SBIR Final Report

Solid-State Thyratron Replacement

Topic Identifier: 31(b)

Award: DE-SC0011292
Diversified Technologies, Inc.
35 Wiggins Avenue
Bedford, MA 01730

December 12, 2017

Dr. Ian Roth
Principal Investigator

Table of Contents

1.0 Executive Summary

The thyratron has been used as a switch in pulsed-power applications for almost a century. In the last 20 years, as a result of developments pioneered at Diversified Technologies, Inc. (DTI), most new applications have transitioned away from thyratrons as solid-state switching technology has become available. As the cost and capabilities of solid-state modulators has improved, their adoption has grown rapidly. With the continued evolution and reliability of solid-state pulsed-power systems, virtually all new pulsed-power systems are designed around solid-state capabilities.

As a result, the thyratron market is in decline—as newer solid-state modulators are deployed and older thyratron systems are taken out of service, the demand for thyratrons has diminished significantly. In response, several vendors have gone out of business or stopped manufacturing thyratrons, further diminishing their availability. For existing systems, however, solid-state modernization typically requires replacement of the entire modulator system, which is expensive. For some applications and facilities, such as at SLAC, the cost of a wholesale upgrade to solid-state modulators is prohibitive, leading to the desire for a true thyratron replacement. A solid-state switch which performs at the level of a thyratron, and could be built for a price that could be recouped in 3 – 5 years of enhanced reliability, would allow replacement of only the thyratrons in existing modulator systems, with rapid payback of the initial investment.

Under this SBIR, DTI developed a solid-state switch as an alternative to legacy thyratron equipment. Our Phase II objective was to make a solid-state thyratron replacement that would provide equivalent or better performance, much higher reliability (at least a 20 year lifetime, compared to a thyratron's two-year lifetime) and would sell for ~3x the cost of a thyratron, or less than $40k. We were successful in building a solid-state switch which could reliably function as a thyratron replacement. The unit was designed to directly replace the thyratrons currently being used at SLAC's Linac Coherent Light Source (LCLS), and was built in a tank that was small enough to fit into the existing thyratron cabinet, providing a true form-fit-function replacement path. We tested the switch at the full operating specifications: 48 kV, 6.3 kA, and 1 μs risetime. We also demonstrated a peak-to-peak pulse jitter of 1.5 ns, which is five times shorter than is typical for thyratrons. This lower jitter would improve the performance of the LCLS beam. The predicted reliability is more than 80 years, which is 40 times greater than a thyratron.

Unfortunately, due to several design issues uncovered during construction and testing, the SBIR contract funds were expended before the switch could be installed at SLAC as originally planned, and tested at 120 Hz, the full pulse rate. The final design was more complex, and required significantly more devices for reliable operation, than we expected after Phase I. As a result, we were not successful in demonstrating that this switch could be built at a price comparable to a thyratron – it would be roughly an order of magnitude higher price, even in quantity.

We do believe that on a long-term basis (10 – 20 years), the solid-state approach would provide a lower life cycle cost, but this is beyond the upfront investment feasible at SLAC. The primary limiting factor to the switch cost is determined by the dI/dt capability of IGBTs, which is improving with each new generation, but is not yet fast enough to make this transition cost-effective. Subsequent to this work, thyristors with a very high dI/dt have been claimed to be developed by Silicon Power. It is possible that these could be used in a much less expensive switch.

2.0 Report Organization

This report is arranged in four parts:

- Phase I
- Phase II (Year 1)
- Phase II (Year 2)
- Phase II (Year 3)

3.0 SBIR Phase I

These SBIR Phase I activities occurred in the period of February 18, 2014 – November 17, 2015.

3.1 Identification & Significance of the Opportunity

As solid-state modulators are deployed in new accelerators, and older thyratron systems are taken out of service, the demand for thyratrons has diminished significantly. However, in many existing systems, it is considered less expensive in the short-term to retain thyratrons, many of which need to be replaced annually, rather than upgrade the system. **A solid-state switch which performs at the level of a thyratron, and could be built for a similar price, would allow replacement of thyratrons in existing modulator systems as they are depleted, with rapid payback of the initial investment.**

Despite their reputation, thyratrons are not cheap when lifetime costs are considered. The L-4888 thyratrons used at SLAC must be replaced approximately every two years, at a cost of $13,000 each, plus costs for installation and regular maintenance to adjust the reservoir heater voltage. **Replacing these thyratrons with solid-state switches that last 20 years or more without maintenance would provide significant savings – over $1M per year.** Furthermore, a solid-state switch will also improve performance, reducing pulse-to-pulse jitter and protecting the klystron from faults. **The solid-state switch's extremely long life and reliability presents an attractive opportunity for many thyratron-based systems.**

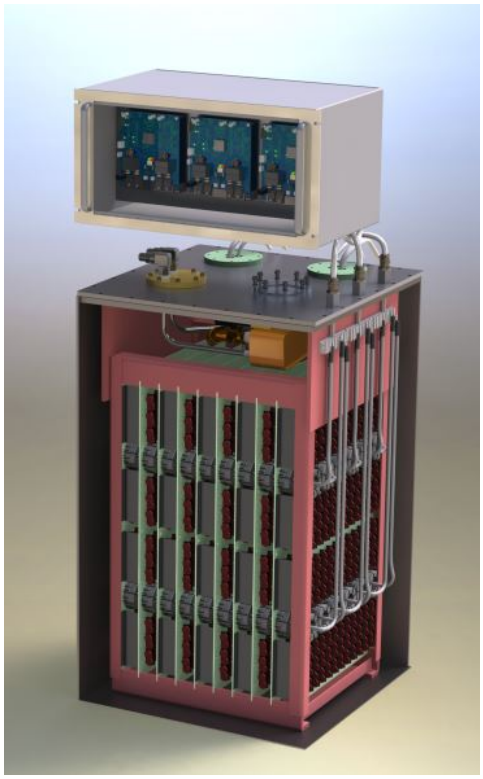

Figure 1. DTI's proposed replacement switch for the L-4888 thyratron used at SLAC. The switch operates at 48 kV, and 6.3 kA. The switch and controls fit in the same location as the legacy thyratron assembly.

In Phase I of this SBIR, Diversified Technologies, Inc. (DTI) successfully demonstrated the technology for a solid-state thyratron-replacement switch assembly (Figure 1 and Figure 2). We have demonstrated full operating capability, arc handling for both individual IGBTs and a representative matrix of IGBTs, and peak-to-peak pulse jitter of 1.5 ns, which is five times shorter than typical thyratron performance which will help stabilize of the LCLS beam.

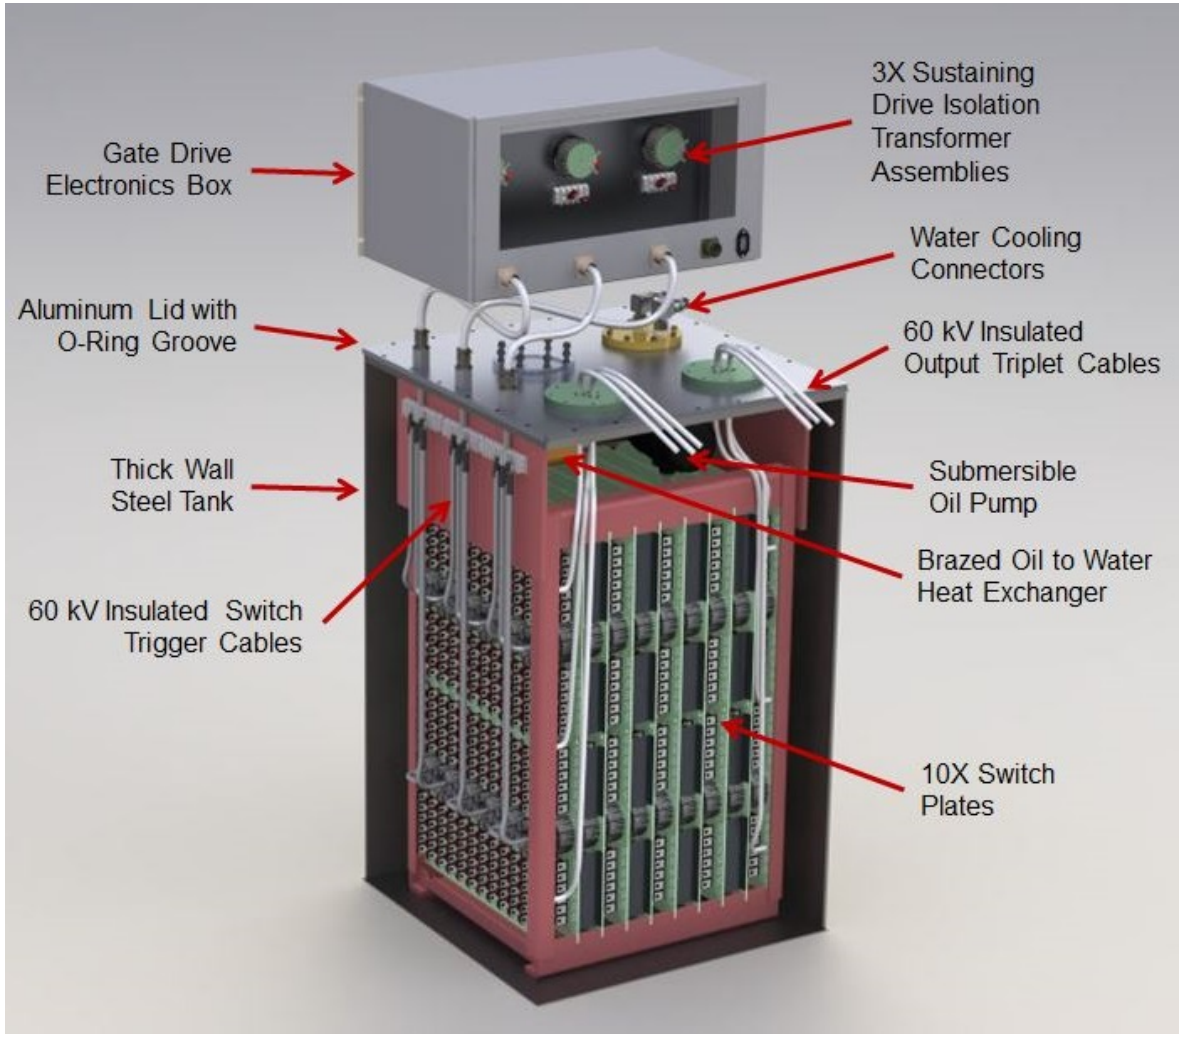

Figure 2. Rear view of DTI's solid-state thyratron replacement switch assembly as proposed. The system contains 10 circuit boards in series, each with 288 IGBTs, and is rated for 48 kV, 6.3 kA peak. Final layout of all interconnections, especially high voltage input and output cables, was optimized in Phase II. Submersible pumps and heat exchangers are standard elements of DTI systems.

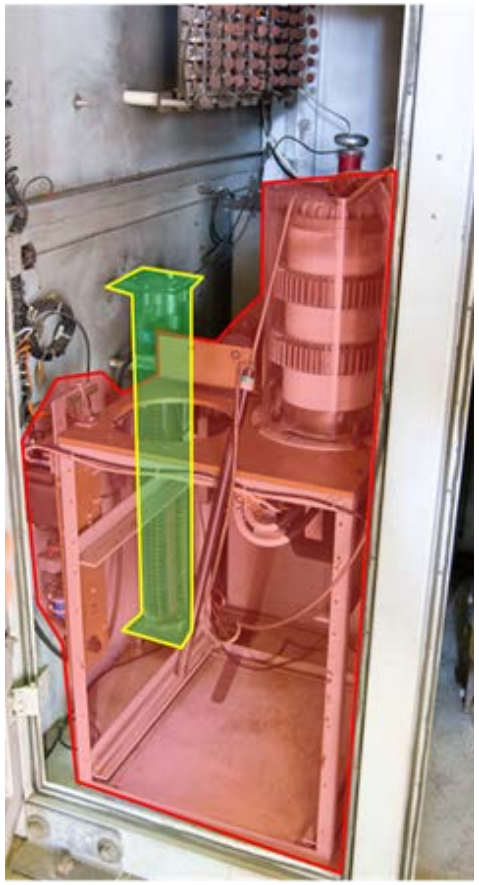

Figure 3. Photograph showing front view of existing SLAC cabinet and area of the cabinet to be replaced. Red highlighted thyratron mounting structure; will be removed entirely. Green highlighted voltage divider, to remain in the cabinet.

3.2 Test and Measurement Results

The Phase I effort aimed to develop an affordable, reliable, solid-state replacement for the klystron modulator thyratrons at SLAC. Explicitly, DTI's primary objective was to develop a solid-state switch meeting SLAC's pulse requirements, fitting into the existing SLAC modulator cabinet, having an expected lifetime of 20 years or longer, and costing well under $100k. The Phase I effort was highly successful, and indicated that each of these objectives was achievable.

In Phase I we evaluated a number of IGBT switching devices, and selected the one with high current capability, short-circuit tolerance, and low cost. We verified that the temperature of this IGBT in operation will be low, giving high reliability. We also tested a series array of 60 devices with the number of faults that they would see in 20 years of operation; there was no sign of damage. We built and tested a 10 × 10 array of IGBTs, which demonstrated operation in both series and parallel, with peak-to-peak jitter less than 1.5 ns. We performed an initial design of the switch assembly, showing that it fits in the modulator cabinet at SLAC. Finally, we estimated the selling price of the switch as designed with the possibility of even lower cost if the current per device can be increased.

3.2.1 Device Selection

The switch developed in Phase I uses IGBTs, which have a fast risetime and can carry substantial currents. Other possible solid-state devices are thyristors and FETs. Thyristors can carry higher currents, but if the current rises quickly (as it does at SLAC), it tends to concentrate in a small region, which can cause a short. FETs, the other alternative, switch rapidly, but cannot carry as much current as IGBTs.

We considered 22 different IGBTs for the switch. The criteria for evaluation were:

- Voltage: 1200 V to give fast turn-on
- Current: >200 A pulsed
- Packaging: TO-247 or TO-264 (largest standard discrete)
- Short-Circuit Tolerance
- Low Cost

Our evaluation concluded that the optimum device, with high current capability, short-circuit rating, and lowest cost is the NGTB25N120LWG from On Semiconductor. Table 1 shows the devices we considered.

Table 1. IGBT Comparison. Peak current based on DTI bench-test measurements, 17 V on gate.

Manufacturer	Device	Ic max (A)	Short-rated	Cost ($)	Comments
Fairchild	FGL40N120ANDTU	550	+	5.56	
	FGH40T120SMD_F155	300	0	3.77	shorted in test
	FGL40N120ANTU	-	+	5.53	no diode
	HGTG18N120BND	-	+	3.59	
Int'l Rectifier	IRG7PSH73K10PBF	700	+	11.15	no diode
	IRGPS60B120KDP	430	+	9.76	
	IRG7PH50K10D	-	+	4.67	
Infineon	IKW40T120T2	240	+	4.70	
	IGW60T120	320	+	4.41	
	IGW40N120H3	-	+	4.57	no diode
	IKW40T120	-	+	5.47	
	IGW40T120	-	+	3.97	no diode
On Semi	NGTB50N120FL	500	+	5.10	
	NGTB40N120IHL	500	+	2.85	
	NGTB25N120LWG	***500***	***+***	***1.76***	***selected***
IXYS	IXYX120N120C3	> 800	0?	15.81	no diode
MicroSemi	APT85GR120L	> 900	+	9.31	no diode
	APT40GR120B	-	+	4.60	
STM	STGW40H120DF2	-	+	2.51	

3.2.1.1 Single-Device Current and Conduction Voltage Measurements

Measuring current is straightforward with a Pearson current monitor. The conduction voltage is somewhat more difficult to measure. It cannot be accurately measured with just a scope probe, since the voltage across the IGBT drops from hundreds of volts (when the IGBT is open) to a few volts (when the IGBT is closed). If the scope sensitivity is set high enough to accurately measure the conduction voltage, the scope will not recover from the high initial voltage in a reasonable time. Alternatively, if the scope sensitivity is set low enough for the initial high voltage, the sensitivity will be too low to accurately measure the conduction voltage.

Instead, the conduction voltage is measured using the simplified test circuit shown in Figure 4. When the IGBT is open, most of the voltage is dropped across the 20 kΩ resistor, and the zener diode limits the scope voltage to 22 V. When the IGBT closes, the voltage across it drops to a few volts, and the zener diode behaves as an open circuit. The 1N4148 diode blocks the charge on the zener capacitance, so the scope can accurately follow the voltage across the IGBT from 22 V down. The RC response time of the circuit, limited by the 20 kΩ resistor and the 17 pF capacitance of the scope probe, is 340 ns.

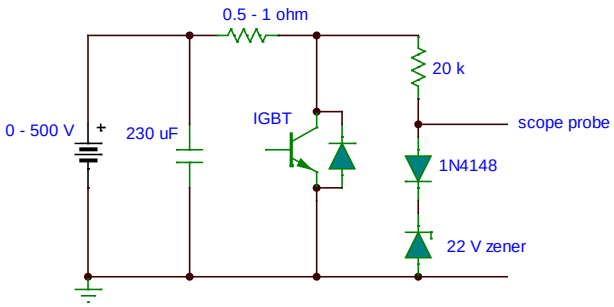

Figure 4. Test circuit to measure conduction voltage; this eliminates overdrive recovery problems with the oscilloscope.

The current and conduction voltage were measured with 17 V on the gate (< 20 V for reliability, >15 V for high current). Figure 5 shows the waveforms for the selected IGBT at 300 A. The steady-state conduction voltage is 5.5 V. The maximum practical current for this device is 500 A, at which the conduction voltage is 8 V. The present switch design uses a device current of 260 A. In future we would examine increasing the current per device, which decreases the number of IGBTs, and so reduces the cost.

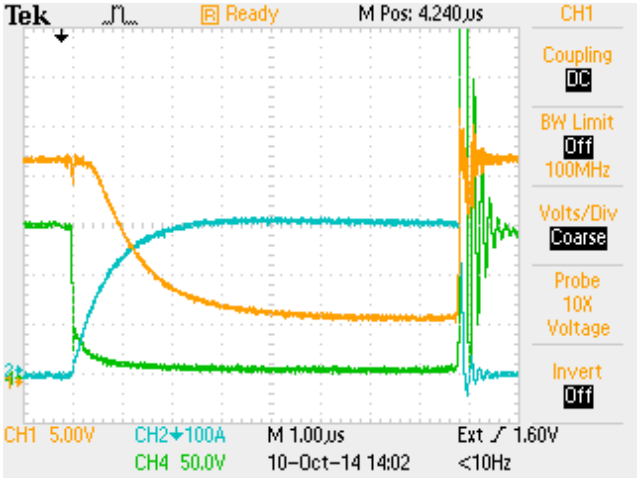

Figure 5. Conduction voltage of a NGTB25N120LWG IGBT at 300 A.
Yellow: Conduction voltage with measuring circuit, 5 V/div. The steady-state conduction voltage is 5.5 V.
Blue: Current, 100 A/div. The peak current is 310 A.
Green: Conduction voltage with scope probe directly across collector-emitter, 50 V/div. Voltage across the IGBT with a fast response time, but lower sensitivity than the yellow trace.

3.2.2 Series-Array Testing

Once the device was selected, 60 devices were connected in series, and tested to demonstrate series operation and ruggedness to arcs at a current of 250 A. The tests were:

- Resistive load, inductance corresponding to normal operation (Figure 6).

  This test demonstrated that the voltages share evenly in a series stack. It also showed that the current rises to its peak value in under 2 µs, as required.

- Arc, inductance corresponding to a klystron arc (Figure 7).

  We performed 300 tests with the switch stack, which corresponds to the number of klystron arcs the switch would experience in 20 years of operation. Based on other DTI work, the IGBTs are expected to survive many more arcs; our IGBT switches are routinely used for klystron conditioning.

- Arc, inductance corresponding to a cable short (Figure 8).

  There were 10 tests for the switch stack. This is more than the number of cable shorts the switch would see in 20 years of operation.

In general, an arc is the most stressful condition in switch operation because the IGBT must open at very high current. Rapid interruption of the current causes substantial transient voltages. The main concern we had was that in these arc tests, the IGBT gates might be damaged. However, there was no sign of this. The trace in Figure 9 shows a typical gate voltage after the arc tests were performed. The gate is charged in 2 µs, and then the drive pulse is removed. Note that the gate voltage trace is flat. Had there been any damage to the gate, there would have been a noticeable decrease in the gate voltage over the next 175 µs.

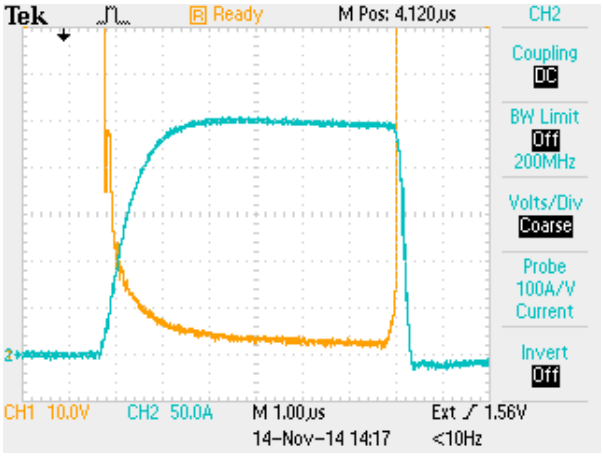

Figure 6. Resistive load, 250 A current.
Yellow: Voltage across 60 IGBT switch stack, 1 kV/div. Note that the voltage settles to 300 V, which agrees with the single-device conduction voltage measurement
Blue: Current, 50 A/div. The peak current is 250 A.

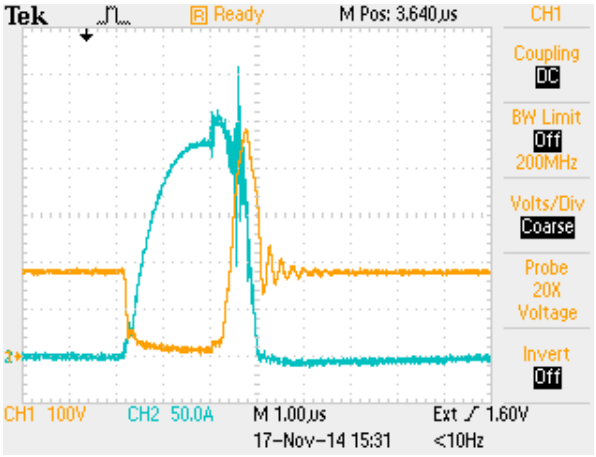

Figure 7. Arc, inductance corresponding to klystron fault; 230 A current.
Yellow: Voltage across 60 IGBT switch stack, 10 kV/div.
Blue: Current, 50 A/div. The peak current before the arc is 230 A.

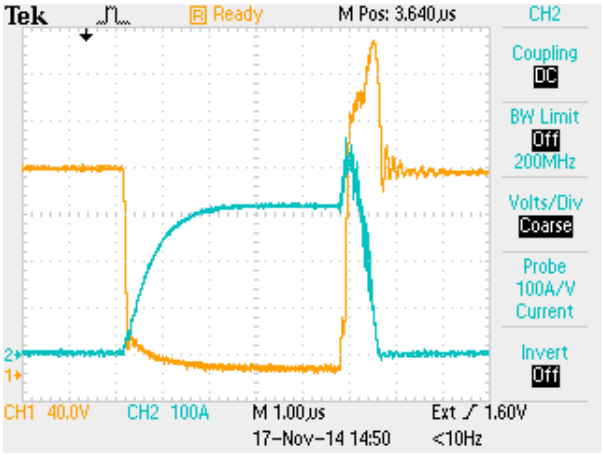

Figure 8. Arc, inductance corresponding to cable fault; 320 A current.
Yellow: Voltage across 60 IGBT switch stack, 4 kV/div. Peak reverse voltage is 28 kV.
Blue: Current, 50 A/div. The peak current is 320 A.

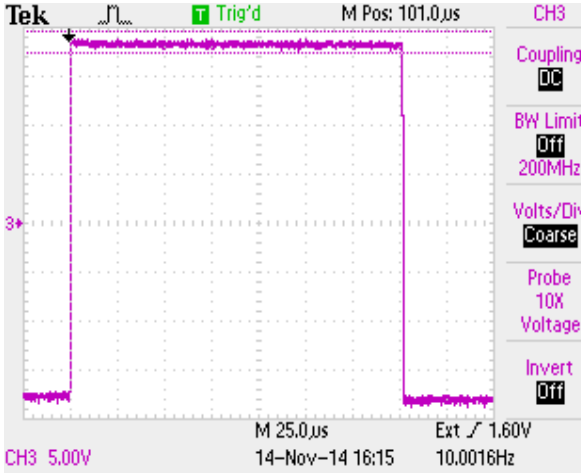

Figure 9. Gate voltage after arc tests; damage would have resulted in a decrease in gate voltage due to leakage current.

To reiterate: the IGBTs withstood a higher number of arc tests than the switch is expected to experience over 20 years of operation at SLAC, and they can be expected to survive far more. We would perform extensive tests in the future to further prove the reliability, as discussed in detail in Section 3.3.2. However, the voltages across the device and across the gate do not exceed their rated values, and the junction temperature will be very low compared to recommended normal commercial operation (see Section 3.3.2.2), so switch reliability should be very high.

3.2.3 Series – Parallel Operation

Operating with both series- and parallel-connected devices is critical to this design. Diversified Technologies Inc. has the proprietary technology to do this, and has fielded well over 500 systems over the past 20 years – including the 500 kV, 530 A NLC modulator delivered to SLAC in 2007[1]. To demonstrate this functionality, an array with 10 parallel by 10 series devices was built (Figure 10) and tested to 2.6 kA (Figure 11).

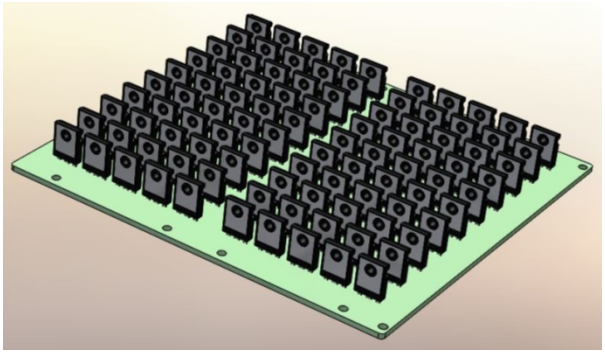

Figure 10. Test board with 10 series by 10 parallel IGBTs; used for data shown in Figure 11.

3.2.4 Jitter

The jitter was measured on the switch with 60 IGBTs. The jitter was derived from the timing difference between the voltage across the switch, and a pulse produced by the trigger generator.

Typical timing variation is shown in Figure 12. Most of the pulses are within 1 ns of each other, and maximum variation is 1.5 ns. This is significantly less than the jitter for thyratrons, typically 5-10 ns. **The lower jitter will increase the HV stability in the accelerator.**

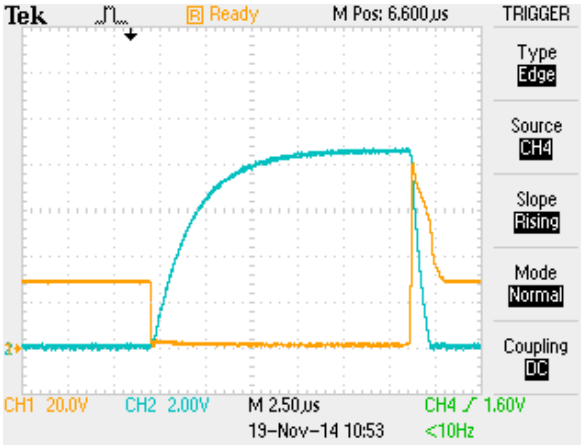

Figure 11. Yellow: Voltage across 10 × 10 array, 2 kV/div. Peak reverse voltage is 8 kV, which is under the 12 kV/device limit (10 devices at 1200 V).

Blue: Current through 10 × 10 array, 600 A/div. The peak current is 2.6 kA.

3.2.4.1 Fault Response

There are two options for the switch to handle a load fault: the switch can either stay closed or open. If the switch stays closed, it needs to carry the full fault current for the 6 µs pulse duration. While the switch is rated for short-circuit operation, this will deposit a significant amount of energy in the devices, and could limit the reliability.

The other option is for the switch to open. In this case it carries the fault current for only 700 ns, and deposits much less energy in the devices. A potential concern with this option is that as the switch opens, the current turns off and on rapidly, which produces a large transient voltage in the emitter-lead inductance (see Figure 13). This transient voltage will couple onto the gate, though it is filtered by the gate resistor and the gate capacitance. The question is whether the gate voltage exceeds its 20 V rating. While it would

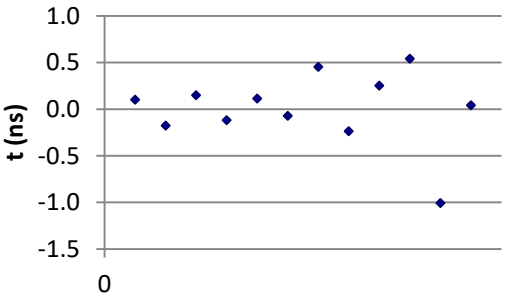

Figure 12. Times between turn-on of the switch and trigger-generator pulse. Most of the times are within 1 ns of each other; the maximum variation is 1.5 ns.

1 Work done under SBIR contracts DE-FG02-99ER82776 and DE-FG02-03ER83637.

be desirable to measure the voltage on the gate directly, this cannot be done since the scope lead is necessarily across the emitter-lead inductance. However, the voltage on the gate can be simulated in SPICE, where the measured gate voltage is filtered by the gate resistor and the capacitance of the gate (While the gate capacitance varies with gate voltage, it is a constant for voltages beyond device turn-on, which is the case as the voltage approaches the 20 V limit.).

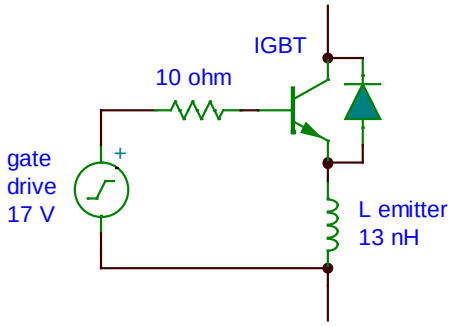

Figure 13. Simplified gate drive connection to IGBT, showing emitter lead inductance.

The simulation results are shown in Figure 14. The measured voltage is the lower waveform; note that the spikes exceed 100 V. The filter gate voltage is the upper waveform. The voltage is under 20 V, except for a 50 ns transient of 23 V. We do not expect that this transient will limit the reliability. However, we would run extensive arc tests in future to confirm that this is the case.

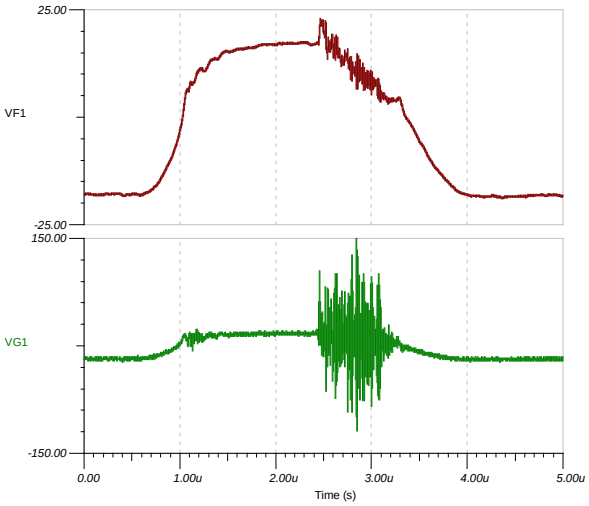

Figure 14. Top: Measured voltage, with filtering calculated for the gate resistor and the gate capacitance. Bottom: Measured voltage between gate and emitter leads during a fault.

Another potential issue is the modulator operation after the switch opens, leaving the PFN partially charged. The next pulse will be at a lower voltage. This is because the resonant charging system rings up into the PFN, increasing the PFN voltage by twice the difference between voltages on the power supply and the initial PFN voltage. The power supply is at 24 kV; if there is any voltage on the PFN, it will ring up to less than the normal 48 kV. All the subsequent pulses will be at normal voltage. We do not expect that this will be a concern, but we will discuss the issue with SLAC. If having a low voltage is a concern, the switch could remain closed in a fault.

3.3 Switch Design

3.3.1 Overview

The preliminary design of the circuit board (“switch plate”) used in the switch is shown in Figure 15. The plate has 4 groups of IGBTs, each with 6 devices in parallel and 12 devices in series.

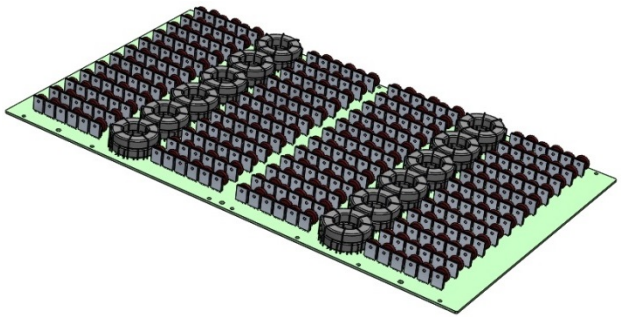

Figure 15. Proposed IGBT circuit board. The board has 4 groups of 12 series by 6 parallel IGBTs; the ferrite toroids couple the drive pulses that control the IGBTs. There are ten of these boards in the thyratron-replacement switch.

In the thyratron-replacement switch, the current per device is 260 A. At this current, and the 6 μs pulse width and 120 Hz frequency at

SLAC, the IGBTs dissipate only 1 W per device. At this power, the IGBT back surface provides sufficient cooling in oil, so the IGBTs do not need heat sinks and can be closely packed.

Since there are 2880 devices in the switch assembly, the total power dissipation is 2880 W. To remove this power, there is an oil-circulation pump and an oil-to-water heat exchanger, which are located in the oil above the switch stack. This is standard practice at DTI.

The IGBTs are gated on and off by pulses coupled by ferrite toroids. The drive wires (shown in Figure 2) pass through the center of the toroids, and have 60 kV solid insulation to prevent breakdown between the drive and the high voltage on the switch. An additional 0.5 inches of oil separates the ferrite secondary winding at high voltage and the solid insulated trigger wire at ground potential. The wires are connected in pairs to minimize the inductance effects of the drive loop.

The complete switch assembly is designed to fit within the 22” wide × 22” deep × 37” high dimensions of the modulator tank (Figure 16). The switch has ten switch plates, and 120 rows of devices in series. This is 20% more than the minimum required, so multiple devices must fail before the switch shorts. Note that IGBTs always fail short, allowing a series-string with margin to maintain full functionality until such a time as the margin is exceeded (“graceful degradation”).

The switch is immersed in oil to conduct heat from the IGBTs and insulate the high voltage. The oil spacings are designed conservatively, with a bulk field strength of 30 kV/in, and a tracking field strength of 15 kV/in.

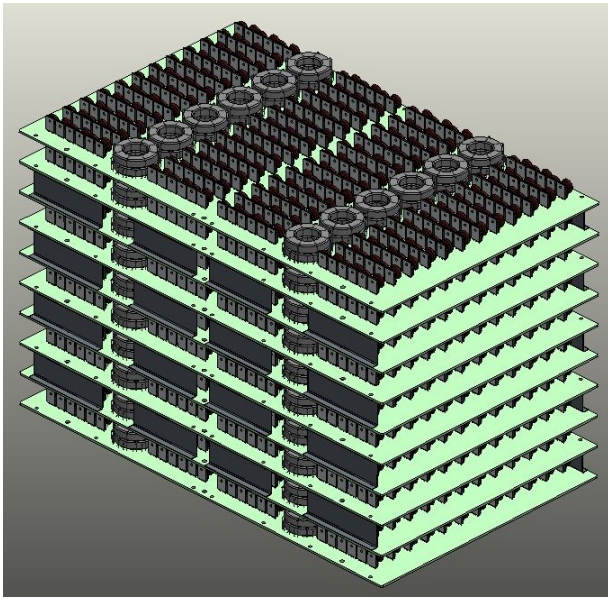

Figure 16. Proposed full switch assembly composed of ten 288 IGBT circuit boards in series, for 2880 total devices. The boards will be held in place by a simple card cage in the oil tank seen in Figures 1 and 2.

3.3.2 Reliability

To reduce the cost, the switch should use as few IGBTs as possible, consistent with reliability. As the number of IGBTs is reduced, either the voltage or current per device will increase correspondingly. The voltage is near its maximum value, and cannot be increased further. However, we anticipate that the device current can be increased above the present 260 A design value, perhaps to as much as 500 A. The limit is set by reliability.

The IGBT lifetime is limited by the voltages on the device and the gate, and the temperature, both steady-state and its variation. We have designed the thyratron replacement switch taking these limits into account; thus, we calculate the mean time to failure to be well over 79 years. The following sections discuss these limits.

3.3.2.1 Device Voltage Limit

A key issue with IGBT operation is that the device voltage cannot exceed the maximum rated value, which is 1200 V for the device chosen here. Since the devices are series-connected in a stack, if the voltage is not divided evenly, in operation the devices will successively fail until the entire stack shorts. While series-connection is generally not straightforward, DTI has a proprietary technology to enable this, and a long history of successfully building series-connected IGBT switches.

3.3.2.2 Steady-State Temperature Limit

The peak current limit specified on the data sheet is for a time of 1 ms. Since the 6 μs pulse here is much shorter than that, the current can be higher than the data-sheet value without the temperature rising excessively. The peak current is determined by the power dissipation, which depends on the conduction voltage. The conduction voltage measurement is described in Section 3.2.1.1. We established that the maximum current for the selected device is 500 A. The design presently uses 260 A; this may be overly conservative and we will reexamine this design in the future.

The dissipated power in the IGBTs is predominantly due to the conduction loss (see Section 3.2.1.1):

$P = 5.5\text{ V conduction} \times 260\text{ A} \times 6\ \mu\text{s} \times 120\text{ Hz} = 1.0\text{ W}.$

The switching losses are small, and can be ignored in comparison to the conduction loss. There is little turn-on loss, since the system inductance keeps the current low during turn-on. There is also no turn-off loss, since the pulse is terminated by the pulse-forming network.

The coefficient for thermal resistance temperature increase in circulating oil is 25°C-cm^2/W, and the metal backing of a TO-247 package has an area of 2 cm^2. Hence the device temperature rise is:

$\Delta T = K_{thermal}\ P / A = 25°\text{C-cm}^2/\text{W} \times 1\text{ W} / 2\text{ cm}^2 = 13°\text{C}.$

This gives a device junction temperature of:

$T = 35°\text{ water} + 5°\text{ heat exchanger} + 1°\text{ junction-case} + 13°\text{ rise} = 54°\text{C}.$

The manufacturer specifies a failure rate of 10 FIT (failures in time, the number of failures in 10^9 hours) at 55°C. This is only 1°C different from the calculated temperature, so we can use the specified failure rate of the device to calculate the switch lifetime. This gives the switch an estimated mean time to failure greater than 79 years (see Section 3.3.2.4).

3.3.2.3 Temperature-Variation Limit

Temperature variation can also affect reliability. Since the temperature variation here is only 6°C, however, it does not limit the reliability. The discussion of this follows.

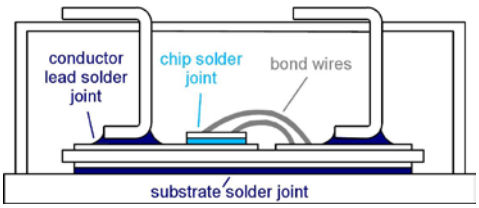

Figure 17. Potential failure locations in an IGBT module.

An ABB application note[2] states that the failures due to temperature variation occur at the joints; at the bond wires, chip, insulating substrate, and conductor leads (Figure 17). The critical joint that needs to be considered here is at the bond wires: temperature cycling produces an immediate plastic deformation at the bond wire joint, independent of the cycle time. For all the other joints, the cycle time needs to be in seconds to have an effect; this is much longer than the 8 ms temperature cycling produced at 120 Hz.

We can calculate the transient temperature rise as due to conduction loss, since the switching losses can be ignored:

$$\Delta T_{transient} = V_{conduction}\ I\ R_{thermal\ transient}$$

$$= 5.5\ \text{V} \times 260\ \text{A} \times 4\times10^{-3}\text{°C/W} = 6\text{°C}.$$

For these numbers, $V_{conduction}$ was measured (see Section 3.2.1), and $R_{thermal\ transient}$ is from the IGBT data sheet. The ABB application note gives the data[3] for lifetime associated with joints at the bond wire versus temperature variation; this is plotted in Figure 18. The bond wire lifetime for the 6°C variation and 57°C maximum junction temperature is off the plot, and it is clear that the lifetime will be longer than the 8×10^{10} pulses in a 20 year lifetime at 120 Hz. Hence, failure due to temperature variation is not a concern here.

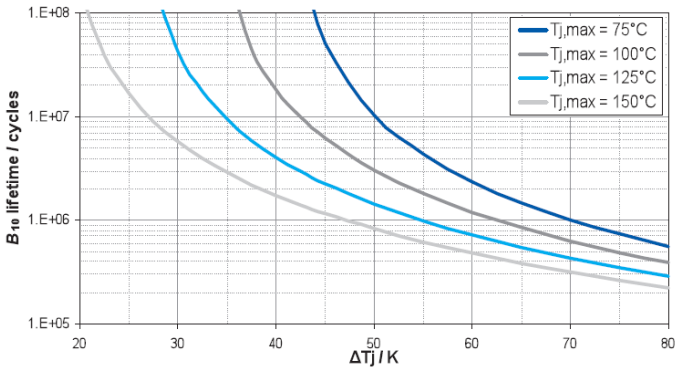

Figure 18. Number of cycles for 10% failures due to a failure at the wire bonds for ABB IGBTs. For the case here (6℃ temperature rise, 57℃ maximum temperature) the lifetime will be very long compared to the 8×10^{10} pulses in 20 years estimated at SLAC.

3.3.2.4 Reliability of Switch Array

We calculated the reliability of the switch array to be well over 79 years. The calculation of this follows.

The failure rate of a single device, as given by the manufacturer, is 10 FIT (failures in time, failures in 10^9 hours). Note that since this is a very small number, the failure rate for multiple devices is approximately additive.

2 Load-cycling capability of HIPak IGBT modules , Figure 1, http://search-ext.abb.com/library/Download.aspx?DocumentID=5SYA%202043-04&LanguageCode=en&DocumentPartId=&Action=Launch.

3 Figure 12 in the application note.

To get an upper bound on the failure rate of the entire switch, which has 2880 devices, let's consider the case where all 24 devices in a plate row (see Figure 15) are connected together, rather than being in 4 groups. A single failure in one of the devices shorts the entire row, so this increases the failure rate over a single IGBT by a factor of 24. In addition, the switch array has 120 of these rows connected in series, giving another factor of 120 to the failure rate. However, the switch has 20 more columns than necessary, so there must be 20 device failures for the switch to fail; this reduces the failure rate by a factor of 20. The net failure rate of the switch array is then

F = 10 FIT / device × 24 × 120 / 20 = 1400 FIT.

This gives a lower bound on the mean time to failure of

MTTF = 1 / F = 1 / (1400×10^{-9}/hr) = 6.9×10^{5} hr × 1 year / 8770 hr = 79 years.

The actual MTTF will be longer than this. This is because the switch has 4 groups of 6 devices in parallel, rather than a single group of 24 devices. If the failures in the switch were evenly distributed, the failure rate would decrease by a factor of four, since each IGBT failure shorts out only 6 devices instead of 24. However, the switch failures will not be perfectly distributed, so the reduction in failure rate will be somewhat less than a factor of four.

3.4 Reliability of the Switch Assembly

The other components making up the switch assembly will contribute to the composite failure rate. The reliability of these components was not assessed during the Phase I, but will be quantified in the future. Based on over 500 high-voltage, solid-state systems delivered by DTI in the last 19 years, we anticipate reliability to be extremely high. This reliability comes from conservative design, with substantial margins for the device voltages and temperatures.

Anecdotally, there are two DTI systems which run nearly continuously, allow us to report their results, and provide the best operating data. One is at the Sondrestrom Research Facility in Kangerlussuaq, Greenland, operated by SRI International. This modulator drives a radar for ionospheric research, and delivers pulses at 130 kV, 300 A, and 1 kHz. This modulator has been in continuous service since October 2003, and has had only a single minor failure (a gate drive board, which was replaced in less than one hour) over this time. The other modulator is for Iotron, a commercial customer in Canada doing radiation processing. The modulator delivers 150 kV, 100 A, up to 300 μs pulses at a pulse rate of up to 300 Hz (limited to 750 kW average power) on a 24/7 basis. This modulator has been in service since August 2007 without any failures.

3.5 Accelerated Testing

The calculated switch lifetime is so long that it cannot be established in a few months of testing under normal operating conditions. To accelerate the failure rate, lifetime testing would be done at an elevated switch temperature, and the higher failure rate produced can be scaled to the lower failure rate at normal operating temperature over a longer time. The failure rate scales as an Arrhenius relation,

$$F = F_0 \exp\left[E_{activation} \left(1/T_0 - 1/T \right) \right]$$

where the temperatures are in °K; the constants F_0 (the failure rate at temperature T_0) and $E_{activation}$ are from manufacturer's data.

The manufacturer of the chosen device gives an IGBT failure rate of 10 FIT for IGBTs at 55°C, and an activation energy of 0.7 eV. (Note that 55°C = 328°K, 150°C = 423°K, and 1 eV = 11,600°K.)

The device failure rate at 150°C is then

F (/hr) = 10×10^{-9} exp [0.7 × 11,600 × (1/328 - 1/423)] = 10×10^{-9} × 260 = 2.6×10^{-6} /hr.

The average number of device failures in 60 days of testing is then

N = 2.6×10^{-6} / device-hr × 2880 devices × 24 hr/day × 60 days = 11.

This is not enough failures to cause the switch stack to fail, but is sufficient to extrapolate to the failure rate for the entire switch. The device failures will be quickly and easily measured with the monitor described in Section 3.5.1.1, below.

There are several ways to increase the device temperature for accelerated testing.

- The flow of cooling water can be restricted, increasing the temperature of the ambient oil and so of the IGBTs.
- The pulse frequency can be increased, using a resistive load. However, this will require both substantial power from the supply, and high power dissipation in the load.
- The switch array could be divided into two halves, and run as a half-bridge to minimize the power required, as shown in Figure 19.

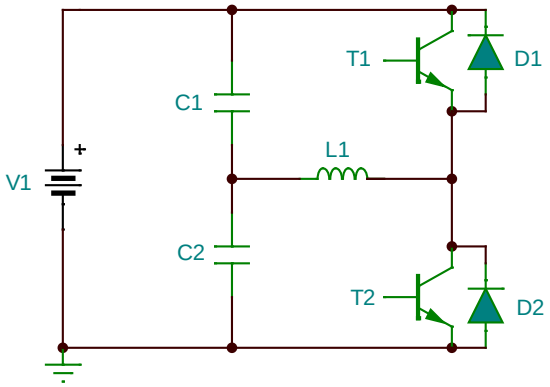

Figure 19. Half-bridge topology for accelerated life testing. Though the switch current is a triangle wave rather than the square wave in modulator operation, this test requires only average power to compensate for losses, and can thus run at high repetition rate for accelerated life testing.

The half-bridge circuit operates as follows. Initially T1 (which represents half the switch stack) closes, and current builds up in the inductor. When the current is at the peak value desired, T1 opens, and the inductor discharges into C2 through D2. When the current has ramped down to zero, T2 (the other half of the switch stack) closes, then opens. The sequence then repeats, transferring charge back and forth between C1 and C2. While this has the disadvantage that the current waveform is a triangle instead than a square wave for the actual modulator, its substantial advantage is that the test requires little average power. Since there is dissipation only in the switch and in the resistance of the inductor, not in a load, the test can run at a high repetition rate for accelerated life testing.

3.5.1.1 Monitoring Number of Failed Devices

The switch stack is built with redundant devices, so that a number of them can short while still permitting operation of the switch (the devices always fail as a short). Knowing how many devices have failed gives confidence in the reliability of the switch, and allows scheduling of switch maintenance.

The concept for the device failure monitor is shown in Figure 20. Each IGBT has a high-impedance resistor across it. The more IGBTs that are shorted, the lower the resistance of the switch array. Hence, the number of good devices remaining can be inferred by knowing the applied voltage, and measuring the leakage current when the IGBTs are off.

In the actual switch, there will be a resistor across each group of IGBTs in parallel, rather than across each device.

This monitor is particularly useful in testing, since it can determine the number of good devices in a few seconds, in-between pulses. It has been used on several previous high-value systems produced by DTI.

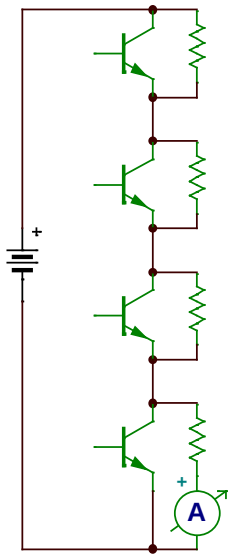

Figure 20. IGBT failure monitor. The number of failures is inferred from the total resistance (switch leakage current) when the switch is open.

3.6 Gate Drive

In the Phase I proposal we assumed an ultra-fast gate drive would need to be developed to reduce the switching losses. However, this is not necessary since the switching losses are small compared to the conduction loss. Instead, the gate drive for the thyratron-replacement switch can use existing DTI technology.

3.7 Evaluation for Outside Applications

Either thyratrons or IGBTs can be used in line-type modulators, since the fixed-width pulse from a PFN does not require opening under load. We anticipate that a switch that opens during an arc will increase the klystron life, since the fault current is interrupted sooner, and less energy will be deposited in the klystron.

Other applications, such as high power (utility-scale) conversion and modulation, demand opening under load. DTI believes the technology developed under this SBIR will be immediately extensible to these applications, and anticipates the switch boards to cost less than a third of our existing switching technology.

4.0 SBIR Phase II (Year 1)

These SBIR Phase II activities occurred in the period of March 6, 2015 – January 21, 2016.

4.1 Summary

The primary objective for Diversified Technologies, Inc. (DTI) in Phase II SBIR was to design and build a solid-state thyratron-replacement switch that meets the performance specifications for the Stanford Linear Accelerator, given in Table 2.

Table 2. Modulator requirements for Stanford Linear Accelerator

Parameter	Requirement
Voltage	48 kV
Current	6.3 kA
Pulsewidth	6 µs
Frequency	120 Hz
Fault Current	18.7 kA
Jitter	5 ns

Meeting the electrical performance requirements is straightforward for DTI. The major issue for this project is to meet these requirements at a life-cycle cost competitive with thyratrons, which need to be replaced every 8 - 18 thousand hours. Thyratrons cost $13 k; hence the replacement switch needs to sell for approximately $30 - 40k (or less if possible) to achieve the investment payback required at SLAC.

We began this project with the idea that the switch would be based on DTI's standard technology. Soon after we started this project, we realized that the switch does not have to open, even in the event of a short, since the peak current and pulse duration is limited by the pulse-forming network. We changed the switch to use a resistor-capacitor-diode (RCD) snubber, which allows the devices to operate at nearly twice the voltage of a switch with DTI standard technology; this would significantly reduce the cost of the switch.

We spent a fair amount of time working on this approach. However, we had not considered the consequences of a short in the output cable. The issue is that the peak current into a short is three times the normal current, so the IGBTs could go into current limiting. Current limiting is made even more likely in a cable short, since the *dI*/*dt* increases the voltage on the emitter lead of the IGBTs, reducing the gate voltage. If the IGBT goes into current limiting, the system inductance drives current into the snubbing capacitors, which will charge and over-volt the IGBTs. To ensure that the switch does not go into current limiting, we could put more devices in parallel. However, it is not certain how many devices would be needed.

While it is possible that a switch with RCD snubbers will have fewer devices, it is risky. Because of this risk, we reverted to the standard DTI technology, which is highly robust for shorts.

4.2 Accomplishments

To produce the thyratron-replacement switch we performed the tasks described in this section. Progress is as of the end of Year 1.

4.2.1 Task 1: Confirm Specifications with SLAC

We confirmed the initial performance specifications, which have not changed. They remain at 48 kV, 6.3 kA, 6 µs pulse width, 120 Hz, 18.7 kA fault, and 5 ns jitter. SLAC has sent us drawings of the switch cabinet, which we have used to model the switch assembly under Task 3. Water cooling is available.

4.2.2 Task 2: Small-Scale Device Testing

In Phase I we had identified the NGTB25N120LWG as the prime candidate device. We reexamined the devices, and found one other worth considering: the NGTB20N120LWG. It has the same rated peak current as the NGTB25 part, 200 A, and the same gate charge, suggesting that it might actually be the same device. While the price is lower, the difference is small: $1.68 (in quantity 1000) compared to $1.85. This would mean a cost difference of about $500 for the complete switch. We decided that this difference is not large enough to be worth investigation.

4.2.3 Task 3: Switch 1 Design, Build, and Test (at DTI)

This task is the major design effort in this work; the other switch designs, Tasks 4 and 5, will be modifications of this design. Most of the work in this project to date has been associated with this task. The switch design is constrained in size, since it needs to fit in the existing SLAC modulator cabinet, and in price, which cannot be much more than about $30 k (twice the cost of a thyratron) for SLAC to consider it for purchase. These constraints make the design significantly more difficult than if this was a ‘clean sheet’ switch design.

4.2.3.1 Snubber

There are generally two types of switches: those that close under load, but cannot open (such as a thyratron or SCR), and those that can also open (such as an IGBT). A switch that opens under load has two advantages: it can produce variable pulse widths, and the switch can interrupt a fault current when the RF tube shorts. All of DTI’s modulators to date have been built with switches that open.

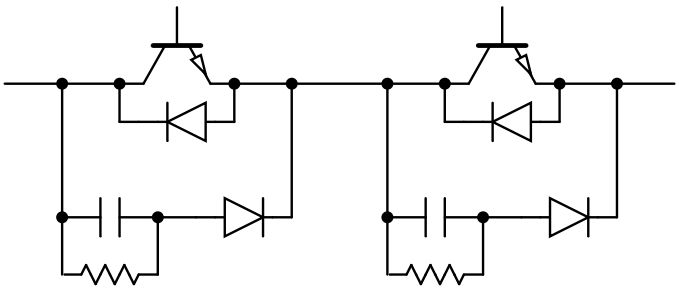

Figure 21. Two series transistors with resistor-capacitor-diode snubbers.

The SLAC modulators use thyratrons and pulse-forming networks (PFNs). The pulse width is fixed, and the klystrons can easily survive a short, so opening under load is not required. Hence, the thryatron-replacement switch could be built either with an opening capability, or without one (if the switch itself can handle the full fault current).

There are tradeoffs for both of these approaches. The DTI opening switch design is rugged, and a low-risk approach to the thyratron-replacement switch. The tradeoff is that the IGBTs in the opening switch operate at a small fraction of their rated voltage – as low as 40% – in order to withstand the voltage spike upon opening, so that the switch built with this technology would need more devices in series (as many as 108 IGBTs in series).

If the switch does not require opening capability, then the IGBTs could use resistor-capacitor-diode snubbers (Figure 21) to handle the variation in turn-on time of the different IGBTs. The snubber capacitor carries current from the other devices until its IGBT turns on. This current is small, since it starts from

zero (Note that snubbers do not work well during turn-off, since the snubber needs to briefly carry the full IGBT current.).

A switch with snubbers has the strong advantage because the devices can operate at a higher fraction of their rated voltage, compared to a switch with opening capability, and so it would have fewer devices in series (as few as 64, or nearly half as many individual IGBTs required for the opening switch).

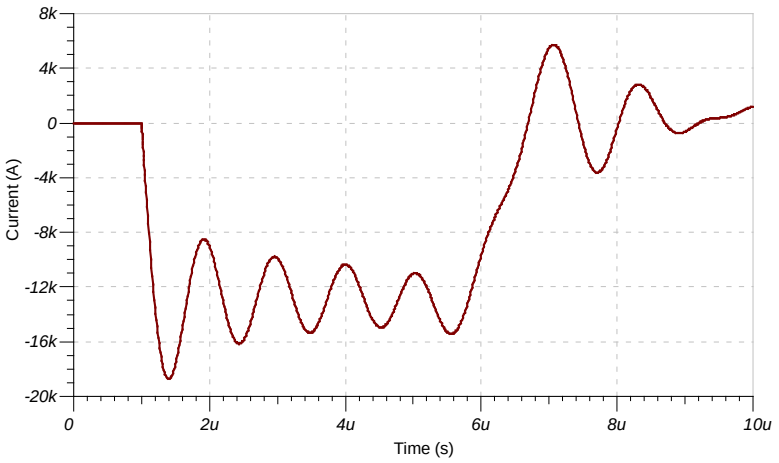

Figure 22. SPICE simulation of worst-case fault current into a cable short. The peak current is 19 kA.

However, a switch with snubbers cannot go into current limiting. Once the current rises above the IGBT limit, the snubbing capacitors would carry the remaining current. The voltage across the capacitors would quickly rise, exceeding the IGBTs device rating, so the IGBTs would short. To prevent the snubbed switch from going into current limiting, it needs more devices in parallel than an opening switch. The worst case is a short in the output cable (Figure 22), where the peak current is 19 kA. The problem is exacerbated by the current rise during a fault. The *L dI*/*dt* voltage drop in the emitter lead of the device reduces the gate voltage, which in turn reduces the limiting current. It is not clear whether the snubber switch will have fewer devices than a standard switch. However, **the snubber approach is risky.**

Because of this, we are using DTI's standard technology.

4.2.3.2 IGBT Voltage and Current

The IGBTs in the opening switch will operate between 400 and 450 V per device, including additional devices in series for margin. At 444 V per device (the expected value) the 48 kV switch needs at least 108 devices in series. By making a circuit board with 16 devices in series (7.1 kV rating), the switch would need seven boards in series, giving a total 112 series devices.

The current per device is limited by the temperature. We can estimate the temperature from the waveforms shown in Figure 23; note the peak current is 320 A, and the voltage is across 60 devices in series. Integrating the power (the product of voltage and current) gives the energy dissipated in the IGBT, shown in Figure 6. The sharp rise in energy at 5 µs is due to the IGBT opening in this test; in normal operation the IGBT does not open. The pulse is 4 µs long. To estimate the energy coupled during the 6 µs pulse, the 10 mJ energy in 4 µs scales to 15 mJ. Hence, the power dissipated at 120 Hz will be:

P (320 A) = 15 mJ × 120 Hz = 1.8 W.

From this, we can calculate the temperature rise of the IGBT in oil:

ΔT = 1.8 W × 50°C-cm^2/W (for still oil) / 2 cm^2 (area of a TO-247 package) = 45°C.

For an ambient oil temperature of 25°, the IGBT junction temperature is only 70°C, which supports a long switch lifetime.

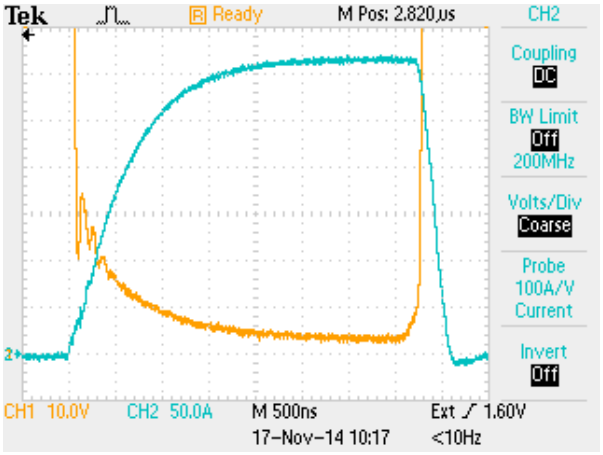

Figure 23. IGBT operation at 320 A. There are 60 devices in series and a resistive load.

Yellow: Voltage across IGBT stack, 1 kV/div.

Blue: Current, 50 A/div.

The switch will initially have 32 IGBTs in parallel. This gives a current per device of only 200 A, rather than 320 A used for the 70°C estimated temperature. In future testing, we will consider reducing the number of parallel devices to further lower the cost of the switch.

4.2.3.3 Fault Tolerance

The switch consists of a large series-parallel array of IGBTs. The switch needs to be redundant, so that individual devices can fail without adversely affecting the switch operation.

To do this, the switch has more devices in series than the minimum required. As a result, when a device shorts, the voltage on the other devices increases, and the switch continues to function.

The switch is divided into four parallel strings, each taking one-quarter of the current. This is because when a device shorts, it carries the total

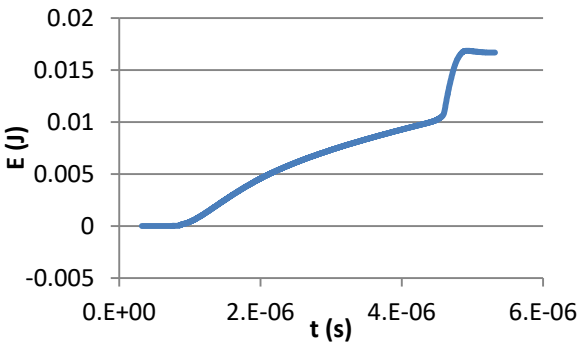

Figure 24. Energy per device at a current of 320 A. The peak is 10 mJ just before the switch opens.

current of the devices it is in parallel with – and this current needs to be small enough that the shorted device does not rupture. There are two limits: RMS and pulsed. The RMS current per string is:

$I_{RMS} = I_{peak} \times \sqrt{(t\,f)} = (6.3\text{ kA} / 4) \times \sqrt{(6\ \mu\text{s} \times 120}$ This is less than the 50 A steady-state current rating of the IGBT, and so it will be fine.

The pulsed limit is that the peak fault current should not melt the two aluminum bond wires in the IGBT, which are 15 mils in diameter.

The temperature rise in the bond wires is given from: $E = m\,C_p\,\Delta T = I^2\,R\,t$

For a current of 4.8 kA (one quarter of the 19 kA peak fault current) and a time of 6 µs; this temperature rise is only 29°C, so the bond wires will not melt.

Another reason to divide the current into multiple parallel strings is reliability. If an IGBT gate shorts, it shorts out the other devices in the string – which pushes the design towards a larger number of parallel strings. However, adding strings means more windings on the gate-drive transformers, which makes the circuit board more complex. As a compromise, we have settled on the current being divided into four parallel strings.

4.2.3.4 Circuit Board Layout

A candidate circuit-board arrangement is shown in Figure 25. The two halves of the board are independent. Since the switch is made with two boards in parallel, each board half carries one-quarter of the total current.

The IGBT gates are driven by toroidal transformers in the center of the board. The primary windings for these transformers are made by a single-turn cable that passes through the center of the toroid. The cable is insulated, and is rated for the 48 kV switch voltage. The secondary windings are single turns, which minimize the inductance, giving a fast rise time for the gate voltage.

The board operates at 7 kV and 3.1 kA, so the thyratron replacement switch needs seven boards in series, and two boards in parallel.

We are in the process of fitting the board design into the SLAC modulator cabinet. The final board design will be similar to this, with transformers centered between rows of IGBTs; the board size and layout details are in progress.

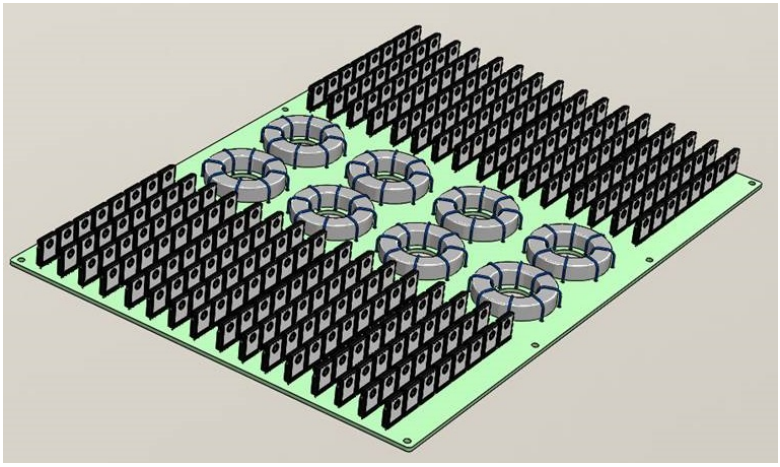

Figure 25. Candidate circuit-board layout. The IGBTs are on the left and right sides of the board, and gate-drive transformers are in the center.

The previous board design is shown in Figure 26; it has three levels of transformers. The first-level transformer is at the top of the board; it is driven by insulated cables that pass through the center of the toroids. It drives the six E-I transformers at the second level. In turn, these second-level transformers drive the eight third-level transformers that distribute the voltage to the IGBTs. The IGBTs had resistor-capacitor-diode snubbers.

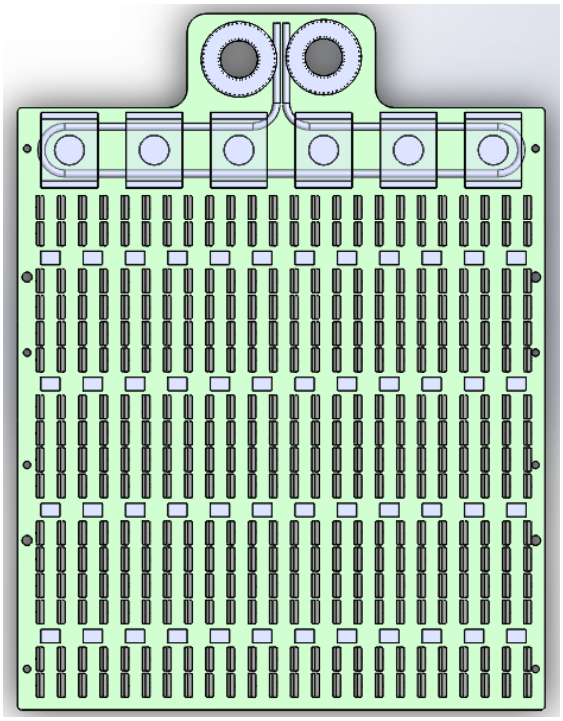

Figure 26. Previous circuit-board layout.

This design had the problems of substantial leakage inductance in the first- and second-level transformers which slows the rise, and insufficient voltage hold-off in the third-level transformers in the event that one of the IGBTs shorts.

4.2.3.5 Gate-Drive System

The gate-drive system design is in progress. It uses one or two driver boards that are the latest iteration in a series of boards that have been in use at DTI for over a decade. The board uses IGBTs rated at 650 V and can deliver pulses at 300 A. The boards drive multiple insulated cables which pass through the toroids on the circuit-board. There are eight cables in the design shown in Figure 25. These cables isolate the driver boards from the 48 kV switch. Having multiple cables reduces the net load inductance, and so gives a fast turn-on for the IGBTs. The cables pass through ferrite toroidal transformers (seen in Figure 25), which couple the drive to the IGBTs.

A concern for this design had been that the IGBTs in the switch plate are not the same distance from the toroids, so the drive inductances are not the same; creating gate voltages with different rise times. However, the effective capacitance of the IGBT gates is small, so the difference is insignificant.

To develop the basic scaling of the gate-drive system, we made a simple model that treats the IGBT gates as capacitances (Figure 27). The components were combined in series and parallel to give the driver

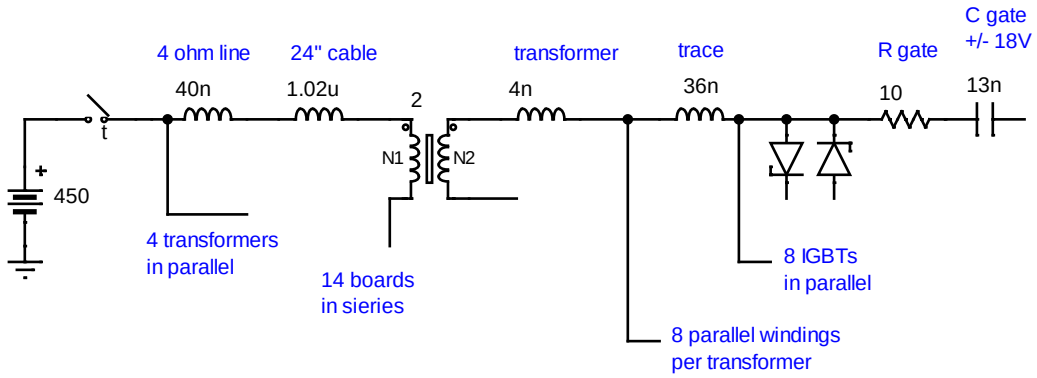

Figure 27. Simple model of the gate-drive system.

current, and the sum of the gate voltages on the IGBTs (Figure 28). The peak gate-driver current is 430 A, and the gate voltage swings from full off to full on in 420 ns. This time is roughly twice as fast as present DTI drives.

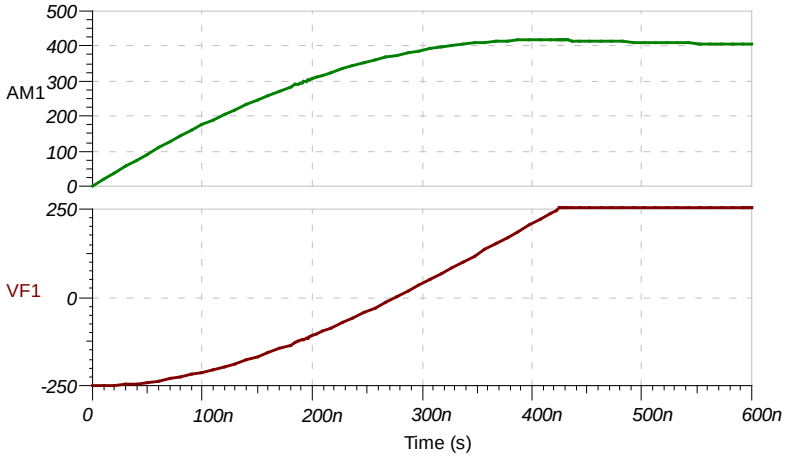

Figure 28. Gate driver current, and sum of IGBT gate voltages. The peak current is 430 A, and the gate voltages rise from full-off to full-on in 420 ns.

4.2.4 Task 4: Switch 2 Design, Build, and Test (at SLAC)

We modeled the SLAC modulator cabinets, using drawings supplied by SLAC, and made an initial design of the switch with the resistor-capacitor-diode snubber. This switch is shown in Figure 29, and is

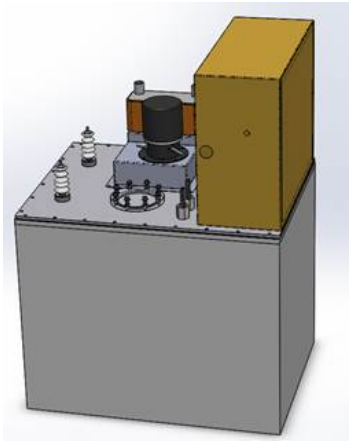

Figure 29. Initial mechanical design of the switch, with resistor-capacitor-diode snubber.

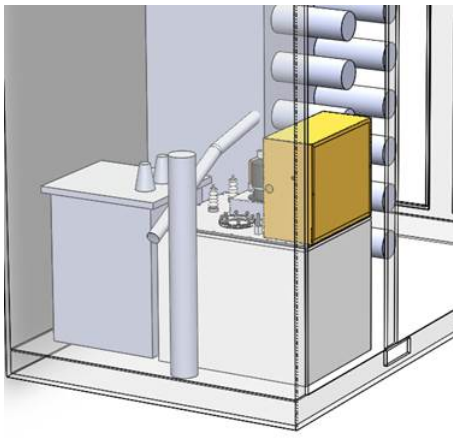

Figure 30. Initial switch design, placed in modulator cabinet.

shown in place in Figure 30. The largest component is the gray oil tank, which holds the switch plates. The yellow box contains the gate-drive boards and the controls. On the top of the tank, at the left, are the high voltage feed-throughs, in the center of the tank are the oil circulation pump and heat exchanger. This switch had a resistor-capacitor-diode snubber, which gave a small volume. However, we have set this approach aside, and are redesigning the switch for the standard DTI configuration, which has much lower risk.

We combined Tasks 2 and 3, so that we will build one switch.

4.2.5 Task 5: Switch 3 Design, Build, and Test (at SLAC)

No activity.

4.2.6 Task 6: Cost Reduction Study

Reducing the switch cost has been part of the design effort. As well as minimizing the number of components, we are considering how to minimize the cost of assembly. This means making the gate circuits with surface-mount components which can be automatically assembled, and using single turns on the gate drive transformers, which minimizes the level of manual labor that will be required to build each thyratron replacement switch.

4.2.7 Task 7: Production and Commercialization Planning

The immediate motivation for this SBIR is to develop a thyratron-replacement switch with the lowest cost possible, optimized for use at SLAC.

A second motivation for this SBIR is to use the thyratron-replacement switch for commercial applications. This switch is made with many discrete IGBTs connected in series and parallel, unlike DTI's present design, which uses IGBT modules in series. The commercial switch will have a faster rise, and will be cheaper as well. We have identified a number of commercial applications, and expect to build commercial products/prototypes in parallel with the version built for SLAC.

The principal difference between the two switches is that the SLAC switch does not need heat sinks, since the duty cycle is only 0.07% (6 μs × 120 Hz), while the commercial version needs heat sinks because its duty cycle is much larger (Both switches are immersed in oil, which gives much better cooling than air.).

There are many heat sinks we can choose from; four of the eight candidates are shown in Figure 31. To choose the heat sinks for the commercial switch, we cannot use the thermal resistance values listed on the data sheets, since these are specified for air, and the switch is immersed in oil. Instead, we measured the heat sink performance as follows.

Figure 31. Four heat sinks. From left to right these are the DTI standard heat sink; C247-025-1AE for the commercial switch, C40-058-AE, and 4-101007U.

A thermocouple was potted into the center of a 3/4" square copper interface block 0.06" thick. In each test, this block was placed between one resistor and the heat sink. Ten ohm TO-218 package resistors

were used for heating, either singly or in series. In each test, approximately 1.73 A was applied, with a 17.3 V drop for single resistor tests, and a 34.6 V drop for dual tests. This provided 29.9 W per device. The heat sink was immersed in still Diala AX transformer oil. Data was logged to a laptop using Omega OM-USB-TC with USB interface.

The temperature rise is listed in Table 3, along with the size and cost of the devices. The C247-025-1AE has a good combination of cooling, volume, and size; the clip gives simple mounting to the IGBT. We will probably use this for the commercial switch. This heat sink gives better cooling than the DTI sink, and is significantly smaller. Another heat sink, the C40-058-AE, gives the best cooling; however, it is substantially larger than the C247 sink. Yet another candidate, the 4-101007U, is smaller than the C247 sink, but it requires screw mounting, while the C247 sink uses a clip, which makes assembly easier.

P/N	N Devices	Cost / Device	Vol / Device (cm^3)	ΔT (°C)	Vol ΔT product
C247-025-1AE	**1**	**$3.82**	**16.2**	**42.2**	**683**
DTI	1	~$7.00	30.9	49.2	1520
C40-058-AE	2	$3.22	40.8	32.9	1053
4-101007U	1	$7.50	11.4	40.8	467
MA-102-55E	4	$1.70	17.3	60.8	1053
MA-301-27E	2	$4.57	31.0	45.8	1419
MA-302-55E	4	$2.37	31.6	56.9	1795
R2A-CT4-38E	1	$3.10	41.1	53.1	2183

Table 3. Heat Sink Characteristics.

4.3 Changes/Problems

We began this project with the idea that the switch would be based on DTI's standard technology. Soon after we started this project, we realized that the switch does not have to open, even in the event of a short, since the peak current and pulse duration is limited by the pulse-forming network. We changed the switch to use a resistor-capacitor-diode (RCD) snubber, which allows the devices to operate at nearly twice the voltage of a switch with DTI standard technology; this would significantly reduce the cost of the switch.

We spent a fair amount of time working on this approach. However, we had not considered the consequences of a short in the output cable. The issue is that the peak current into a short is three times the normal current, so the IGBTs could go into current limiting. Current limiting is made even more likely in a cable short, since the dI/dt increases the voltage on the emitter lead of the IGBTs, reducing the gate voltage. If the IGBT goes into current limiting, the system inductance drives current into the snubbing capacitors, which will charge and over-volt the IGBTs. To ensure that the switch does not go into current limiting, we could put more devices in parallel. However, it is not certain how many devices would be needed.

While it is possible that a switch with RCD snubbers will have fewer devices, it is risky. Because of this risk, we reverted to the standard DTI technology, which is highly robust for shorts. With this decision, we expect progress to occur quickly. We will reconsider the snubber later, if budget permits.

5.0 SBIR Phase II (Year 2)

These SBIR Phase II activities occurred in the period of January 21 – September 14, 2016.

5.1 Summary

The primary objective for Diversified Technologies, Inc. (DTI) in this Phase II SBIR was to design and build a solid-state thyratron-replacement switch that meets the performance specifications for the Stanford Linear Accelerator, given in Table 4.

Table 4. Modulator requirements for Stanford Linear Accelerator

Parameter	Requirement
Voltage	48 kV
Current	6.3 kA
Pulsewidth	6 µs
Frequency	120 Hz
Fault Current	18.7 kA
Jitter	5 ns

Meeting the electrical performance requirements is straightforward for DTI. The major issue for this project is to meet these requirements at a life cycle cost competitive with thyratrons, which need to be replaced every 8 – 18 thousand hours. Thyratrons cost $13 k; hence the replacement switch needs to sell for approximately $30 k (or less if possible) to achieve reasonable life cycle cost savings for SLAC.

5.2 Accomplishments

In Year 2, to produce the thyratron-replacement switch we performed the tasks detailed in this section.

5.2.1 Task 1: Confirm Specifications with SLAC

The control interface has now been defined, and the specifications are complete.

5.2.2 Task 2: Small-Scale Device Testing

The device used in the switch is the NGTB25N120LWG, as we had proposed. We considered arc testing with 60 devices in series and one device in parallel, but eliminated this because of the test described in the next section.

5.2.3 Task 3: Switch 1 Design, Build, and Test (at DTI)

This task is the major design effort in this work; the other switch designs, Tasks 4 and 5, will be modifications of this design. Most of the work in this project to date has been associated with this task.

5.2.3.1 Switch Circuit Board

We designed the circuit board with DTI's standard switch configuration, since it is less risky than one with a resistor-capacitor-diode snubber. A circuit board stuffed with components is shown in Figure 32. The IGBTs are on the left and right sides of the board. The eight IGBTs in each row are connected in parallel; the rows are then connected in series with each other.

The gates of the IGBTs are driven by the four gray toroidal transformers in the center of the board. The secondary windings are the black wires on the toroids. The primary windings are not shown; these are made with single-turn cables that pass through the center of the toroids. These cables are insulated, and rated for the 48 kV switch voltage.

The switch has are more devices in series than the minimum required, so multiple devices can fail short without affecting the switch operation. The switch has monitors that measure the number of devices that have failed, so that maintenance can be scheduled as necessary. These monitors are voltage dividers, made with the blue resistors next to the toroids.

The board operates at 8 kV and 3.1 kA, so the thyratron replacement switch will be made with six boards in series, and two boards in parallel.

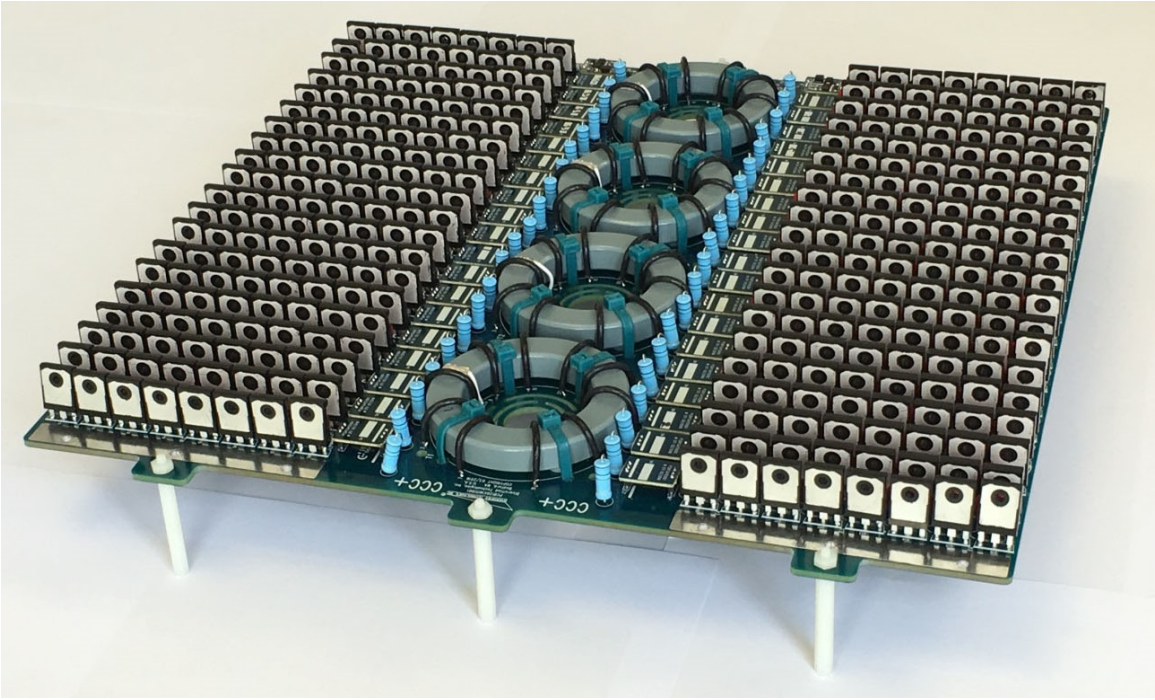

Figure 32. Switch circuit-board. The IGBTs are on the left and right sides of the board, and gate-drive transformers are in the center. The complete switch will use 3840 IGBTs

5.2.3.2 Circuit Board Construction

We built and checked the twelve circuit boards needed for a modulator.

5.2.3.3 Gate-Drive System

We built a gate drive board, and used it to drive one of the four drive loops in the system, using six switch plates. A set of waveforms is shown in Figure 33, taken with 250 V on the gate driver, and no applied voltage on the IGBTs. The gate voltage is the yellow trace, on Channel 1; it rises from - 19 V to + 20 V in 900 ns. The gate drive current is the blue trace, Channel 2; the peak current is 64 A. Hence the peak gate drive current with all four drive loops will be four times this, 256 A. This is close to the 300 A pulse current rating of the IGBTs in the gate drive – hence the gate drive will need two gate drive boards.

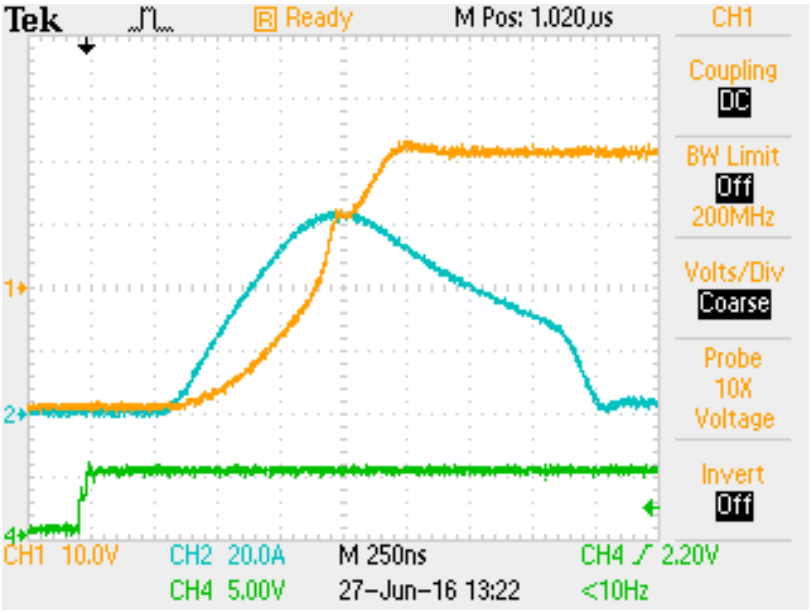

Figure 33. Gate driver current, and sum of IGBT gate voltages. The peak current is 430 A, and the gate voltages rise from full-off to full-on in 420 ns.

Devices Turned Off at Zero Current

We tested an IGBT switch plate into a resistive load, and turned the switch off as it was carrying current; this broke some IGBTs. We believe this was due to the variation in turn-off times of devices in parallel; the last device to open carried the current of the other devices, which was more than the device could handle. As a result, we will only turn the switch off when the current has completely discharged. This means that the switch will stay on during a fault, and the energy in the pulse-forming network will dissipate in the damping resistors.

Switch Testing, and Current Sharing Between the IGBTs

We tested the switch operation without opening the IGBTs while they carry current. To do this, we used the CLC circuit shown in Figure 34. This circuit works as follows. The capacitor on the left is initially charged. The switch is closed, and the left capacitor discharges ("rings") into the right capacitor. The right capacitor then rings back into the left capacitor through the diode. As current is ringing back, the switch is opened (with no current flowing through it). Finally, the remaining voltage on the right capacitor discharges through the resistor.

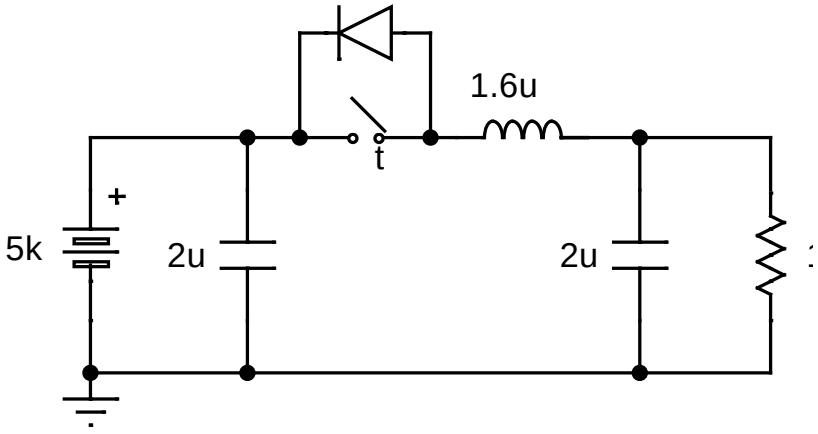

Figure 34. CLC circuit used to test the current sharing of the IGBTs.

The IGBTs readily operate in this configuration, without failure. Typical waveforms are shown in Figure 35, for a charge voltage of 5 kV, and a peak current of 3.4 kA.

A design issue for the switch is that the current should share evenly between the IGBTs, so that they have similar temperatures and none get hot, which could eventually cause a failure.

To determine the current sharing, it is not feasible to make direct measurements, since there are 320 devices in a plate. Instead, we took thermal pictures of the switch after operation. Some examples are shown in Figure 36.

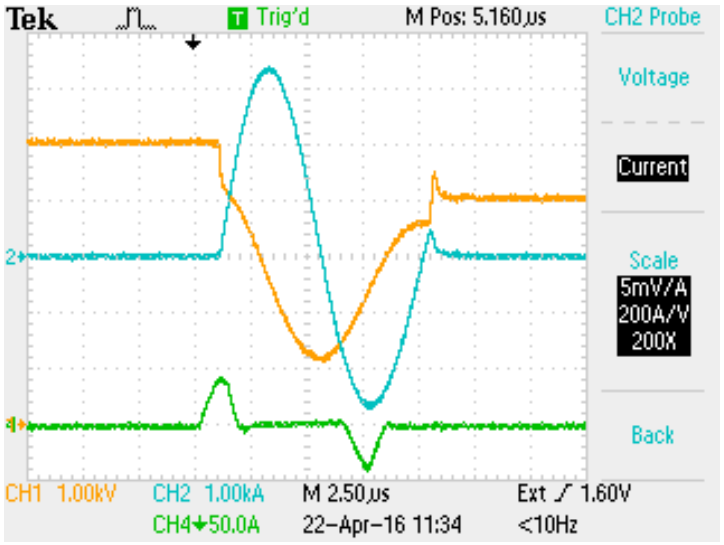

Figure 35. Waveforms for CLC test circuit.

Ch 1 (yellow) V on charged capacitor, 1 kV/div

Ch 2 (blue) Current, 1 kA/div

Ch 4 (green) Gate drive current, 50 A/div

The picture at upper left shows the uneven heating that occurs with the current feed at the top; there is less inductance through the upper IGBTs than through the lower ones, so the current concentrates near the top. The picture at the upper right has the current feed in the middle of the plate, so the heating is fairly even. The picture at the lower left shows uneven heating in two of the columns. The IGBTs in a column all have the same gate voltage, so this is the origin of the uneven heating. We are currently working on this issue.

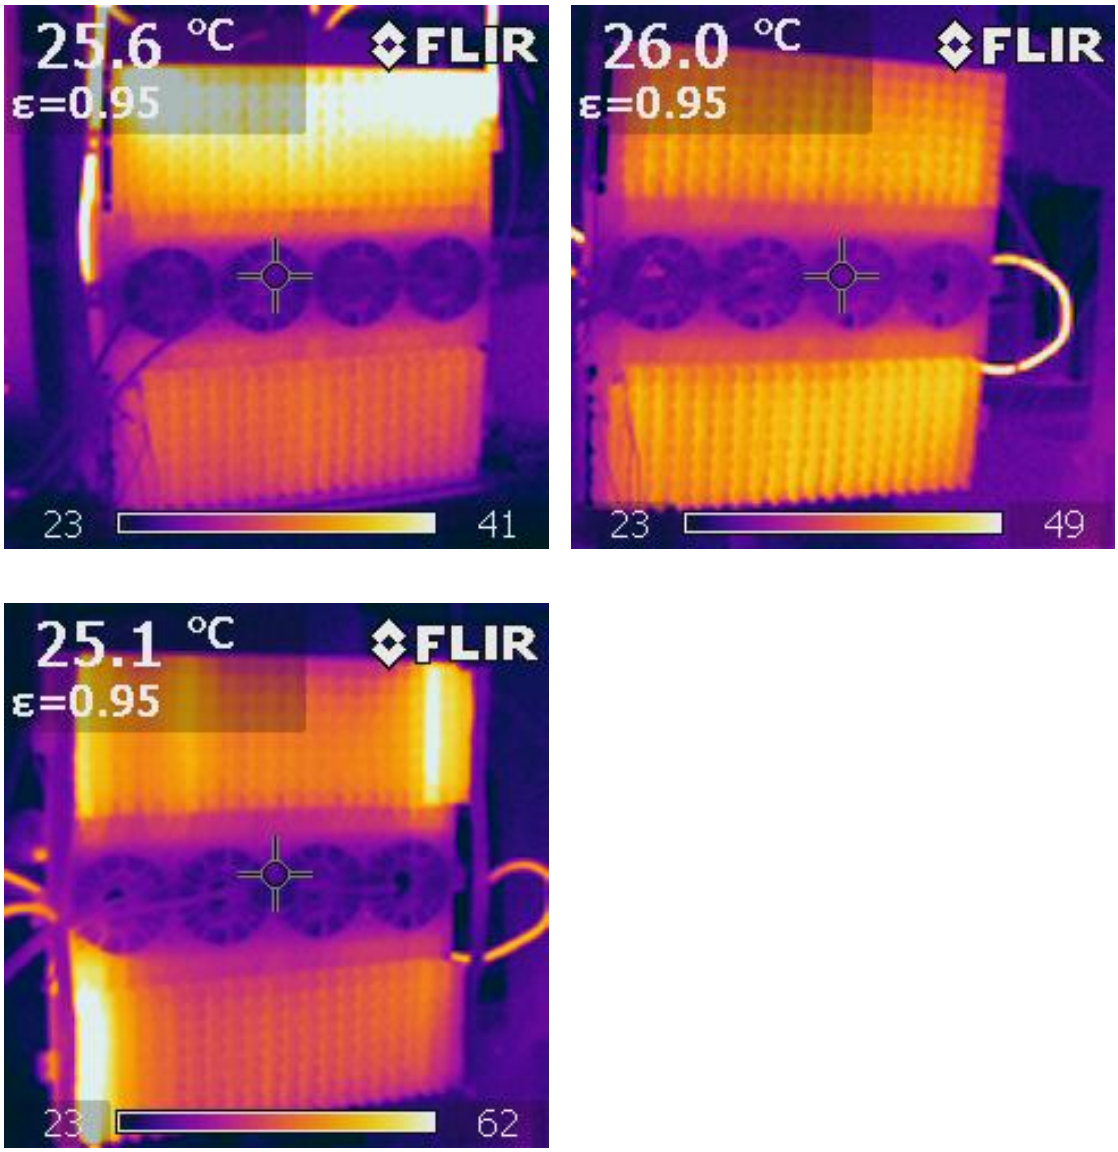

Figure 36. Thermal images of a switch plate after operation for several minutes.

Upper left. Uneven heating produced by a current feed at the top

Upper right. Even heating produced by a current feed in the center

Lower left. Uneven heating of two switch columns.

5.2.4 Task 4: Switch 2 Design, Build, and Test (at SLAC)

The mechanical design uses the IGBT circuit board. The layout is shown in Figure 37. Most of the volume is taken up by the green switch plates, which are in an oil tank (cut away in this drawing). The gate drive is in the gold box to the right. At the top of the tank are the oil pump and oil-water heat exchanger.

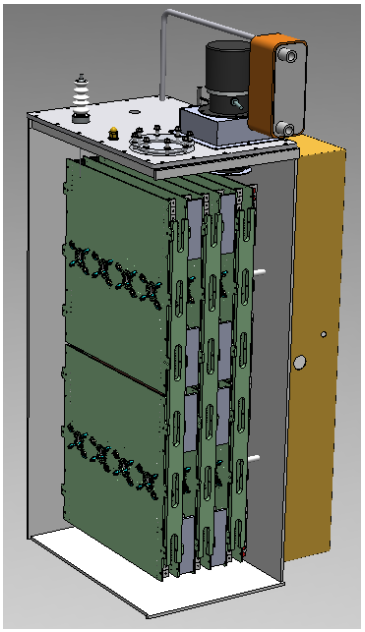

Figure 37. Present mechanical design of the switch.

5.2.5 Task 5: Switch 3 Design, Build, and Test (at SLAC)

No activity.

5.2.6 Task 6: Cost Reduction Study

No activity.

5.2.7 Task 7: Production and Commercialization Planning

During Year 2 of the Phase II effort, DTI presented our progress on this SBIR at several relevant conferences to drum up interest on the technology (listed below). At these conferences our progress was presented via poster or oral presentations. In addition, for each of these conferences a short paper on our progress was submitted for publication in the conference proceedings. As a result of these presentations, we received a number of inquiries regarding the technology.

- 2016 International Vacuum Electronics Conference – Oral Presentation
- 2016 International Conference on Particle Accelerators – Oral Presentation
- 2016 International Power Modulator and High Voltage Conference – Poster Presentation

6.0 SBIR Phase II (Year 3)

These SBIR Phase II activities occurred in the period of September 14, 2016 – September 29, 2017.

6.1 Uneven Heating of IGBTs

The previous Phase II report revealed that some IGBT columns heated unevenly, as shown in the thermal images (Figure 36). This heating did not always occur on the same columns. It was apparent that the heating was caused by low gate voltages, but it was not initially clear why just a few columns heated. The previous report also described that the IGBTs near the input and output connections were hot; see Figure 36. The following discusses the fixes for these.

A simplified diagram of the gate circuit is shown in Figure 38. (Only one secondary winding on the coupling transformer is shown, instead of ten, and there is only a single gate resistor and IGBT shown on the secondary winding, while in actuality there are eight) The gate driver couples a pulse to a 1:1 transformer. When the secondary voltage is large enough, the switch closes, coupling the pulse into the gate of the IGBT. The inductance is the stray inductance in the IGBT emitter lead.

The stray inductance caused problems at the high *dI*/*dt* here. During the current rise, the inductive voltage drop in the IGBT emitter lead acts to reduce the gate voltage. The initial voltage drop at full current is:

$$V_{inductive} = L_{emitter\ lead}\ dI/dt$$

$$= 13\ \text{nH} \times \pi/2 \times 6.3\ \text{kA}/1000\ \text{ns}/32\ \text{devices}$$

$$= 4.0\ \text{V}.$$

This stray inductance can cause a problem if the gate drive pulse is too short. The gate voltages had been initially set to the full 20 V when the switch was unpowered; an example of this is shown in Figure 39.

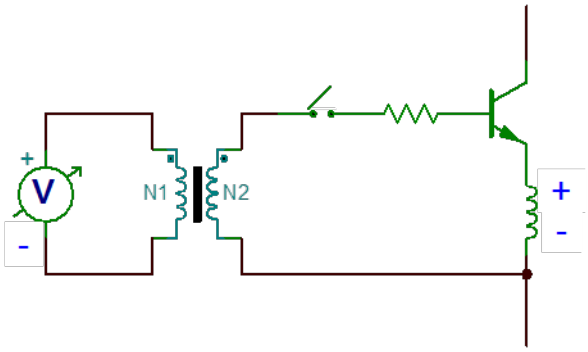

Figure 38. Simplified diagram of gate drive circuit. The inductance shown on the emitter lead is the stray value, about 13 nH; the voltage polarity is for increasing current.

While it would be desirable to measure gate voltages when the switch is powered, this can only be easily done for the lowest-voltage IGBT, which is referenced to ground; the other IGBTs float at high voltage. Instead, we measured the voltage on a gate drive transformer using a loop of insulated wire as shown in the Figure 40 waveforms. The large oscillations on the transformer voltage make the IGBTs resistive, which dissipates power. This problem was fixed by extending the duration of the gate drive pulse over the entire current rise, rather than just a fraction of it.

A second problem was due to the switch in the gate circuit. This switch disconnects the gate drive when the drive voltage is too low, preventing stray voltages from turning on the switch. However, the two voltage-divider resistors that control the switch had been swapped, so the switch was turned on all the time. This allowed the gate voltage to discharge through the transformer. The fix was to correctly specify the divider resistors.

A third problem, the heating near the input and output connections, was solved by using large current-carrying plates. These plates allow the current to spread before it reaches the IGBTs. The plates wrap around the switch in order to fit in the volume required. The fix was first tested with plates that were made of perforated metal so that the IGBTs were visible.

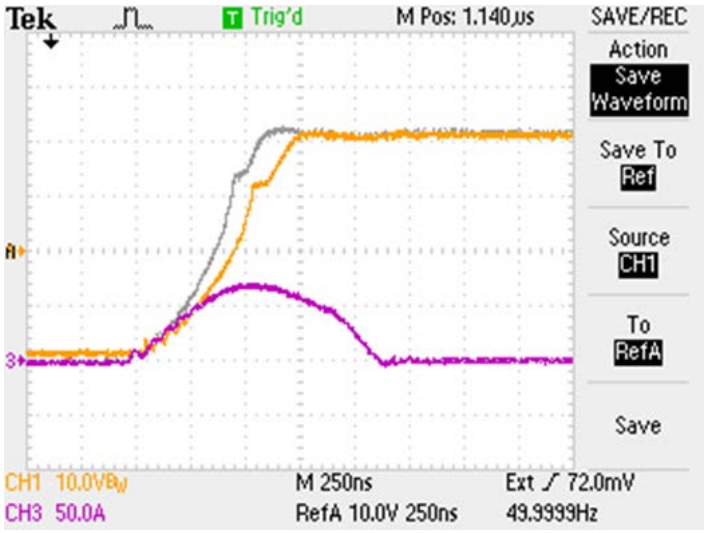

Figure 39. Waveforms with switch unpowered

Ch 1 (yellow) IGBT gate voltage, 10 V/div

Ch 3 (purple) Gate drive current, 50 A/div

Reference (gray) Another gate voltage, 10 V/div

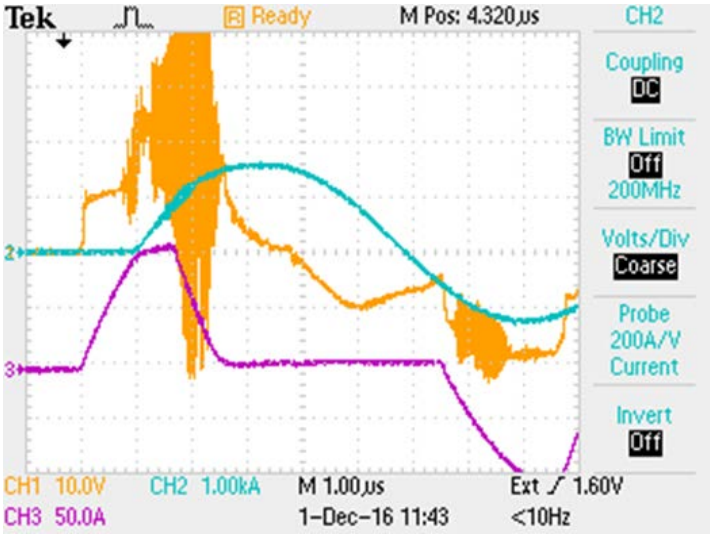

Figure 40. Waveforms associated with uneven IGBT heating due to a too-short gate drive pulse.

Ch 1 (yellow) Gate transformer voltage 10 V/div

Ch 2 (blue) Switch current, 1 kA/div

Ch 3 (purple) Gate drive current, 50 A/div

Figure 41 shows a thermal image of a test with switch circuit boards stacked two high, as they will be in the switch. The current in the boards is quite uniform. The darker region at the upper left of the image is due to the camera. The bright images to the left of the switch are the connection wires and a damping resistor.

Voltage and current waveforms for this test are shown in Figure 42. The gate drive current was 150 A in each loop. Note that the gate transformer voltages have a large oscillation at the start of the current, but very little 1 µs later, when substantial power dissipation would occur.

We then decreased the gate drive current to 75 A per loop; the waveforms for this drive current are shown in Figure 43. Note that the gate transformer voltages start to oscillate 1 µs into the current pulse. While this may be associated with increased dissipation in the switches, the dissipation was not enough to decrease the peak current.

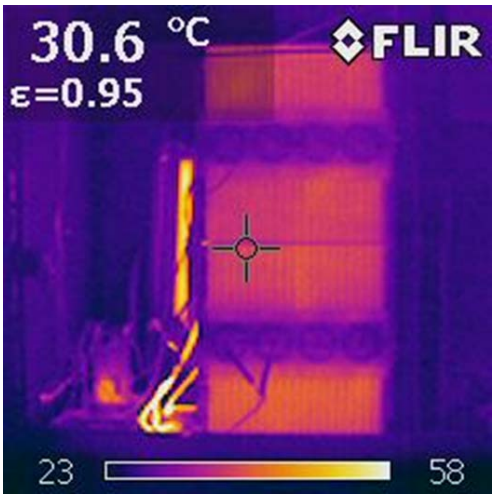

Figure 41. Thermal image showing uniform current distribution in the switch circuit boards stacked two high.

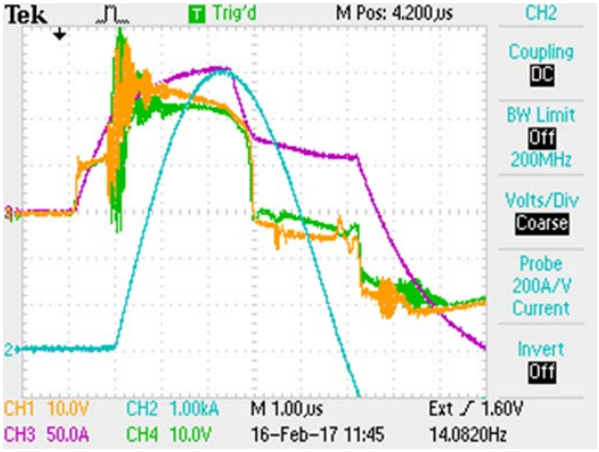

Figure 42. Waveforms for 150 A drive current.

Ch 1 (yellow) Gate transformer voltage 1 10 V/div

Ch 4 (green) Gate transformer voltage 2 10 V/div

Ch 2 (blue) Switch current, 1 kA/div

Ch 3 (purple) Gate drive current, 50 A/div

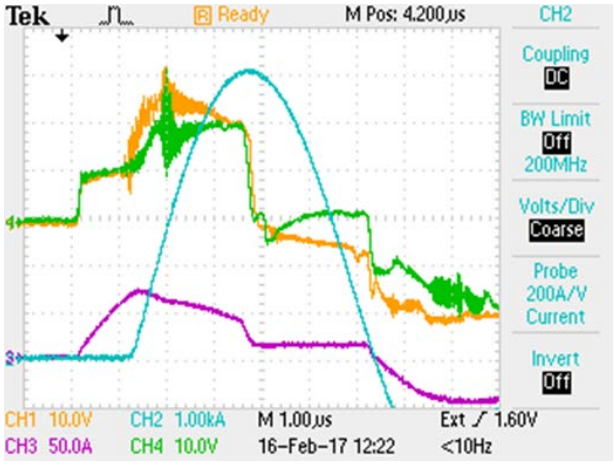

Figure 43. Waveforms for 75 A drive current.

Ch 1 (yellow) Gate transformer voltage 1 10 V/div

Ch 4 (green) Gate transformer voltage 2 10 V/div

Ch 2 (blue) Switch current, 1 kA/div

Ch 3 (purple) Gate drive current, 50 A/div

6.2 Switch Design and Construction

With these problems solved, we designed and constructed the switch. Figure 44 is a SolidWorks view of one side of the switch, showing the current-connecting plates. Figure 45 shows the other side of the switch.

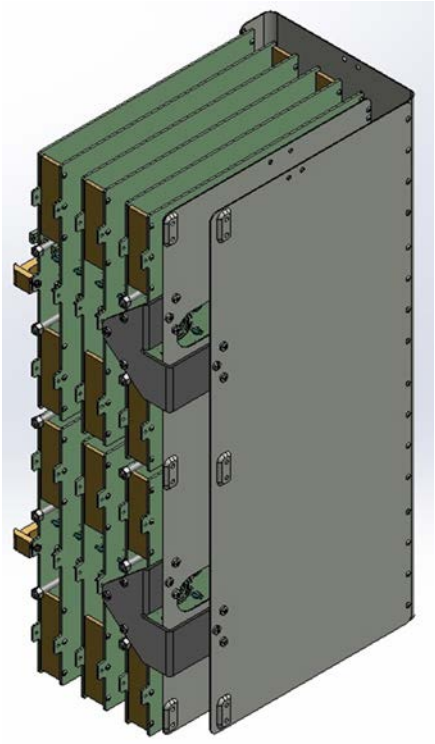

Figure 44. Switch assembly showing current-connection plates at the front. The input and output wires are connected to the left side of these plates; the IGBTs are connected to the right side. The plates allow the current to spread out, giving an even current distribution to the IGBTs.

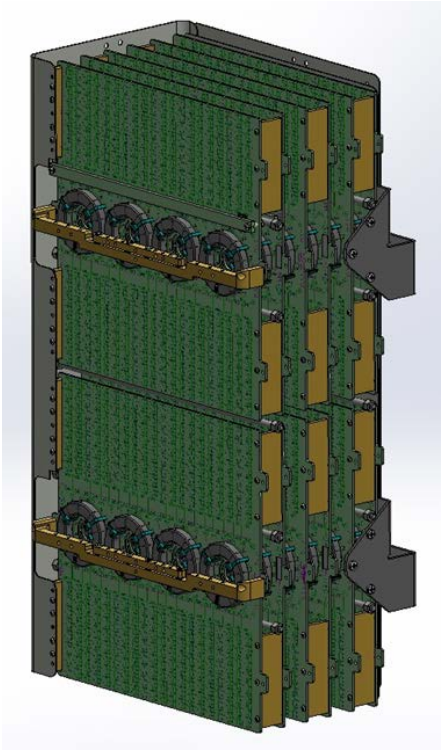

Figure 45. SolidWorks view of the switch assembly. The green sheets are the circuit boards. (The electrical components on these boards are not shown, since they make the file too large.) The switch has twelve circuit boards, two in parallel (upper and lower) by six in series. The yellow bars on the right connect the circuit boards to each other. The toroids are the gate drive transformer cores.

Figure 46 shows the switch that was built. On the right of the top plate are the electrical feedthroughs: the HV feedthrough is the tall tube with corrugations, and the LV feedthrough is the flat square next to it. On the left is a box that allows for oil expansion; it contains the oil-water heat exchanger and pump.

Figure 46. Thyratron replacement switch (as built)

We tested the switch at full current and rate-of-rise. A current trace is shown in Figure 47. Additional waveforms are shown in Figure 48.

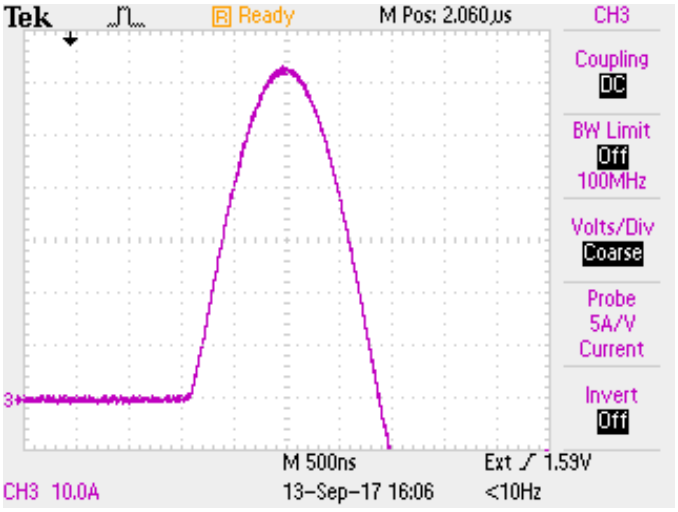

Figure 47. Full operating current, 6.2 kA

Ch 3 (purple) I switch,1 kA/div (not 10 A, as listed)

Sweep speed 500 ns/div

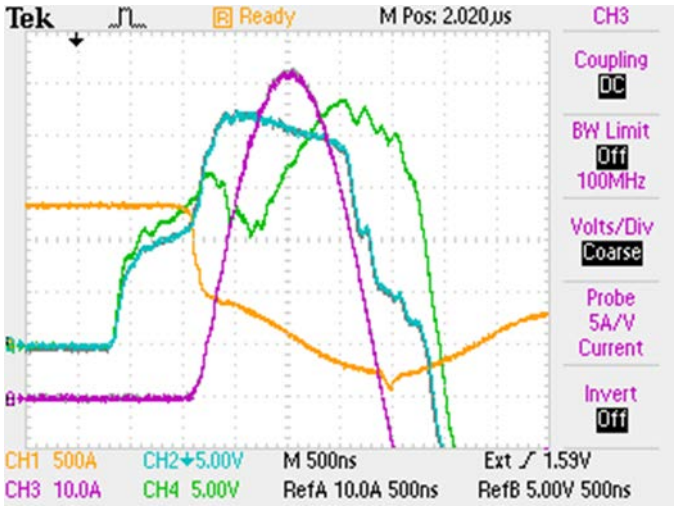

Figure 48. Full operating current, 6.2 kA.

Ch 4 (green) Voltage on first (outer) gate transformer on lowest-voltage plate, filtered, 5 V/div.

Ch 2 (blue) Voltage on second (inner) gate transformer on lowest-voltage plate, filtered, 5 V/div.

Ch 1 (yellow) V switch, 10 kV/div (not 500 A). Max is 27 kV.

Ch 3 (purple) I switch, 10 kA/div (not 10 A). 6.2 kA max.

For comparison, the voltages on the gate drive transformers are shown in Figure 49.

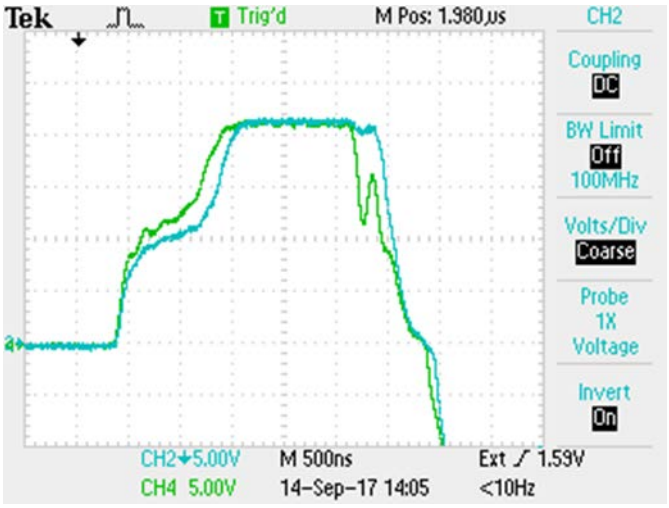

Figure 49. Voltages on two drive transformers on the lowest-voltage plate; no voltage applied to system.

Ch 4 (green) Voltage on first (outer) transformer on lowest-voltage plate, 5 V/div

Ch 2 (blue) Voltage on second (inner) transformer on lowest-voltage plate, 5 V/div

In Figure 49, the voltages on the inner and outer transformers are significantly different. The question is whether this is a real difference, or due to pickup. We believe that this difference is due to a capacitive pickup, since the signals on the outer transformers (first and fourth) are similar. If the signal were due to an inductive pickup, the signals on the first and third transformers would match. However, we are not certain that this is the cause.

We successfully ran the switch at full current for an hour at 2.5 Hz (9000 pulses); no components in the switch failed. These tests were done at a 27 kV charge voltage. We also demonstrated the switch hold-off at the full voltage, 48 kV, for five minutes without failure.

Due to the expenditure of the SBIR funds in other tasks, the Phase II effort was unable to support installation and testing at SLAC. We briefly considered continuing this effort under DTI funding, but the expected cost of the thyratron replacement switches precluded SLAC's ability to purchase them if we had been successful. Without a longer-term path to payback DTI declined to make this investment, ending the Phase II effort prior to testing at SLAC.

6.3 Production and Commercialization Planning

As in Year 2 of the Phase II effort, in Year 3 DTI presented our progress on this SBIR at three relevant conferences to drum up interest on the technology (listed below). At these conferences our progress was presented via poster or oral presentations. In addition, for each of these conferences a short paper on our progress was submitted for publication in the conference proceedings. As a result of these presentations, we received a number of inquiries regarding the technology.

- 2016 Euro-Asian Pulsed Power Conference – Oral Presentation

- 2016 Linear Accelerator Conference – Poster Presentation
- 2016 North American Particle Accelerator Conference – Poster Presentation

In addition to presenting at conferences, DTI submitted a press release on this switch, through our advertising partner, Venmark International. These press releases are submitted to relevant online electronics and technology publications and are readily found via searching for key terms online.

6.4 Future Tests

We recommend these tests if the switch were to be run further:

- To resolve whether the differing gate-drive transformer voltages are real or due to pickup, the sensor wires would have multiple turns around the transformer cores, rather than one. This would increase the signal, but not capacitive pickup.
- The IGBT temperatures would be measured to see if the losses were consistent across multiple devices.
- The switch operation would be demonstrated at the conditions for a klystron arc (12 kA peak) and an arc in the output cable (18kA peak). It is plausible that the switch would handle these, since we have demonstrated that an individual IGBT could carry 500 A; the 32 devices in parallel could carry 16 kA, if the current shared evenly.
- The switch would be tested at the full operating frequency, 120 Hz.

6.5 Cost Issues

The objective of this work was to develop a long-life solid-state switch that could replace the SLAC thyratrons. These thyratrons typically have a two year lifetime, and have a cost (including maintenance) of about $13,000. For this switch to be commercially viable, it would need to sell for about $40,000 or less in quantity.

However, the selling price for building the next switch would be high, and estimated at $150 - $200k in small quantities. Without substantial investment, it seems unlikely that the price in quantity would come down to the $40,000 objective to make it affordable for SLAC. There were two factors that contributed to this high price. The first was the 3,840 individual IGBTs required in this design. While these had a low cost, $1.70 each, the labor to assemble the circuit boards was substantial, since the devices are in through-hole TO-247 packages. To reduce labor cost, the devices could be wave-soldered rather than hand-soldered. However, hand-stuffing of the devices would still be required, since surface-mount parts, which could be machine-installed at a lower cost, are not available for devices with the high power ratings required.

The second factor driving the high price is that the switch was mechanically complicated, and the parts cost was substantial. This should be addressable in a next-generation design, but would require significant investment that does not appear recoverable in the near term.

7.0 Conclusion

Overall, we consider this SBIR to be a qualified success. We were able to demonstrate full thyratron switching capabilities, in a similar size unit, allowing a drop-in solid-state replacement for thyratrons. This success may have been simply premature. The costs and capabilities of existing semiconductor

devices don't yet support a viable product. These costs and capabilities have been improving significantly over time – over the past 20 years, we have observed a process similar to 'Moore's Law', which has doubled the peak switching capability per dollar every generation (3 - 4 years). When we began this effort, we believed that the cost cross-over would be achievable within the SBIR timeframe, and that might have been true if our original design concept proved out. Having completed this effort, the crossover still appears to be close, but not currently achievable. In the next 3 -5 years, we believe an affordable thyratron replacement switch will be possible - but the question is whether any meaningful demand for such a switch will exist, as thyratrons continue to be replaced in pulsed power systems. The reality appears that existing pulsed power systems may soon face a choice between complete modulator upgrades or termination of operations.