NORMAL-CONDUCTING RF STRUCTURE TEST FACILITIES AND RESULTS *

C. Adolphsen
Stanford Linear Accelerator Center, Stanford University, Stanford CA 94309 USA

Abstract

The designs for a next-generation linear collider based on normal-conducting rf structures require operation at gradients much higher than those in existing linacs. For the NLC/JLC 11.4-GHz structures, the design unloaded gradient is 65 MV/m, which is about four times that of the 2.9-GHz SLAC Linac. The CLIC proposal using 30-GHz structures requires an even higher gradient, 170 MV/m. Both the CLIC and NLC/JLC groups are aggressively pursuing programs to develop structures that operate reliably at these gradients and also have acceptable efficiencies and transverse wakefields. Much progress has been made in the past few years, and this paper reviews the programs, test facilities and results from this research.

INTRODUCTION

During the past six years, the major R&D on normal-conducting accelerator structures has been in support of two linear collider initiatives [1]. The SLAC, FNAL and KEK based NLC/JLC design proposes 11.4-GHz rf technology as a 'modest' extrapolation of that of the 2.9 GHz SLAC Linac. The CERN-based CLIC design aims at multi-TeV operation, proposing a higher rf frequency, 30 GHz, at the limit to which conventional copper fabrication techniques can still be used. The desire for a higher operating frequency results from the associated cost benefits of having lower rf energy per pulse (hence fewer rf sources) and a higher operating gradient with reasonable structure efficiency (hence a shorter linac).

Higher frequencies have several drawbacks including the large transverse wakefields generated in the structures by the beam. The structure designs are strongly influenced by the need to limit emittance growth caused by short-range wakefields. In particular, it tightly constrains the allowed minimum average structure iris radius (a), which would otherwise be reduced to improve efficiency. The average radii currently being considered by both groups are about 17% of the rf wavelength (λ). This constraint has had implications for the structure high-gradient performance, as will become evident below.

The long-range transverse wakefields also need to be suppressed, and both designs use a combination of dipole-mode detuning and damping for this purpose. For the NLC/JLC structures, moderate damping is achieved by coupling the cells to four circular waveguides (manifolds) that run parallel to the structure. The CLIC design achieves stronger damping with four terminated, transverse waveguides attached to each cell. In each approach, there is enhanced pulsed heating near the openings in the cells that is a concern for high gradient operation.

The greatest challenge related to the choice of high rf frequency has been in achieving the design gradients. Both groups propose constant-gradient, traveling-wave structures where the upstream ends of the structures operate near the unloaded gradient. The design unloaded gradient for NLC/JLC is 65 MV/m (52 MV/m loaded), and for CLIC, it is 170 MV/m (150 MV/m loaded). The rf pulse length is tied to the frequency choice through the Q's of the structure cells. For efficiency, the pulse length is several times the structure fill time, and the fill time is roughly equal to the attenuation time, $Q/\omega \sim \omega^{-3/2}$. For NLC/JLC (CLIC) the rf pulse length is 400 ns (130 ns) and the structure fill time is 110 ns (30 ns).

Both groups had similar initial experiences in developing structures to meet these goals. In the mid-to-late 1990s, each group built and commissioned a test facility aimed at demonstrating rf system performance (for CLIC, the CLIC Test Facility II, or CTF II, at CERN, and for NLC, the NLC Test Accelerator, or NLCTA, at SLAC). During this time, high-gradient operation was not a major concern since earlier results showed >100 MV/m gradients were achievable at X-band, and that higher gradients seemed possible at higher frequencies [2]. However, when the test facilities came into full operation in 1999-2000, it became clear that the structures tested at that time, which required higher input powers and longer pulse lengths than in the earlier tests, would not meet performance requirements due to breakdown-related damage.

In the following sections, the NLC/JLC and CLIC high-gradient programs, facilities and test results are reviewed with the emphasis on the past year's findings [3].

NLCTA STRUCTURE TESTING

The first structures installed in the NLCTA linac were NLC prototypes built in part to test wakefield detuning and damping. They are 1.8-m long (206 cells), traveling wave ($2\pi/3$ phase advance per cell), nearly constant gradient (the group velocity varies from 12% to 3% c), and have an a/λ = 18% aperture. At the completion of NLCTA commissioning in 1997, four such structures had operated concurrently at unloaded gradients of 40-45 MV/m, close to the 50 MV/m requirement for the 0.5 TeV NLC design at that time. During 1997-99, the two linac power sources were upgraded to allow 70 MV/m operation, which had become the 0.5 TeV and 1 TeV NLC/JLC

* Work Supported by DOE Contract DE-AC03-76F00515.

design unloaded gradient. It was during this period that performance limitations due to breakdown were fully realized, which initiated the high-gradient structure design program that is ongoing today. Since 2001, power from the two linac rf stations has been used to test 21 structures (up to four at a time), with over 10,000 hours of operation logged at 60 Hz [4].

The important structure performance metrics are (1) the time it takes to process to an unloaded gradient about 10% higher than nominal, (2) the breakdown rate at the nominal unloaded gradient and (3) the damage incurred during normal operation. With the large number of NLC/JLC structures, the processing time per structure should be less than a few hundred hours so they can be realistically conditioned during the five-year production period envisioned for a linear collider. Structure damage is characterized by the change in net phase advance through the structure and a few-degree shift per year would likely be acceptable, although much less is expected with stable operation.

An acceptable trip rate has been defined as one that would rarely (once a year) deplete the planned 2% reserve of NLC/JLC power sources assuming a 10-second recovery period after a trip (such periods have been shown feasible). For the 60-Hz operation at NLCTA, this requirement translates to < 0.1 trip per hour for a 60-cm structure. To be conservative, the structures are qualified at the 65 MV/m unloaded gradient with 400-ns square rf pulses. The corresponding NLC/JLC trip rate would likely be lower since the pulse length is effectively shorter (due to the 100-ns ramp for beam loading compensation) and the gradient lower (20% on average due to beam loading).

NLC/JLC TEST RESULTS

The past year saw the transition from testing experimental structures (so called T-Series structures) to testing those designed for use in the NLC/JLC (so called H-Series structures). The former had been built to examine how performance depends on structure length and group velocity. This study was motivated by the pattern of damage observed in the 1.8-m structures: the high group velocity (> 5% c) portions incurred extensive pitting and phase change with operation above about 50 MV/m, while the low group velocity portions remained relatively unscathed.

Seven T-Series structures were built with different lengths (20, 53 and 105 cm) and initial group velocities (5% c and 3% c). By mid-2002, six structures had been tested and showed that breakdown-related damage decreased significantly at lower group velocity and that structure length had little effect on performance. However, the breakdown rates in the input and output coupler cells were noticeably high. At 70 MV/m in the 3% c structures, the coupler rates were generally > 0.3 per hour while for the other 59 cells combined, they ranged from < 0.1 per hour to 0.3 per hour.

An autopsy of the input couplers revealed melting on the edges of the waveguide openings to the cell and extensive pitting near these edges and on the coupler iris. It was subsequently realized that the waveguide edges see large rf currents and the associated pulse heating can be significant if the edges are sharp. By design, the edges in the T-Series structures were much sharper (76-µm radius) than those in the 1.8-m structures (500-µm radius) where this problem was not seen. Simulations showed that the pulse heating for the T-Series structures was in the 130-270 °C range, well below the copper melting temperature, but high enough to produce stress-induced cracking, which can enhance the heating.

To see whether reducing the pulse heating would help, a structure was built using an input coupler design with lower peak currents (a ‘mode-converter’ type [5]) and an output coupler with larger radius (3 mm) edges. For the regular cells, a previously tested T-Series design (53 cm, 3% c) was used. This structure performed very well, with no enhancement of the coupler breakdown rates relative to the other cells. For the full structure, the breakdown rate at 90 MV/m with 400-ns pulses was about 1 per 25 hours. All structures have since been made using similar couplers.

Although the T-Series structure results were encouraging, their average cell iris radii (a/λ = 13%) are too small in that the resulting wakefield is unacceptably large. Thus, the next step was to develop NLC/JLC-compatible versions. The first task was to increase the iris size while keeping the low group velocity and maintaining a reasonable shunt impedance. This required both thickening the irises (from about 2 mm to 4 mm) and increasing the phase advance per cell (from 120º to 150º). The structures were also detuned and their lengths were optimized for efficiency. The resulting designs (H-Series) require higher input power relative to comparable T-Series structures (about 50% more at 3% c). This increase is more than expected from just the larger iris diameter because the thicker irises also reduced the shunt impedance. To date, six such structures have been tested at high gradients.

The first test was of a pair of structures, one 60 cm, 3% c (designated H60VG3), and the other 90 cm, 5% c (designated H90VG5). These structures were built before the coupler problem was remedied and have the sharp-edged design. Also, the pulse heating is higher for the H-Series structures due to the larger stored energy. During processing, coupler breakdowns did indeed limit the gradients to values below those achieved with the T-Series structures. In addition, at short pulse lengths where the coupler events did not dominate, the processing rate was much slower than that for the T-Series structures. The thicker H-Series irises may be a contributing factor since the high-field surface area is about twice as large. Nonetheless, the H60VG3 structure did eventually process to 72 MV/m, and excluding the input coupler, the breakdown rate was < 1 in 10 hours at 65 MV/m.

The next four structures processed to 70 MV/m much faster than the first two, and in one case, faster than most of the T-Series structures. None of these structures showed an enhanced coupler breakdown rate. However, at gradients > 70 MV/m with 400-ns pulses, the processing was slowed considerably by what are aptly called 'spitfests.' These are a series of breakdowns close to each other in location and in time (in many cases, they occur during the ramp-up sequence in power and pulse width that follows a breakdown). The damage caused by the initial breakdown in these sequences is likely creating additional breakdown sites. Such events were observed in the 1.8-m structures at lower gradients, and in the T-Series structures at higher gradients. In all cases, they become noticeable at roughly the same structure input power level (60-80 MW).

Two of the four structures are H60VG3 designs built by the FNAL structures group, which are the first ones they produced for testing [6]. Problems with the brazing furnace used for assembly caused a slight oxidation to the copper, which was removed from the second structure using a hydrogen bake. At NLCTA, the first structure would not process above 70 MV/m due to the large number of spitfests. The second structure performed much better, with a breakdown rate of about 1 per hour at 65 MV/m during a relatively short run. Future structures will be brazed in a custom-designed vacuum furnace at FNAL.

Another of the four structures is a 90-cm, 3% c design (H90VG3), which was built to test the feasibility of longer structures. It requires 30% more power for the same average gradient as H60VG3. When processed, it reached 75 MV/m with 400-ns pulses before being limited by spitfests. At 65 MV/m, the breakdown rate was about 1 per hour, but the spitfest nature of the events was still apparent. About 70% of the breakdowns occurred within 5 minutes of the previous one, although there were periods of up to 25 hours without any breakdown. Fig. 1 shows the dependence of the trip rate on gradient and pulse width after processing. At a fixed breakdown rate, the gradient scales as the -1/6 power of pulse width. By the end of the run, the breakdown rate had decreased to about 1 in 10 hours at 61 MV/m (less than 1 in 6 hours is required for this longer structure).

Roughly half of the breakdowns at 65 MV/m occurred in a few cells (three near the front and one in the middle). After the run, these 'hot' cells were examined using a boroscope. In the interior cell, a 100-µm piece of material that was pitted by breakdown was seen on the outer wall (a similar-looking contaminant in an earlier structure turned out to be aluminum). No obvious foreign material was seen in the upstream cells and a more thorough autopsy is underway.

The last of the four structures is a H60VG3 design whose middle six cells have manifold slots for wakefield damping. They were included to test whether pulse heating near the slot openings leads to breakdowns such as those in the sharp-edged couplers. To reduce this possibility, the slot openings were rounded to minimize the heating. A 43 °C temperature rise is expected at 65 MV/m with 400-ns pulses, which is believed to be safe based on the sharp-edged coupler operation experience.

This structure processed fairly quickly (≈ 50 hours) and achieved 80 MV/m with 400-ns pulses before spitfests began to limit further gains. During this period, there was no obvious breakdown rate difference between the slotted and non-slotted cells. At lower gradients, a few hot cells became apparent, including a slotted one. These cells accounted for a sizeable fraction of breakdowns, as shown in Fig. 2. Since only one of the six slotted cells was hot and hot cells occur elsewhere, the breakdown enhancement in the slotted cell is probably not slot related.

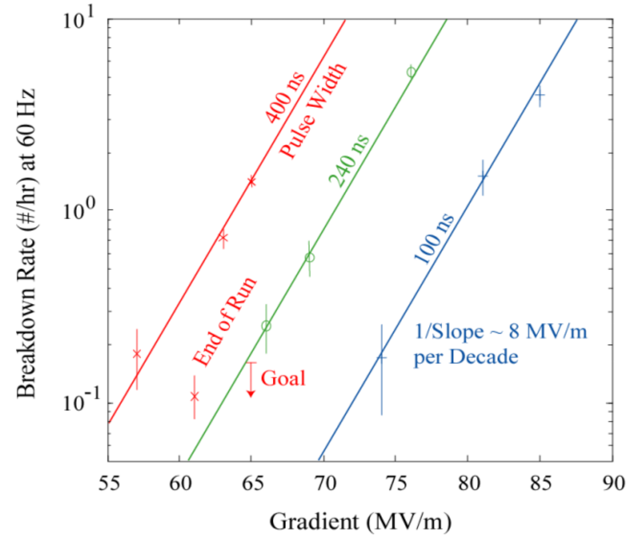

Fig. 1: Breakdown rate versus gradient and pulse width for the H90VG3 structure.

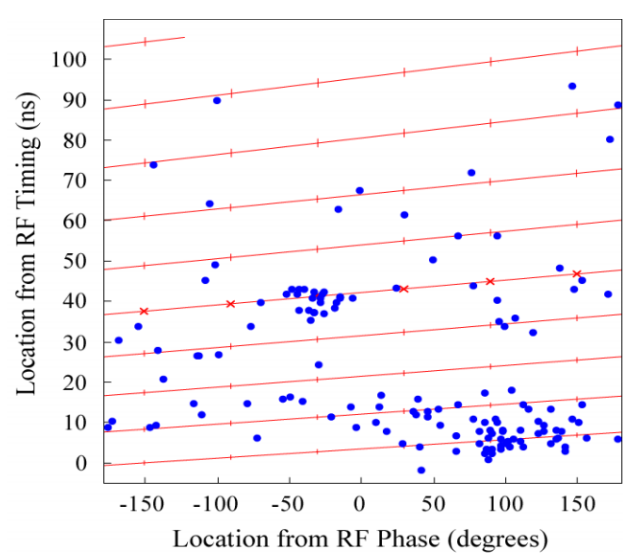

Fig. 2: H60VG3 breakdown locations (dots) at 65 MV/m inferred from reflected and transmitted RF. The cross hatches along the lines are cell locations (crosses are slotted cells) with the first upstream cell at the lower left.

During a 32-day run at 65 MV/m, the average trip rate was 0.21 per hour. Large day-to-day fluctuations occurred (up to 0.7 per hour) from ‘flare-ups’ in the hot cells. The spitfests were much reduced at this gradient, with 25% of breakdowns occurring within 5 minutes of the previous one. At 60 MV/m, the breakdown rate was well below 1 in 10 hours. This structure will be autopsied as well.

NLC/JLC SUMMARY AND OUTLOOK

The H-series results show that in structures with essential NLC/JLC features, reasonable processing times are possible and that breakdown rates close to those required can be achieved (damage rates also appear acceptable). Nonetheless, more operating overhead is desired, and a new version of the 60-cm structure design was recently adopted that has a smaller average iris radius ($a/\lambda = 17\%$) than the present structures ($a/\lambda = 18\%$). Decreasing the radius increases efficiency and will likely improve high-gradient performance due to the lower (7%) input power required. The transverse wakefield is about 20% larger, which is not expected to impact significantly the control of emittance growth in the NLC/JLC linacs.

The main goal within the next year is to operate eight such structures at 65 MV/m in the NLCTA. Tests of experimental structures will also continue, including ones aimed at achieving significantly higher gradients. One test will be of CERN-built, X-band structures with molybdenum and tungsten irises (part of the CLIC program described below). Another test will be of standing-wave structures with rounded-edged couplers. Earlier versions were limited to about 60 MV/m gradients, which may have been the result of their sharp-edged couplers [7].

At a more fundamental level, many questions remain such as the origin of breakdown and the large (×10) variation in processing rates. Surface analyses such as SEM, EDS, and AES have not found contaminants linked to breakdown sites, and residual gas analyses have not revealed any unexpected impurities. Also, the link between field emission and breakdown is unclear since the size of emitters would have to be very small (« micron) to explain the low dark currents (≈ 1 mA) that have been measured. Finally, the relationships among pulse heating, surface melting and breakdown are poorly understood. Not all couplers with high pulsed heating broke down so other conditions may be required, such as some minimum surface electric field. These and other questions will continue to be explored in future structure, waveguide and single-cell tests.

CLIC STRUCTURE DEVELOPMENT

The CLIC Test Facility II (CTF II) basically consists of two parallel beam lines (drive and probe) separated by about a meter [8]. It was designed when the CLIC linac required shorter rf pulses (12-18 ns) and lower gradients (80 MV/m). The drive-beam injector produces 16-ns long trains of 48 bunches with several nC per bunch, and the probe-beam injector generates single, 0.6-nC bunches. Both beams are first accelerated to about 45 MeV in individual S-band linacs. The drive beam is then de-accelerated in transfer structures to generate 15 ns, 30-GHz rf pulses. This power is fed locally to 30-GHz structures in the probe linac to accelerate the probe bunch.

When CTF II commissioning concluded in 1999, the probe linac contained five CLIC Accelerator Structures (CAS), each constant impedance, 28-cm long (86 cells) with a $2\pi/3$ phase advance, 8.6 % c group velocity, and an $a/\lambda = 0.2$ iris aperture. A constant-impedance design was used to simplify cell fabrication, but has the disadvantage that only the upstream cells ‘see’ the highest gradient (the gradient decreases by 20% along the structure). Moreover, the single-feed power coupler used in the CAS design enhances the surface fields on the coupler irises across from the waveguide opening by about 40%.

The basic results from the CAS testing were that processing plateaued after a 250-MV/m peak surface field was attained on the input coupler irises, and that rf breakdown severely eroded (≈ 100-μm deep) the high field surfaces, independent of how the irises were prepared. The corresponding first cell gradient was about 60 MV/m with the 15 pulses, far from the 170 MV/m required with 130-ns pulses.

These results led to the consideration of iris materials with melting temperatures higher than copper (1100 °C). For the structure to be practical, a reasonable conductivity (> 1/3 of copper) was required. Tungsten (W) was the first material tested due to its very high (3400 °C) melting temperature (its use in high-voltage switching applications also made it a plausible choice). In the initial tests, only the CAS input coupler iris material was changed (a clamped-on coupler was used for this purpose). The tungsten irises sustained higher surface fields (320 MV/m), limited by breakdowns elsewhere (mainly on the copper output coupler irises after a 250 MV/m surface field was reached). This success prompted a new structure design to exploit the higher possible surface fields.

For the couplers, a mode-converter design was used that has a 15% lower peak surface field than in the neighboring cells [9]. For the regular cells, the surface to on-axis field ratio was reduced from 2.8 to 2.2 by making the irises thicker and their radii smaller ($a/\lambda = 0.175$). The smaller radius is acceptable for CLIC, but it cannot be reduced much further because of wakefield problems. The success of the clamped-on couplers led to the adoption of a fully ‘clampable’ design where four bolts run through the structure to hold it together. The cells are made of copper and include a 10-mm OD slot to accommodate snugly fitting, interchangeable irises. With the smaller structure group velocity (4.6% c), the structure was shortened to 30 regular cells (10 cm) to keep the gradient attenuation at about 20% with all-copper cells.

Three such structures have been tested to date, one with ground tungsten irises, one with ground molybdenum (Mo) irises, and for comparison, one with standard, brazed copper cells [10]. When processed, the all-copper structure performed as expected, with little gain after a 250 MV/m surface field was reached on the first cell iris (later visual examination showed some erosion at the highest field regions of this cell). With the lower surface to on-axis field ratio, the first cell gradient was 110 MV/m at this limit, almost twice the CAS result.

The gradients achieved in the other structures were limited by the allowed processing time. After about three million pulses, the gradient reached in the first cell of the W structure was 150 MV/m (340 MV/m surface), and in the Mo structure, it was 193 MV/m (426 MV/m surface). As a candidate iris material, molybdenum was chosen because it is relatively easy to machine and has good HV hold-off properties at DC. Its melting temperature (2600 °C) is midway between copper and tungsten, so its faster processing is somewhat surprising. However, this may be a consequence of the Mo iris heat treatment (900 °C vacuum firing), which was done to prevent the high out-gassing experienced with the non-heat-treated W irises. When disassembled, the high-field regions on both the W and Mo first cell irises showed surface-layer melting but no obvious erosion; a more thorough examination is underway.

While the Mo structure results showed that CLIC-like gradients are possible, many challenges remain to achieve a 'CLIC-ready' structure. First, an acceptable breakdown rate at 170 MV/m needs to be shown at the 130-ns design pulse width in longer (≈ 80 cell), constant-gradient structures. Also, if Mo or W irises are ultimately used, their surface area must be made smaller or else the Q loss would be prohibitive. If only the high field regions (> 50% of maximum) of the irises are made with these materials, a modest Q loss could be achieved (5%), although the fabrication of such tipped irises may be difficult. Finally, four radial waveguides need to be added to the outer region of each cell for wakefield damping. Pulse heating near the waveguide openings to the cells is of particular concern, given the high rf frequency and required field levels. In the latest design aimed at minimizing the heating, a 120 °C pulse temperature rise would occur under CLIC operating conditions, while a temperature rise below 60 °C is desired.

As for future testing, CTF II was decommissioned in October, 2002, and its components are being used for CTF 3. This facility will produce longer (130 ns), higher current (35 A), higher frequency (30 GHz) drive beams using the type of combiner rings envisioned for CLIC. High-power testing should resume in spring 2004 using the 5-A, 3-GHz injection beam for CTF 3. In the meantime, versions of the clamped structure scaled to X-band will be tested at NLCTA with Mo and W irises.

SUMMARY

The results from three generations of NLC/JLC structure tests show the need to operate at power levels (60-80 MW at 400 ns) below the regime where damage occurs. At this limit, the peak iris surface fields vary significantly: about 110 MV/m in the 1.8-m structures, 140 MV/m in the H-Series structures and up to 195 MV/m in the T-Series structures. In contrast, operation of the copper CLIC structures with 15-ns pulses appears to be surface-field limited (at about 250 MV/m). At this limit, the input power differs significantly for the two structure designs tested (23 MW and 35 MW). This insensitivity to power compared to the NLC/JLC structures may be due to the higher surface fields, shorter pulses and the much higher breakdown rates at which they were operated.

In regard to demonstrating a feasible design, the NLC/JLC group has operated structures with essential linear-collider features that basically meet the 65 MV/m performance requirements (at 60 MV/m, breakdown rates well within spec have been achieved). A lower-power structure design has been recently adopted that should prove more robust. The CLIC group has attained over 300 MV/m surface fields on W and Mo irises with 15-ns pulses. Although CLIC-level gradients were achieved in a Mo iris structure, significant challenges remain to demonstrate acceptable performance with 130-ns pulses in longer, damped, efficient structures.

REFERENCES

[1] International Linear Collider Technical Review Committee Second Report, SLAC-R-606, 2003.

[2] J. Wang and G. Loew, "RF Breakdown Studies in Copper Electron Linac Structures," SLAC-PUB-4866, March 1989.

[3] For an expanded version of this paper, see SLAC-PUB-9906.

[4] For earlier results, see C. Adolphsen, et al., SLAC-PUB-8573 and SLAC-PUB-8901, and J. Wang et al., TH464, LINAC 2002, August 2002.

[5] C Nantista et al., "Novel Accelerator Structure Couplers," these proceedings and Ivan Gonin et al., "Coupler Design for NLC/JLC Accelerating Structures," these proceedings.

[6] N. Solyak et al., "Development of X-Band Accelerating Structures at Fermilab," these proceedings.

[7] V. Dolgashev et al., "Status of X-band SW Structure Studies at SLAC," these proceedings.

[8] H. Braun, "Experimental Results and Technical R&D at CTF II", EPAC2000, June 2000.

[9] I. Syratchev, "Mode Launcher as an Alternative Approach to the Cavity-Based RF Coupler of Periodic Structures," CERN PS/RF/Note 2002-013.

[10] W. Wuensch et al., "A Demonstration of High-Gradient Acceleration," these proceedings.