misidentification of only 0.007 for pions.

The CLEO system uses three 36-cm thick steel plates, while our Muon Magnet endplate, which is 30-cm thick, serves as the first absorber layer and 1-m thick concrete slabs provide the second and third layers. (Concrete is being used for budgetary reasons.) Assuming a concrete density of 2.5 g/cm^3 each concrete slab has a nuclear interaction length about 30 % larger than the 36-cm thick steel plates used for CLEO. Thus our identifier has an overall thickness that is slightly larger than the CLEO detector. Simulations with PISA show no significant difference between the behavior of our concrete absorbers and steel absorbers of the same interaction length. The only disadvantage to concrete is that the large thickness of the identifier causes a loss in acceptance angle.

Muon Identification Algorithm

The main difference between the situation in CLEO and ours is that we must contend with a much larger backgound of hadrons. Muon identification in the CLEO detector is done by tracking particles from the vertex into the identifier and selecting tracks that range out in the absorber layers at the depth expected for a muon of the momentum determined by the tracking chambers. This scheme works well in the low background environment of CLEO and with the narrow range of muon momenta characteristic of that experiment. The Discriminant Analysis technique developed to discriminate between muons and pions and described in the PHENIX pCDR, works well in the wide momentum range environment of RHIC for many sampling layers. However, PISA simulations show that this technique becomes unsatisfactory with only three absorber layers. The probability of misidentification for pions is larger than 0.15 for pion energies below 3 GeV. We have therefore turned our attention to a scheme similar to the one used at CLEO, namely to track particles through the muon identifier using the momentum and hit multiplicity in the streamer tubes to select muon.

Our technique starts with a track in the deepest layer of the identifier, and searches for hits produced by the same track in forward detector layers. This is still under development; so far we have obtained results on tracking particles through the muon identifier. The particle identification efficiencies we have obtained will improve with the track information from the muon arm tracking chambers.

Muon Identifier Efficiency

The simulations performed with PISA determine how well muons can be identified in the range of particle momenta and the large background expected at RHIC. Simulations of central events with the ORNL HGMC code using the UA1 p_T spectrum show that approximately 18 pions and six muons per event reach the muon identifier. To simulate this background we have used PISA to generate events consisting of 32 particles as follows:

- A $\mu^+\mu^-$ pair (dimuon) with lab momenta corresponding to emission from the decay of a particle of prescribed rapidity, p_T, and mass. For a given pair-mass the values of y and p_T are chosen to put both muons into the muon identifier. The angle ϕ is chosen randomly for each event. We generate test dimuons in this way to study the tracking efficiency for dimuons in our acceptance.

- An additional eight muons and 22 pions with y and p_T from a simple event generator providing a uniform distribution in y between $y = 1.0$ and 2.5, and p_T according to $\exp^{-p_T/0.4}$.

The p_T distribution was deliberately chosen to give a hard spectrum. In order to save computation time, the central magnet is omitted. The muon magnet is included and the magnetic fields are turned on. The y distribution covers all angles between the muon magnet "piston" and "lampshade". In addition, the initial trajectory of each particle is tested to ensure that θ is between 10° and 35°. This technique, coupled with the omission of the central magnet, guarantees that 32 particles per event enter the muon identifier. The numbers we have used in our simulations are a factor of 2–3 larger than predicted by the HGMC code, to deliberately complicate the tracking. Three figures of merit have been evaluated in this study, namely:

1. The efficiency for correctly identifying a dimuon.
2. The single muon loss rate.
3. The single pion misidentification rate.

The tracking algorithm searches for isolated hits, beginning in the deepest layer of the muon identifier. It has been tested for a muon identifier with 10×10 cm^2 pads, so "isolated hit" implies a minimum transverse distance of 10 cm between hits. The program follows the centroid of a track from the back to the front of the identifier. In each plane all hits in a square array of nine adjacent pads are assumed to belong to this track.

The lower energy muons present the most difficulties, because they may not traverse all three layers of the identifier. Therefore, we have studied dimuons with masses of 2.0 and 2.5 GeV, pair rapidities from 1.2–1.6, and p_T from 0–1 GeV/c. Single muon candidates are selected by searching for tracks with an isolated hit in the second and third planes of the identifier. Muons are identified as the subset of those tracks with one and only one hit per layer in the forward detector layers. For low mass dimuons, often only one of the muons has sufficient energy to reach the third layer of the identifier, and the other ranges out before reaching the second layer.

The dimuon tracking efficiency is defined as the number of dimuons accurately tracked through the identifier divided by the total number generated by PISA. Note that "accurately tracked" indicates the tracks are identified as muons because they produce only one hit in each plane. It should also be noted that, for a substantial fraction of dimuons with 2 GeV mass at these rapidities, one particle ranges out before reaching the second plane. So, for low mass we count the dimuon as correctly tracked and identified if one muon produces only one hit and the other reaches the third plane with only one hit per plane. With these conditions we find tracking efficiencies of 0.91 to 0.99.

The muon loss rate is given by the fraction of the eight additional muons in each event that are misidentified as pions because they produce more total hits than the number of planes

Table 10.6: Particle identification performance of the muon identifier for events with one dimuon with invariant mass M_{inv}, eight single muons and 22 pions single pions. All "muons" must reach at least the second μID plane. Refer to main text for further definitions of entries.

	M_{inv} = 2.0 GeV	M_{inv} = 2.5 GeV
Dimuon Tracking Efficiency	91%	99%
Muon Misidentification Ratio	4%	6%
Hadron Misidentification Ratio	4%	1%

they traverse. The fraction of misidentified muons is between 0.03 and 0.06 for muons with energy greater than 1 GeV at the vertex.

The fraction of misidentified pions is the number of pions tagged as muons because they produce only one hit per layer. Note that some of these may actually be muons from pion decay. This fraction is quite large, ranging from 0.1–0.19. This result is disappointing and is similar to the results found for pion misidentification with the Discriminant Analysis using only three absorber layers. Close inspection of the simulation results shows that the misidentified pions are almost all associated with single, isolated hits in the first detector layer making them indistinguishable from the lowest energy muons which do not reach the second detector plane. These pions can cause a combinatorial background with the higher energy muon which is cleanly identified. However, if we require that tracks reach at least the second layer, increasing the soft muon energy threshold somewhat, the fraction of misidentified pions drops to between 0.01 and 0.04. These results are summarized in Table 10.6.

We should emphasize that these studies do not yet use information from the tracking chambers. By requiring that the isolated hits in the first detector plane match up with a clearly defined particle trajectory emanating from the vertex, we will be able to eliminate many of these as possible muon candidates. It is also important to emphasize that, for dimuons with mass greater than 2 GeV in our acceptance, the kinematics cause both muons to have energies sufficient to reach the second or third detector layer. In fact, our muon identifier, similarly to the CLEO identifier, is optimized for detection of dimuons with pair-masses around that of the J/ψ mass. For this mass range one muon ranges out in the third absorber layer and the other ranges out in this layer or beyond, making it possible to use the range-energy information in identifying the source of the dimuon. For dimuons from ϕ decays with both muons entering the muon arm, only about 60% of the muons produce any hits at all. Of these only about 60% produce hits in the second or third detector layers. This represents a rather large effective decrease in acceptance in this mass range. The acceptance is in fact peaked at the most forward rapidities covered by the muon arm, as seen in the acceptance plots in Fig. 10.4.

Aside from the expected poor performance for very low masses at $y \leq 1.5$, the results on dimuon identification and tracking are satisfactory. Realistic backgrounds do not cause the

Table 10.7: Signal to background ratios for the most important vector meson resonances in the muon detector. The table is based on 2×10^9 central events each containing 10,000 charged pions and was done with the ORNL HGMC code with the UA1 p_T spectrum. Calculations done by PISA using HIJET as primary pion source and with additional kinematical cuts on the accepted tracks can lead to S/N improvements for each of the vector mesons by factors of 3–6.

Resonance	Signal/Background
ρ	1.2×10^{-2}
ϕ	0.8×10^{-2}
J/ψ	2.7×10^{0}
ψ'	2.1×10^{-1}
Υ	2.9×10^{2}
Υ'	7.0×10^{1}

occupancy of the muon identifier to increase sufficiently to cause confusion in tracking and identifying muons.

10.4.4 Integrated Performance

Probably the single most important aspect of the feasibility of the muon detector is the question of the vector meson signal-to-background ratio (S/N). The primary tool for identifying vector meson resonances is to create the invariant mass of all possible unlike-sign muon pairs in each event. This process, however, also creates a combinatorial background below each resonance which is one of the main experimental problems for all multiparticle measurements in the RHIC environment. Since the whole feasibility of the vector meson measurements detector depends critically on the S/N ratio between the resonance and underlying background yields, we have performed these calculations redundantly using four different approaches. PISA calculations were done independently at Vanderbilt and LANL, whereas specialized High Gain Monte Carlo (HGMC) codes were used at Tennessee and ORNL. The final results of all these calculations agree to within a factor of 2–3 for the final S/N's. The differences between the results of the codes does *not* mean that the calculations are incorrect, but does indicate the theoretical uncertainty in these calculations arising from the underlying assumptions. Note that the background discussed in this section is the combinatorial background unless stated otherwise. This implies that, for example, Drell-Yan production is not considered a background in this section.

Figure 10.31 shows the invariant mass spectrum of all unlike-sign muon pairs corresponding to 10^8 central Au + Au events. In contrast to the calculation shown in Fig. 3.6, which used the ORNL HGMC, this calculation was done using PISA with HIJET as the primary pion source. The calculation includes the following resonances: ϕ, J/ψ, ψ' and Υ. In ad-

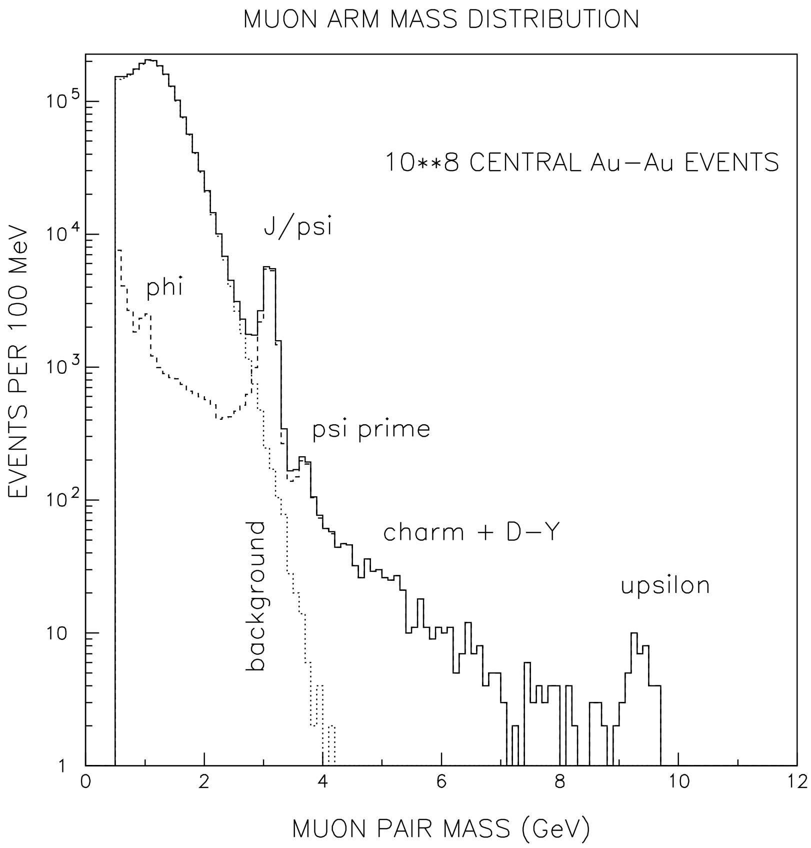

Figure 10.31: Dimuon invariant mass spectrum based on PISA calculations with no kinematic cuts on the accepted muons.

dition $D\overline{D}$ and DY production is included. Vertex cuts were imposed to reduce the muon background originating primarily from kaon decay and, to a much smaller extent, from pion decay. Finally, only muons with momentum larger than 1 GeV/c after the absorber were accepted. The actual simulation calculated the background for 10,000 central Au + Au collisions and the complete background spectrum corresponding to 10^8 events was obtained by using event mixing. The HGMC codes, which obtain a factor 10–100 higher statistics than PISA within the same CPU time, verified that the event mixing procedure causes no systematic error by running a few high statistics cases without event mixing and comparing to equivalent event mixed calculations.

Figure 10.31 shows how both the J/ψ and the ψ' resonances are clearly above the background. Table 10.7 lists the S/N-values for the six dominant vector meson resonances. The values quoted in this table probably represent more conservative estimates than values extracted from PISA calculations with HIJET. The main reason is that the HIJET p_T spectrum is softer than the UA1 spectrum used in the HGMC calculation.

As will be explained in more detail below, improvement in the S/N for particular reso-