ALEPH 89-115
PHYSIC 89-45
R. Miquel et al.
28.6.1989

VAX implementation of Monte Carlo simulation of two photon processes

R. Alemany, V. Gaitan, M. Martinez and R. Miquel
Universitat Autònoma de Barcelona

June 23, 1989

Abstract

Problems with the VAX implementation of the Monte–Carlo event generators for two photon processes DIAG36N and DIAG36ZN are described. After studying the reason for the overflows and underflows occurring, some methods are explained to cure them and have a version running on VAX. Differences in the results of the QED and electroweak versions of the programs (DIAG36N and DIAG36ZN, respectively) are shown.

1 Introduction

The set of programs DIAG36N, DIAG36ZN (both specially designed for no-tagging configurations) has been written by P.H. Daverveldt and include the calculations made by F.A. Berends, P.H. Daverveldt and R. Kleiss [1]. These programs are able to compute the cross section for the process

$$e^+e^- \longrightarrow f_1\bar{f}_1 f_2\bar{f}_2$$

where f_1 and f_2 are any charged fermion (leptons or quarks at the parton level) using a calculation which includes all the Feynman diagrams describing the process at the tree level. DIAG36N takes into account only the QED diagrams whereas DIAG36ZN takes also into account the diagrams with Z^0 exchange. Therefore, they give a reliable cross section for processes such as

$$e^+e^- \longrightarrow e^+e^-\mu^+\mu^-$$

and

$$e^+e^- \longrightarrow e^+e^-e^+e^-$$

which have a huge total cross section, in fact around 200 nb and 10^7 nb respectively when the beam energy is about 50 GeV. Moreover, at LEP energies they may be an important background to other reactions of interest.

Originally these programs where adapted by their author to run on IBM computers and due to differences in the way double precission variables are defined in VAX and IBM Fortran they cannot run directly on VAX.

Our goal is obtaining a single compatible program for VAX and IBM computers. The problem is that standard double precission variables on VAX have less exponent range that on IBM in spite that actually they have one digit more in the mantissa (so that in fact, on VAX, better precission is reachable); for this reason often the program crashes giving errors

of underflow and overflow. However, as we know that the needed precision for the calculation lies well inside the range of the VAX variables, then it must be possible, writing in the proper way the intermediate operations, to fix the problem.

There is an important difference between standard double precission VAX variables and IBM ones : while the last ones allow the exponent to be bigger, in absolute value, than 60, for VAX the magnitude cannot be smaller than approximately $0.29\text{x}10^{-38}$ or bigger than approximately $1.7\text{x}10^{38}$ for the D-floating implementation. The number of digits that preceed the exponent is 16 in D-floating, while IBM computers have less significant digits.

Since VAX computers have two implementations for double precission variables, D-floating and G-floating, one can think naively that the problem will dissapear using the G-floating implementation in the FORTRAN compiler instead of the standard D-floating one. In fact, the G-floating implementation has less significant digits and more exponent range. Nevertheless, in many steps of the programs it is very important not to lose precision on the number of significant digits (due to important cancellations ocurring in the calculations) so that the direct use of the G-floating implementation doesn't solve the problem either. Another inconvenient in the G-floating implementation is that it increases the computing time by a very considerable amount.

Therefore, the solution is to rearrange the arithmetic operations according to the double D-floating precission variables on VAX . The processes mentioned above are the most important and problematic ones because we are working close to the limit of VAX variables accuracy and exponent.

The main reason of all this kind of problems is that the denominator of some propagators can be extremely small (the difference from zero coming only from terms of the order $O(m_e^2/s)$) because the outgoing particles go at very small angle. For instance, in the reaction

$$e^+e^- \longrightarrow e^+e^-\mu^+\mu^-$$

the electron and positron tend to go at very small angle, and this region dominates clearly the total cross section. That implies problems in the DIAG2, DIAG4, DIAM subroutines. (see appendix 1).

2 Modifications

The programs modified are DIAG36N, DIAG36ZN. The subroutines and functions which have been modified necessarily are the following : DIAG2, DIAG4, MCA, MCB, DIAM, GETRID and GRAAF.

The different methods used to solve the problem are :

a) Exchanging the order of variables in a statement.

It is possible that the intermediate steps of some multiplication and division expressions overcome the exponent range of VAX D-floating double precission variables while the final result does not. Exchanging the order of the variables we obtain a new expression that doesn't lead to any under-overflow.

For instance in the subroutine MCB the expression:

$$EBOB1 = 8D0 * FF * DD * EDF/[(EE + SDELTA)*$$

$$(EDF - DF - SDELTA) * (EDF + DF + SDELTA)]$$

is substituted by:

$$EBOB1 = 8D0 * FF/(EE + SDELTA) * DD$$

$$/(EDF - DF - SDELTA) * EDF/(EDF + DF + SDELTA)$$

That is, for calculational purposes, the following mathematical expression:

$$EBOB1 = \frac{8F_F D_D E_{DF}}{[(E_E + S_{DELTA})(E_{DF} - D_F - S_{DELTA})(E_{DF} + D_F + S_{DELTA})]}$$

is replaced by:

$$EBOB1 = \{\frac{8F_F}{(E_E + S_{DELTA})}\}\{\frac{D_D}{(E_{DF} - D_F - S_{DELTA})}\}\{\frac{E_{DF}}{(E_{DF} + D_F + S_{DELTA})}\}$$

Obviously the order followed to do this operation is different and that allows to avoid under-overflows.

b) Change of variables.

In this case we suppress certain variables used in the programs just for intermediate calculations and define instead new ones free of under-overflow.

For instance in the subroutine MCB the calculation of WEIGHT(2) is performed originally in the following way:

$$XMBP1 = 512.D0/(T2 * W2 * DD3 * DD4 * WEA1 *$$

$$WEA2 * WEA3 * WEA4 * WEA5 * WEA6)$$

$$WEIGHT(2) = XME(2)/[(XMBP1 + XMBP2) * WAP(2)]$$

and, instead, the following way can be used to avoid the under-overflow problem:

$$XMBD1N = 512.D0/DD3/DD4/W2$$

$$XMBD1D = WEA1 * WEA2 * WEA5 * T2$$

$$XMBP1E = WEA4 * WEA6$$

$$WEIGHT(2) = (XMB1 * T1 * 1.D - 15 + XMB2 * T2 * 1.D - 15 + XMB3 * W2)/WAP(2)$$

$$/(XMBP1N * XMBP2D + XMBP1D * XMBP2N * XMBP1E * WEA2)$$

$$*XMBP1D * XMBP2D/T1/T2/W2 * XMBP1E * WEA2$$

In this new expression there is also a substitution for the value of XME(2) and a similar trick for XMBP2 and, in addition, contains some factors of rescaling (commented below).

c) Rescaling.

In some cases it has been necessary to multiply and to divide in the proper places for a given scale factor to preserve the exponent range in intermediate operations.

For instance:

$$DELTAA = 2 * DSQRT(((4 * Y * W2/(XKH + 2 * XKV) + 2 * XKV * VK -$$

$$-2 * XKV * DEPA) * 1.D10) ** 2 + 4 * XKV * DEPA * (XKH + 2 * XKV) * 1.D20 * VK)/1.D10$$

That is, matematically speaking, the operation performed reads:

$$\frac{2\{[(\frac{4YW_2}{X_{KH}+2X_{KV}} + 2X_{KV}V_K - 2X_{KV}D_{EPA})\mathbf{10^{10}}]^2 + 4X_{KV}D_{EPA}(X_{KH} + 2X_{KV})\mathbf{10^{20}}V_K\}^{\frac{1}{2}}}{\mathbf{10^{10}}}$$

For this rescaling, in some of the subroutines, it has been useful to define a new variable used as an scale factor (SCAL) which takes different values according to the kind of process.

d) Alternative method to compute the scalar product.

The programs had already two different methods to compute the propagators in subroutine GRAAF but it isn't enough to handle the required accuracy in some cases.

Thus we introduce a third method that is valid when the two others give problems. Our method allows to obtain more precission for the variables since it works with a complement to one. The modification consists in computing the scalar product in the following way:

Let $P = (P_0, P_1, P_2, P_3)$ and $Q = (Q_0, Q_1, Q_2, Q_3)$, then:

$$(P+Q)^2 = P^2 + Q^2 + 2(Q_0(P_0 - 1/Q_0) - P_1Q_1 - P_2Q_2 - Q_3(P_3 - 1/Q_3))$$

The four-momentum P_μ and Q_μ can be any of the ingoing or outgoing particle four-momentum which are rescaled to the energy beam. For this reason the 0 and 3 components can be very close to 1 and, in this case, this way of rewriting the scalar product gives the largest precission.

e) Additional changes.

In addition to the methods to avoid under-overflow problems explained above, it has been necessary the modification of the subroutine GETRID to fix an existing bug in DIAG36ZN.

The subroutine GETRID determines which of the supressed spin configurations will be omitted during the summation of the exact matrix element squared over all the possible spin configurations. And it also selects for a given event which diagrams are likely to give the largest contribution. This speeds up the calculation of the final weight considerably.

For the selection of the more important diagrams, the program looks at the combination of propagators appearing in each diagram, looks for the minimum of those combinations and keeps in the calculation all the diagrams whose propagators are larger than the minimum by a factor chosen by the user with the variable PROC (whose default is 10^6).

In GETRID (not in the actual calculation of the matrix elements), the propagator for the Z^0 diagram is written as

$$WEZ = WE - AMZ2$$

where $AMZ2$ is M_Z^2 and WE is the q^2 in the Z propagator. When we are sitting on top of the Z^0 resonance it gives 0, and then only the diagrams where this propagator appears will be taken into account. This a very bad approximation, since, even on top of the Z peak, for most of the process, the QED diagrams still dominate.

To correct for this, we have taken into account the full Z propagator

$$\frac{1}{q^2 - M_Z^2 + iM_Z\Gamma_Z}$$

and, therefore, where we find in the program :

$$WEZ = WE - AMZ2$$

it is changed to

$$WEZ = DSQRT((WE - AMZ2) * *2 + AMGZ)$$

After this modification, we do not find any problem even when running the program at the pole.

3 Results

We present the results which were obtained by the use of the programs mentioned before. First of all, the total cross sections are presented for four-lepton production processes and we have reproduced the results published in the doctoral thesis of P.H. Daverveldt [2].

In table 1 the total cross sections for the production of $e^+e^-e^+e^-$, $e^+e^-\mu^+\mu^-$, $\mu^+\mu^-\mu^+\mu^-$ and $\mu^+\mu^-\tau^+\tau^-$ final states are listed for beam energies equal to 50 GeV (where the mass of Z^0 is taken to be 88.615 GeV — the result obtained from $\sin^2\theta_W = 0.23$ using the tree level relations between the parameters of the Standard Model) and equal to 46 and 50 GeV (where the mass of Z^0 is 92 GeV). The first case has been computed just in order to cross-check the results obtained with the modified program with the ones obtained by P.H. Daverveldt [2].

The table is divided into two parts. In the first part all lowest order QED diagrams are calculated. In the second one the lowest order cross section as predicted by the electroweak theory is calculated. The total cross section for $e^+e^-e^+e^-$ (and to a lesser extent for $e^+e^-\mu^+\mu^-$) is very large, due to the fact that these final states are predominantly produced by multiperipheral scattering. This implies that the vast majority of the $e^+e^-e^+e^-$ events will be invisible, all the particles being produced at very small angles with respect to the beams.

We can see that in the $e^+e^-e^+e^-$ and $e^+e^-\mu^+\mu^-$ processes the corrections due to Z^0 are quite negligible. On the contrary, for the process $\mu^+\mu^-\mu^+\mu^-$ and $\mu^+\mu^-\tau^+\tau^-$ the corrections due to Z^0 are by far much more significative. For these processes, when we are sitting on top of Z^0 peak the effect is larger than some GeV above (see table 1) and this can be understood just taking into account the contribution of the so-called annihilation diagrams, which dominate and tend to reproduce the Z^0 line shape. (See ref. [1] for a definition of the annihilation diagrams).

For every Monte-Carlo event we determine the number of tracks coming out at an angle (with respect to the e^+e^- beams) larger than some θ_0. In table 2 and table 3 we give the visible cross sections for the four-lepton production processes : $e^+e^-e^+e^-$, $e^+e^-\mu^+\mu^-$, $\mu^+\mu^-\mu^+\mu^-$ and $\mu^+\mu^-\tau^+\tau^-$ where we apply a value of 25 degrees for θ_0. The table 2 takes into account all lowest order QED amplitudes whereas the table 3 does the complete electroweak calculation.

Looking at these tables we can observe the following:

a) $e^+e^- \longrightarrow e^+e^-e^+e^-$

In spite that the total cross section doesn't exhibit any difference as commented above, we can see that the visible cross section is slightly sensitive to the effect of the Z^0 peak. The effect is only visible when we are sitting just on top of the Z^0 resonance and seems to be a little bit more significative in the case where two electrons are tagged than in the case in which just one electron is seen.

The size of the effect logically will depend strongly on the angular cut applied since there is a strong correlation between minimum angle and Q^2. Low angle (correlated with low Q^2) corresponds to QED dominance whereas large angle (high Q^2) tends to select EW effects.

b) $e^+e^- \longrightarrow e^+e^-\mu^+\mu^-$

Unlike the previous case, this process don't show any difference in the total cross section nor in the visible cross sections in the tables we have computed. Nevertheless if very strong detection cuts are applied (very small regions of the available phase space are selected) then the EW effects might be enhanced in such a way that they become really important. This is the case in which a strong muon invariant mass cut is applied for instance.

c) $e^+e^- \longrightarrow \mu^+\mu^-\mu^+\mu^-$

For this process the visible cross sections, like the total ones, in QED and EW are sizabily different for all kind of observed final states. The case in which this difference is more relevant (and also the one with larger cross section) is when four μ's are observed : then the EW cross section is almost a factor 80 larger than the QED one.

The case in which two muons are detected has a very tiny cross section so that actually this is a harmless background for most of the interesting processes. Anyway it is clear that there is about a factor 20 difference between the EW and the QED calculations.

d) $e^+e^- \longrightarrow \mu^+\mu^-\tau^+\tau^-$

Like in the previous process, in this case the visible cross sections and the total ones in QED and EW are clearly different. This is specially true when the four particle produced are detected. In this case not just the total cross section is the larger one but also the difference between the EW and the QED calculations becomes the larger one (about a factor 70 for the angular cut we are using)

4 Conclusions

The use of some computational tricks to avoid the problem of underflows and overflows after a careful study of the programs, has enabled us to prepare a computer version of the two-photon Monte Carlo programs by P.H.Daverveldt able to run on VAX as well as on IBM and more stable numerically speaking than the previous ones in any case.

The programs, called DIAG36N and DIAG36ZN, are available on request from ALEMANY@EB0UAB51 and GAITAN@EB0UAB51 and allow directly the calculation of the cross section of any leptonic two photon process using a complete tree level QED calculation (DIAG36N) or a complete tree level EW one (DIAG36ZN).

These programs have been checked producing about 3 x 10^6 events for each channel without any kind of problems on VAX and IBM computers. Furthermore, we have also checked, running the program on an IBM, that the results are exactly the same as those obtained with the original version, after solving the problem with the Z propagator in the routine GETRID.

A study of the results obtained running the programs near the Z^0 peak shows that the use of the EW versions (obviously slower) is extremely necessary when final states without electrons are to be produced whereas can be safely avoided for most of the detection configurations when the final state contains electrons using the QED versions, which are much faster. In this last case, the study of very restricted regions of phase space characterized by high momentum transfer should be taken with caution since depending on the cuts, the size of EW corrections becomes important.

Appendix

In this section we will outline the program structure. A flow chart of the program is shown in fig. 1.

The whole procedure used to generate events consists of three steps. During the first step the initialization is performed. Partly this is done by the user, who has to specify various input parameters such as the process desired, the beam energy, the masses and charges of the particles involved, the α_i and β_i parameters, the cuts which must be applied to the events and finally the number of rejection algorithms used in the program.

A rejection algorithm can be used twice. First it is used to obtain events distributed according to the sum of the matrix elements squared of all the contributing subgroups of Feynman graphs. (At this stage the interference between different subgroups is not included). Later on it is used to obtain unweighted events reproducing the complete matrix element squared distribution.

The arbitrary input parameters α_i and β_i enable the user to tune the 'strengths' of the four subgenerators in order to obtain the best possible efficiency of the program.

For example, if the user suspects that the cuts applied to the events favor the annihilation group and that the other diagrams can be treated as only corrections, he or she may use only the annihilation subgenerator for the event generation by setting α_m, α_b and α_c equal to 0. Of course, as the contributions from the other diagrams grow bigger, the fluctuations in the weight of the events will increase accordingly.

The rest of the initialization is performed by the subroutine START. There, among others, the approximate total cross sections of the groups are determined. The next step consists of the event generation and the weight calculation. Finally, FINISH calculates the mean weights, exact cross sections and event statistics.

Basic subroutines in the programs and their function

Name	Description
MCA	Event generator for the multiperipherical diagrams.
MCB	Event generator for the bremsstrahlung diagrams.
MCC	Event generator for the conversion diagrams.
MCD	Event generator for the annihilation diagrams.
START	This routine performs the necessary initialization.
FINISH	This routine calculates the exact cross section.
DIAM	This routine determines the complet matrix element squared.
DIAG2	This routine determines the matrix element squared of a gauge invariant subset of two Feynman diagrams.

DIAG4	This routine calculates the interference between the subsets of Feynman diagrams which contribute to the bremsstrahlung and annihilation subgroup.
SPINOX	This performs the initialization for DIAM.
GROUP	This routine defines the various groups of diagrams.
PERMUT	This labels the four-momenta in each group.
CHOICE	This routine selects the Feynman diagrams which contribute to the process chosen. It also puts in the charge factors when we produce quarks in the final state.
GRAAF	This routine calculates the amplitude of a diagrams for a given four-momentum and spin configuration.
ZZ	This function represents the contraction between two corrents which are used to built the amplitude in GRAAF.
GETRID	This routine helps to speed up the calculations of DIAM.
RNF100	Multiple random number generator which shuffling which generates most variables.
RNDM	Simple multiplicative random number generator.
HISTO1	This routine fills histograms. The entry HISTO2 prints the histograms.
CHANGE	This routine interchange two four-momenta.
DOT	This routine calculates the dot product between two four-momenta.

Acknowledgements

Two of us (R.A. and V.G.) would like to thank M.C. Delfino and A. Pacheco for their help with technical problems on the VAX. One of us (R.M.) wish to acknowledge the help of S. Orteu when trying to run the program on our IBM 3090.

References

[1] F.A. Berends, P.H. Daverveldt, R.Kleiss,
Phys. Lett. 148B (1984) 489,
Nucl. Phys. B253 (1985) 421, 441
Comp. Phys. Comm. 40 (1986) 271, 285, 309 .
[2] P.H. Daverveldt, Ph.D. Thesis, University of Leiden, Holland (1985).

FIGURE 1

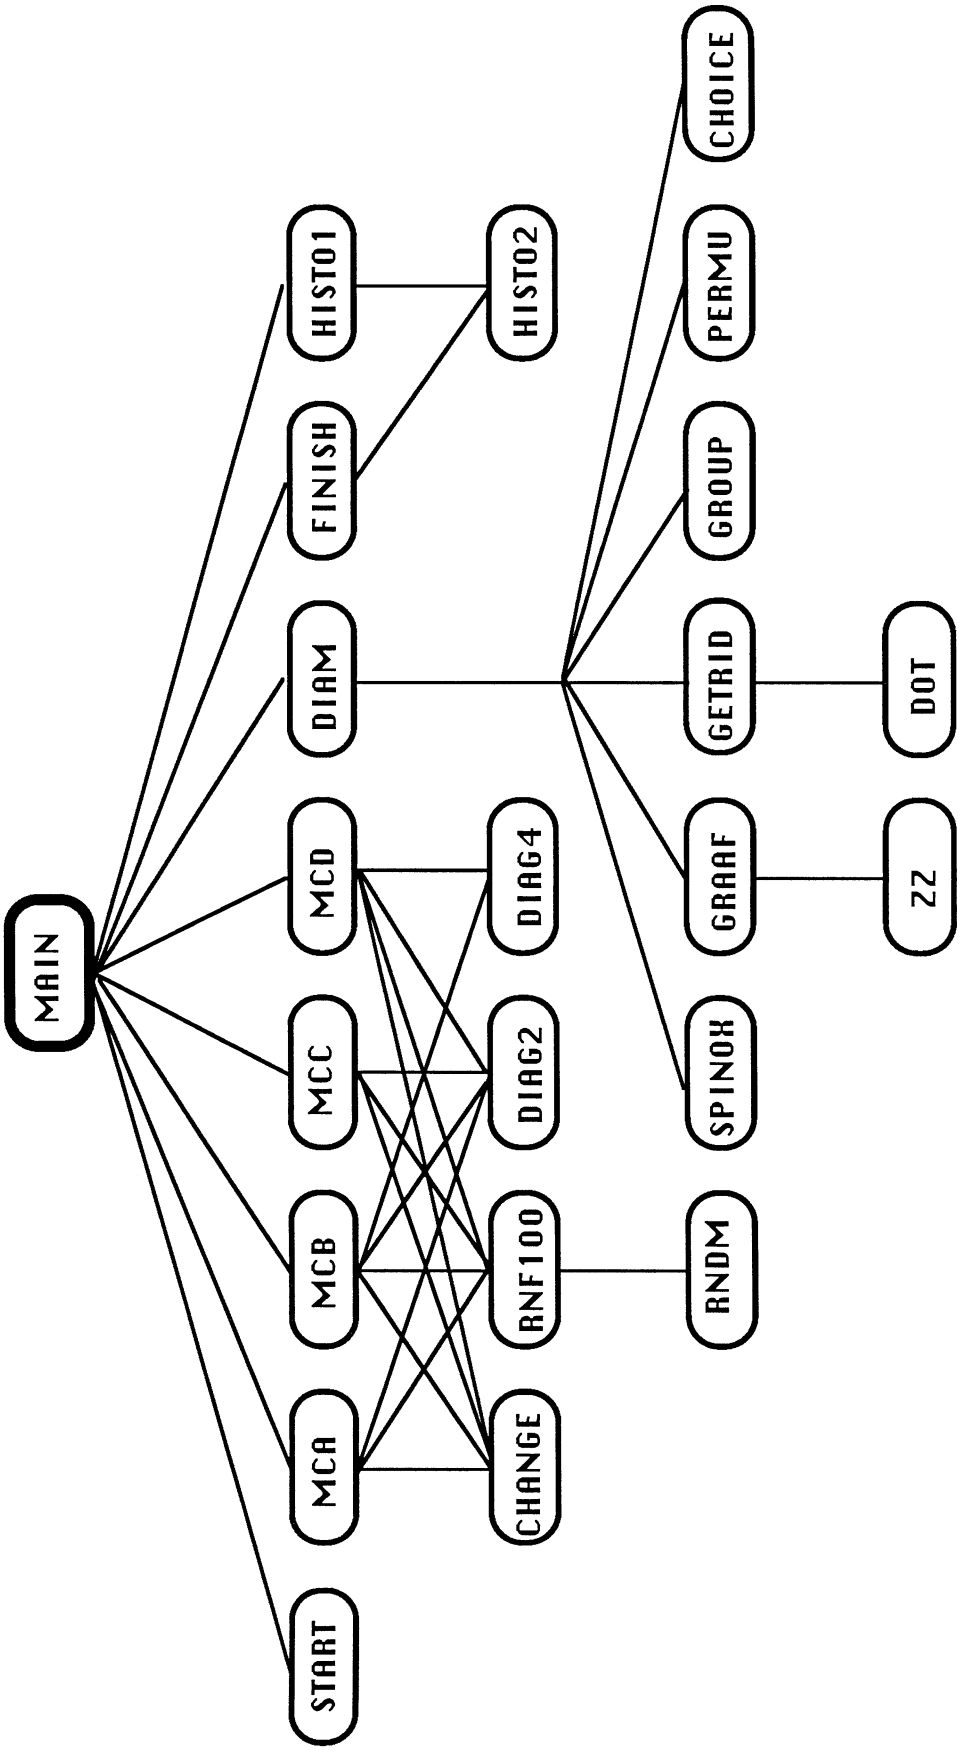
MAIN
START
MCA
MCB
MCC
MCD
DIAM
FINISH
HIST01
HIST02
CHANGE
RNF100
DIAG2
DIAG4
RNDM
SPINOX
GRAAF
GETRID
GROUP
PERMU
CHOICE
ZZ
DOT

TABLE 1

TOTAL CROSS SECTION

CONTRIBUTIONS OF ALL LOWEST ORDER QED DIAGRAMS IN NANOBARNS

FINAL STATE	46 Gev	50 Gev
eeee	$0.1438(9)\ 10^{8}$	$0.1472(8)\ 10^{8}$
ee$\mu\mu$	197.3(2)	204.2(2)
$\mu\mu\mu\mu$	$0.2035(5)\ 10^{-4}$	$0.1805(5)\ 10^{-4}$
$\mu\mu\tau\tau$	$0.1768(5)\ 10^{-4}$	$0.1591(4)\ 10^{-4}$

CONTRIBUTIONS OF ALL LOWEST ORDER ELECTROWEAK DIAGRAMS IN NANOBARNS

FINAL STATE	Mz = 92 GeV		Mz = 88 GeV
	46 GeV	50 GeV	50 GeV
eeee	$0.143(1)\ 10^{8}$	$0.144(2)\ 10^{8}$	$0.142(2)\ 10^{8}$
ee$\mu\mu$	198.5(8)	204.4(8)	204.3(7)
$\mu\mu\mu\mu$	$0.70(1)\ 10^{-3}$	$0.111(2)\ 10^{-3}$	$0.88(1)\ 10^{-4}$
$\mu\mu\tau\tau$	$0.652(7)\ 10^{-3}$	$0.112(2)\ 10^{-3}$	$0.91(1)\ 10^{-4}$

table 2

tracks visible	46 Gev QED cross section (nb)	50 Gev QED cross section (nb)
	$e^+e^- \rightarrow e^+e^-e^+e^-$	
e	$2.48(8)\ 10^{6}$	$2.40(6)\ 10^{6}$
ee	$1.07(5)\ 10^{6}$	$1.10(4)\ 10^{6}$
	$e^+e^- \rightarrow e^+e^-\mu^+\mu^-$	
e	0.60(3)	0.57(3)
μ	53.1(3)	54.6(3)
$\mu\mu$	24.6(2)	25.1(2)
	$e^+e^- \rightarrow \mu^+\mu^-\mu^+\mu^-$	
μ	$3.35(4)\ 10^{-6}$	$3.00(2)\ 10^{-6}$
$\mu\mu$	$3.17(3)\ 10^{-6}$	$2.65(2)\ 10^{-6}$
$\mu\mu\mu$	$1.66(2)\ 10^{-6}$	$1.37(2)\ 10^{-6}$
$\mu\mu\mu\mu$	$6.51(5)\ 10^{-6}$	$5.37(4)\ 10^{-6}$

table 2 (cont.)

tracks visible	46 Gev QED cross section (nb)	50 Gev QED cross section (nb)
	$e+e- \rightarrow \mu+\mu- \tau+ \tau-$	
μ	$0.61(1)\ 10^{-6}$	$0.57(1)10^{-6}$
τ	$2.39(2)10^{-6}$	$2.17(2)10^{-6}$
$\mu\mu$	$0.50(1)\ 10^{-6}$	$0.46(1)10^{-6}$
$\tau\tau$	$2.27(2)\ 10^{-6}$	$2.06(2)10^{-6}$
$\mu\tau$	$0.64(1)\ 10^{-6}$	$0.55(1)10^{-6}$
$\mu\mu\tau$	$0.88(1)\ 10^{-6}$	$0.80(1)10^{-6}$
$\tau\tau\mu$	$1.12(1)\ 10^{-6}$	$1.01(1)10^{-6}$
$\mu\mu\ \tau\tau$	$6.67(5)\ 10^{-6}$	$5.9(4)10^{-6}$

table 3

tracks visible	Electroweak cross section (nb)		
	Mz = 92 GeV		Mz = 88 GeV
	46 GeV	50 GeV	50 GeV
	e+e- → e+e-e+e-		
e	$2.7(1)\ 10^{6}$	$2.5(1)\ 10^{6}$	$2.4(2)\ 10^{6}$
ee	$1.22(7)\ 10^{6}$	$1.1(1)\ 10^{6}$	$1.1(1)\ 10^{6}$
	e+e- → e+e- μ+μ−		
e	0.6(1)	0.6(1)	0.58(8)
μ	53(1)	55(1)	55.6(8)
μμ	24.1(6)	25.0(6)	24.6(5)
	e+e- → μ+μ−μ+μ−		
μ	$2.1(2)\ 10^{-5}$	$0.58(5)\ 10^{-5}$	$0.56(4)\ 10^{-5}$
μμ	$0.59(3)\ 10^{-4}$	$3.22(7)\ 10^{-4}$	$0.271(5)\ 10^{-4}$
μμμ	$0.76(3)\ 10^{-4}$	$1.30(4)\ 10^{-4}$	$0.105(3)\ 10^{-4}$
μμμμ	$5.10(9)\ 10^{-4}$	$0.509(9)\ 10^{-4}$	$0.361(7)\ 10^{-4}$

table 3 (cont.)

tracks visible	Electroweak cross section (nb)		
	Mz = 92 GeV		Mz = 88 GeV
	46 GeV	50 GeV	50 GeV
	$e^+e^- \rightarrow \mu^+\mu^- \tau^+ \tau^-$		
μ	$12(1)\ 10^{-6}$	$1.6(2)\ 10^{-6}$	$1.6(2)\ 10^{-6}$
τ	$10(1)\ 10^{-6}$	$4.0(5)\ 10^{-6}$	$3.5(4)\ 10^{-6}$
$\mu\mu$	$35(2)\ 10^{-6}$	$5.1(3)\ 10^{-6}$	$3.5(2)\ 10^{-6}$
$\tau\tau$	$25(2)\ 10^{-6}$	$27.2(7)\ 10^{-6}$	$24.5(5)\ 10^{-6}$
$\mu\tau$	$6.7(9)\ 10^{-6}$	$0.8(1)\ 10^{-6}$	$1.0(2)\ 10^{-6}$
$\mu\mu\tau$	$31(2)\ 10^{-6}$	$3.6(2)\ 10^{-6}$	$3.5(2)\ 10^{-6}$
$\tau\tau\mu$	$55(2)\ 10^{-6}$	$11.1(4)\ 10^{-6}$	$8.3(3)\ 10^{-6}$
$\mu\mu\ \tau\tau$	$460(7)\ 10^{-6}$	$53(1)\ 10^{-6}$	$39.6(7)\ 10^{-6}$