On the Engineering of an Inexpensive, Portable X-Ray Machine

A REALLY LONG PAPER BY ADAM MUNICH

2/27/2011

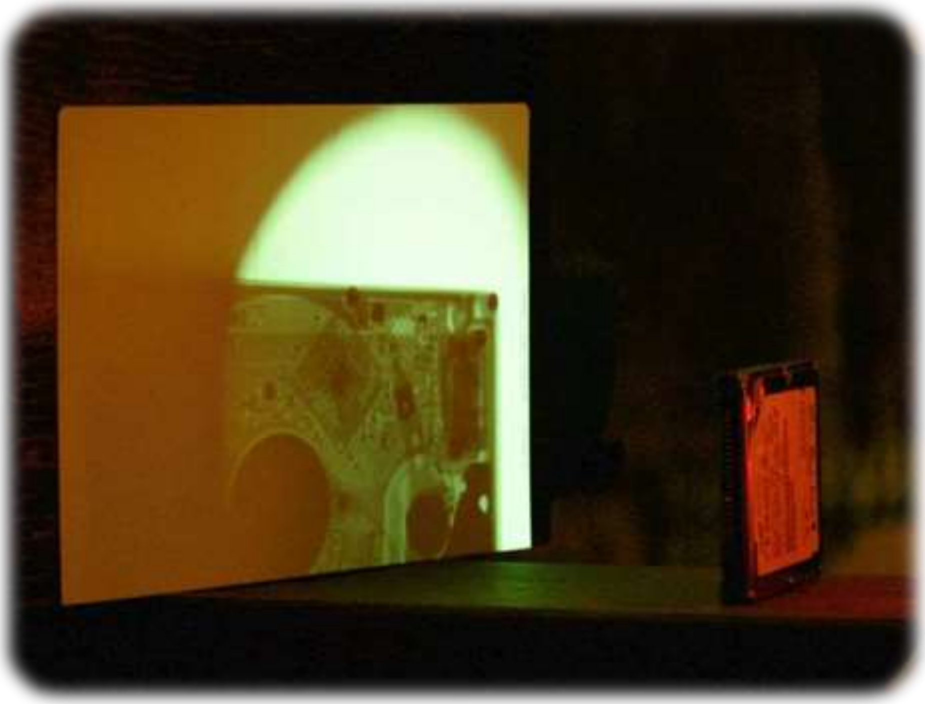

Part I - Foreword

Introduction

Throughout certain parts of the world, access to electricity can be scarce, or in some areas simply inconsistent. To some such a problem would be viewed as just an inconvenience, but to those who have suffered physical trauma a wait of hours or possibly days to get a diagnosis can prove to be debilitating, perhaps even fatal.

Often in these parts where electricity is scarce money is scarce as well. According to a 1998 survey by World Bank Development, 50 percent of the world lives on a yearly income of less than $850. Adjusted for the average 3% inflation this figure would equate

to roughly $1000 in today's economy. When a typical price for even a low-budget, used x-ray machine is $6,000, it becomes quite obvious that much of the underdeveloped world could never afford such an amenity even if they could reliably supply power to the device.

Inspiration

This project was inspired after IRC chat with a Pakistani man who complained of rolling blackouts in his country, and a man in Mexico who complained of the time it took to get a Radiograph of his foot after a motocross accident. How I manage to meet these people will remain one of life's greatest mysteries, yet during the remainder of the chat I pondered both thoughts. Compared to Mexico, Pakistan is medically, a well-equipped country. Mexico on the other hand has a fairly reliable power grid, but in certain parts lacks expensive medical supplies.

But what can a doctor do in countries which lack both amenities?

What is an injured soldier supposed to do when they are stationed hours away from a base camp?

Following the thought up with a Google search, there appeared to be only one battery powered x-ray machine on the market; a power-drill style machine dubbed the "Nomad Pro". That said, the $8,000 unit is not rugged, nor is able to charge without use of a proprietary charging system which accepts line power only. This device, while novel, fails to solve both problems!

After some more research it was concluded that such a machine, a portable, inexpensive x-ray generator; something which could help the diagnosis and treatment of millions of people around the world does not exist. I found it odd that, despite the dramatically falling costs of every component in the electronics industry, the vast majority of x-ray generators still remain expensive, non-portable items. With an annual market expenditure of 12 billion dollars there certainly was demand enough for x-ray equipment.

Being an inventor, it was felt that I had the duty to solve this issue. The task at hand was simple in theory: to build a $400 portable x-ray machine. In practice however, it proved to be quite a complicated task which tested skills in everything from electrical, physical, mechanical and to a certain extent, nuclear engineering.

Part II – Designing the Radiating Head

Choosing a Tube

The first step required of course, was to obtain a suitable Coolidge tube. While the prices of these tubes are often kept purposefully inflated by manufactures, lately a flood of tubes from china has reduced the average asking price by nearly fivefold.

Ideally the tube would be of the stationary anode variety, have a tungsten target and be rather small physically. After about one day of searching through the various manufacturers' websites, an ideal tube was found on the website of the Hangzhou Wandong Electron Co. Their model XD-12 not only fulfilled all of these requirements but was priced at a reasonable $75.

Before we delve further into the design and assembly though, it would be wise to first explain how such an x-ray tube operates.

Figure 2 – Visual detail of the same tube.

Figure 1 -- Dimensions of the XD-12 Coolidge tube.

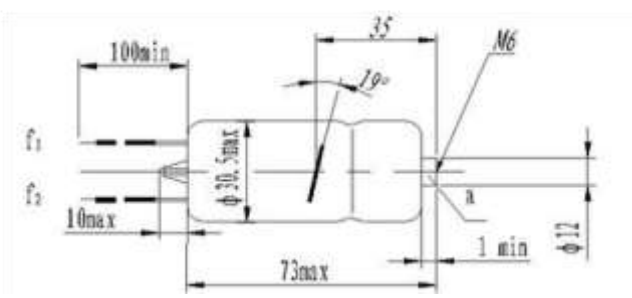

Coolidge Tube Operation

Essentially, a Coolidge tube is a thermionic diode optimized for both high voltages and high powers. Like a thermionic diode, all elements are contained in a glass envelope which has been evacuated to the hardest vacuum reasonable.

A Coolidge tube's method of operation deviates not too far from that of a diode. The heater is given a bit of current to warm it to incandescence, where the now-hot tungsten cathode boils off a cloud of electrons while simultaneously focusing them into a beam. These electrons are then attracted to the positively biased anode and move towards it at a very high speed.

Upon arrival at the anode, the high energy electrons lose energy through collisions with the metal atoms. Most of these electrons will do little more than heat the anode, but about 2% will generate x-rays in a process called *bremsstrahlung.*

Figure 3 – Operation of a Coolidge tube

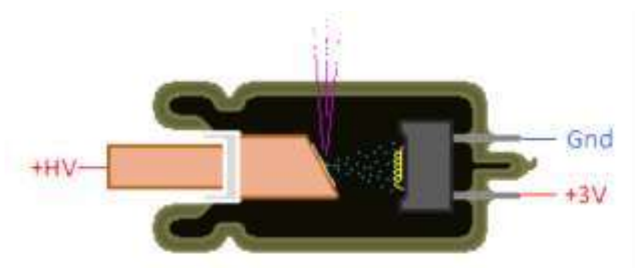

Bremsstrahlung

Literally translating to 'braking radiation', *bremsstrahlung* is the process where a high speed electron 'brakes and sling-shots' around an atom's nucleus, dumping its kinetic energy on one photon. An imparting electron might have an energy in excess of 60keV so some very energetic photons are be made; photons which fling off into space and become the xrays we all know and love.

What determines the energy of the x-rays produced is the amount of voltage present on the anode. It's quite simple actually; more voltage means more electron attraction, and more attraction results in a faster electron beam. Faster electrons are able to make higher energy photons, and thus "harder", higher energy x-rays would be generated.

Bremsstrahlung is a continuous spectrum of radiation akin to a "white light" source. Since most electrons graze a few atoms before they have the chance to sling-shot, they often lose some energy before they make any x-radiation. A whole range of xray energies is thus produced.

The maximum energy that an x-ray can have is limited to the energy of electron producing it, itself directly proportional to the voltage applied on the tube's anode. Often this energy is measured as kilovolts peak, or kVp. In reality the majority of the xrays produced are low energy, 'soft' x-rays, but these are greatly attenuated by the tube's glass wall.

Characteristic Production

Characteristic or k-line production is the second mode in which an electron might produce an x-ray. In this method, electrons knock others out of an atom's lowermost shell and leave a hole which must be filled. This unstable arrangement is then promptly made stable by electrons from higher shells who jump down to fill the hole, emitting an x-ray photon during the journey.

Tungsten k-shell electrons have a binding energy of 69.5keV, so in order to kick these out your impacting electrons must have energies greater than 69.5keV. Typically one would need to give the anode a bit more than 72kV to accomplish this, hence the standard 75kV x-ray tube.

After a k-shell electron gets the boot, its hole will immediately be filled by an electron from tungsten's l-spectrum shell; binding energy 10.2 keV. The difference between these two energy states; 69.5 keV and 10.2 keV gives us the characteristic tungsten x-ray energy of 59.3 keV. A molybdenum anode would produce two peaks, one at 19.7keV and the other at 17.6keV.

Interestingly, this process can be used to identify elements based on their k-lines. By bombarding a sample with electrons and measuring the output spectra, an x-ray florescence analyzer can determine what elements a compound contains. Oh how I wish I had one of those...

Figure 4 – *Bremsstrahlung* and characteristic production

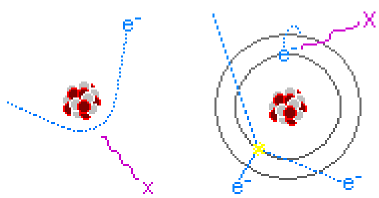

Figure 5 – An x-ray tube's emission

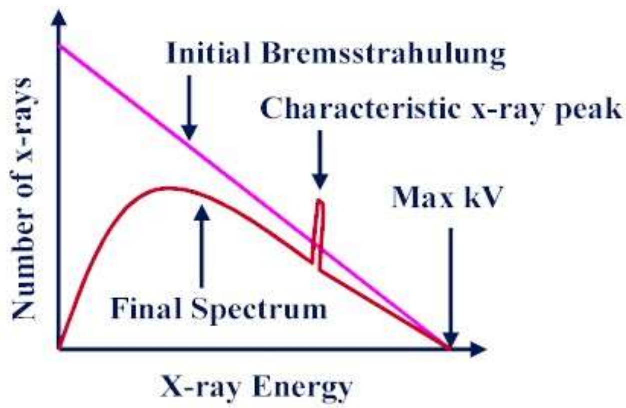

Anode Current

While anode voltage controls the 'hardness' of the x-rays, the voltage applied to the tube's heater gives independent control of the total x-ray flux. A hotter cathode boils off more electrons, reducing the impedance and allowing for a larger current to flow through the tub. One must be careful though, as too much power will damage the tube's anode via excessive heating. Ideally, a medical x-ray generator would provide a short, high intensity burst of radiation to reduce the 'shutter speed' and overall dose absorbed by the patient.

Conveniently, the chosen XD-12 tube comes with some graphs to help one set parameters for the design of a machine: figure 6 displays the relation between heater voltage and anode current.

Figure 6 – Heater voltage v Anode current

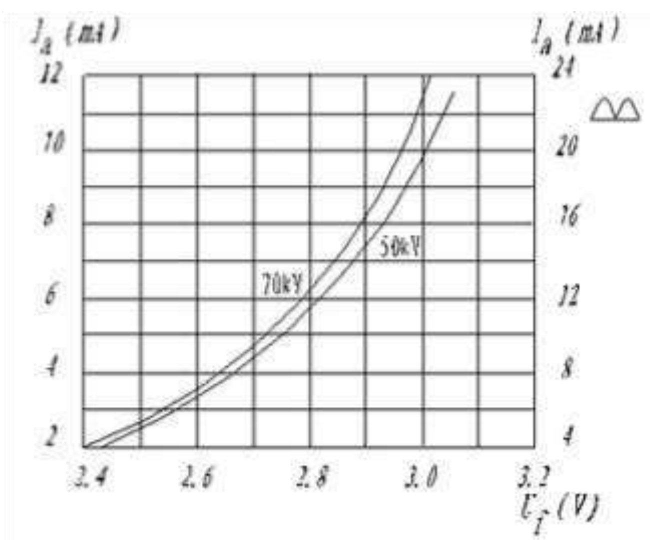

From what the above curve tells, one must apply 2.6V to the heater in order to allow a decent 3mA to flow through the tube's anode.

While 3mA might not sound like a lot of current, at 75,000V that is a respectable 225 watts of power. Assuming a typical 3% efficiency, this would equate to 6.7W of x-ray energy out. When one thinks about how a 100W light-bulb emits on average, 4 watts of visible light, it becomes quite clear that the tube will emit quite a hefty amount of radiation; certainly enough to expose film of any speed.

Thermal Limitations

A critical value which must be known is the anode's heat storage capacity. The tungsten/copper anode has a limited thermal conductance, thereby limiting its ability to dissipate the great heat generated by the focused electron beam. To cope with this Coolidge tubes often are run at a duty cycle, limited by both the operating power and the anode's heat storage capacity. In the case of this XD-12 tube, the anode's heat capacity is 7kJ.

Fortunately, manufacturers of both x-ray machines and x-ray tubes are mandated by federal law [CFR Title 21] to provide anode heating and cooling curves for their devices.

Figure 7 – Thermal curves

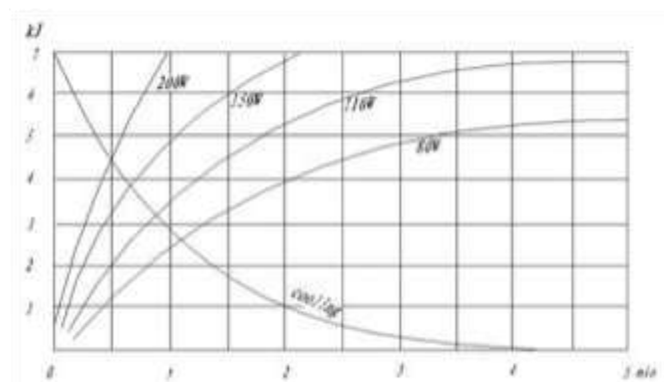

It is evident then, that operating this tube at a power of 225W would limit exposure time to a bit less than 1 minute with a 5 minute cool-down period. Of course x-ray exposures are never actually 1 minute long, but instead often consist of multiple short exposures. Provided the machine is not abused, 225W would not be an unreasonable power to operate this generator at.

Generating an extra high tension

There are many methods of generating high voltages including but not limited to Tesla coils, induction coils, Marx generators, Van-De-Graff generators and Cockroft-Walton cascades. Even some unorthodox methods such as piezoelectric and pyroelectric crystals exist.

While all of these methods have their merits, the Cockroft-Walton voltage multiplier would be the circuit of choice for this project. Properly designed cascades are capable of transforming large powers with comparatively little loss, and their lightweight and small stature make them well suited for a portable x-ray generator.

Figure 8 – A two stage CW multiplier

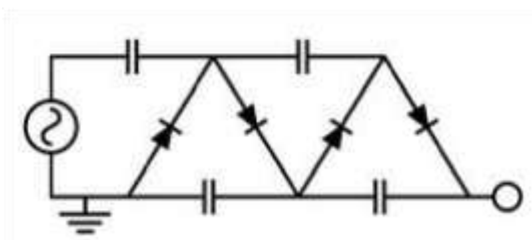

By feeding an alternating current into the circuit one can sacrifice cycles and current in return for a doubled, tripled or even quadrupled DC voltage.

CW theory of operation

All and all the circuit's mode of operation is rather simple, as it is nothing more than a cascade of Greinacher voltage doublers. Please refer to figure 9.

On the negative alternation, the bottommost plate of C1 is charged to -10kV via D1. Afterward, the positive alternation puts C1 in series with another 10kV creating a total potential difference of 20kV, which is shared with C2 via D2. This 20kV can then be discharged, or another cascade may be added to create 30kV, or another to make 40 kilovolts.

In reality it takes several more cycles for the stack to reach its full potential due to parasitic resistances limiting what would otherwise be very high currents.

Figure 9 – A single CW stage

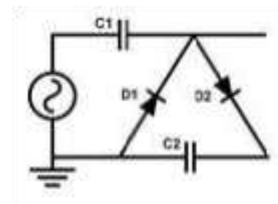

Designing a CW

The maximum voltage which may exist between ground and stage n can be predicted using the following formula:

$$Eout = Ein * \sqrt{2} * n$$

-where -

E_{in} *is the RMS input voltage*

n *represents the number of stages in the multiplier*

There is nothing overtly complex about this math; it's simply the peak DC input value multiplied by the number of stages in the multiplier stack. Of course, this is only the theoretical output voltage. Large high voltage capacitors are expensive and bulky, so in most cases we are stuck using small capacitors –and their high X_C.

This problem is compounded by the fact that many capacitors are placed in series in a multiplier. The high impedance nature of this circuit allows any sort of load to pull down the voltage substantially. This pulldown can be so profound that time was spent developing a formula to estimate the voltage drop under various conditions.

$$Edrop = I \div (f * C) \ * \left(\frac{2}{3} * n^3 + \frac{n^2}{2} - \frac{n}{6}\right)$$

-where-

I is the load current in Amperes

f is the frequency in Hertz

n represents the number of stages

C is the stage capacitance in Farads

The large influence of *n* in the latter half of the equation tells that using as few stages as possible in a multiplier will help minimize voltage drop. Fewer stages would mean fewer series-strung capacitors, and thus a lower X_C. Likewise, $I \div (f * C)$ tells us that both higher frequencies and larger capacitances will reduce the voltage drop under load. In both scenarios X_C would be reduced. A practical multiplier would thus contain no more than 5 stages, operate at a very high frequency and have capacitors whose values are not less than 500pF.

Designing *the* CW

Making do with what could be found on eBay the multiplier for this machine will consist of 4 stages, themselves consisting of a pair of series-connected 1.5nF, 15kV capacitors and another pair of 15kV ultra-fast diodes. The operating frequency will be set to 70kHz to keep capacitive losses in the transformer's windings to a minimum, yet that still should be high enough to push 3mA through without too much voltage drop.

$$Edrop = \frac{3mA}{75{,}000 * 7.5x10^{-10} * \left(\frac{2}{3} * 4^3 + \frac{4^2}{2} - \frac{4}{6}\right)}$$

$$Edrop = \frac{0.003}{0.0028125}$$

$Edrop$ = 2,666V

This 2.6kV drop is certainly reasonable, and may be compensated for by increasing the input voltage by an additional 650 volts.

However there is one minor problem; this formula is for the most part, useless. While the theoretical voltage drop should be 2.6kV, it will most likely be on the order of 6 to 12kV. Although this is significantly higher than predicted is not impossible to compensate for by upping the frequency and input voltage a bit.

Figure 10 – The completed cascade

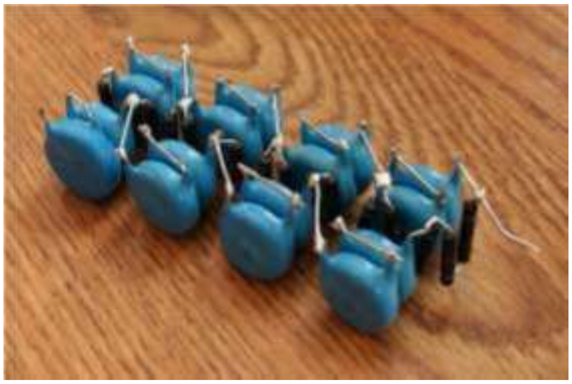

DC to HV AC Conversion

A CW multiplier still requires a moderately high voltage, high frequency input. Obviously, this cannot be obtained from a stack of 5,000 AA cells, so it must be generated somehow. The most logical method would of course be a forward mode flyback converter, preferably one which switches on the zero voltage or zero current crossing point to minimize losses.

In order to obtain 225W at 12kV we need 54mA, and assuming a 100% efficiency we'd likewise need to draw 6.25A from a 36V source. Of course 100% efficiency is unobtainable, but 10A at 36V is still not unreasonable for a battery to supply.

A choice oscillator for this task is the current-fed Royer Oscillator; an LC resonant zero-voltage switching circuit. Although a Hartley or a Colpitts oscillator could be used, both are not ZVS and would thus burn significant amount of power when the MOSFET switches through its linear region. A Zero Voltage switching oscillator on the other hand does not.

Figure 11 – Mazilli Oscillator

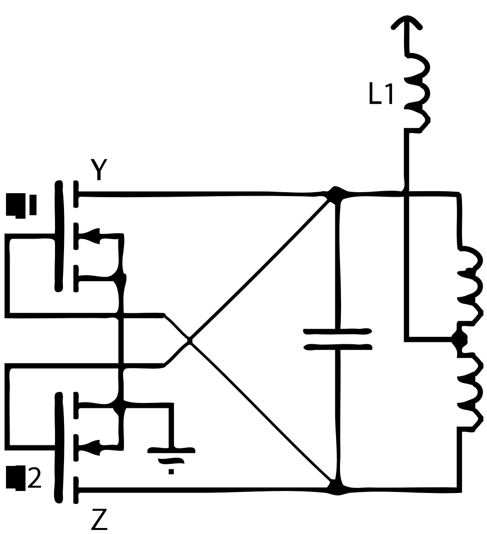

[Assuming conventional current flow]

When power is applied to the circuit, current begins to flow through L1 and into the MOSFETs' drains via the center tapped load coil. Simultaneously, this voltage appears on both gates and begins to turn the MOSFETs on. Since no two FETs are alike, one turns on faster than the other and drags down the voltage on the opposite MOSFET's gate. One FET is now latched on while the other is off. The tank capacitor prevents the circuit from staying in this state as the LC resonance causes a sinusoidal reactance in the circuit; reactance which will 'flip' the MOSFETs' states and feed more current into the tank. This oscillation keeps going until power is removed, or some instability such as a saturated load inductor latches one MOSFET on and explodes the circuit.

Though this circuit would work in theory, such an oscillator proves to be very unreliable without some modifications. Instead of connecting the gates directly to the LC tank, it'd be wiser to instead leave them normally pulled up via a pair of resistors, where the LC tank would alternately ground the gates via ultrafast feedback diodes. This method ensures that no stray currents lock up the gates and kill the circuit.

Of course MOSFET gates will not tolerate a VGS in excess of 30V, so it is a wise idea to use 18V zener diodes to protect the gate from such excess voltages. 10K discharge resistors ensure that no stray charges are left on the gate while it is being pulled down by a feedback diode.

Figure 1 - A practical Mazilli oscillator

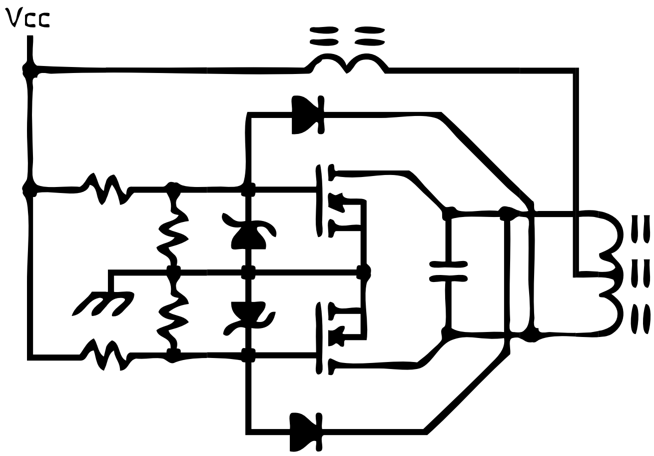

As it is nothing more than LC oscillator, a Mazilli's operating frequency may be derived using the parallel LC resonance formula:

$$f = \frac{1}{2\pi * \sqrt{LC}}$$

-where-

L is the tank inductance in Henries.

C is the tank capacitance in Farads.

Parallel resonant LC oscillators often have poor tank power factors, with the Mazilli being no exception. When drawing 8 amps the circulating tank current may exceed 50A! This can become a problem if the DC resistance of the load coil is high. To combat this one may wish to use a larger capacitor than inductor, since a metalized polyester capacitor's parasitic resistance tends to be smaller than a large coil's.

$$f = \frac{1}{2\pi * \sqrt{0.00000068 * 0.0000075}}$$

f = 70,474Hz

To achieve 70kHz with 680nF of tank capacitance we'd need roughly 7.5µH of tank inductance, or about 5 turns of wire on an average ferrite core. Since 5 turns of copper wire has for the most part an insignificant series resistance I^2R losses in the coil will be minimized. A trouble may arise if the coil voltage raises high enough to saturate the transformer core, but an air gap should solve any issues.

Since commercially produced HV transformers are hard to come by in small quantities, the transformer for this prototype had to be handmade. Special attention was paid to prevent saturation that could otherwise occur at 15 volts per turn; the core was chosen to have a low permeability, a large cross sectional area, a small magnetic path and a 0.5mm air gap.

Figure 13 – A homemade transformer

Switch-Mode Power Conversion

The X-Ray tube requires a low voltage, high current power supply to energize its heater; something which must be derived from the 14V bus. A linear regulator would burn away unacceptable amounts of power so a switch mode regulator must be used instead.

The heated cathode does not require an isolated power supply, so a simple buck topology may be used to provide the regulated voltage it requires. Buck converters are the simplest of the switch-mode topologies. Essentially the circuit measures the voltage across a capacitor, and attempts to maintain a set voltage by varying the current supplied to that capacitor. Often this is done by varying the duty cycle on a rapidly switching MOSFET.

While that would work in theory, in practice the very high currents experienced by the MOSFET would increase losses to intolerable levels. In the real world an inductor is placed in series with the MOSFET to 'average' this current flow.

This solution creates another problem though. Quickly interrupting an inductor's current flow would create damaging high voltage spikes which could destroy the MOSFET. Typically this is solved by putting a diode in antiseries with the inductor, but that would create a very lossy circuit. Instead, the diode is placed in antiseries with the load.

Figure 14 – A buck converter

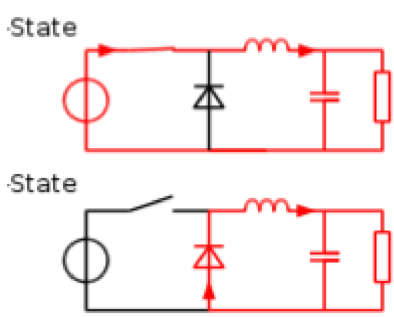

When the switch is closed, current flows through the load and filter capacitor via an inductor. When the specified voltage is reached on the capacitor, the switch is opened. Both the inductor and capacitor then deliver power to the load via a schottky flyback diode until the voltage falls enough to once more turn on the MOSFET. This happens thousands of times per second.

While it might be a bit more inexpensive to build a buck converter out of discrete components, using a Simple Switcher from National Semiconductor was found to both increase reliability and reduce real estate. Perfect for the job was the LMZ12003 buck converter IC, with an inbuilt inductor. Averaging 94% efficiency in a surface mount package, this IC in my opinion is a feat of semiconductor engineering like no other!

Figure 15 – Typical LMZ12003 application

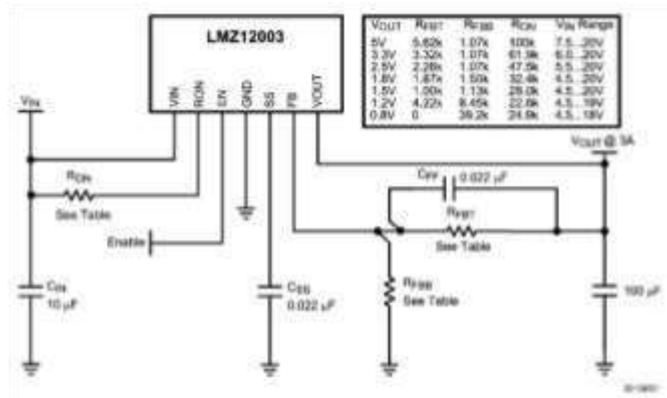

Fortunately there is not much engineering to be done using this IC, apart from setting the feedback reference via a voltage divider connected to the output. This feedback pin feeds into an inbuilt comparator which the onboard oscillator uses to set the MOSFET's duty cycle.

Figure 16 – Prototype buck converter

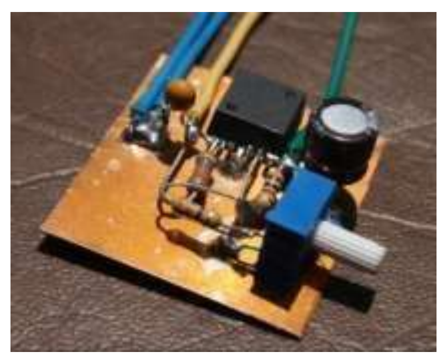

Assembly of the Radiating Head

Atmospheric air has a dielectric strength of 1.1 million volts per meter despite what Wikipedia argues. This translates to 11kV per centimeter, or 6kV with pointed electrodes. Since reliably insulating 75kV using air would require a distance larger than 10cm, making the device compact is near impossible.

That is of course, if one does not use insulating oil. Most oils have a dielectric strength 4 times that of air and eliminate the corona losses which would otherwise

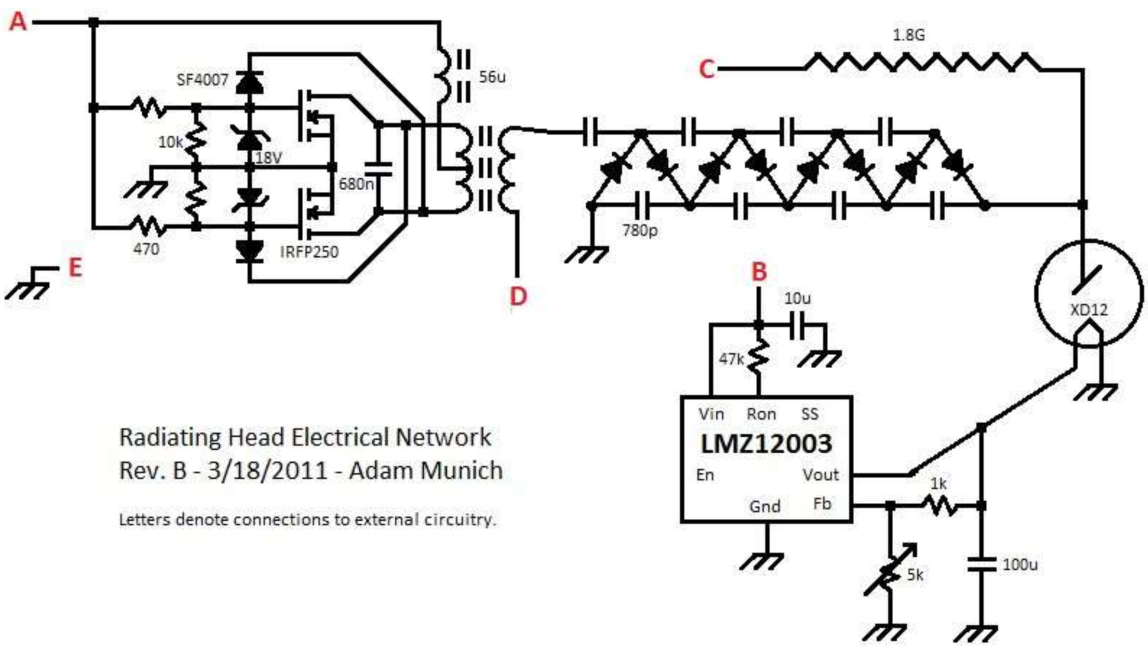

occur in an open air design. This reason, coupled with increased thermal conductivity is the reason why nearly all x-ray machines insulate all of their high voltage components with oil, and why mine will follow suit.

Since injection molding was far out of budget most of the assembly had to be made with commercially available components. Instead of a custom-molded enclosure a junction box was used to house the EHT components. Figure 17 displays the junction box which houses the Coolidge tube, its lead shield, the voltage multiplier and a 1.8 billion ohm resistor to measure the anode voltage. A 90kV this resistor will leak the 50uA needed to fully deflect a galvanometer.

Figure 17 – The ‘hot box’

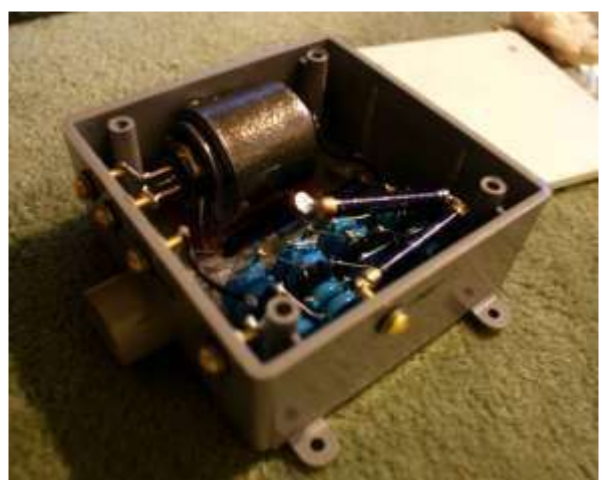

The thickness of the box’s wall will attenuate the xrays somewhat so there likely won’t be any low energy rays escaping generator. An aluminum window will be fitted in later designs to solve this problem, as aluminum is a lot more x-ray transparent than polystyrene is. However, x-rays with energies higher than 35keV should still be able to escape the ‘hot-box’.

Limited by what I could fabricate with a lone Dremel, everything was fitted into a wooden craft box and painted to look nice. Later prototypes will be built with 3D printed cases.

Part III – The Control Box

Designing a User Interface

With the radiating head designed and built it was time to build a second device to both power and control it. Typically, an x-ray machine follows the following operating procedure:

1. The technician chooses an exposure time and kVp.
2. The tube's heater warms up for a short period.
3. High voltage is turned on and x-ray photo is taken.

It's a process that must be replicated in this prototype machine. While the circuitry could be as simple as a 555 timer and a few MOSFETs, such a circuit would be rather inaccurate and for the most part inferior in today's digital world. The safest, most economical method would be to base the circuit around a micro-controller. Not only would a microcontroller be reliable, but it would have the added benefit of allowing for easy modification later on. An ATMEGA368 programmed with the arduino bootloader was the controller of choice.

It was suggested by a friend that I use a set of nixie tubes for the machine's display. Although an HD44780 interface LCD would have been the most logical option, the challenge of designing a multiplexed nixie display was one too enticing to pass up. The next prototype will certainly use an LCD however.

Figure 19 – Multiplexed nixie display

In the end though, this proved to be quite reliable and accurate enough for 'real use'. Since the x-ray tube's anode current is fixed, its impedance will also be fixed. With its nonvariable impedance the anode's voltage will remain constant time and time again, so a calibrated dial is all that was needed to set an accurate kVp.

Figure 20 – Detail of the atmega and voltage regulators.

Exposure time is to be set using three knobs: one for tens of seconds, one for the seconds and one for the 0.1 seconds. In practice this makes it simple to adjust the exposure time from 0.1 to a software-limited 60.0 seconds. A 1 minute exposure is completely useless for any medicinal purposes of course, but such long exposure times would be required if the generator were used to inspect dense, inanimate objects.

A 10 second countdown ensues after pressing the exposure button, allowing both time for the technician to ready themselves and adequate time to fully heat the cathode. After the countdown high voltage is applied for the set exposure duration and the x-ray burst is produced. The system then resets.

The power source is a pair of lithium polymer batteries; one 4 cell pack named Hans and a 6 cell pack who goes by the name of Frantz. Hans provides the +14.4V rail while the series-strung Frantz supplies the +36V that's fed to the linear regulator. Both voltages are monitored by the microcontroller, and if one happens to dip below the safe discharge threshold for lithium cells, the machine shuts down.

Figure 21 – Hanz and Frantz

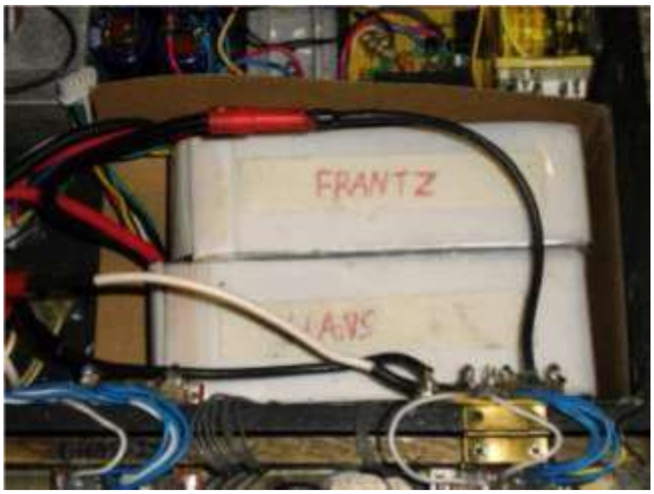

Using standard-size lipo packs gives both an economical and user-friendly advantage to the machine. Such R/C lithium batteries can be had near anywhere for a mere $30 and could be purchased from many different sources if need be. Chargers for these batteries exist which will operate on nearly every standard voltage; even the 12V that could be obtained from a solar panel if that is the only power source available.

The control box is kept separate from the radiating head by means of a detachable octal connector cable. This arrangement not only allows the head to be properly positioned, but also provides distance between the x-ray source and its operator. Distance is of course one of the best ways to shield from any ray source which follows inverse square law. X-rays are not exempt from this law, so by doubling the distance between the source and the target the x-ray flux drops by a factor of 4.

That said, little to no stray radiation would escape the Coolidge tube's lead jacket. Any exposure which the operator might receive would be limited to x-rays which are scattered and reflected off of objects in the beam. Even then, such a minute exposure could be all but nulled by 0.5mm leaded vest.

Figure 22 – The prototype

Microcontroller Code, Rev B.

The following annotated code is the sequence the ATMEGA368 has been programed with.

```
// Define compiler constants -------------------------------------------------------------
#define HVsupp 2       // HV supply relay on D2.
#define heater 4       // Tube heater relay on D4. High = heater on
#define expose 10      // Expose button is on D10
#define speaker 1      // Speaker on A5
#define battcheck A5   // Button to check oscillator voltage on D1
#define nixTens 13     // Tens-Place nixie multiplex switch on D13
#define nixOnes 12     // Ones-Place nixie multiplex switch on D12
#define nixTenths 11   // Tenths-Place nixie multiplex switch on D11
#define potTenth 2     // Tenths potentiometer on A0
#define potone 1       // Ones potentiometer on A1
#define potTen 0       // Tens potentiometer on A2
#define Vbatt 3        // 4-Cell voltage measurement on A3
#define Vbatt2 4       // 10-Cell voltage measurement on A4

// Initialize variables on a global scope -----------------------------------------------
int tenths = 1;
int ones = 0;
int tens = 4;
int bytesread = 0;
int val = 0;
int voltageA = 0;
int voltageB = 0;
int avg[5];
unsigned int storeTenths = 0;
unsigned int storeOnes = 0; unsigned
int storeTens = 0; unsigned int
storeOnes2 = 0;
byte index = 0;
byte numb = 0; byte
digit = 0;
boolean digA = 0;
```

```
boolean digB = 0;
boolean digC = 0;
boolean bitA = 0;
boolean bitB = 0;
boolean bitC = 0;
boolean bitD = 0;
boolean doTheExposure = false;
boolean safetyCountdown = true;
boolean hasExposed = false;
boolean safetyDone = false;
boolean doIt = false;
 int clockCount = 0;

// Initialize pins ---------------------------------------------------------------------
void setup(){
    pinMode(HVsupp, OUTPUT);
    pinMode(3, OUTPUT);
    pinMode(4, OUTPUT);
    pinMode(5, OUTPUT);
    pinMode(6, OUTPUT);
    pinMode(7, OUTPUT);
    pinMode(8, OUTPUT);
    pinMode(9, OUTPUT);
    pinMode(expose, INTPUT);
    pinMode(battcheck, INTPUT);
    pinMode(speaker, OUTPUT);
    pinMode(nixTens, OUTPUT);
    pinMode(nixOnes, OUTPUT);
    pinMode(nixTenths, OUTPUT);
    digitalWrite(heater, LOW);
    digitalWrite(HVsupp, LOW);
}
```

```
// Main Routine ---------------------------------------------------------------------------
void loop(){
   while(digitalRead(battcheck) == true){ // While button is pressed check regulator voltage
         voltageA = map(analogRead(Vbatt), 10, 1024, 0, 500);
         voltageB = map(analogRead(Vbatt2), 10, 1024, 0, 500);
         numb2digit(voltage);
         printNixie();
   }

 // Stop by and grap the potentiometer readings
     storeTens = tens;
     storeOnes = ones;
     storeTenths = tenths;
     tenths = map(analogRead(potTenth), 10, 1000, 0, 9);
     ones = map(analogRead(potone), 10, 1000, 0, 9);
     tens = map(analogRead(potTen), 10, 1000, 0, 9);
     if(storeTens != tens || storeOnes != ones || storeTenths != tenths)
         tone(speaker, 200, 50);

    // Check if the expose button is pressed
    if(digitalRead(expose) == true){
       wait(1000); // For at least one second
       if (digitalRead(expose) == true){
         doIt = true;
       }else{
         doIt = false;
       };
     };


    // If it was pressed for 1 second then start the 10 second countdown
 if (doIt == true){
    safetyCountdown = true;
    storeTens = tens; // Copy the potentiometer values to the countdown variables
    storeOnes = ones;
```

```
    storeTenths = tenths;
    tone(speaker, 200, 1000);           // 200Hz beep on the speaker
    printWait(2000, 2, 0, 2);           // Wait 2 seconds
    digitalWrite(heater, HIGH);         // Turn on the tube heater
    printWait(100, 1, 0, 0);
    while (true) {                      // Get stuck in a loop
      storeOnes2 = ones;
      printNixie();
      countdown();                                  // Call the countdown routine, should take 1 second
      if (storeOnes2 != ones) tone(speaker, 80, 500);     // Beep every second
        // If the countdown is complete, set a variable true and break the loop
      if (tens == 0 && ones == 0 && tenths == 0){
        safetyDone = true;
        break;
      };
    };
  };

  if (safetyDone == true){              // If the safety countdown is done then start making x-rays
     tone(speaker, 80);                 // Turn on warning buzz and wait a sec
     wait(1000);
     digitalWrite(HVsupp, HIGH);        // Turn on x-rays
     tens = storeTens;                  // Copy the stored pot values back to the active variables
     ones = storeOnes;
     tenths = storeTenths;
     doTheExposure = true;              // Tell the other routine to count down
   };

  if (doTheExposure == true) {          // Here's the 'other routine'
     clockCount = 0;
     while (true) {                     // Countdown and break when all numbers reach 0
        printNixie();
        countdown();
        if (tens == 0 && ones == 0 && tenths == 0) break;
     };
```

```
      digitalWrite(HVsupp, LOW);             // Then turn off the HV supply
      wait(500);
      noTone(speaker);                 // And the buzzer
      hasExposed = true;
      doIt = false;
      safetyDone = false;
      doTheExposure = false;
      for (int x = 0; x < 4; x++){      // Superfluous 'Food is ready' beeps
            tone(speaker, 200, 500);
            printWait(500, 0, 0, 0);
            delay(500);
      };
     digitalWrite(heater, LOW);
   };
   printNixie();                         // Update the nixie display, in case none of those routines ran
};

// Subroutines ------------------------------------------------------------------------------------------------------------
void countdown(){        // The countdown routine. Counts down whatever is on the 3 number vars
   switch(clockCount){
         case 100: case 200: case 300: case 400: case 500: case 600: case 700: case 800: case 900: case 1000:
         -- tenths;
         break;
   };
  if (tenths == -1){          // Reset tenths if < 0
      tenths = 9;
      clockCount = 0;
      -- ones;
      if (ones == -1){        // Reset ones if < 0
            ones = 9;
            -- tens;
            if (tens == -1){   // Reset tens if < 0
                  tens = 9;
            };
      };
```

```
 };
 clockCount += 2;
};

//----------------------------------------------------------------------------------
void wait(int n){    // Waits 'n' milliseconds, doesn't blank nixies
    n = n / 2;
    for (int x = 0; x < n; x++) printNixie();
};

//-----------------------------------
void printWait(int n, int A, int B, int C){    // Same as above but also prints a number
     n = n / 2;
     tens = A;
     ones = B;
     tenths = C;
     for (int x = 0; x < n; x++) printNixie();
};


//-----------------------------------
void printNixie(){ // Prints the pre-stored numbers to the tubes, takes exactly 2mS
          toTheTubes(tenths, 0);
          toTheTubes(ones, 1);
          toTheTubes(tens, 2);
};

//-----------------------------------
void toTheTubes(byte numb, byte digit){           // Sends an actual byte to the tube's decoder IC
     switch (numb) {      // Convert the digit to binary-coded-decimal
          case 0: bitA = 0; bitB = 0; bitC = 0; bitD = 0;  break;
          case 1: bitA = 1; bitB = 0; bitC = 0; bitD = 0;  break;
          case 2: bitA = 0; bitB = 1; bitC = 0; bitD = 0;  break;
          case 3: bitA = 1; bitB = 1; bitC = 0; bitD = 0;  break;
          case 4: bitA = 0; bitB = 0; bitC = 1; bitD = 0;  break;
```

```
        case 5: bitA = 1; bitB = 0; bitC = 1; bitD = 0;  break;
        case 6: bitA = 0; bitB = 1; bitC = 1; bitD = 0;  break;
        case 7: bitA = 1; bitB = 1; bitC = 1; bitD = 0;  break;
        case 8: bitA = 0; bitB = 0; bitC = 0; bitD = 1;  break;
        case 9: bitA = 1; bitB = 0; bitC = 0; bitD = 1;  break;
    };
    switch (digit) {                    // What anode should be turned on?
        case 0: digA = 1; digB = 0; digC = 0; break;
        case 1: digA = 0; digB = 1; digC = 0; break;
        case 2: digA = 0; digB = 0; digC = 1; break;
     };
    digitalWrite(8, bitA);              // Write the BCD data to the decoder IC
    digitalWrite(7, bitB);
    digitalWrite(6, bitC);
    digitalWrite(5, bitD);
    digitalWrite(nixTens, digA);        // Send a square wave pulse to the proper anode transistor
    digitalWrite(nixOnes, digB);
    digitalWrite(nixTenths, digC);
    delayMicroseconds(666);             // Wait 2/3 of a millisecond. 3 calls == 2mS
}

//-----------------------------------
void numb2digit(int number){            // Split a 3 digit number into 3 digits by modulo division
    tens = number / 100;
    ones = (number / 10) % 10;
    tenths = number % 10;
}



//-----------------End---------------
```

Part IV - Testing

Equipment

Dosimetry of the x-ray beam was measured with a cold war era JAN-5979 Geiger-Muller tube.

Essentially a GM tube is a neon lamp which is biased just below its breakdown voltage. When a highenergy particle enters the tube it invokes an electron avalanche after ionizing a path the low pressure fill gas. Since the GM tube is connected to a high impedance power source, this avalanche causes a voltage drop which can be amplified and registered as a 'count'. About 75 microseconds later the avalanche is extinguished by a quench gas, and the tube is ready for another count. Typically the quench agent is a 1000ppm fill of either methane or a similar halogenate gas.

Methane tubes degrade over time, and old ones are usually unusable for measurement. Halogen tubes such as the JAN-5979, on the other hand last forever. In fact, this model was so reliable that it is still produced today, under the identity LND-7225.

Figure 22 – Geiger Muller Tube

While particle radiation such as beta and alpha can be detected directly, electromagnetic radiation such as gamma and x-radiation won't cause an electron avalanche in the tube; at least, not directly. In order to cause an avalanche one of these photons must first knock an electron out of the tube's wall, and that's not exactly an efficient process. Typically a Geiger tube will respond to only 3 percent of the total x or gamma radiation present, though this is increased to about 5 percent in a tube that has a thicker steel wall like the 5979.

Beam Dosimetry

Initial tests were performed at an energy of about 70kVp. The milliammeter reading of 2.4mA suggested that a total of 168W was delivered to the x-ray tube, a small but not insignificant amount of power. Assuming 2.5% efficiency this equates to 4.2W of x-ray energy, certainly not a negligible sum considering that is the amount of light a 100W Edison lamp emits! Assuming the tube operated for 1 second, 4.2J of x-ray energy would be emitted.

Figure 23 – The counter ready for testing

Results:

As expected, the near-field radiation was intense enough to fully saturate the Geiger tube. Thus, the dose-rate in the x-ray beam was measured to be 'pretty damn high'.

Leakage Dosimetry

Although the beam's x-ray flux was impossible to quantify using the tools at my disposal, leakage radiation was still measurable. The ambient background radiation was averaged to be 64 counts/minute, so this was subtracted to prevent skewing of the data. Below is a diagram of both the experiment and test results.

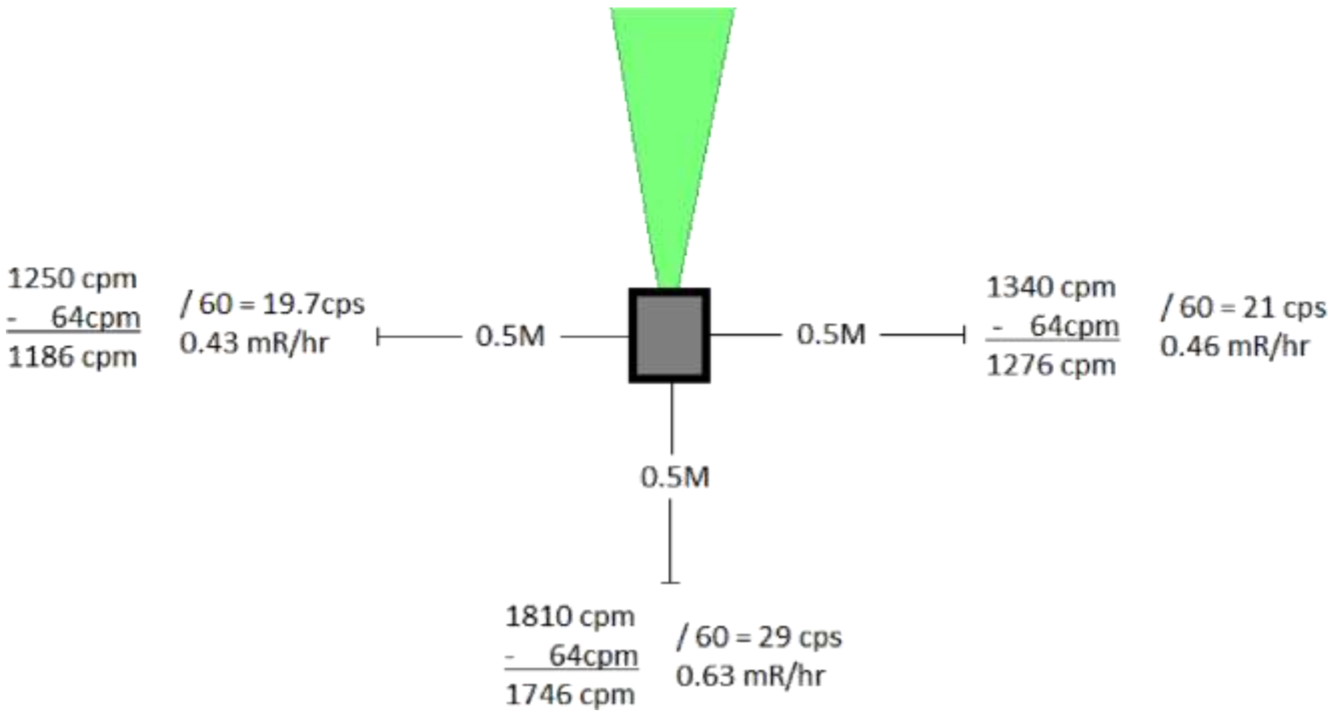

The gamma/x-ray sensitivity of a JAN-5979 Geiger tube is 45 counts/second per mR/hr. On average, leakage from the generator was about 50 mR/hr, or 500 µSv/hr in the related Si unit. 500µSv/hr is a rather high dose-rate as far as leakage radiation goes, especially when the background was measured to be 0.2 µSv/hr.

Strangely, most of this the radiation seems to be coming not from the machine but from the beam itself. I suspect the culprit is Compton scattering, coupled with small reflections off the trees in the uninhabited forest at where the beam was directed. The fact that the Geiger tube gives an increased response when its mica window is pointed toward the beam further supports this hypothesis, since scattered radiation is more likely to have a lower energy than directly emitted radiation.

Radiography Tests

It is now time to test the most important aspect of this machine; its ability to take a radiograph. Because of the risk associated with exposing living tissue to ionizing radiation, these tests will be performed using inanimate, highly detailed objects.

The first subject will be a 2.5 inch hard disk drive. The small parts and high detail in such a hard drive will provide a good estimate of the x-ray machine's resolving power.

Figure 25 – First light

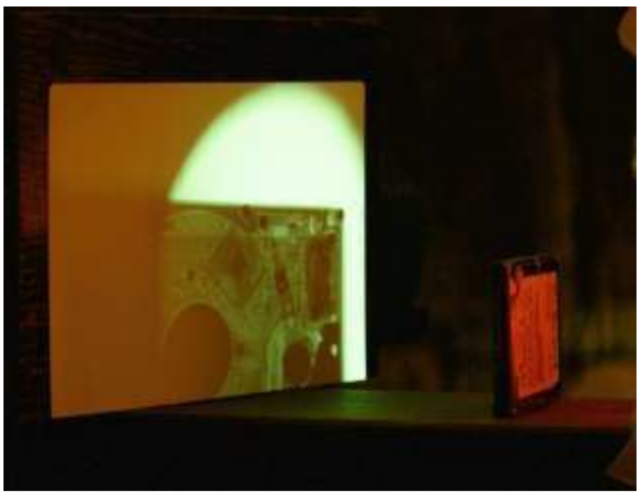

Figure 26 – Pitiful results...

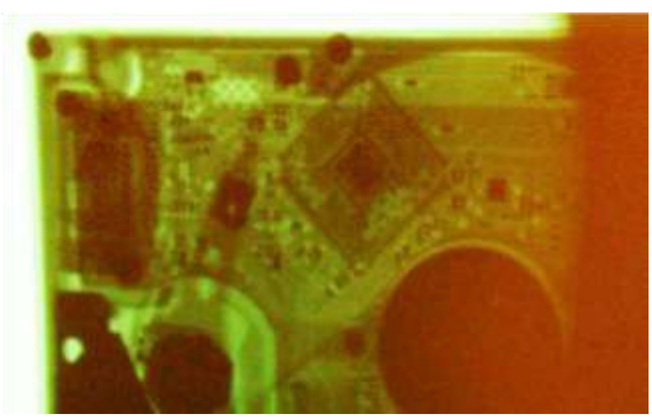

Although it looked promising from far away, up close the radiograph was quite disappointing. I doubt it was due to the machine itself; but rather the $50 camera used to take the photos.

Attempt #2

For the next test a higher quality camera was used, this time on a 3.5 inch hard drive. A 'normal speed' intensifying screen also substituted the 'ultra-rapid' one used during the last test. Although the slower screen is a bit dimmer, the thinner substrate should cause less blur and distortion.

Figure 27 – Second attempt at radiography

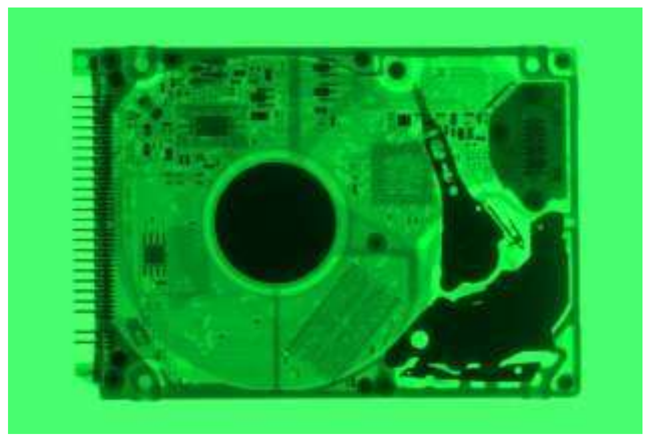

There, much, much better! This picture was taken using a DSLR camera set at f/8, exposure time 11 seconds. X-Ray film would need a much shorter exposure, but with Kodak going out of business it's hard to find Dektol and rapid fixer at my local camera shop, let alone paper film.

Figure 28 – Detail of HDD radiograph

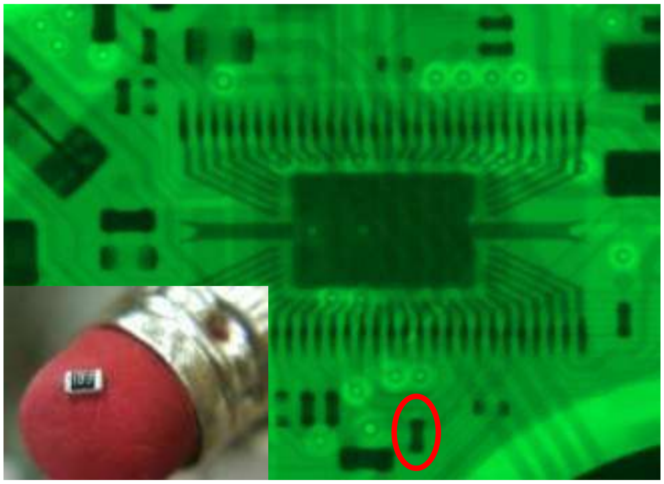

This encircled component is a 0603 SMT resistor, a 0.06" tall by 0.03" wide high-tech speck. The resolution of this radiograph is quite good, certainly comparable to what a commercially made machine can produce. Not bad considering my budget was about 1/20th of something from GE Medical.

The space below hosts a few radiographs, each displaying the hidden wonders of household objects with astounding acuity.

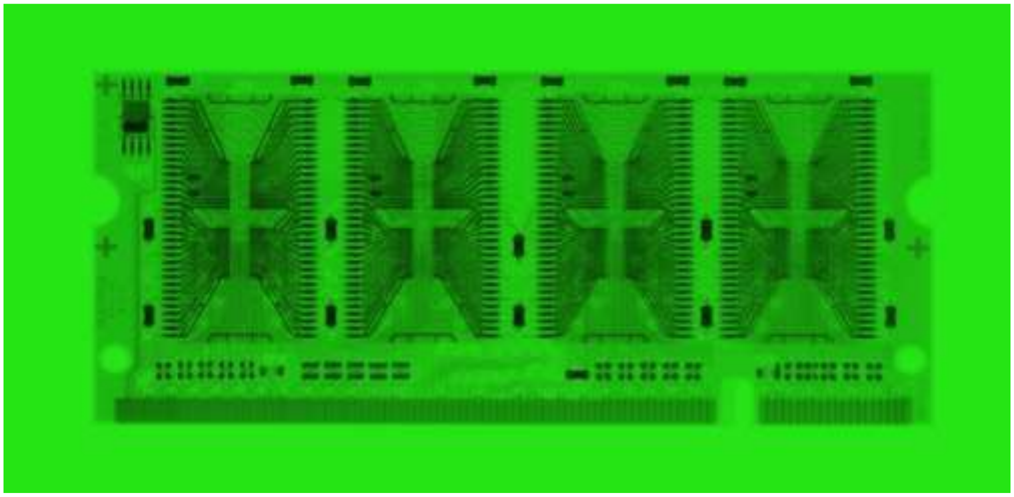

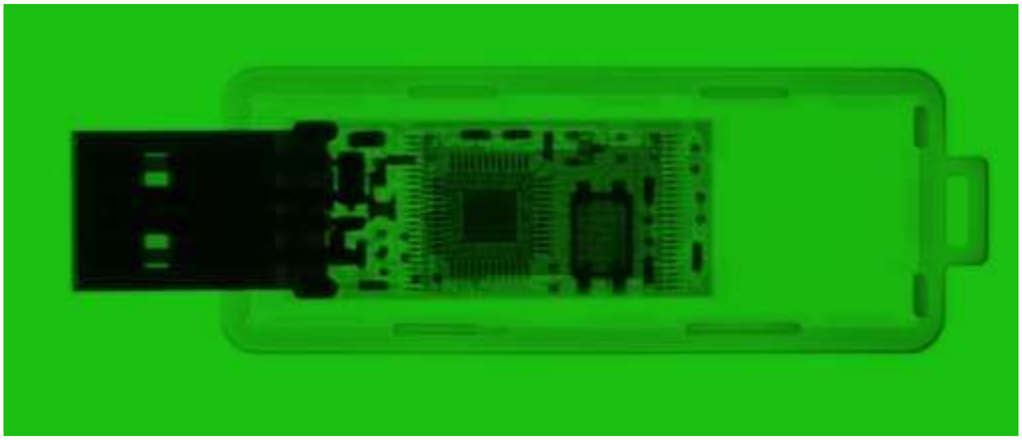

Part V - Conclusion

Portability

Was it possible to build a portable x-ray machine? Yes, it was. Not only was it possible to build a portable x-ray generator, but it was possible to build useable x- ray generator. Despite its small stature the generator has proven itself capable of taking sharp, high-contrast radiographs; perfect for that medical unit stationed at a tent in Afghanistan.

Cost

Though I've now proven it possible to build a portable x-ray machine, could it be made for less than $400? Although an army unit might have a pretty penny to spend, that village hospital in Africa is likely a bit short on funds…

- 1 Coolidge tube $75
- 16x 15kV 1.5nF capacitors $45
- 16x 15kV ultrafast diodes $20
- Ferrite Transformer $15
- Electrical components $65
- LiPo batteries $60
- Mechanical things $50
- LiPo charger $30

Grand total $360

Of course this all assumes labor is free, something which any businessman will tell you is false. In reality the cost of producing such a machine would likely be $600, where it would then be sold for $1000.

I suppose this is just a problem with x-rays; producing them has always been an inefficient, clumsy and expensive process. While the Coolidge tube was a major advancement in x-ray technology no real game-changers have come along since the late 1940's. Sure, there has been some research into a pyroelectric generation method by AMPTEK, but this technology has proven incapable of producing the large x-ray flux needed to take a medical radiograph.

There's been research into triboelectric production via scotch tape, but the Coolidge tube still has that process beat in almost every area. Despite the huge advances in LEDs in recent years that technology can't be applied to x- ray production, unless of course somebody discovers a semiconductor with a 70,000V bandgap!

Suppose all the amenities were dropped though. Rather than having it adjustable, the kVp would be fixed at 60kV. Rather than having an internal battery, the generator would be powered by a 24VDC input. Perhaps that 24V might come from two car batteries. Rather than having a nixie or LCD display, exposure duration could be set by a simple knob. Could it then be possible to build an x-ray machine on a budget?

- 1 Coolidge tube $75
- 12x 15kV 1.5nF capacitors $33
- 12x 15kV ultrafast diodes $15
- Ferrite Transformer $15
- Electrical components $20
- Mechanical things $20

Grand total $177

From the looks of it, yes, it would be possible! Phase 2 of this project will be to do just that –to skim down one machine to bare minimums, and to design a second machine, capable of imaging any object you can throw at it.

Until this is done, I bid you farewell!