FINAL DESIGN OF THE SLAC P2 MARX KLYSTRON MODULATOR

M.A. Kemp[ξ], A. Benwell, C. Burkhart, R. Larsen, D. MacNair, M. Nguyen, J. Olsen
SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., M/S 49
Menlo Park, California, United States

Abstract

The SLAC P2 Marx has been under development for two years, and follows on the P1 Marx as an alternative to the baseline klystron modulator for the International Linear Collider. The P2 Marx utilizes a redundant architecture, air-insulation, a control system with abundant diagnostic access, and a novel nested droop correction scheme. This paper is an overview of the design of this modulator.

I. PROJECT OVERVIEW

A. The International Linear Collider

The International Linear Collider (ILC) is a proposed high energy physics accelerator [1]. It is composed of two main linacs with a total of 560 L-band RF stations. Each station drives 22 superconducting cavities, which have a fill-time of ~800 us. The beam consists of 2625 individual bunches spaced 333 ns apart. An RF station consists of a modulator which drives a single klystron. Parameters for this modulator are shown in Table 1. The baseline ILC modulator is the Fermi bouncer modulator [2].

The Marx program at SLAC National Accelerator Laboratory (SLAC) was undertaken to provide an alternative topology to with the characteristics of high-availability, simple maintenance, and low cost. The P1 Marx was designed, built, and is currently undergoing lifetime testing into a Toshiba multi-beam klystron [3]. Both building upon the lessons learned in the P1 development and employing some fundamentally different concepts, the SLAC P2 Marx (P2 Marx) is a follow-up effort to further develop the Marx concept as a viable alternative to the baseline modulator [4-7].

Table 1. ILC Klystron Modulator Parameters

Output Voltage	120 kV
Output Current	140 A
Pulse Width	1.6 ms
Pulse Repetition Frequency	5 Hz
Average Power	134 kW
Output Pulse Flat-top	±0.5%
Energy Deposited into Klystron During a Gun Spark	<20 J

B. SLAC P2 Marx Design Goals

There are several points of emphasis for the P2 Marx design. First, the modulator must be compatible with the ILC two-tunnel design. In this scheme, the modulator and klystron are located within a service tunnel with limited access and available footprint for a modulator. Access to the modulator is only practical from one side. Second, the modulator must have high availability. Robust components are not sufficient alone to achieve availability much higher than 99%. Therefore, redundant architectures are necessary. Third, the modulator must be relatively low cost. Because of the large number of stations in the ILC, the investment needed for the modulator components is significant. High-volume construction techniques which take advantage of an economy of scale must be utilized. Fourth, the modulator must be simple and efficient to maintain. If a modulator does become inoperable, the MTTR must be small. Fifth, even though the present application for the modulator is for the ILC, future accelerators can also take advantage of this development effort. The hardware, software, and concepts developed in this project should be designed such that further development time necessary for other applications is minimal.

II. MODULATOR TECHNOLOGY

In general, technology advancements in modulators might be grouped into several categories: switching, topologies, controls and diagnostics, energy storage, and packaging. These five technology areas must be focused on to achieve the design goals mentioned above. In each following section, these technologies are highlighted and are tied to the design goals of the modulator.

A. SLAC P2 Marx Topology

1) Overview

Advancements beyond the current state-of-the-art in Marx modulator topologies are necessary to achieve the project's goals of high availability, low cost, and allowance for a portable design. The Marx topology is inherently modulator; it is an array of many, low-voltage components. The large numbers of small components allow an economy-of-scale to be taken advantage of during procurement, which aids in achieving the goal of low-cost. The P2 Marx cell is designed such that each cell

* Work supported by the US Dept of Energy under contract DEAC02-76SF00515.
ξ email: mkemp@slac.stanford.edu

provides a square output pulse. By doing this, the Marx can be made truly redundant (a separate, single-point failure compensation mechanism is not relied upon). This redundancy can be taken advantage of to achieve the goal of high availability.

There are many different ways in which a Marx modulator can be constructed to achieve the parameters listed in Table 1. The design process involved with selecting the parameters utilized in the P2 Marx are detailed elsewhere [7,4]. The cell parameters are summarized in Table 2.

Table 2. P2 Marx cell parameters.

Number of Cells	32
Minimum Cells Needed to Provide Full Output Pulse	30
Bus Voltages	-4kV, -1.1kV
Nominal Cell Voltage	-3.75 kV
Peak Cell Power	560 kW
Average Cell Power	4.48 kW
Duty Cycle	0.008
Main Storage Capacitor Droop	<20%

2) Compensation Scheme

The simplified schematic of a single Marx cell is shown in Fig. 1. A basic Marx cell consists of a storage capacitor and an output switch. In Fig. 1, this is shown by C1 and Q1. When discharging into a resistive load, the output voltage droops with an RC decay. To achieve a flat-top pulse, a buck converter is placed in series with the main storage capacitor. The droop correction (or "regulation," "compensation") scheme utilized in the P2 Marx has been presented in detail elsewhere [1, 8].

In brief, the compensation portion of the cell functions by chopping Q3 at the PWM frequency. The LC output filter smooths the ripple and has a slow ramp up in voltage during the duration of the pulse. The voltage across the filter capacitor adds with the main storage capacitor voltage droop to produce a flat output pulse.

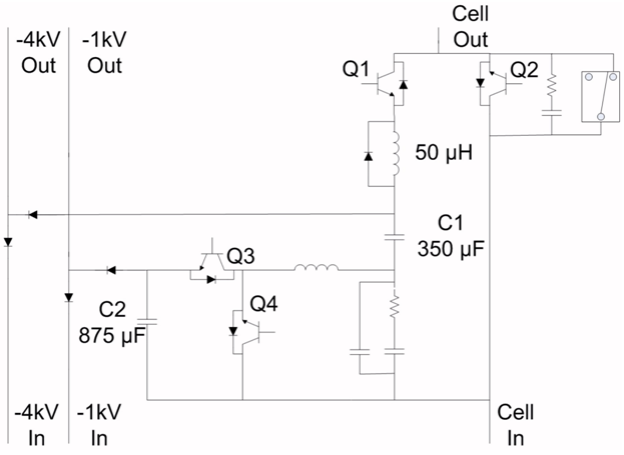

Figure 1. Simplified schematic of one P2 Marx cell.

After the pulse, energy stored in the filter capacitor is recovered back to C2.

For the ILC klystron modulator, the specification is for a voltage flat-top of ± 0.5%. The size of the LC PWM filter in-part determines the characteristic ripple of the cell. The maximum ripple occurs when the Q3 switches with a 50% duty cycle and is given by

$$\Delta V \cong \frac{V_{C2} - V_{Cf}}{8 C_f L_f f_{PWM}^2} \qquad (1)$$

where ΔV is the peak-to-peak voltage ripple in volts, V_{C2} is the voltage across C_2 in volts, V_{Cf} is the voltage across the PWM filter capacitor in volts, C_f is the PWM filter capacitor in Farads, L_f is the PWM filter inductor in Henrys, f_{PWM} is the PWM switching frequency in Hz. With the values utilized in the P2 Marx cell, the calculated maximum cell ripple is ~8.5V. At -3.75kV nominal cell voltage, this corresponds to a ripple of ±0.12%.

Figure 2 shows the measured output waveform for one cell into a 29 Ω resistive load. The waveform was captured with the integral cell diagnostics. The ADC used has a resolution of 12-bits with a sample rate of 1 MS/s. At this scale, the ripple is not perceptible. At very low values of ripple, an offset box might be utilized to better resolve features on top of the waveform [8]. In the case studied here, the internal offset of the scope is appropriate to resolve the top of the pulse, as will be demonstrated next.

To demonstrate the concept of the Marx, a two-cell Marx was constructed. Each cell was charged to -3.75kV, and therefore the output pulse was -7.5kV. Fig. 3 illustrates a zoomed-in view of the top of the waveform. Note that at this stage of the project, the correction algorithm utilized is still under development and further optimization [9]. Nevertheless, the characteristic ripple of the PWM filter can be viewed. As shown, the ripple is about 20-25V_{p-p}. This corresponds well with the calculated value of 8.5V*2=17V. Also shown are the tolerance levels of ±0.2%. This falls well within the ILC specification of ±0.5%.

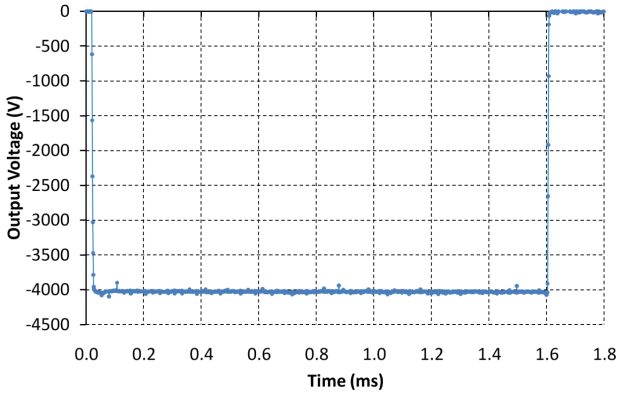

Figure 2. Measured output voltage for one cell operating into a resistive load. Data are taken with the in-cell voltage diagnostic.

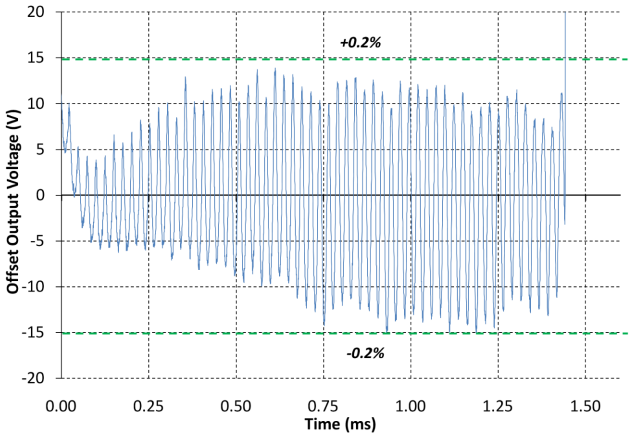

Figure 3. Measured flat-top of the output voltage for two cells. A constant offset equal to the average output voltage is applied to the data. The green dashed lines represent the ±0.2% tolerance levels. PWM switching for the two cells is **in phase.**

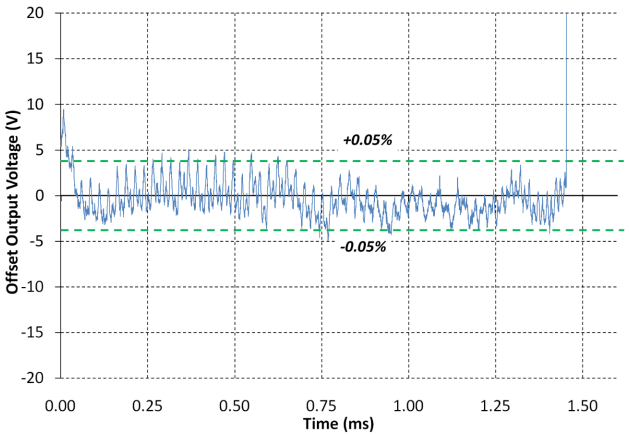

Figure 4. Measured flat-top of the output voltage for two cells. A constant offset equal to the average output voltage is applied to the data. The green dashed lines represent the ±0.05% tolerance levels. PWM switching for the two cells is **180° out of phase.**

There are at least three ways to fundamentally reduce the output ripple of the cell. First, as shown in (1), the filter capacitance and/or inductance may be increased. The trade-off of doing this is increased cost, mass, and, potentially, heat dissipation. Second, one could increase the PWM frequency. This would increase device losses and increase the IGBT junction temperature change during a cycle, which may lead to reduced device lifetimes. Third, one can take advantage of the characteristic of the Marx as an array of many cells, shown next.

In the present implementation, all switch timing is generated by the Application Manager (detailed in section C). In addition, all cells have the same PWM frequency. Therefore, the cells can have a staggered turn-on sequence which results in a phase shift of the ripples generated by the cells. Since the output voltage is determined by the sum of all the cell voltages, the cell ripples can add to cancel out some portion of the output ripple.

This was demonstrated on a small scale with only two cells. One cell is triggered 180° after the first cell. Since the PWM frequency is 40 kHz, the second cell is delayed 1/(2*40kHz)=12.5 µs. Figure 3 shows the measured output waveform when the PWM switching for the two cells is in phase. Figure 4 shows the output waveform when the cells are out of phase. Indeed, the ripple of one cell partially cancels the other cell. With this method, an order of magnitude reduction in ripple is achieved compared with the ILC specifications. Future iterations of the cell design might take advantage of this characteristic and have relaxed constraints on the PWM filter.

As an aside, it is interesting to note that the above technique does not correct for an overall slope on the output nor correct for large artifacts in the waveform. For example, in the present implementation of the correction algorithm, there is a slight undershoot at the front of the pulse. Addition of the two cells does not remove this deviation from the flat-top, as shown at the front of the waveform in Fig. 4. Also, cells are tuned separately, and do not necessarily have the same switching duty cycle profile throughout the pulse. Nevertheless, independently tuned cells still appear to cancel each other out well. Further investigation is underway to fully take advantage of this.

B. Switching

Differing from the approach taken with the P1 Marx, in the P2 Marx, the maximum cell voltage was derived from the highest-voltage commercial IGBTs available. A characteristic of the Marx topology is that it arrays lower voltage components into one high voltage output. Arrays of devices within the Marx cell, especially the switching elements, do not take advantage of this inherent feature of the Marx topology. This is not to say that a Marx cell can not successfully contain arrayed components. However, control, protection, and diagnostic access are simplified with single modules controlling the output. The P2 Marx 4kV cell operating voltage, the 1.1 kV correction portion

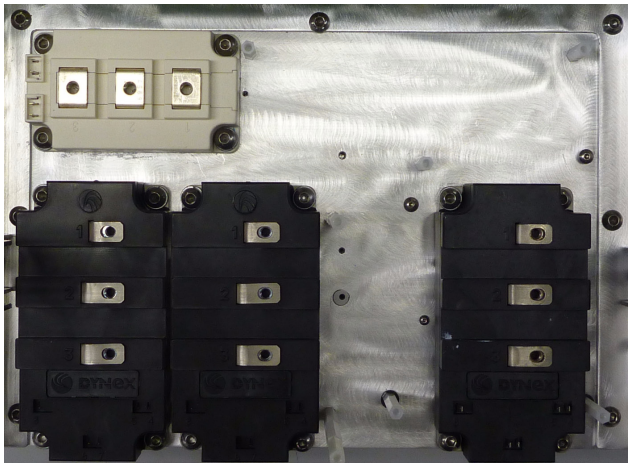

Figure 5. Photograph of the power semiconductors mounted on the air-cooled heat sink.

operating voltage, and the 40 kHz PWM frequency were chosen by considering the requirement for a 10^5 hr lifetime and the effects of bond-wire liftoff and cosmic ray effects (two common failure mechanisms in IGBTs) [7].

Typically, pulsed power modulators have demanding requirements on switching; many times operating "outside the datasheet." Because the rise-times, peak power, and packaging requirements of the P2 Marx are not beyond the state-of-the-art of presently available solid-state devices, COTS power semiconductors were chosen and all operation is within manufacturer's specifications. Fig. 5 is a photograph of the devices mounted on a bonded-fin heatsink.

Each P2 Marx cell protects itself. Overvoltage, over current, and over di/dt faults are detected and suppressed at the gate drive level. The second-level protection is accomplished by the hardware and application managers. These faults are typically much slower. For example, over-temperature, enclosure interlocks, water flow, etc. The gate drive protection scheme is detailed in a previous paper [10].

C. Control System

The control system of the Marx follows a hierarchical scheme [11]. At the first level, the gate drives provide fast detection and response to faults. They also have the capability, utilizing an optional diagnostic card, to provide diagnostic access to signals at the IGBT potential. The second level, the hardware manager (HM), implements the control algorithm, controls timing within the cell, acquires diagnostic signals, and is the interface between the gate drives and the third level, the application manager (AM). This third level performs high level operations such as cell coordination, interlock handling, start up and shut-down, power supply control, and modulator-level diagnostics.

A diagram of the control system components is shown in Fig. 6. Each cell contains one HM which communicates with the application manager level via a gigabit Ethernet fiber and a single trigger fiber. The HM contains twelve, 12-bit, 1 MS/s ADCs which monitor voltage and current

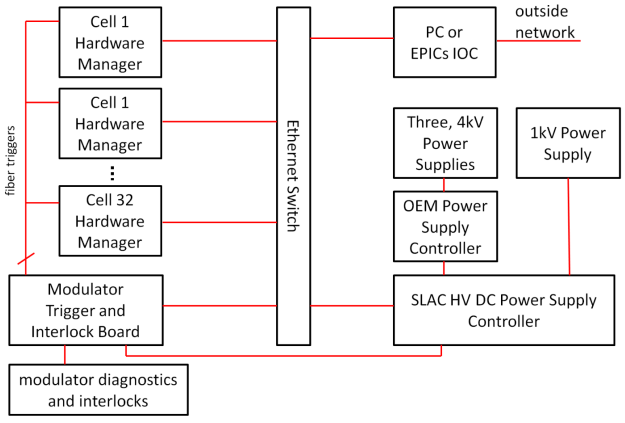

Figure 6. Diagram of the P2 Marx control system.

signals within the cell. These are transmitted to the AM for processing.

The application manager level consists of several components. A gigabit fiber optic switch routes communication between the components. A trigger and interlock handling board provides trigger signals to the cells, records modulator level diagnostics (output voltage, output current, etc.), and integrates interlocks. A SLAC-built power supply controller controls the OEM power supply banks. Finally, a PC (or, in the future, potentially an EPICs IOC) runs the software for processing diagnostics, provides the user interface, and handles high level applications.

The FPGA card utilized in the hardware manager, the trigger board, and the power supply controller was designed at SLAC. In addition to being portable between platforms within the modulator, various systems at SLAC also take advantage of this control card. The Link Node and upgraded Modulator Klystron Support Unit (MKSU 2) utilize the same basic architecture.

D. Energy Storage

Energy storage technology is important for the P2 Marx design goals of low cost and ease of maintenance. In the P2 Marx, the storage capacitors account for about 6% of the total materials cost of the modulator. However, in longer pulse-length applications, the capacitors become a large portion of the total cost. For example, for a 25 ms long pulse Marx, the capacitors would be as much as 25% of the total cost [12]. Also, capacitor technology is important for achieving ease of maintenance. Smaller, lighter capacitors improve ease of handling the Marx cells.

There are five different energy storage devices in each cell. There are two inductors: the di/dt limiting inductor and the PWM inductor. A design for these inductors was completed utilizing magnetic cores; however the benefits of smaller size and decreased copper loss did not outweigh the negatives of increased cost, mounting complexity, and risk of core saturation when compared to an air-core design. These elements were designed at SLAC.

There are three energy storage capacitors in the design; all are of the self-healing metalized film type. For the PWM filter, COTS film capacitors were chosen. For the 1kV and 4kV main storage capacitors, a custom capacitor is used. Because the size, weight, and lifetime of the capacitors heavily influence the cell design and performance, extra scrutiny was placed upon the design of these elements.

An initial study was conducted on the polypropylene film chosen for the cell [13]. Although it has been shown that many factors affect aging in self-healing capacitors (temperature, reverse voltage magnitude, reverse voltage frequency, etc.) [14], discussions with the manufacturer indicated that field stress on the film would be the primary aging mechanism in our operating parameter range. Accelerated lifetime testing was performed with

film electric field stress as the aging parameter. Results indicated that a field stress of 260 V/μm for this film would result in a 10^5 hour lifetime.

However, during life testing of the P1 Marx modulator, premature end of life was reached on the main storage capacitors. These capacitors had a maximum field stress of ~190 V/μm. Experiments are ongoing to fully understand this aging mechanism. Preliminary results indicate that ac effects are dominant. The majority of previous studies in HV film capacitors have focused on filter applications (such as in a DC power supply) or in complete discharge applications (such as in a pulse forming network). The Marx applications reside in-between these parametric regimes. Further, the accelerated aging seen for the P1 capacitors, which discharge to 60% of their charge voltage (40% droop), does not occur under the P2 operating conditions, 20% droop. Additional research is needed to optimize energy density, discharge cycle, and construction materials to achieve a given lifetime.

E. Packaging

Packaging and construction techniques must be emphasized in the Marx design to achieve the goals of a portable design and ease of maintenance. Modular components within cells, an easily scalable cell layout, and standardized design criteria allow for the ILC Marx design to be simply transferred to other needs. Also, efficient and accessible packaging is necessary for simple and quick maintenance.

1) Cell Layout

There have been two major versions of the P2 Marx cell. The first version was used to demonstrate the droop correction concept and perform component-level testing. The second version was designed to integrate efficiently with a modulator, minimize cell volume, and incorporate high energy density capacitors. Also, the shape was chosen to be long and short to allow for simple access to cell components. Cells weigh less than 50 lbf, so they can be hand-carried by maintenance staff.

The cell components are divided in to several modular sections. These modules then attach to a single sheet metal piece. Module to module connections are made by pre-fabricated cable harnesses. All components were fabricated commercially and were assembled into cells at SLAC. A strict quality control regiment of checking of outside-manufactured components and verification of SLAC-assembled modules was followed. So far, there have been no component failures resulting from poor assembly practices.

2) Modulator Layout

During the course of the design phase of the project, there have been several iterations on the scheme utilized for the overall modulator layout. The final scheme is shown in Fig. 7. The modulator consists of five columns of cells. Two columns have seven cells, the other two have six. The voltage output from each cell is illustrated in Fig. 8. The dimensions of the cells and the modulator are given in Table 3.

Table 3. Modulator size parameters.

Height	93.5" (2.37m)
Width	109" (2.76 m)
Depth	60" (1.52 m)
Single Cell Volume	1.63 ft^3 (0.046 m^3)
Total Cell Volume	52.2 ft^3 (1.48 m^3)
Modulator Volume	354 ft^3 (10.0 m^3)
Cell Power Density	2.57 kW/ft^3 (90.8 kW/m^3)
Modulator Power Density	0.379 kW/ft^3 (13.4 kW/m^3)

There are several factors which influenced the design choices made in the modulator layout. First, the layout is constrained to allow access only one side. Second, there are several components within the cell which constrain the cell shape. This determines the maximum/minimum

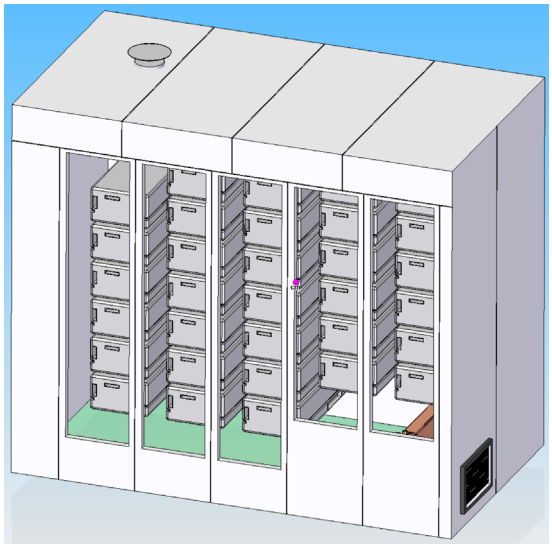

Figure 7. Sketch of the overall modulator layout of the Marx.

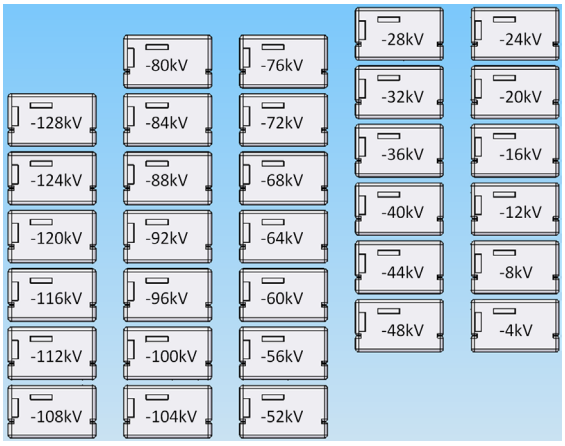

Figure 8. Arrangement of cell output voltage within the Marx enclosure.

realistic cell height and width. Third, high voltage considerations determine the spacing and mounting necessary to hold the cells.

3) High Voltage Effects

There are several advantages to utilizing air insulation rather than oil. First, access and therefore maintenance is straightforward. The modulator does not need to be lifted out of an oil tank (or the oil removed and tank access granted) to repair or remove components. Second, the choice of components becomes less restrictive. Many electronic components have premature failure associated with long-term degradation in oil. Finally, secondary containment and specialized handling of hazardous materials is not required in an all-air modulator.

However, there are drawbacks with air insulation. First, the specific heat of air is lower than oil. Therefore, if air instead of oil is used also for cooling, higher volumetric flow rates are necessary to remove the same amount of heat. Second, the dielectric strength of air is lower than oil. A larger volume is required to insulate with air rather than oil. Third, in some situations, high voltage points in air can result in the creation of ozone. This can lead to long-term degradation effects in, for example, cables or printed circuit boards.

To properly handle high voltage effects, carful design is necessary. At the onset of the P2 Marx project, we revisited the modulator high voltage design criteria. A review of published fundamental phenomena, UL standards, and examples from industry was conducted [15]. Design criteria were chosen for the project and adhered to on all system components: 1) The maximum bulk electric field magnitude in air is lower than 18 kV/cm, 2) The creepage along insulators is less than 2.7 kV/cm, 3) The electric field at all triple points is less than 5.0 kV/cm.

All critical areas of the system were simulated in Maxwell 3D [16]. Some gap distances, conductor radii, and insulator placements were determined by an iterative optimization routine which coupled a Matlab [17] script with Solid Edge [18]. This routine aided in minimizing the modulator volume while adhering to the design criteria.

Each Marx cell mounts inside of an electrostatic shield. This effectively isolates the cell such that, regardless of its position within the 120 kV Marx, it only "sees" a maximum of 4kV. In addition, the shield grades the fields such that the electric field magnitude outside the cell is kept to a minimum. It was found using this shield scheme, only two additional small external field shapers were required. Therefore, the sheet-metal shields can be cost-effectively produced with high-volume manufacturing techniques. Figure 9 shows a photograph of a cell inside the shield. Also shown is the electric field simulation of this shield with a prototype field shaper. Compliance with the design criteria mentioned previously can be easily visualized by adjustment of the scale in the plot, as shown in Fig. 9c. For this part, the design was revised in the trouble areas of the shield corners.

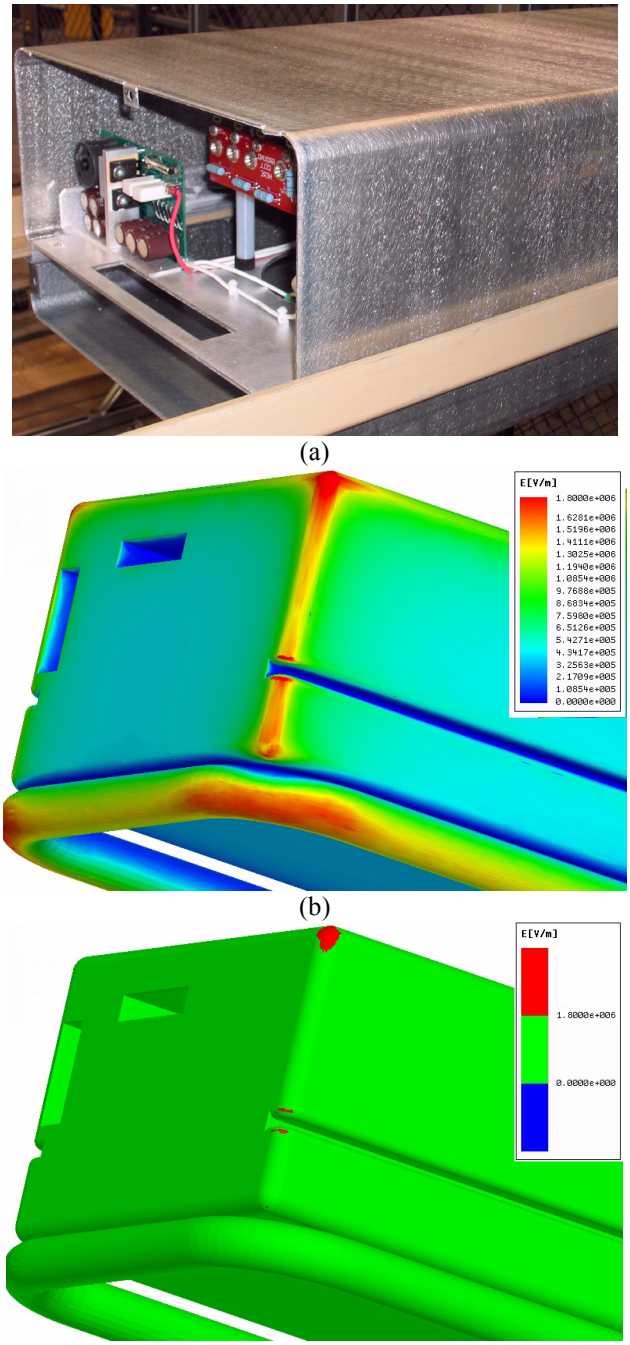

Figure 9. (a) Photograph of a cell inside a shield. The shield front is omitted. (b) Simulated electric field magnitude on outside of a cell shield with full scale and with a (c) "pass/fail" scale.

4) Heat Removal Scheme

Several issues impact the choice of the heat removal scheme. The minimum efficiency of the Marx is estimated to be >93% [1]. This necessitates the removal of 10.1 kW of heat from the modulator cabinet, or an average of 315 W from each cell. Because of the modest cell heat load, to reduce the number of connections to the cells, and to bypass the issue of water cooling at elevated voltages, the Marx cells are air-cooled.

Table 4. Calculated cooling system parameters.

Input Water Temp.	30 °C (86 °F)
Calculated Modulator Heat	10.1 kW
Water Flow Rate	0.38 l/s (6 gpm)
Water ΔT	6.4 °C (11.5 °F)
Air Volumetric Flow Rate	708 l/s (1500 CFM)
Air Velocity Over Heat Sinks	1.9 m/s (75 in/s)
Avg. Modulator Air Temp.	47.5 °C (118 °F)
Avg. Heat Sink Temp.	62 °C (144 °F)

In the two-tunnel ILC design, modulators are located within a tunnel. For ease of tunnel cooling, the vast majority of heat from the modulator shall be transferred to cooling water. The Marx cooling scheme is illustrated in Fig. 10. The cell heat sinks, the modulator air/water heat exchanger, and the air blowers are all COTS.

Using manufacturers' data sheets, the estimated water temperature, and the estimated modulator efficiency, the temperatures associated with the air and heatsink can be calculated. These values are shown in Table 4. A primary value of interest is the maximum junction temperature in the power semiconductor devices. The device with the largest power dissipation in the modulator is the charge switch in the cell closest to ground potential. This switch is expected to dissipate ~0.22 kW. The datasheet value of junction-heatsink thermal resistance is ~80 °C/kW. Therefore, the expected junction temperature is (0.22 kW) * (80 °C/kW) + (62 °C) = 75.2 °C which is much less than the device maximum junction temperature of 125 °C.

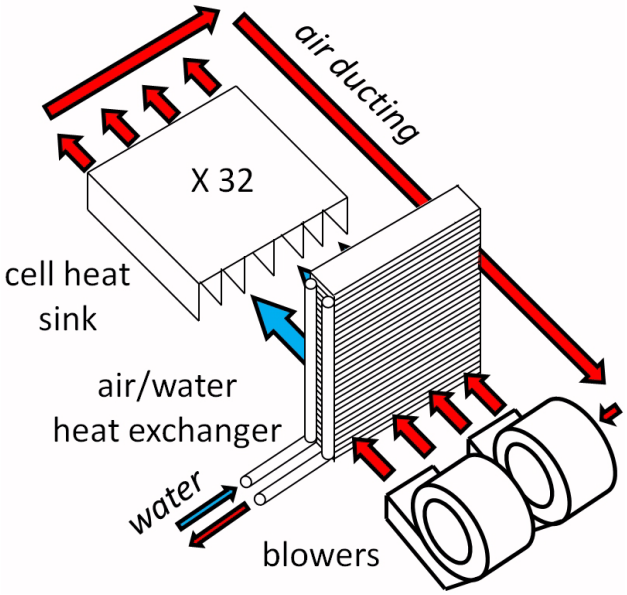

Figure 10. Marx cooling scheme.

A feature of the P2 Marx cooling scheme is that it is conservative. The added cooling capacity potentially allows for operation outside the ILC parameter range. In addition, the extra headroom allows for redundancy in the air blowers. One blower can fail and the modulator can continue providing the specified output pulse.

5) Maintenance Scheme

The general maintenance philosophy used with the P2 Marx is to minimize modulator downtime. Generically, repair time is typically determined by the sum of 1) time to determine there is a fault, 2) time for maintenance staff to travel to the modulator, 3) time to lock out the modulator and control hazards, 4) time to diagnose the source of fault, 5) time to repair the fault, and 6) time to put the modulator back in operation. Portions 3) - 5) can be minimized with a modular system.

The lock out scheme is simplified by two factors. Under certain conditions, zero voltage verification can be achieved with the on-cell diagnostics. In addition, the capacitors have easily accessible shorting terminals at the front of the cells. The control system records the source of each fault. Instead of repairing the cell, a pre-tested spare cell is brought to be modulator to replace the disabled cell. Down time is spent quickly replacing a cell rather than repairing in-place.

However, the easiest way to minimize downtime is to keep the modulator from becoming disabled in the first place. Stated previously, the Marx has N+2 redundancy, meaning two cells can become disabled and allow the Marx to provide the full output pulse. In addition, the capability to implement prognostics has been implemented [19]. In some cases, before a condition becomes disabling, it can be detected. This way, a cell can be replaced during a scheduled down time. Finally, a classification scheme is in place to categorize types of faults. Not all out-of-tolerance conditions disable a modulator. Even in the event of a major event such as a klystron arc, the modulator can terminate the present pulse and be ready to fire by the next pulse.

III. PRESENT STATUS

As of this publication, the design for all modulator mechanical components has been finalized and all parts are either being assembled or are under construction. The cells are undergoing testing on a 2-cell test stand, shown in Fig. 11. In this test stand, cells are tested up to full peak and average power. In addition, they are fault tested into both hard and soft shorts.

The remaining project effort is in three areas: 1) finishing construction of the modulator, 2) testing and troubleshooting of the full modulator, and 3) finishing the controls hardware and software. For the controls

Figure 11. Photograph of the two-cell test stand.

elements, one more platform, the modulator trigger and interlock board, is still to be developed. Otherwise, the majority of effort is needed in the area of software development. Presently, a PC-based program controls the 2-cell test stand. This must be revised in order to meet the requirements of a 32-cell modulator.

All components for the full modulator enclosure will be in-house by mid-September. Testing of the full modulator will progress as components become available and controls development allows.

IV. ACKNOWLEDGEMENT

The authors acknowledge the outstanding support and contributions from Patrick Shen, Alfred Viceral, Chris Adolphsen, Zen Szalata, Chuck Yee, John Hugyik, Joe Olszewski, Wanda Gorecki, Chaofeng Huang, David Anderson, Richard Cassel, Koen Macken, Phillip Nguyen, Tony Beukers, and the SLAC Test Facilities staff. This work is supported by the US Department of Energy under contract DEAC02-76SF00515.

V. REFERENCES

[1] N. Phinney, *et al.*, "ILC reference design report volume 3- accelerator,"arXiv:0712.2361, 2007.

[2] H. Pfeffer, *et al.*, "A long pulse modulator for reduced size and cost," 21st Power Modulator Symposium, 1994.

[3] C. Burkhart, *et al.*, "P1-Marx modulator for the ILC," 2010 IEEE Power Modulator Conference, Atlanta, GA, 2010.

[4] M.A. Kemp, *et al.*, "Status update on the second-generation ILC Marx modulator prototype," 2010 Power Modulator Conference, Atlanta, GA, 2010.

[5] M.A. Kemp, *et al.*, "Design of the second-generation ILC Marx modulator," LINAC 2010, Tsukuba, Japan.

[6] K. Macken, *et al.*, "Towards a PEBB-based design approach for a Marx-topology ILC klystron modulator," PAC'09, Vancouver, BC April 2009.

[7] K.J.P. Macken, *et al.*, "IGBT PEBB Technology for Future High Energy Physics Machine Operation Applications," APEC 2011.

[8] C. Huang, *et al.*, "Development of modulator pulse stability measurement device and test results at SLAC," 2011 Pulsed Power Conference, Chicago, IL.

[9] D. MacNair, *et al.*, "SLAC P2 Marx control system and regulation scheme," PAC'11, New York, March 2011.

[10] M.N. Nguyen, *et al.*, "Compact, intelligent, digitally controlled IGBT gate drivers for a PEBB-based ILC Marx modulator," IPAC'10, Kyoto, Japan, May 2010.

[11] K. Macken, *et al.*, "A hierarchical control architecture for a PEBB-based ILC Marx modulator," in *Proc. IEEE Int. Pulse Power Conference*, Washington, DC, June 2009

[12] M.A. Kemp, *et al.*, "The ILC P2 Marx and application of the Marx topology to future accelerators," PAC'11, New York, March 2011.

[13] M.A. Kemp, *et al*, "Lifetime tests on a high ohms/square metalized high crystalline polypropylene film capacitor with application to a Marx modulator," in *Proc. IEEE int. Power Modulator and High Voltage Conference,* May 23-27, 2010, Atlanta, Ga.

[14] C.W. Reed and S.W. Cichanowski, "The fundamentals of aging in HV polymer-film capacitors," *IEEE Trans. Dielectrics and Electrical Insulation*, vol. 1, pp. 904-922, Oct. 1994.

[15] A. Benwell, "High voltage design review pertaining to breakdown and flashover," SLAC Power Conversion Department internal note, Feb. 2010.

[16] Maxwell 12. Canonsburg , PA: Ansys, 2008.

[17] Matlab 7.4, Natick, MA: Mathworks, 2007.

[18] Solid Edge 20. Munich, Germany: Siemens, 2006

[19] A. Benwell, *et al.*, "A prognostic method for scheduling maintenance on the P2-Marx modulator," 2010 IEEE Power Modulator Conference, Atlanta, GA, 2010.