DIRECT PHOTON PRODUCTION*, **

By T. Ferbel

Department of Physics and Astronomy, University of Rochester, Rochester, NY 14627, USA

(*Received October 4, 1985*)

First, I summarize the physics issues involved in studies of prompt photons in hadronic collisions at large transverse momenta (p_T); next, I review the latest comparisons between data and predictions from QCD; finally, I discuss some of the features of a second-generation experiment at Fermilab (E706) that is designed to study the flavor dependence of direct photon production at beam energies of 600–900 GeV.

PACS numbers: 13.85.Qk

Introduction

Extensive recent studies of e^+e^- collisions, of deep-inelastic scattering of leptons, and of hadroproduction at large p_T has made it difficult to doubt that Quantum Chromodynamics (QCD) has great relevance to strong interactions. Results on jet production at colliders, for example, have shown that: (*i*) scattering of confined constituents accounts for all the essential features of the data, (*ii*) nucleon structure functions measured at lower energies and low Q^2 values, when evolved to larger Q^2, have the x-dependence that is inferred from data at high energy, and (*iii*) two kinds of constituents appear to be present within hadrons: quarks and gluons. The production of direct photons at large p_T, at rates comparable to pions, provides particularly cogent support for the existence of confined gluons [1].

Figure 1 depicts the two leading-order QCD graphs thought responsible for the production of direct photons in hadronic collisions. The Compton graph is dominant in nucleon-nucleon reactions at all p_T, while the annihilation diagram can be important in meson-nucleon or in antinucleon-nucleon collisions, especially at large $x_T = 2p_T/\sqrt{s}$ values. From Fig. 1, it is clear that triggering an experiment on a direct photon is equivalent

* Research supported by the Department of Energy.

** Presented at the XXV Cracow School of Theoretical Physics, Zakopane, Poland, June 2–14, 1985.

to selecting gluon physics: (*i*) probing gluon structure of hadrons in the case of the Compton diagram, or (*ii*) measuring gluon fragmentation in the case of the annihilation graph.

The advantages that direct-photon production provides over jet studies for examining QCD phenomenology are primarily due to: (*i*) the ability to isolate specific mechanisms that contribute to any given process, and (*ii*) the simplification in the unfolding of the dynamics that is realized as a result of the fact that the produced photon is an observable object that couples in a point-like way to the electrical charge of a quark.

Unfortunately, as for all reactions measured at finite Q^2 values, the leading order graphs in Fig. 1 cannot correspond to the full perturbative-QCD story. Bremsstrahlung

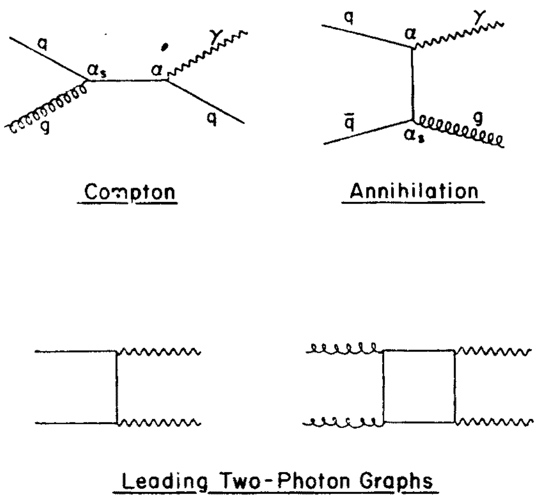

Fig. 1. Leading-order QCD graphs that contribute to direct-photon production in hadronic collisions. Also shown are the two important diagrams for di-photon yields

in quark-quark or quark-gluon scattering (or gluon-gluon in higher order) can also provide direct photons. Higher-twist contributions, that have a softer p_T dependence, can provide observable direct-photon signals in certain unique regions of phase space (large x_T). Most important to understand are the higher-order QCD contributions. The complete calculation to order α_s^2 for pp collisions [2] indicates that an effective "*k*-factor" of about 2 is appropriate for $\sqrt{s} \sim 50$ GeV and photon p_T's of $\sim$5–10 GeV/*c* [3]. Although the *k*-factor is, clearly, not just an overall multiplicative factor, and depends on the definitions used for the Q^2-scales, the predicted p_T-dependence of the direct-photon cross section is nevertheless relatively insensitive to the ultimate choices used for Q^2 [2]

A great amount has been learned about direct-photon production over these past ten years. Broadly speaking, most experimenters would agree that photons are produced copiously in hadronic collisions, and that the γ/π^0 ratio in pp and πp reactions for $\sqrt{s}$ values of 20–60 GeV is small ($\lesssim 0.05$) at $p_T \lesssim 3$ GeV/*c* and is substantial ($\gtrsim 0.5$) at $p_T \sim 10$ GeV/*c*. Most phenomenologists would agree that the yield of direct photons in pp collisions is consistent to within a factor of $\sim$2 with expectations from QCD.

Comparison of data with QCD

Recent preliminary measurements from the NA-24 collaboration are shown in Figs. 2 and 3 [4]. The pp results are similar to those obtained previously at the ISR and at Fermilab; however, the $\pi^{\pm}p$ data are of far better quality than available before and can provide an excellent check of QCD, in that the dominance of the annihilation graph in $\pi^{-}p$ reactions requires the ratio of the γ-yields in $\pi^{-}p/\pi^{+}p$ to increase with x_T. A preliminary test of this

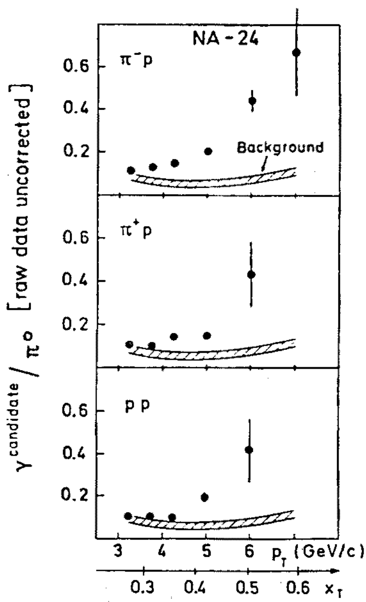

Fig. 2. The ratio of the yield of single-photon showers to that of π^0's. The contribution to the γ/π^0 ratio from known background sources is shown as the shaded band. The results from NA24 at CERN are preliminary

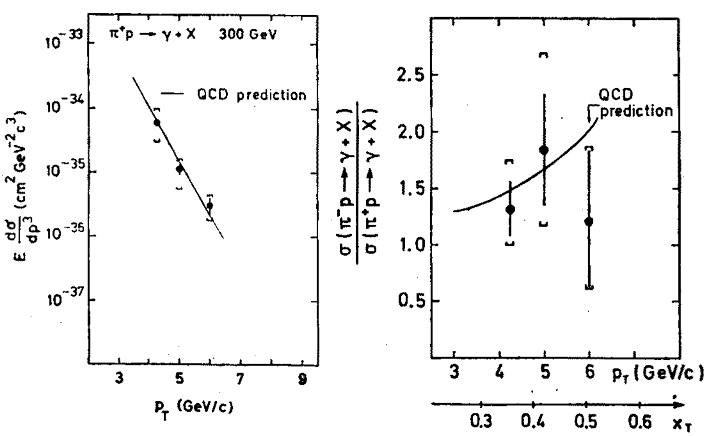

Fig. 3. The yield of direct photons in the preliminary $\pi^{+}p$ data of NA24. The ratio of yields in $\pi^{-}p$ and $\pi^{+}p$ collisions is also shown in the figure. The QCD calculation is that of Aurenche et al., (Ref. [1])

prediction is also shown in Fig. 3. At the present level of statistics, the data do not provide a very stringent test of the phenomenology. Nevertheless, the absolute yields of direct photons are consistent with observations in previous experiments and with QCD.

Azimuthal correlations between direct photons and accompanying hadrons are also consistent with previous measurements and with phenomenology. Namely, prompt photons are emitted without accompanying particles, while π^0's are emitted with other hadrons. This is to be expected because photons emerge isolated from the primary collisions (Fig. 1),

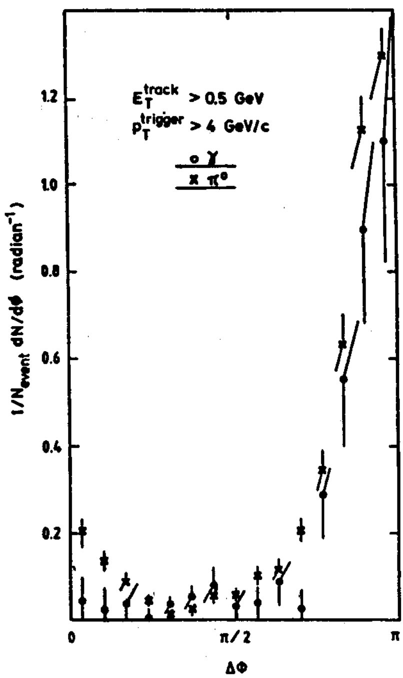

Fig. 4. Azimuthal correlation between direct photons and any accompanying hadrons, observed in CERN experiment NA24 (preliminary). Similar data, for π^0 triggers, are shown for comparison

while π^0's are just leading remnants of constituents and are therefore accompanied by other hadrons from the fragmentation of the common constituent. The new data are displayed in Fig. 4 [4]. Particle multiplicity on the away-side of the trigger is larger than on the same side, and similar for photon and π^0 events.

There is evidence from the ISR [1] that the away-side charge correlation, defined as the ratio of number of positively to negatively charged particles recoiling from direct photon (or from π^0) triggers, is similar to that expected from the dominance of the Compton graph for direct photons. That is, for direct-photon triggers, the recoiling constituent is generally an up quark (charge $+2/3$), while for π^0 triggers the recoil can be a gluon, an up or a down quark. Thus, qualitatively, there should be relatively more positive particles recoiling from direct photons than from π^0's. Although the ISR data are limited by poor statistics, they provide, nevertheless, confirmation of this effect.

From comparisons of yields at Fermilab and ISR, there is also some evidence that the p_T dependence of direct-photon production is comparable to that measured for jets at similar energies [1]. This provides support for a point-like source for direct photons.

Choice of technique

Measuring prompt-photon production is a rather challenging problem, particularly because of the presence of background, mainly from π^0 decays. There are two ways to proceed. One is the statistical conversion method, as depicted in Fig. 5, and the other is the conceptually simpler direct measurement, shown in Fig. 6.

The conversion method relies on the presence of a relatively coarse-grained calorimeter of high detection efficiency. Let us assume that the two decay photons from any π^0 are close enough in space so that the π^0 is observed as a single electromagnetic shower (complete

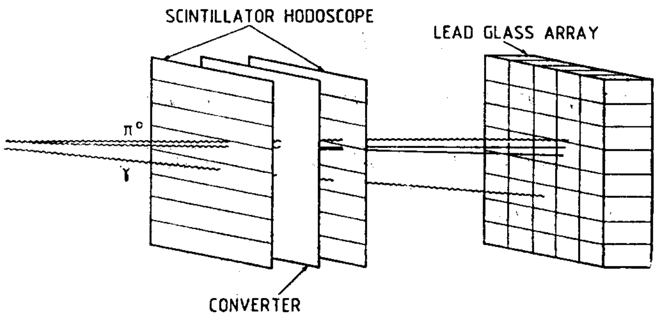

Fig. 5. Sketch of statistical conversion method for measuring direct photons

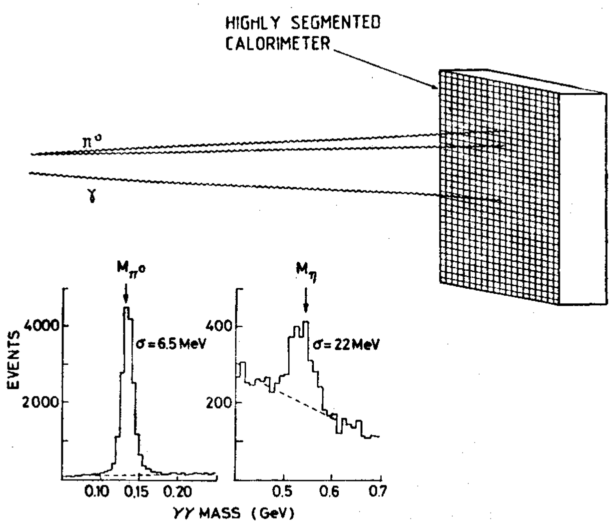

Fig. 6. The direct method of measuring prompt photons. Any two photons observed in the detector can be paired to see if π^0 or η signals are present

overlap of two γ-showers) in the calorimeter. If there is some mixture of γ's and π^0's produced in the hadronic collision, then, from the presence or absence of conversion signals (e^+e^- pairs) in the scintillation counter just downstream of the converter material (radiator), we can determine this γ/π^0 mix. For a single photon, the probability that the γ does not convert in the converter ("the non-conversion probability") is:

$$P_\gamma(\mathrm{NC}) \sim \exp\left(-\tfrac{9}{7}x\right), \tag{1}$$

where x is the thickness of the converter in units of radiation lengths. The non-conversion probability for a π^0 is just:

$$P_{\pi^0}(\mathrm{NC}) \sim [P_\gamma(\mathrm{NC})]^2 \sim \exp\left(-\tfrac{18}{7}x\right) \tag{2}$$

(we are neglecting the weak energy dependence of the photon conversion probability). Thus, for N_γ-photons and N_{π^0}–π^0's produced in the target, defining the photon fraction as:

$$f_\gamma = N_\gamma/(N_\gamma + N_{\pi^0}), \tag{3}$$

we expect the following non-conversion probability for our converter:

$$P_{\mathrm{EXP}}(\mathrm{NC}) = f_\gamma P_\gamma(\mathrm{NC}) + (1-f_\gamma)P_{\pi^0}(\mathrm{NC}). \tag{4}$$

This can be inverted and solved for f_γ:

$$f_\gamma = \frac{P_{\mathrm{EXP}}(\mathrm{NC}) - [P_\gamma(\mathrm{NC})]^2}{P_\gamma(\mathrm{NC})\,[1-P_\gamma(\mathrm{NC})]}, \tag{5}$$

where we have used relation (2).

Now, of course, in any experiment we usually do not know f_γ. However, we can calculate $P_\gamma(\mathrm{NC})$ from the geometry and energy dependence of the conversion probability, and we can also measure the experimental value of the non-conversion probability through observing how often we see hits in the converter for showers observed in the calorimeter (this can be done as a function of p_T, E_T, etc., of the showers). Hence, by measuring $P_{\mathrm{EXP}}(\mathrm{NC})$ we can, in principle, use relation (5) to determine f_γ. The idea of this method is straightforward, its execution, however, is not so trivial. First, when f_γ is small, high statistics are needed to reduce the error on f_γ. (For typical values of $P_\gamma(\mathrm{NC})$ of ~ 0.5, the measurement involves the ratio of two small differences.) The most important difficulty with this kind of method is that the experimental value of $P_{\mathrm{EXP}}(\mathrm{NC})$ assumes that all individual π^0's and γ's can be observed separately in the calorimeter. For example, when showers from several π^0's overlap completely in the calorimeter, there is no way to determine how many particles entered the converter, and the value of $P_{\mathrm{EXP}}(\mathrm{NC})$ for such events is therefore ambiguous. To extract the yield of direct photons, data for $P_{\mathrm{EXP}}(\mathrm{NC})$ must be corrected for overlapping showers from jets, as well as for other neutral hadrons such as neutrons, K_L^0, etc., that do not convert in the converter but can simulate electromagnetic showers in the calorimeter. Such corrections are uncertain because they rely on models of jet decay.

The direct method, sketched in Fig. 6, is cleaner than the conversion technique, and is especially valuable when f_γ is small. The high degree of segmentation needed in the calorimeter to separate the two photons from, e.g., π^0 decay makes the apparatus expensive and difficult to construct. This method provides important checks on data that are not available using the conversion approach. For example, the π^0 decay angular distribution in its own rest frame is isotropic. Any observed departure from isotropy implies the presence of biases in the data. Consequently, measuring the angular decay of π^0's, and comparing

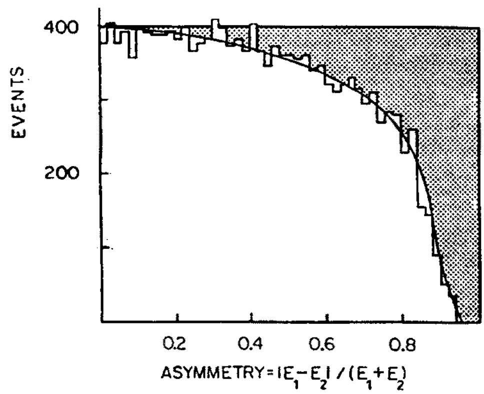

Fig. 7. Energy asymmetry of photons from π^0's observed in E629. The smooth curve is the expectation for this distribution based on a Monte Carlo of the experiment. The shaded region corresponds to π^0's that were not detected, and consequently contributed to the background in the direct-photon signal

that distribution to that expected from a Monte Carlo of the experiment, can provide the background from π^0's that is expected in the direct-photon signal. An example is given in Fig. 7 (from Fermilab Experiment E629) [1]; shown is the measured energy asymmetry in the laboratory for the two photons from detected π^0's. (For ultrarelativistic π^0's, the asymmetry parameter is the same as the cosine of the decay angle of a photon relative to the π^0 line of flight, that is, the helicity angle.) The cross-hatched area corresponds to π^0's that were not detected in the experiment. These π^0's have photons that differ substantially in energy, and, typically, the lower energy photon is lost. These π^0's consequently contribute to the background to the direct-photon signal. Figure 7 shows how many background events to expect (the cross hatched region), and because the Monte Carlo agrees with the data, the correction to the direct-photon yield can therefore be applied with some degree of confidence.

Experiment E706

Fermilab Experiment E706 is a second-generation attempt at extracting precise information from direct-photon production in hadronic collisions in the 500–1000 GeV beam momentum range. It is based on the use of a highly-segmented liquid-argon calorimeter (LAC) that will be able to distinguish γ's from π^0's up to 200 GeV/c in momentum. E706 is a collaboration among physicists from the University of Delhi, Fermilab, Michigan State University, the University of Minnesota, Northeastern University, Penn State University, the University of Pittsburgh, the University of Rajasthan and the University of Rochester. The technical details of the calorimeter have been described in the literature [5].

A layout of the E706 spectrometer is given in Fig. 8. The acceptance is typically $\gtrsim 80\%$ of 4π. The LAC has both electromagnetic and hadronic sections. The latter will be particularly important in measuring the neutral components of jet decays. This is vital because, to provide useful guidance to the phenomenology of QCD, the direct photons must be

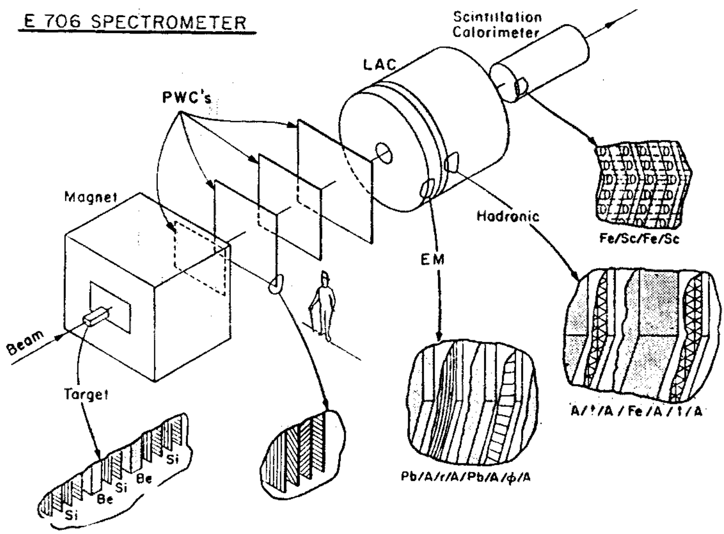

Fig. 8. Layout of experiment E706

measured along with their associated jets. Complete measurements of this kind can be used to reconstruct the x-values of the incident constituents in the collision and thereby provide a direct measure of the x-dependence of the structure functions.

The goals of E706 can be summarized as follows: (*i*) To learn about the Q^2 and z-dependence of gluon fragmentation. (*ii*) To learn about the Q^2 and x-dependence of the gluon structure functions of hadrons. (*iii*) Measure the A-dependence of the direct-photon yield in nuclei and study the character of the recoil jets as a function of A. (*iv*) Using the hadronic part of the LAC (HALAC), study associated jet production at the edge of phase space. (*v*) Compare di-lepton with di-photon yield to extract additional information on gluon structure functions and check calculations for higher-order contributions to these final states. (The leading di-photon graphs are also shown in Fig. 1.)

The $r-\phi$ geometry of the electromagnetic part of the LAC (EMLAC) is particularly well suited for triggering the experiment on photons of large p_T. The acceptance threshold will be set to ~ 5 GeV/c for single photons and to ~ 3 GeV/c for each of two photons (or two electrons). The entire LAC will be located within one large cylindrical cryostat that will be about 6 m high and 5 m in diameter.

"When you get stretched, you come out a little stronger"

As an aside on the engineering complications encountered in a project of this magnitude, I thought that it would be of interest to describe how one particular problem in the design of the EMLAC was tackled.

We wished to construct a device that was reasonably easy to assemble. Because G-10 (circuit board material) is manufactured in only 122 cm widths (and in lengths of less than 366 cm), and the radius of the EMLAC is 150 cm, there was no way to modularize the read-out boards into quadrants. It was therefore decided to use octants as read-out units, but to abut two against each other so that they effectively formed quadrants. The signals from the octants would be read out on outer edges, leaving an essentially seamless region along the radial between them. Lead sheets for the EMLAC were to come in quadrant size. It was important to maximize the area covered by any single sheet of lead so that dead areas, where photons could not convert, would be minimized. It should be pointed out that the coefficient of thermal expansion of lead and of G-10 material are sufficiently different so that, if the EMLAC stack were assembled with the lead pinned rigidly to the G-10 at room temperature, the cooldown to liquid argon would destroy the entire detector because of the difference in the temperature dependence of the expansion of the two materials. The lead is therefore placed, freestanding, within a G-10 frame, which is in turn sandwiched together with the G-10 read-out boards in the stack. Because these G-10 frames produce dead-space for the EMLAC, the larger are the lead sheets, the less dead-space is there in the detector.

At the time the EMLAC was designed, our group had previous experience with manufacture of 2 mm hardened-lead sheets that were about 1 m × 1.5 m in area. We felt reasonably confident that we could obtain 1.5 m × 1.5 m lead sheets, and consequently chose to modularize the EMLAC into quadrants. We realized that industry could not provide lead sheets to meet our flatness specifications. (We were always amused by the fact that the sales departments of all the likely vendors thought they could quite easily provide the lead within specs, while the plant foremen and mill personnel thought that our requirements were absolutely impossible to meet!) Our experience with a previous LAC [6] was that industry could provide excellent thickness tolerance for the hardened 2 mm lead sheet (the hardening is achieved by adding small quantities of calcium and tin to the lead), but, for plate widths of $\gtrsim 50$ cm, flatness over any sizeable area was exceedingly difficult to maintain.

The 1 m × 1.5 m plates for our previous detector arrived at Fermilab looking like the three-dimensional saddle contours protrayed in that classic mathematical physics book of Morse and Feshbach (they reminded me of giant potato chips!). Following an inspired suggestion of Bill Willis (CERN), we took the lead sheets to a roller-leveler at a local steel mill. Through a miracle of metallurgy, the lead sheets went into the rollers looking like giant potato chips and came out on the other side smooth and flat as jig-plates. Remarkable! Well, it is small wonder that, with that sort of memorable achievement in our minds, we felt confident about our new 1.5 m × 1.5 m lead. What we did not anticipate was that this potato-chip effect in the initial lead rolling grew (according to my metallurgist friends) as some non-negligible power of the width of the sheet. We also did not anticipate that larger roller-levelers at steel mills could not be adjusted with precision for sheets that were wider than 122 cm (the universal 4 ft. limitation!). Consequently, when we received our first several sheets of lead from the vendors, they looked far worse that we had ever imagined they would look, and, what is more, we could not flatten them, at all, at any of the steel mills we investigated. We were in a state of utter panic! The design was essentially completed, it was six months before the LAC assembly was to start, and we saw no way to supply the lead sheet for the experiment. (I had nightmares about calling up Leon Lederman and explaining to him that we will need several million dollars for tungsten plate to replace the lead.)

My normal reaction to a panic situation of the kind we faced was to start calling experts (an expert, as my colleague Paul Slattery explains, is someone out of town whom you do not know very well). I spoke with metallurgists, other physicists, engineers, technicians — anyone who would listen. What slowly emerged from all my gossiping was that all the *cognoscenti* thought that lead should not be roller-leveled but rather stretched past its elastic limit to flatness. Well, the problem now was to find a plant that had a plate stretcher, and that would let us use it on lead.

We eventually found a large stretcher at the Allegheny-Ludlum steel bastion outside of Pittsburgh. We still had two remaining problems to solve. One, was that the steel stretcher could only stretch plate that was over two meters in length, and, second, was that, because it was a steel stretching device, it had steel vice jaws that bit off the lead edges before the sheet could be properly stretched. We performed our successful tests on double-length lead sheet ($1.5 \text{ m} \times 3 \text{ m}$) in mid-winter, with the temperature inside the plant hovering near $-10°\text{C}$, and finished all the stretching during one week of the following summer, with the temperature in the plant holding steadily at $+32°\text{C}$. Figure 9 shows three of our young colleagues during a preliminary production run at Allegheny-Ludlum. (Other disasters befell us in the meantime. One particularly memorable annoyance was the firing of the foreman at the lead mill that manufactured our lead sheet. Unfortunately, with the foreman went the alchemy of producing hardened lead ingot. This was noticed when the modulus of elasticity of the new lead was found to be far smaller than that for our old hardened lead. If this were not caught in time, our lead sheets would not have been able to support their weight and could have flowed or buckled in the detector.) Thus, on a relatively happy note, ended a minor chapter in E706 calorimetry.

Advantages of high energy

Having just gone through one of the challenges that faced E706, and having already indicated that there are several groups that have already taken data in the 200–300 GeV beam-energy range, you might wonder whether the new experiment is really worth all the effort. The answer is, of course, a resounding yes! To convince you of this, I show in Fig. 10 the anticipated ratio of yields of direct photons as a function of p_T for two experiments, one at 500 GeV and one at 250 GeV. It is clear that, in terms of absolute yield at high p_T, E706 will have a great advantage over present-day fixed-target experiments. Another advantage is displayed in Fig. 11, where I have sketched the typical γ/π^0 ratio as a function

Fig. 9. Joey Huston (middle), Clark Chandlee and Phil Gutierrez (at right) aligning lead sheet on a plate stretcher at Allegheny-Ludlum

of p_T. The direct-photon cross section increases markedly with increasing p_T, while the background decreases, thereby yielding a far cleaner experiment at the higher than at the lower energy.

The fact that direct-photon measurements become easier at high energies has been illustrated recently in a dramatic fashion by the UA-2 Collaboration at the S$\bar{p}$pS of CERN [7]. Figure 12 displays the cross section extracted for direct photons at $\sqrt{s} = 630$ GeV. Although the conversion method was used by UA-2, the fact that the γ/π^0 ratio in this p_T

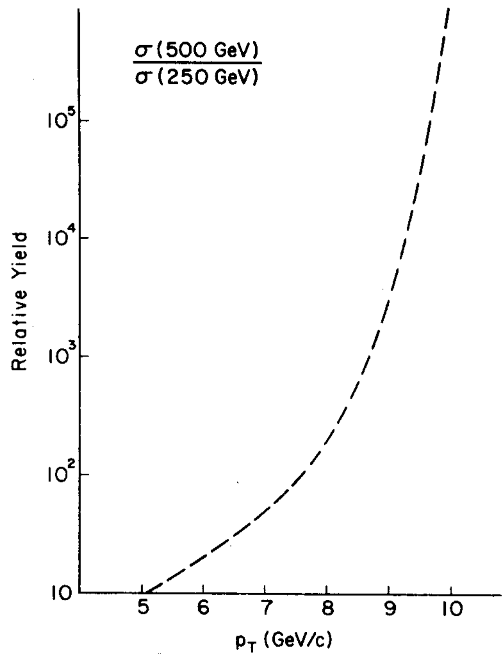

Fig. 10. Anticipated ratio of yields of direct photons at 500 GeV/c relative to 250 GeV/c as a function of p_T

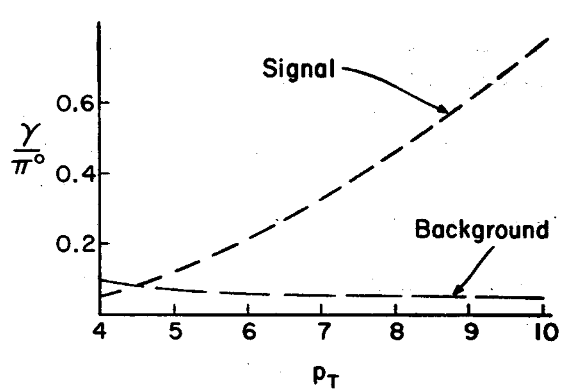

Fig. 11. Typical signal and background for the ratio of the γ/π^0 yields as a function of p_T

range is so large makes the measurement more reliable and less subject to uncertainty from corrections for multi-π^0 overlap. (It should be recognized, however, that an isolation cut was imposed on these data; that is, it was required that the electromagnetic shower not be accompanied by any charged tracks within a cone of $\pm 20°$. This has a small effect on the direct-γ signal, but produces a substantial reduction in the yield of π^0's, η's and other jet fragments.) The calculation of Aurenche et al. [1], corrected to 630 GeV, is seen to agree remarkably well with the data. The isolated γ cross section is, typically, 3 times that of the

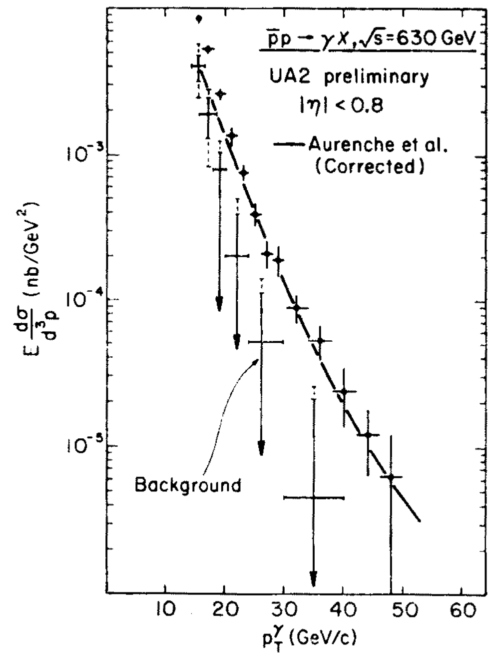

Fig. 12. Direct-photon production observed by UA-2 (preliminary data). The expected γ yield, based on QCD, is shown in the figure

isolated π^0/η production, and, consequently, the background from π^0's or η's cannot possibly account for the large direct photon signal.

At SSC energies, the direct photon yield will be sensitive to quark compositeness [8]. Any departures from predictions of QCD will signify the presence of a new scale. Although there are no operators of dimension six for the process (with amplitude of order $\hat{s}/\Lambda^2$), the shape of the single-photon cross section as a function of p_T (Fig. 13) will still be sensitive to quark substructure.

A final remark on expectations from E706. The point cross section for the Compton graph can be written as:

$$\frac{\pi\alpha\alpha_s}{3s^2}\, e_q^2\, \frac{\hat{u}^2+\hat{s}^2}{\hat{s}\hat{u}}\,, \tag{6}$$

while that for the annihilation graph can be written as:

$$\frac{8\pi\alpha\alpha_s}{9s^2} e_q^2 \frac{\hat{u}^2+\hat{t}^2}{\hat{u}\hat{t}}, \tag{7}$$

where e_q is the charge of the interacting quark, and $\hat{u}^2$, $\hat{t}^2$ and $\hat{s}^2$ are the absolute values of the Mandelstam variables for the colliding constituents. The direct photon yield is just the convolution of these point cross sections with the relevant structure functions. Because gluon coupling is charge-independent, it follows that the yield (to first order in α_s) of direct photons from the Compton graph is the same for π^+p as for π^-p reactions. However, ignoring the sea quarks, the annihilation contribution in π^-p [involving ($\bar{u}$d) quarks interacting with (uud) quarks] is eight times that of π^+p [(u$\bar{d}$) quarks interacting with (uud) quarks]. Thus, to first order, the difference between π^-p and π^+p data is just proportional to the annihilation graph. For a carbon target, with an equal number of up and down quarks, the ratio of contributions (π^-C/π^+C) from the annihilation graph is only four. In E706, we expect that the *difference* in cross sections between π^-C and π^+C γ-yields will amount to 14,000$\pm$250 events for $p_T > 5$ GeV/c. This difference measurement will provide a clean means for studying gluon fragmentation and gluon interaction with nuclear matter (A-dependence). The same sample of data, corresponding to $\sim$1200 hours of running time, will also yield 450$\pm$25 double photon events for masses beyond 10 GeV. This would certainly make the experiment unique.

Figure 14 provides an unfair comparison between last year's data from CERN and the anticipated quality of results from E706. The CERN measurements will certainly improve over the next year or two (they already have for NA24), however, it is clear that E706 will have much to contribute, and particularly, in a regime of x_T that is beyond the

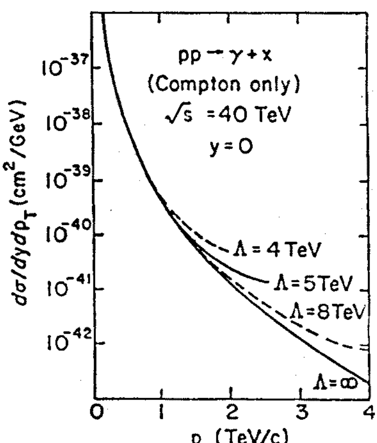

Fig. 13. Predictions for $d\sigma/dydp_T$ based on the dominance of the Compton graph, modified to include effects of compositeness. The parametrization is $\sigma_{\text{Compton}} \times \left[1+\left(\frac{\hat{s}}{\Lambda^2}\right)^2\right]^2$, where $\hat{s}$ is the value of s for the gluon-quark collision ($\sim 4p_T^2$), and Λ is the characteristic mass scale for compositeness

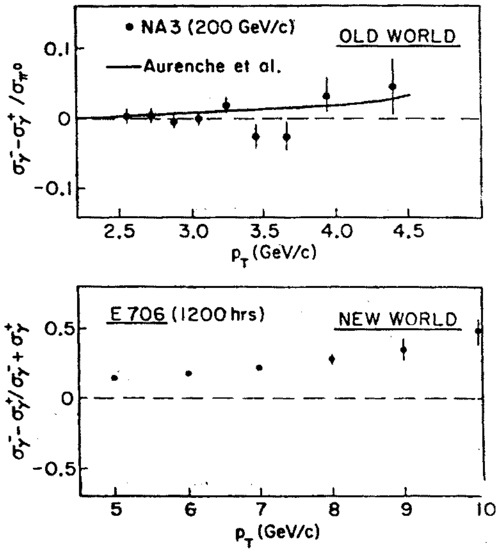

Fig. 14. Difference measurement of direct photon production in π^- and π^+ collisions with nucleons. The "old world" data are from 1984 (Leipzig Meeting), the "new world" data are yet to be run

reach of present-day colliders. Now, if only our calorimeter were guaranteed to work, then our collaboration would have fewer sleepless nights!

I thank the organizers of the Cracow School for inviting me, and for providing the usual warm atmosphere that is so characteristic of Polish hospitality. I also thank Peter Hansen, Klaus Pretzl and Dominique Schiff for comments on this manuscript.

REFERENCES

[1] For a summary of data and phenomenology, see T. Ferbel, W. Molzon, *Rev. Mod. Phys.* **56**, 181 (1984). Since the publication of our review of direct-photon physics, several other interesting papers have appeared in the literature, these include: J. Badier et al., NA3 Collaboration, CERN--EP/84-101 (1984); C. De Marzo et al., NA24 Collaboration, 126 P/ma (1985); E. Berger et al., *Nucl. Phys.* **B239**, 52 (1984); P. Aurenche et al., *Phys. Lett.* **140B**, 87 (1984); A. Contogouris et al., *Phys. Rev.* **D32**, 1134 (1985); and T. Akesson et al., CERN-EP/85-55 (1985). A new calculation of double photon production beyond leading order can be found in P. Aurenche et al., Orsay Preprint LPTHE 85/20 (1985).

[2] See P. Aurenche et al., in Ref. [1].

[3] This is for the standard choice of $Q^2 = p_T^2$. See the phenomenological parametrization of various k-factors for π-nucleon collisions in A. Contogouris et al., of Ref. [1].

[4] This is an update of C. De Marzo et al., of Ref. [1], that was presented by K. Pretzl at the XVI$^{\text{th}}$ Symposium on Multiparticle Dynamics, Israel, June 1985.

[5] F. Lobkowicz et al., *Nucl. Instrum. Methods* **A235**, 332 (1985).

[6] C. Nelson et al., *Nucl. Instrum. Methods* **216**, 381 (1983).

[7] CERN Seminar by J. R. Hansen and the report of P. Hansen at the XVI$^{\text{th}}$ Symposium on Multiparticle Dynamics, Israel, June 1985.

[8] J. Owens et al., Proc. of 1984 Summer Study on the SSC at Snowmass, R. Donaldson and J. Morfin, eds.