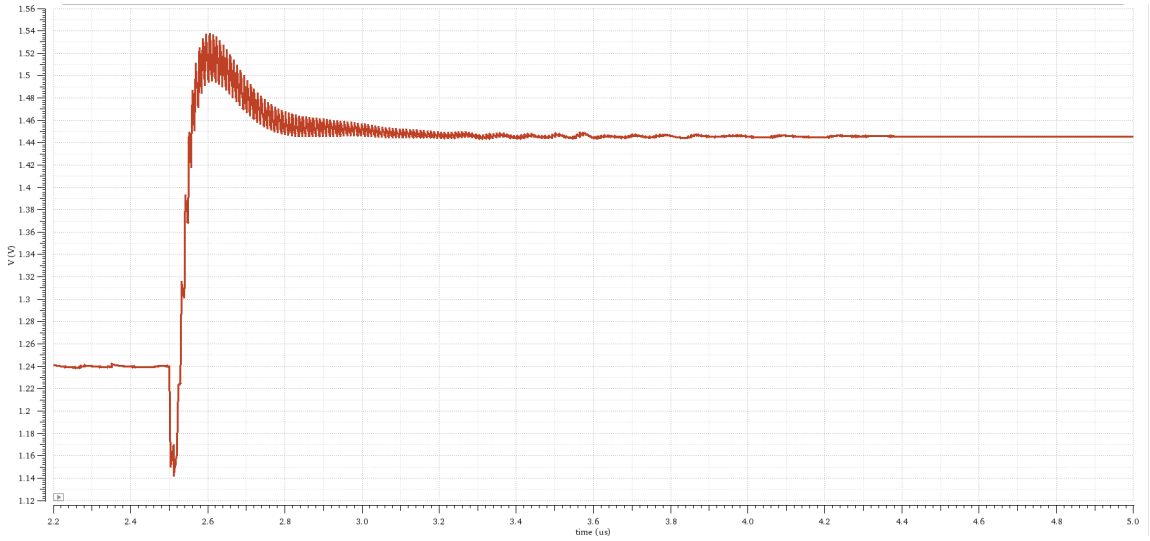

Figure 23: Step Response at 1.80V, 0°, SS Process

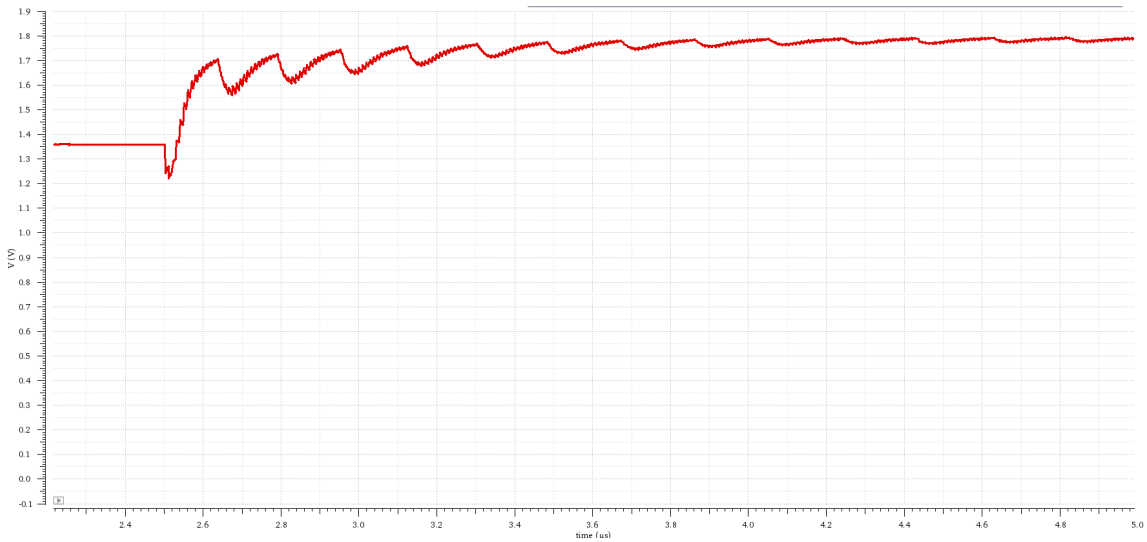

Figure 24: Step Response at 1.80V, 70°, FF Process

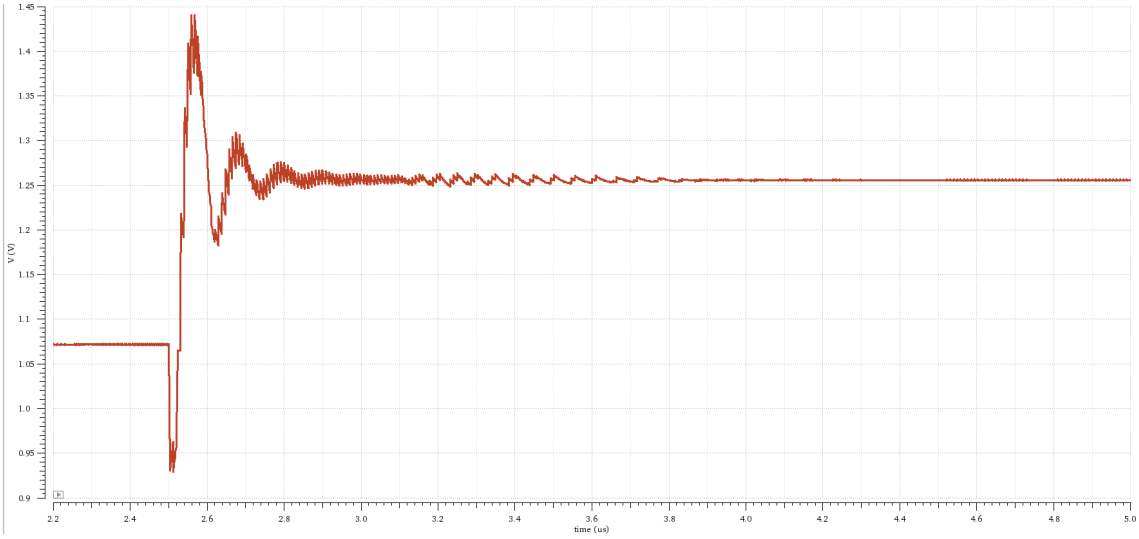

Figure 25: Step Response at 1.98V, 25°, TT Process

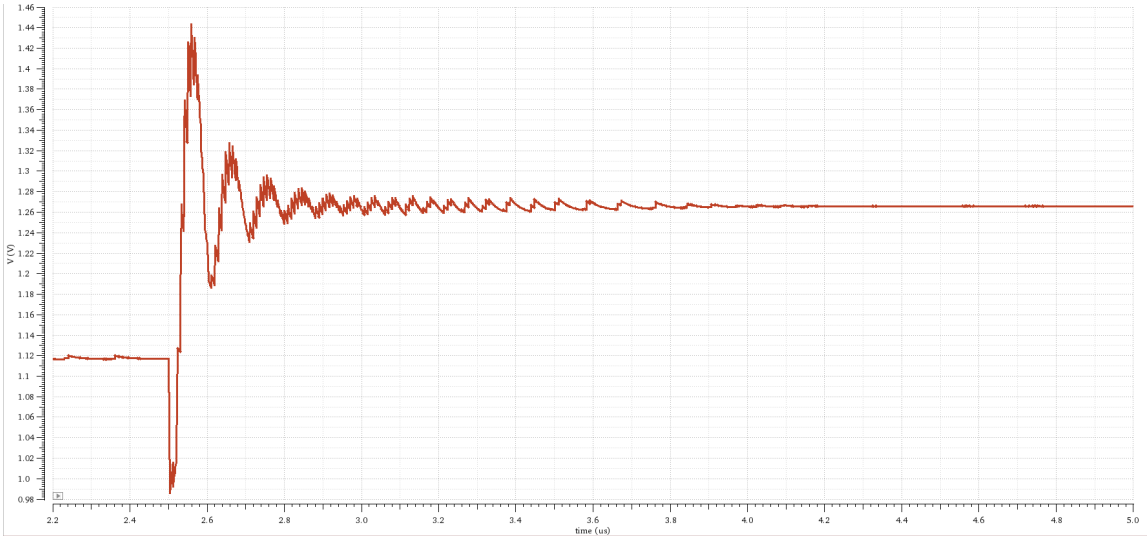

Figure 26: Step Response at 1.98V, 0°, SS Process

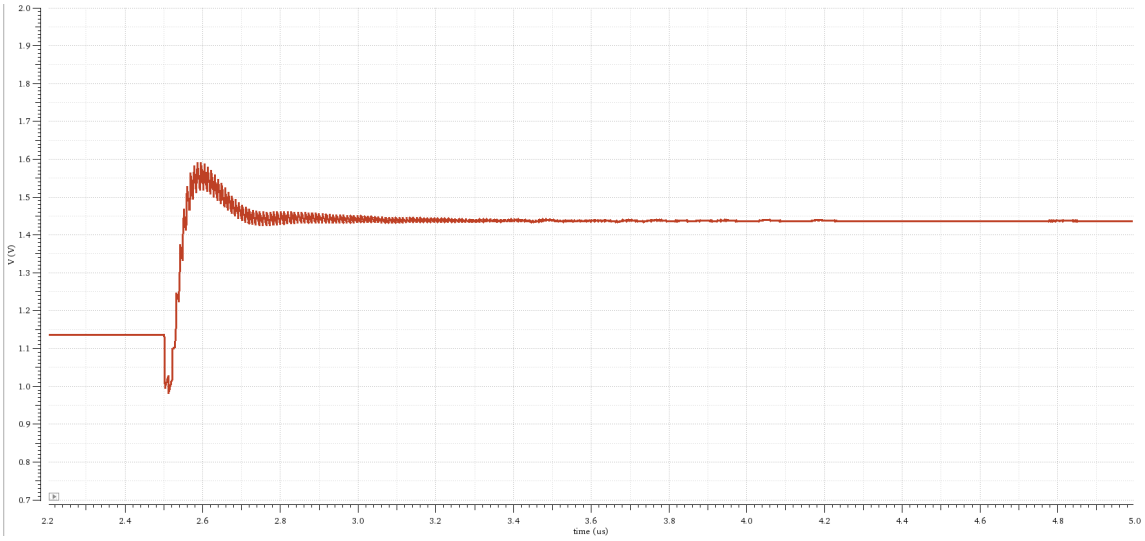

Figure 27: Step Response at 1.98V, 70°, FF Process

Appendix III – Layout Images

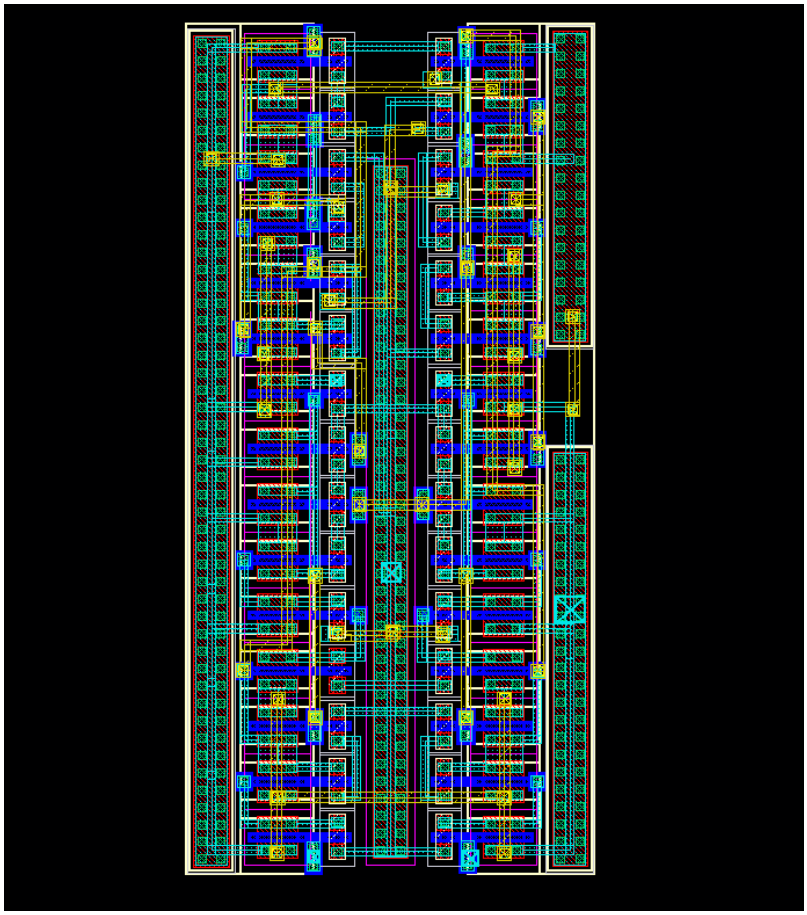

Figure 28: PFD Layout

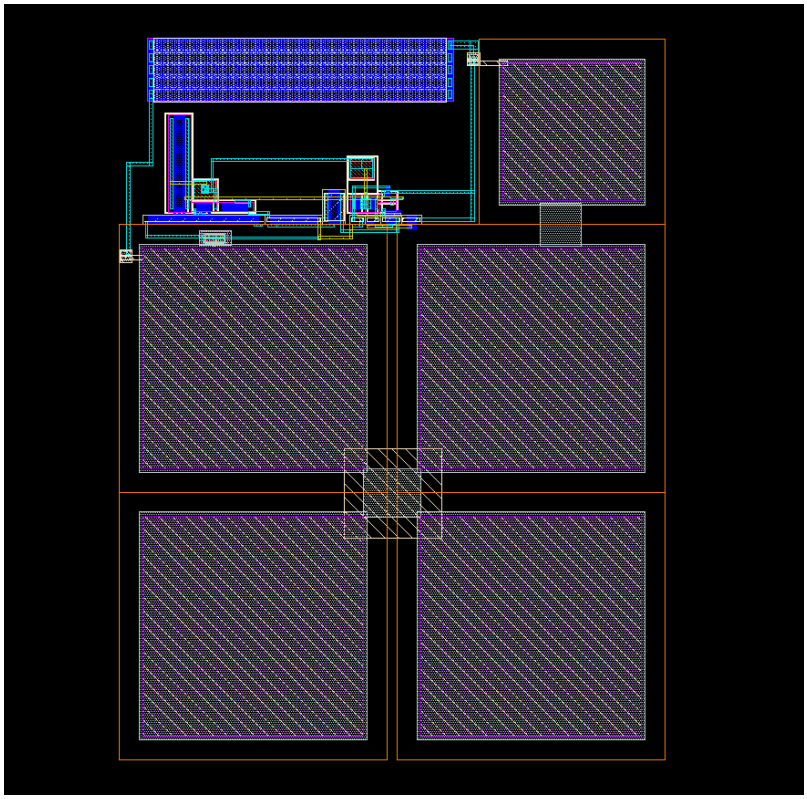

Figure 29: Charge Pump/Loop Filter Layout Layout

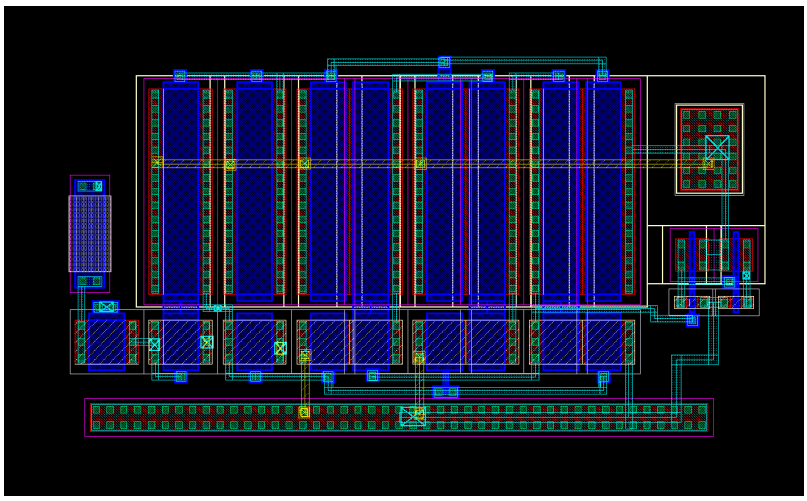

Figure 30: VCO Layout

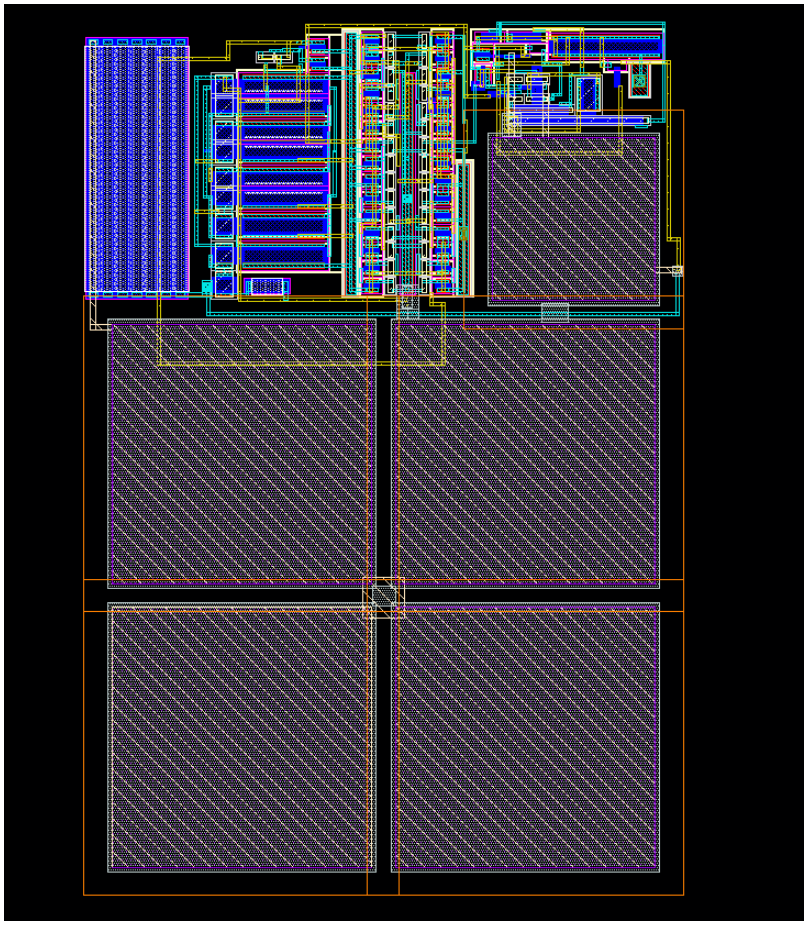

Figure 31: Complete PLL Layout