FFI-rapport 2014/02403

Shock attenuation by porous materials

Jan Arild Teland

FFI Forsvarets forskningsinstitutt

Shock attenuation by porous materials

Jan Arild Teland

Norwegian Defence Research Establishment (FFI)

24 March 2015

FFI-rapport 2014/02403

1325

P: ISBN 978-82-464-2502-3
E: ISBN 978-82-464-2503-0

Keywords

Porøse materialer

Numerisk analyse

Simulering

Sjokkbølger

Approved by

Eirik Svinsås	Prosjektleder
Jon Skjervold	Avdelingssjef

English summary

Shock waves from explosions can do great damage to humans, buildings and other structures. Protective measures may therefore be useful to attenuate the effect of the shock waves. Porous materials have traditionally been considered to be good at shock mitigation and could possibly be placed in front of an object that requires protection from shock waves.

An initial literature survey of various shock attenuation experiments is carried out, showing apparently contradictory results regarding the mitigation effect of different materials. In order to investigate this, a theoretical and numerical study of the shock wave attenuation phenomenon was performed. The study showed that the experimental design largely determines the results. Porous materials can reduce the shock wave amplitudes, but at the expense of longer wave duration. In some cases, protective materials (in particular porous ones) can actually increase the maximum load on the object that is to be protected. In a given case where something requires protection, expert analysis is necessary to ensure that any mitigation measure does not have the opposite effect.

Sammendrag

Sjokkbølger fra eksplosjoner kan gjøre stor skade på både mennesker, bygninger og andre strukturer. Det kan derfor være nyttig med beskyttelsestiltak for å dempe effekten av sjokkbølgene. Porøse materialer har tradisjonelt blitt antatt å være sjokkdempende og kan derfor tenkes plassert foran et objekt som skal beskyttes mot sjokkbølger.

I rapporten gjennomføres en litteraturstudie av tidligere eksperimenter med sjokkdempning. Det viser seg at resultatene varierer sterkt og er til dels motstridende når det gjelder effekten av forskjellige dempematerialer. For å undersøke dette nærmere ble det gjennomført en teoretisk og numerisk studie av fenomenet sjokkdempning. Studien viste at designet på det eksperimentelle oppsettet i stor grad avgjør hva slags resultater man oppnår. Porøse materialer kan redusere amplituden på sjokkbølgen, men samtidig får den lengre varighet. Imidlertid kan bruk av beskyttelsesmaterialer (særlig porøse) i enkelte tilfeller faktisk føre til økt belastning på objektet som skal beskyttes. I et konkret tilfelle hvor noe skal beskyttes, trengs det derfor omfattende kompetanse og grundig analyse for å sikre at man ikke gjør vondt verre.

Contents

1 Introduction

Close-in attenuation of the pressure wave from a detonation is desirable in several applications. One example is within an ammunition storage, where it is important to prevent an accidental detonation of a warhead (or similar) resulting in a full detonation of all the stored objects. Reduced pressure from a buried IED detonated under a vehicle may increase the survivability of the crew. Another application is during EOD operations, where mitigation of the blast can increase survivability for the surroundings during an accidental or provoked explosion.

Many porous materials are thought to exhibit useful properties for shock attenuation and several different experiments by various groups have been carried out to investigate this. Interestingly, taken as a whole, the results of these experiments have been inconclusive and contradictory with regards to the damping properties of the materials. Why do the experiments contradict each other and are porous materials actually useful for shock attenuation? Those are the topics of this report. We will examine the phenomenon of shock attenuation in detail both theoretically and numerically in order to understand what is really going on.

The report is structured as follows. First we review the shock attenuation experiments of different groups, explaining their set-up, idea behind their test and the results obtained. Before attempting to explain these results physically, we first take a step backwards and look at some basic impact physics involving solid objects. We then proceed to numerically study how damping of a shock wave works in a very simple 1D-situation. Interestingly, we shall see that even in the simplest possible case, things are not trivial at all.

Having understood how things work in 1D, we will see how including other “experimental features” complicates the problem and makes it possible to obtain very different results for the shock damping properties of a given material, depending on the set-up of the experiment. All the previously reviewed experiments will then be numerically simulated to demonstrate this. Finally we will apply our achieved knowledge to a situation involving an IED under a vehicle to demonstrate how the burial depth relates to the effect of the IED and discuss whether any countermeasures are possible.

2 Shock attenuation experiments

We shall later see that it is not trivial to actually define what is meant by shock attenuation. For the moment, however, we will proceed by naively thinking of shock attenuation as doing “something” to an incoming shock wave that prevents or minimizes the ensuing damage, either to a structure or a human. A shock attenuating object will here consist of a material placed between the shock wave and the object which requires protection.

How can we test whether a material is good for achieving shock attenuation? Various groups have had different ideas about this and consequently several different experiments have been

defined. In this chapter we will review some of the different shock attenuation experiments which we are aware of.

2.1 NTNU

The first experiments we will mention were performed by NTNU (1). They used a set-up with a large ballistic pendulum exposed to a blast wave. Their idea was that if the pendulum was protected by panels of aluminium foam, this could absorb some of the momentum and stop it from being transferred to the pendulum. Experiments were then carried out to find how much impulse could be absorbed and also investigate how this depended on the foam density and whether adding an aluminium cover plate would have any effect.

The panels had dimensions 68.4 cm x 70 cm and two different PE4 charges were used, 1 kg and 2.5 kg, both placed at a distance of 500 mm from foam panels attached to the ballistic pendulum. The actual set-up is illustrated in Figure 2.1.

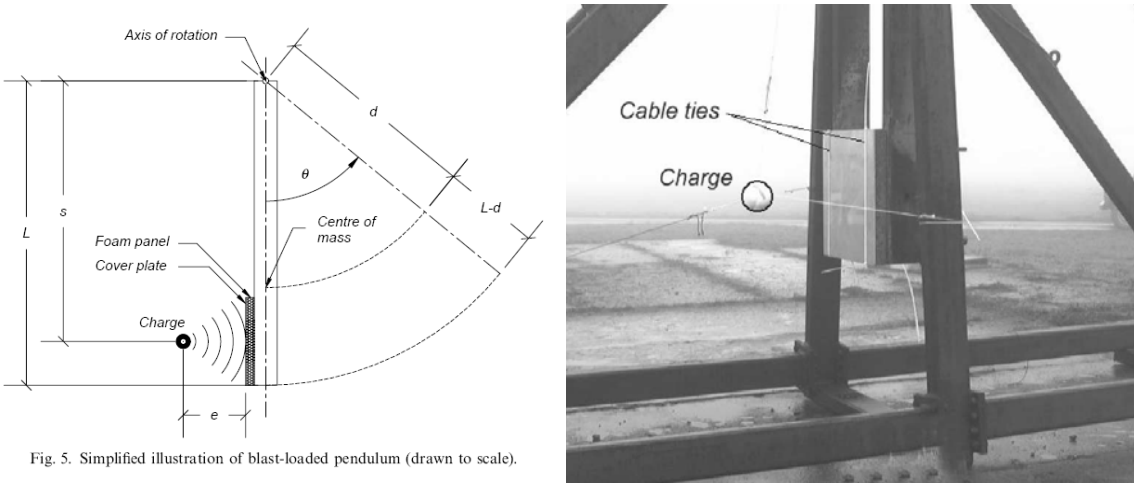

Figure 2.1 NTNU experiment. (Reproduced from (1)).

The experiments were repeated several times and the results were quite conclusive, though unexpected to NTNU. The aluminium foam panels did not decrease the transferred impulse at all. In fact, they were consistently seen to do exactly the opposite, i.e. increase the momentum that was transferred to the pendulum. This happened regardless of foam density and whether a cover plate was present or not. Typically the increased momentum was in the range of 10-20%.

2.2 DSTL

DSTL tried a different approach in their experiments (2-4). The damping material surrounded the charges and it was investigated whether this led to attenuation of the shock wave at a distance away from the charge (compared to the situation of no damping material).

Several experimental series using different charge sizes and attenuation materials were performed. The explosive charges were always spherical and consisted of 20 g – 5 kg PE4.

They used a large range of attenuation materials, including very porous ones (Perlite) as well as sand, glycerine and water.

In their first experiments, the blast wave was measured using blast gauges at different distances from the detonation. Compared to the situation without damping material, a considerable attenuation was found. An example is shown in Figure 2.2 for the case of 40 g PE4 charge with and without surrounding Perlite. The pressure amplitude is clearly mitigated by around a factor of 5 in this case.

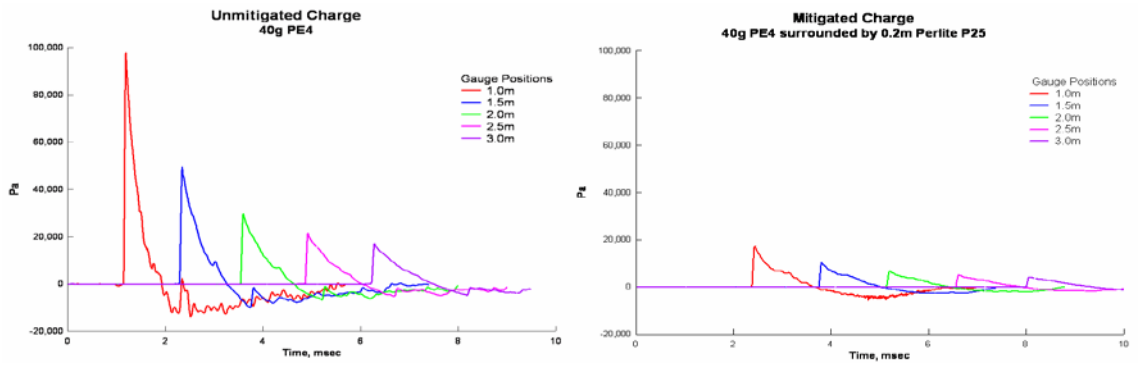

Figure 2.2 Mitigation of blast wave surrounded by Perlite (Reproduced from (4)).

DSTL also noticed a scaling rule where the mitigation seemed independent of charge size as a function of scaled mitigant volume. On comparing materials, DSTL found that sand appeared to give the best damping for a given volume, whereas Perlite was the best for a given mass.

While these experiments seemed to indicate considerable shock attenuation, DSTL noted that such overpressure measurements in air did not account for momentum transferred to the mitigation material, which had the potential to cause significant damage. To investigate this more closely, DSTL performed some experiments using a ballistic pendulum, similar to the NTNU experiment except for the damping material still surrounding the charge. They then found all damping materials to increase the transferred impulse compared with no damping material, which was also confirmed using an "impact gauge" measurement set-up. A comparison of the "attenuation" properties of the different materials, as measured by DSTL, is shown in Figure 2.3.

Thus, the DSTL experiments indicate that the pressure amplitude can be attenuated, but that the impulse is increased.

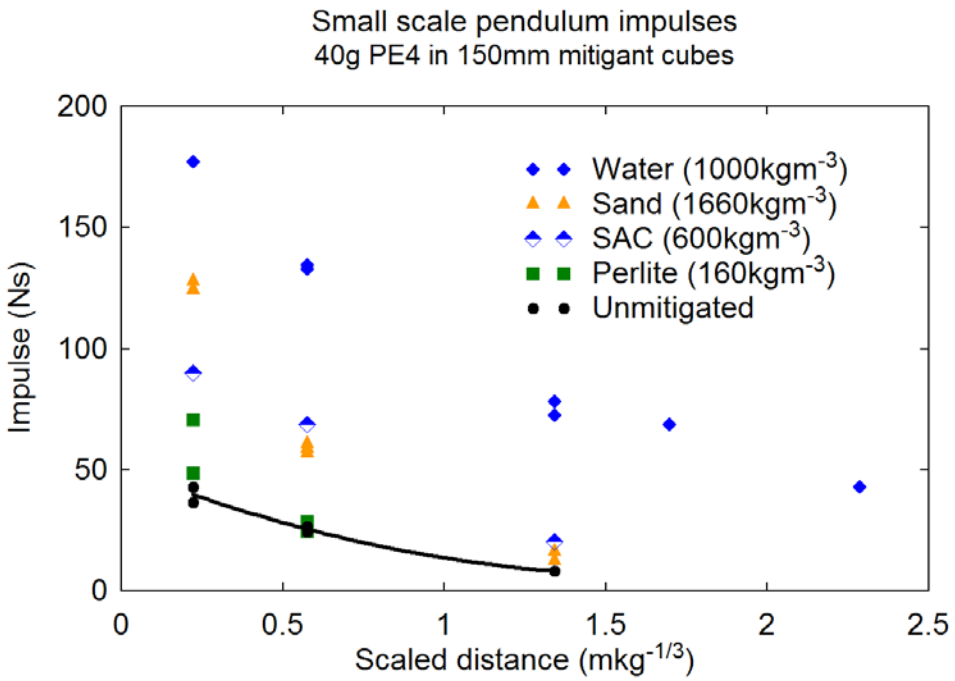

Figure 2.3 Results from DSTL experiments using a ballistic pendulum. (Reproduced from (4)).

2.3 NAVAIR

NAVAIR (USA) used another setup (Figure 2.4) to examine shock attenuation (5). In their experiments, a Pentolite charge with a mass of 175 g was placed directly on top of different attenuation materials of thickness 1-3 inches and detonated. The resulting pressure was then measured using PVF gauges on a PMMA block placed under the attenuation material.

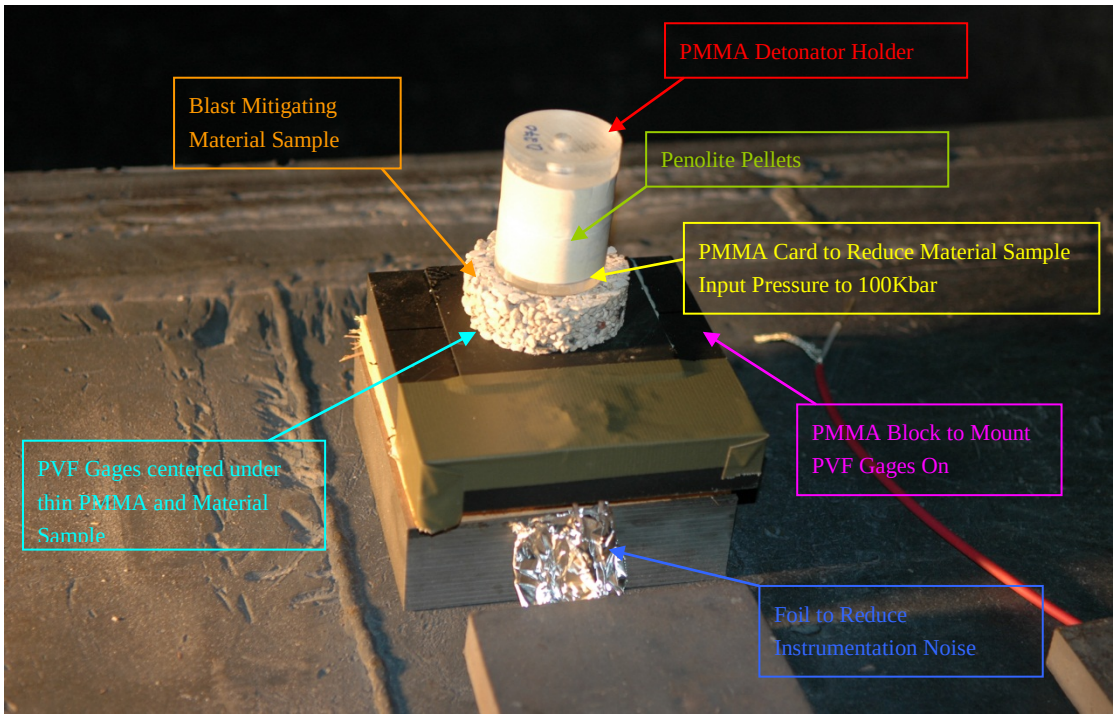

Figure 2.4 NAVAIR setup (Reproduced from (5)).

A range of attenuation materials were tested (mostly porous ones) and the measured pressure as a function of their thickness was determined. No experiments were made without attenuation material, so these experiments only investigated the relative damping properties of the different materials and how they varied with thickness. A summary of the results (grouped into man-made, geological and layered materials) are shown in Figure 2.5. These results seem to indicate that porous materials can indeed attenuate the pressure amplitude from the blast.

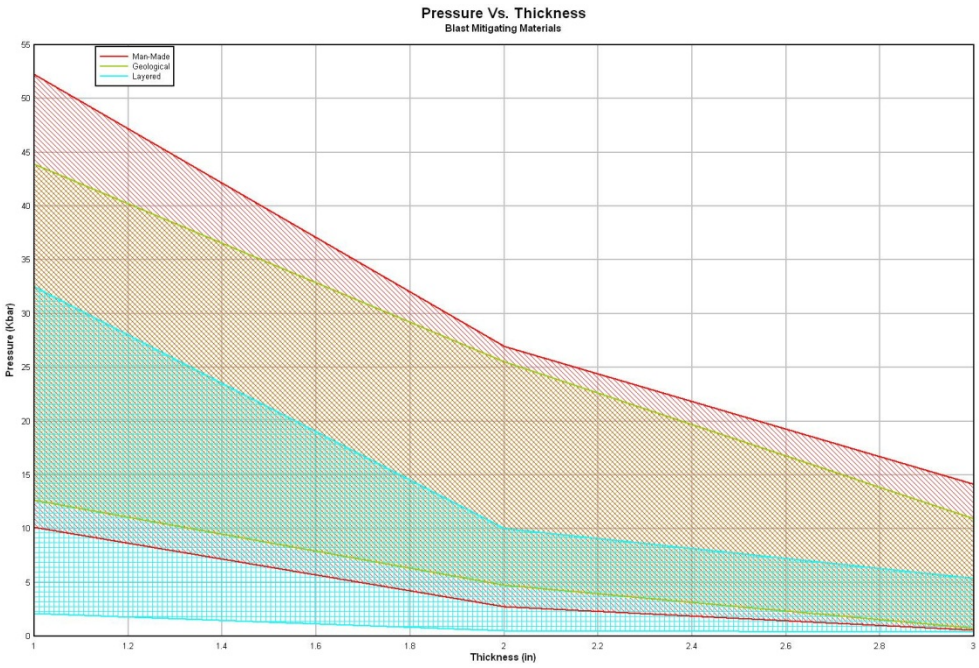

Figure 2.5 Results from NAVAIR experiments. (Reproduced from (5)).

2.4 FFI

In our own experiments at FFI (6-8) another set-up was used (Figure 2.6). This involved detonating a cylindrical TNT charge (radius 25 mm, 129 g) at 80-100 mm distance from a Hopkinson steel bar (radius 25 mm, length 3000 mm). Attenuation materials of different thicknesses were placed between the charge and the steel bar (either close to the charge or close to the bar). Strain gauges were placed at the Hopkinson bar to measure the strain (and calculate the stress) transferred from the explosive. In this way, different damping materials could be compared with each other and with the case of no damping material.

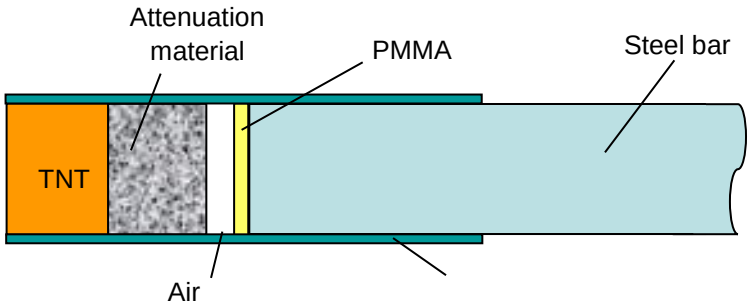

Figure 2.6 FFI set-up

A whole range of different materials were tested including pumice, LECA (coarse and fine), aluminium foam, rubber granules, gravel, wood shavings, saw dust, Glasopor, Siporex and brick. At least two tests were performed for each material and the repeatability of the experiments seemed to be excellent.

In general, most materials seemed capable of shock attenuation. A typical result is shown in Figure 2.7 where wood shavings are compared to the situation without damping material. The shape of the resulting stress pulse in the steel bar is seen to remain largely unaffected, but with reduced amplitude when an attenuation material is present. Since the momentum of the wave is proportional to the time integral of the pressure wave, this seems to imply that both maximum amplitude and impulse is decreased due to the attenuation materials.

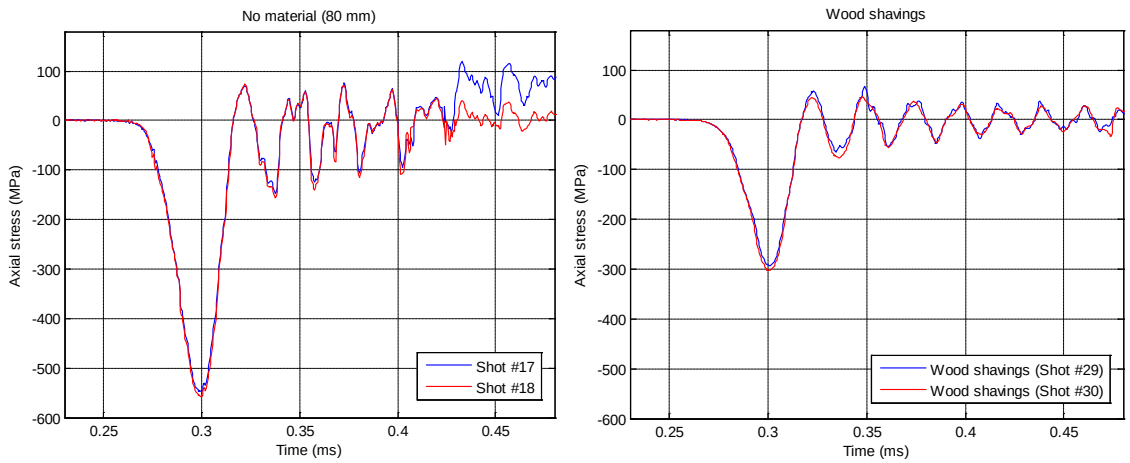

Figure 2.7 Results from FFI experiments (left: no damping material, right: wood shavings)

Figure 2.8 shows the stress amplitude in the bar as a function of the damping material thickness, regardless of material type. There seems to be an almost linear relationship, indicating that the material thickness matters much more than what kind of material is used.

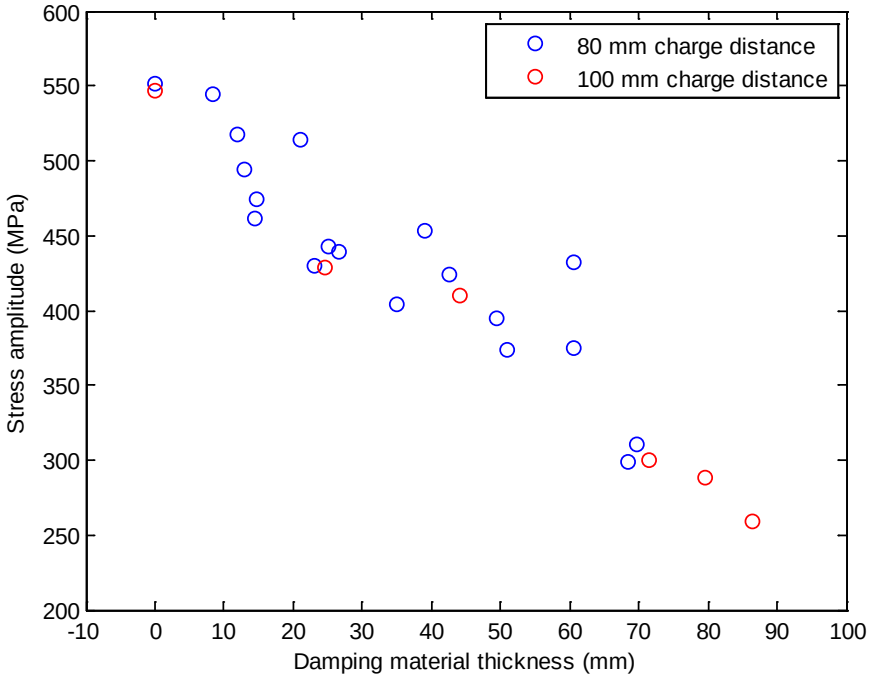

Figure 2.8 Results from FFI experiments

2.5 Summary of experimental results

Having reviewed several different shock attenuation experiments, it is now natural to ask whether porous (and other) materials can really be used for shock attenuation. A naive interpretation of the experimental results would give the following answers (compared with the case of no damping material):

NAVAIR: Yes, the maximum pressure is decreased.
DSTL: Yes, the blast wave is attenuated. But, no, the total impulse is increased.
FFI: Yes, both maximum pressure and total impulse is decreased.
NTNU: No. The impulse and energy is increased.

Thus, the results seem mixed and inconclusive so far. An analysis of the various experiments is needed to reveal the reason for them giving apparently contradictory results and to answer the question of whether shock attenuation is possible.

3 Basic impact physics

We will start our investigation into shock attenuation by looking at some basic physics. Although relatively simple, we shall see that this exercise will provide some enlightening and slightly non-intuitive results that will be very useful in the discussion later on.

3.1 Solid point-like objects

First, let us consider the impact of two solid point-like objects. By point-like we mean that the objects do not deform and that all mass can be considered to be located in one single point.

Let an object of mass m (from now on called the incoming object) impact another object of mass M (called the receiving object) with an impact velocity of v_0. (A reference frame can always be found where one object is at rest, so this is the most general situation). This situation is illustrated in Figure 3.1.

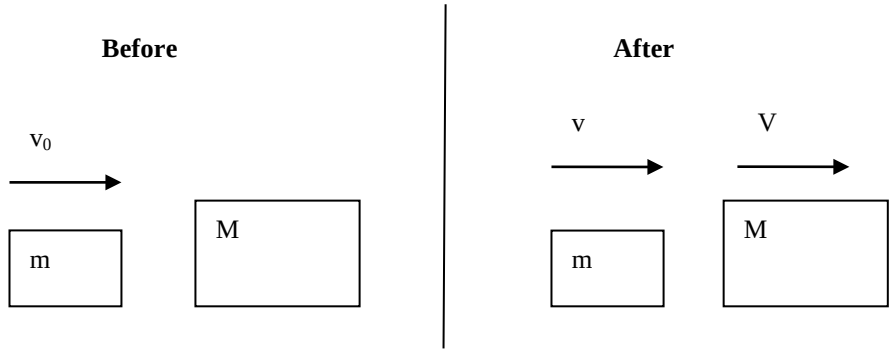

Figure 3.1 Impact between two solid point-like objects of different mass.

The situation after impact can be determined by the laws of conservation of momentum and energy:

$$\begin{aligned} mv_0 &= mv + MV \\ \frac{1}{2}mv_0^2 &= \frac{1}{2}mv^2 + \frac{1}{2}MV^2 + E_i \end{aligned} \tag{3.1}$$

where E_i is the increased internal energy due to the collision.

Equations (3.1) can be solved for the final velocities of the two objects:

$$\begin{aligned} v &= \left(\frac{m}{m+M}\right)v_0 - \sqrt{\left(\frac{Mv_0}{m+M}\right)^2 - 2\left(\frac{M}{m(m+M)}\right)E_i} \\ V &= \left(\frac{m}{m+M}\right)v_0 + \sqrt{\left(\frac{mv_0}{m+M}\right)^2 - 2\left(\frac{m}{M(m+M)}\right)E_i} \end{aligned} \tag{3.2}$$

We note that the final velocities depend on the amount of kinetic energy converted to internal energy E_i. In general, the specific value of this quantity will depend on the material properties of the objects. Still, it is interesting to look at some special cases.

3.1.1 Elastic collision

First, let us assume that no kinetic energy at all is lost in the impact process, i.e. $E_i = 0$. This is called an elastic collision. The expressions for the final velocities then reduce to the following:

$$v = \left(\frac{m - M}{m + M} \right) v_0$$
$$V = \left(\frac{2m}{m + M} \right) v_0 \tag{3.3}$$

To get a feeling for what happens, let us consider a few special cases.

3.1.1.1 Identical masses

First, let the two masses be identical, i.e. $m = M$. This gives us:

$$v = 0$$
$$V = v_0 \tag{3.4}$$

So, for two identical solid point-like elastic objects, the incoming object comes completely to rest after transferring all its momentum to the other object.

3.1.1.2 Small mass impacting huge mass

Still assuming $E_i = 0$, we look at another special case. Say that the object at rest is much more massive than the impacting object, i.e. $M >> m$. A typical example could be a ball impacting a huge wall. If we let M approach infinity, we have:

$$v = -v_0$$
$$V = 0 \tag{3.5}$$

This result corresponds very well with our intuition of a ball bouncing off a wall. But what about conservation of momentum? Kinetic energy is obviously conserved, but the momentum of the ball is reversed while the wall momentum $P=MV$ appears to be zero since the wall is at rest. How can this add up?

In fact, the wall momentum is not zero. Even though the velocity goes to zero, this is compensated by the mass approaching infinity, so mathematically, we have to be quite careful here. Doing everything properly, we find that the momenta p and P can easily be calculated from Equation (3.3). In the M goes to infinity limit, we have:

$$p = -mv_0$$
$$P = 2mv_0 \tag{3.6}$$

which adds up to the initial momentum mv_0. Note that the wall has zero kinetic energy[1] and considerable momentum at the same time!

It is interesting to note that, maybe contrary to our intuition, the receiving object emerges from the collision with a higher momentum than the initial momentum of the incoming object. This is obviously a result of momentum conservation since the receiving object needs a higher forward momentum to compensate for the incoming object being reflected backwards.

3.1.1.3 Huge mass impacting small mass

Finally, let us consider another special case, in which the incoming object has a much larger mass than the receiving object, $m >> M$. The mathematics is straightforward, but it can also be considered as equal to the situation in the previous chapter, just in a difference reference frame. In either case, the results for velocity and momentum are:

$$\begin{aligned} v &= v_0 \qquad p = mv_0 \\ V &= 2v_0 \qquad P = 2Mv_0 \end{aligned} \tag{3.7}$$

Thus, the incoming massive object keeps moving at the same velocity, but pushes the lightweight object ahead at twice the velocity[2].

3.1.2 Non-elastic collision

For a non-elastic collision we need information about the materials to determine the lost kinetic energy. However, without any such knowledge, it is still possible to put an upper limit on the maximum amount of kinetic energy that can be converted to internal energy E_i. This follows from the requirement that the expression under the square root in Equations (3.2) must remain positive (or the final velocities will obtain complex values). The most extreme case is when this expression is zero, which is called a completely non-elastic collision. In this case the converted energy is given by:

$$E_i = \left(\frac{M}{m+M} \right) \left(\frac{1}{2} m v_0^2 \right) \tag{3.8}$$

From Equation (3.8) it is clear that the internal energy will always be smaller than the initial kinetic energy of the incoming object. Further, the final velocities are given by:

$$v = V = \left(\frac{m}{m+M} \right) v_0 \tag{3.9}$$

[1] This is obviously due to the kinetic energy being proportional to the square of velocity, while momentum is proportional to velocity.

[2] Once more, it is tempting to think that this is in conflict with momentum conservation. However, again, this is just an illusion since the incoming object initially has infinite momentum due to the infinite mass. This can easily be shown through a proper mathematical treatment, but here we are only interested in the result.

Thus, if the objects stick together after the collision, maximum kinetic energy is lost. The total momentum is of course still the same as the momentum of the incoming object (as can easily be verified).

If the receiving mass happens to be much larger than the incoming mass (i.e. M approaches infinity), we see that the internal energy approaches the initial kinetic energy of the incoming object. Thus, in this case, the incoming object impacts the receiving object and comes to rest. However, the total momentum is still conserved since the receiving mass is infinite. For the opposite case of the the incoming mass being much larger than the receiving mass, we see that $v=V=v_0$, as could be expected.

It is also interesting to compare the final momentum of the receiving object, depending on whether the collision is elastic or totally non-elastic. We easily see that the ratio is 0.50 for all masses. This is of course because the incoming object keeps moving forward (instead of being reflected backwards) in the non-elastic case, so that the receiving object does not have to compensate for the extra momentum.

3.2 Application to shock attenuation

Let us now analyse shock attenuation using our solid point-like objects. This can be done by considering the same situation as in Chapter 3.1, but with three objects: the incoming object, the "shock attenuation object" and the receiving object, as illustrated in Figure 3.2.

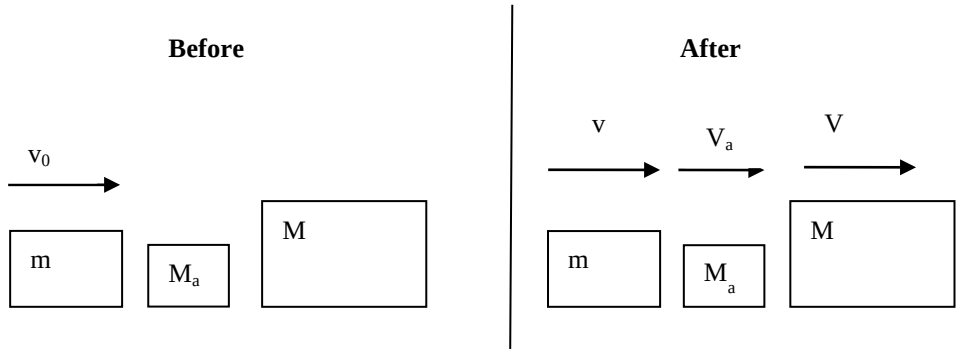

Figure 3.2 Shock attenuation situation with solid point-like objects.

We can analyse this situation as a composition of two impact situations of the kind that were examined in the previous section. First the incoming object impacts the shock attenuation object and then the shock attenuation object impacts the receiving object. From a shock attenuation perspective, the question is: Compared with the incoming momentum, how little momentum is it possible to transfer to the receiving object by tuning the properties of the shock attenuation object?

3.2.1 Elastic shock absorber

Let us first assume that the shock attenuation object has material properties that make the collision with the incoming object elastic. Then the equations of Chapter 3.1.1 can be applied and we immediately note that the shock attenuation object will receive more momentum (in fact, up to twice as much if it has a large mass) than the initial momentum of the incoming object which is reflected backwards.

When the shock absorber proceeds to impact the receiving object, one might be tempted to think that under no circumstances would it be able to transfer less momentum than the incoming object would have done alone. But, imagine that the shock absorber is massive, in which case the situation is composed of the two impacts described in Chapter 3.1.1.2 and 3.1.1.3 respectively. First the shock absorber receives maximum momentum from the incoming object, but after impact it will still move at a very low velocity (due to the huge mass). Then, in the second impact, the receiving object obtains twice the velocity of the shock absorber, which is still a very low velocity. Thus, a massive shock absorber will ensure that very little momentum is transferred to the receiving object.

3.2.2 Completely non-elastic shock absorber

A shock attenuation object with properties that enables some energy to be absorbed in the collision with the incoming object might be even better. From Equation (3.8) we see that the final momentum P_a of the shock absorber after a completely non-elastic collision is given by:

$$P_a = M_a v_a = \left(\frac{M_a}{m + M_a} \right) m v_0 \tag{3.10}$$

The expression in the parenthesis is always less than one, which means that after the collision the momentum of the shock absorbing object will always be less than the initial momentum mv_0 of the incoming object. So, it would seem that by placing a completely non-elastic shock absorber between the incoming and receiving object, we have a way of decreasing the momentum transfer. For a massive shock absorber, we see from Equation (3.10) that the momentum P_a approaches mv_0, the momentum of the incoming object. Thus, in this case very little momentum is transferred to the receiving objects which ends up moving together with the shock absorber at a very low velocity.

3.2.3 Summary

We have seen that to attenuate the momentum for a solid point-like object, it is necessary to have a shock absorber with a huge mass. An example of this could be a steel barrier being placed in front of a person to protect against an incoming football. This seems to be in agreement with our physical intuition, but having a massive shock absorber is not always practical. Does this mean that we might just as well give up on shock attenuation right away? That would be premature. In the real world, it also matters “how” the momentum is transferred, not only how much.

In our simplified analysis using point-like objects, the momentum is transferred immediately, whereas in a real situation the transfer will take some time. The incoming object will generate a force on the shock absorber and ultimately the receiving object. While the total momentum transferred cannot be altered, the duration it takes to transfer the momentum also matters in terms of damage to the receiving object. A huge force that lasts for 1 millisecond may be more damaging than an extremely tiny force which goes on for a year, even though the total transferred momentum is the same in both cases. In fact, this is the idea behind airbags in vehicles. The change in momentum to the passenger is the same with or without airbag, but the change takes place over a longer period of time with the airbag. Similarly, falling from a height and landing on a concrete floor will give more injury than falling from the same height and landing in a haystack, even though the change in momentum is the same in both cases.

3.3 Real solid objects

Let us look at how our analysis must be changed for real solid objects. The situation is still governed by conservation of momentum and energy, but now that the objects are no longer point-like, we must also have local conservation of these quantities. This means that momentum and energy must flow continuously from one location to another. So, instead of all energy being instantaneously transferred from one object to the other, there will be waves set up inside the objects. The propagation of these waves will depend on the material properties and the geometrical shape of the objects. As a consequence, the situation is much more complicated and in general not solvable by analytical methods.

For objects of similar material that are of similar size and shape, the point-like analytical expressions derived in Chapter 3.1 will give a very good estimate, though. For very different objects, stress waves will remain in the objects long after the impact has finished, thereby storing some potential (internal) energy, even for an elastic collision.

3.4 Gas-solid interaction

Finally, let us see how our analysis must be extended when the incoming “object” is a shock wave instead of a solid object. Reflection of shock waves at an interface with a solid object is similar to two solids interacting in that momentum and energy must be conserved at all times. However, typically the solid is too massive to get a very high velocity, which means it obtains very little kinetic energy (since this quantity is proportional to velocity squared). Also, the relationship between energy and momentum is different than for elastic objects. The physics here is quite complicated and everything depends on both the amplitude and shape of the shock wave, so we will not go into detail. However, one result is that a shock wave can be reflected with a momentum of much larger magnitude than the incoming momentum. Consequently the solid can then obtain a larger final momentum than double the incoming momentum of the shock wave. A special case of this is called “confinement”, where a charge is detonated inside some confining structure. This can lead to very high pressures being generated from reflections.

The momentum transfer to the shock absorber leads to shock wave propagation inside the solid objects. When the shock absorber impacts the receiving object, it may both have waves propagating back and forth as well as having a velocity as a whole. The transfer of the shock waves to the receiving object is therefore a very complex situation depending on both the properties of the incoming shock wave, the geometry and material properties of the shock absorber and the geometry and material properties of the receiving object. This means that obtaining an analytical solution is impossible, and that it is far from trivial to determine how a particular situation is going to turn out in terms of the stress wave that eventually is transferred to the receiving object.

Although analytical solutions are ruled out, it is still possible to gain an understanding of how shock attenuation works for real objects and materials. This can be done using numerical simulations and will be the topic of the next chapters.

4 1D-simulations of shock attenuation

The best way to explore the phenomenon of shock attenuation numerically is to start with the simplest possible scenario and then gradually move on to more complex situations. In this chapter we therefore examine shock attenuation in one dimension using ANSYS AUTODYN.

4.1 Set-up

In the 1D-scenario we will detonate an explosive and let the shock wave pass through various damping materials. These damping materials will then interact with a receiving object and we will take the response of this receiving object as a measure of the effect of the relevant damping material. As a receiving object, we will start with a massive steel object, that will only move negligibly.

Thus, our 1D-setup is as follows. We will have an explosive, then air, followed by the damping material in front of a steel bar. To learn about the attenuation properties of the damping material, we will compare the stress inside the steel bar as a function of the damping material properties and layer thickness.

However, for variation of the damping material thickness, there are three possible ways to proceed:

- Constant distance between the explosive and the steel bar . Put damping material at various thicknesses L close to the steel bar. (Figure 4.1a)
- Same as above, except damping material close to the explosive. (Figure 4.1b)
- Constant distance between explosive and damping material. Let the damping material be near the bar, and have different thicknesses L. This means that the explosive will be further away from the steel bar for thick damping materials. (Figure 4.1c)

We will study all of these scenarios to obtain as much information as possible about how shock attenuation works in the 1D-case.

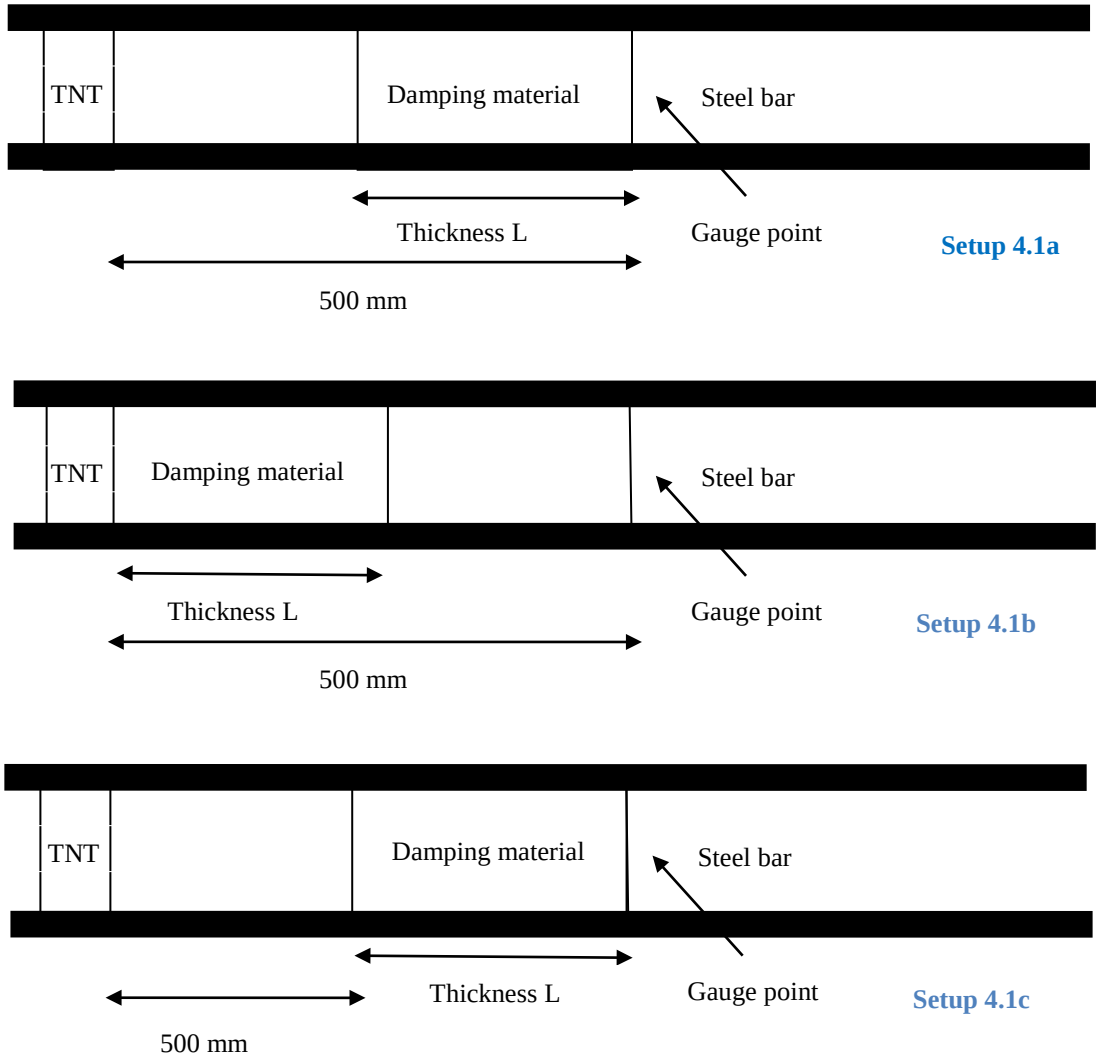

Figure 4.1 Possible 1D-setups

4.2 Material models

We need material models to model the attenuation materials in AUTODYN. It turns out that the Equation of State (EOS) is by far the most important material parameter. The materials are also described by some other parameters, but they will not have much effect in the scenarios that we will study.

In our simulations we will use three materials with very different EOS in order to illustrate how this affects the shock attenuation properties. The material models to be used are as follows:

- Dry sand (From the AUTODYN material library). It has been widely used in simulations with sand and has been seen to give good results (9,10).
- Wet sand (From (11)). This will be an example of a material that is not very porous. It is a special case of a more general sand model, for all kinds of saturations.
- Porous sand (Scaled version of the dry sand model). Very porous material with initial low density. Not been calibrated to any particular material, but may correspond to a very porous material like pumice. Remember that the idea is not to study a given material, but to study the general idea behind shock attenuation.

4.2.1 Loading

The EOS of these three materials are shown in Figure 4.2 and the complete materials models are reproduced in Appendix A. All three are described using the Compaction EOS[3] material model in AUTODYN, but we note that the wet sand is more or less linear and cannot be compacted very much. Physically, this is because the pores of the wet sand are filled with water, which is almost incompressible. On the other hand, the dry and especially the porous sand are very compactable because they contain a lot of empty space. However, we also note that they both eventually reach a limit where there are no empty space (pores) left, where it suddenly becomes much more difficult to compact them further. This is seen as a sudden change in slope for these materials.

[3] Although the AUTODYN theory manual warns against using the Porous EOS (of which the Compaction EOS is an extension) for high pressures and energy absorption, private communication with ANSYS (12) has shown that generally this should not be a problem in our case.

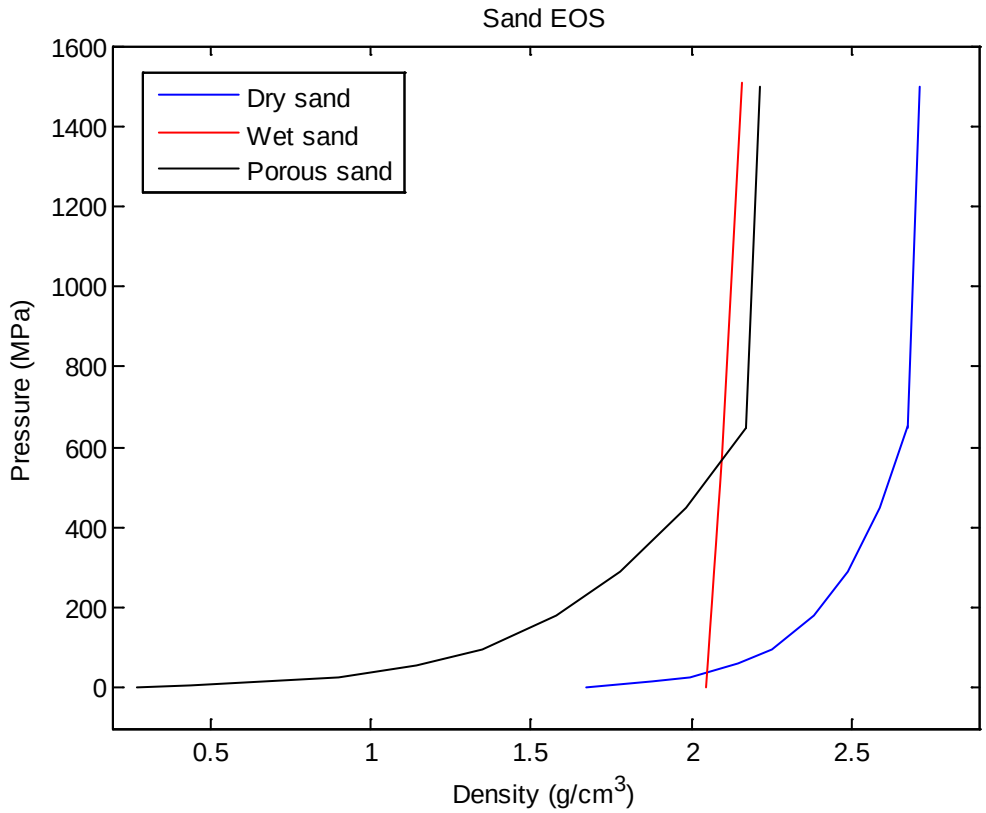

Figure 4.2 EOS of the three materials studied in this report.

4.2.2 Unloading

One important point not shown in Figure 4.2 is what happens to the material during unloading. The compaction is plastic so the material obtains permanent volume change and the unloading will therefore not proceed along the same curve as the loading. If the material is fully compacted, the unloading will go along the new slope. If it is not fully compacted, the unloading will follow a slope that is an interpolation with the fully compacted slope and the totally uncompacted slope. This is shown in Figure 4.3. Note that unless the material is fully compacted, the unloading slope is steeper than the loading slope. Since the propagation velocity of a wave in a given material is related to this slope, it follows that, an unloading wave will travel faster than a loading wave in a porous material, unless the material is fully compacted. This observation will turn out to be crucial in explaining the behaviour of porous materials later on.

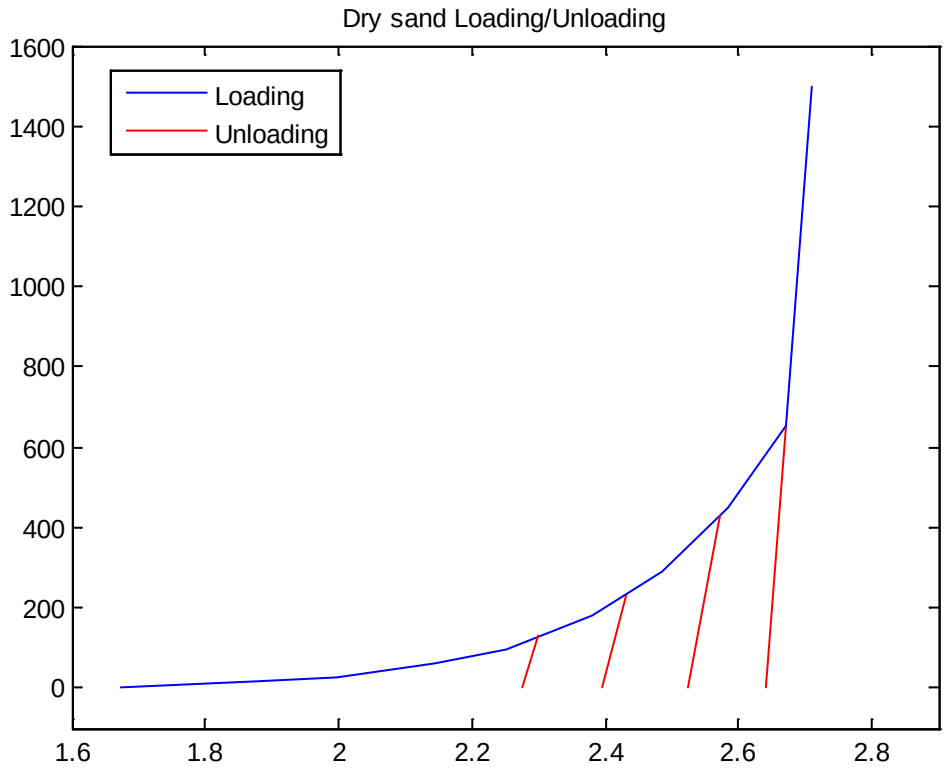

Figure 4.3 The difference between the loading and unloading EOS for dry sand.

4.3 1D confined results

In all simulations we used ANSYS AUTODYN 13.0. The 1D-simulations were performed using the 3D Euler-Godunov solver. The grid size was 1 mm in the axial direction and no boundary conditions were used so that everything was confined in one dimension. In principle the simulations could also have been run in 2D, but since the 3D Godunov solver is more accurate (2^{nd} order) than the 2D Godunov solver, 3D was chosen instead. Since there is only one grid cell in the y and z directions, this did not have much effect on the CPU-time, except taking slightly longer due to the more accurate 2^{nd} order scheme than in 2D.

Initially we used a TNT charge with a 5 mm thickness. (Charge mass has no meaning for a 1D-simulation).

4.3.1 Damping material near bar

We start by looking at the case where the damping material is near the steel bar. The AUTODYN-setup is shown in Figure 4.4. The gauge points are also indicated in the figure. In general there was little difference between the various gauge points.

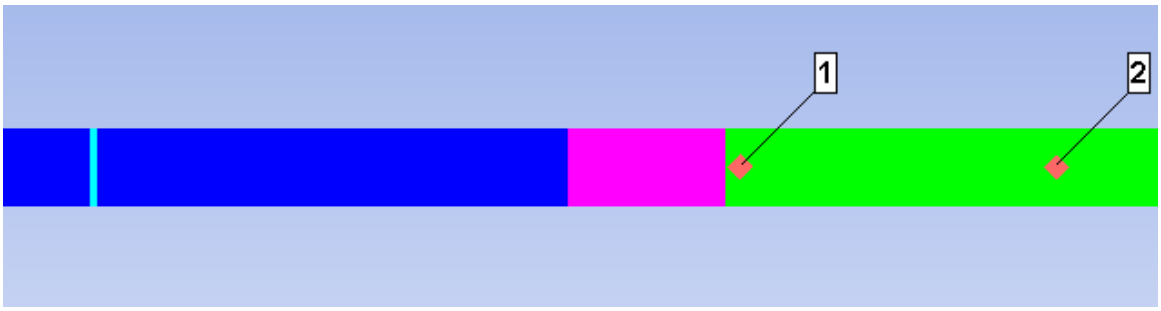

Figure 4.4 AUTODYN setup when the damping material is near the bar. (Charge is cyan, air is dark blue, damping material is pink and steel bar is green).

Starting with dry sand, we ran simulations with this setup for various thicknesses of dry sand. In Figure 4.5 we have plotted the stress as a function of time in gauge #2 for layer thicknesses in the range 100 mm – 500 mm. (Compressive stresses are negative).

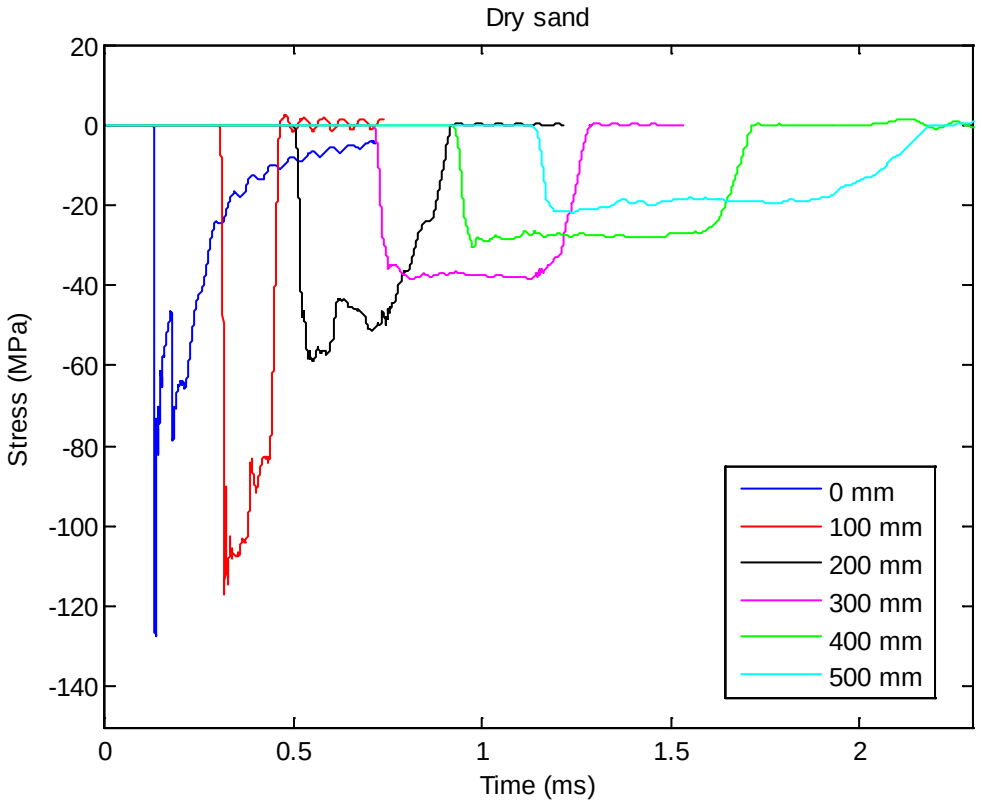

Figure 4.5 Damping material near bar 1D results (Dry sand)

It is worthwhile to study Figure 4.5 very carefully. First, we note that as the damping material layer increases in size, the stress wave arrives later. This is natural since the blast wave travels faster through air than through the damping material. Further, we see that for thicker layers, the amplitude of the stress wave in the bar is reduced, but the wave has a longer duration.

Thus, the hypothesis that sand can attenuate the shock wave amplitude seems to have been confirmed. The amplitude clearly decreases as more sand is placed in front of the steel bar, and for layer thicknesses of 300 mm, 400 mm, 500 mm the amplitude is down to 20-40 MPa from around 130 MPa. This represents a considerable reduction and the pulses are square and quite nice looking, although they have a longer duration, which is in accordance with our discussion in Chapter 3.

But, let us take a closer look at what happens for damping material layers of less than 100 mm thickness. This is shown in Figure 4.6.

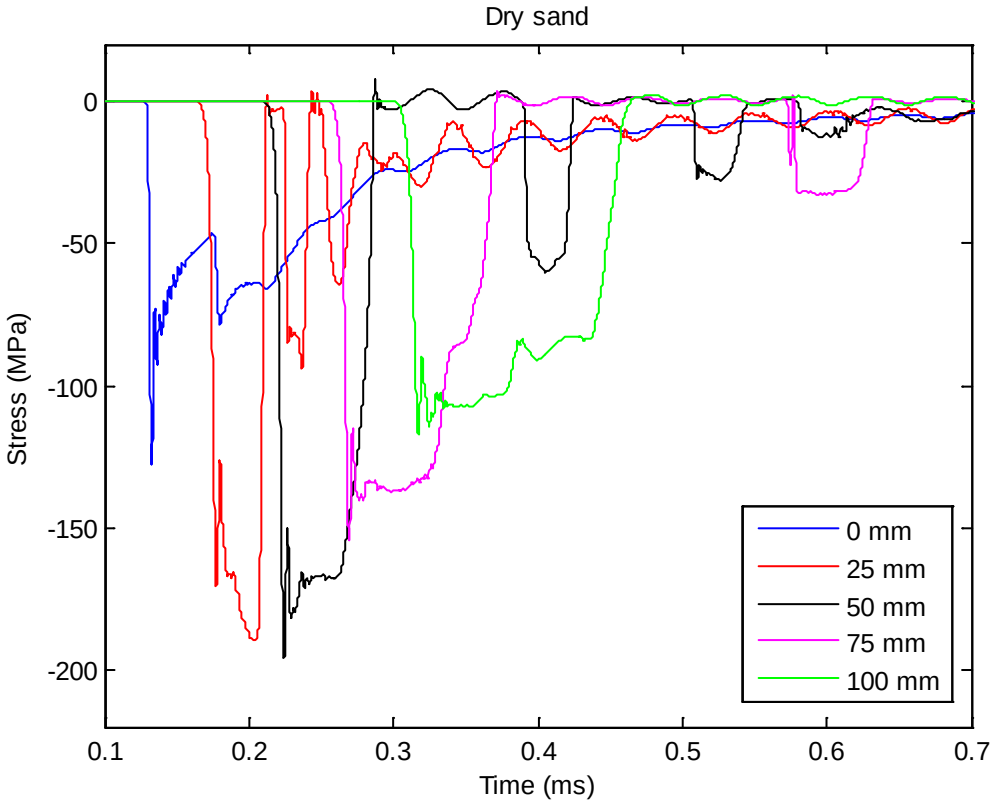

Figure 4.6 Damping material near bar 1D results (Dry sand – thin layers)

These results may come as a little surprise. When the damping material thickness decreases below 100 mm, the amplitude does not immediately converge towards the result for no damping material. Instead it continues to increase and for 75 mm thickness, the amplitude is higher than without damping material. It increases further for 50 mm thickness, but is smaller for 25 mm, so there seems to be a thickness that gives maximum amplitude somewhere between 25 mm and 50 mm.

So, the dry sand "damping" material can both increase and decrease the shock amplitude! This is a very important result which reminds us that shock attenuation is far from trivial. It will now be interesting to see if the porous and wet sand materials exhibit the same behaviour. The results from simulations with these materials are shown in Figures 4.7 and 4.8.

The porous sand does indeed show the same tendency as dry sand, and in fact, it behaves even more extremely. Around 200 mm of damping material is needed to avoid enhancement of the amplitude. For 25 mm and 50 mm the enhancement is enormous.

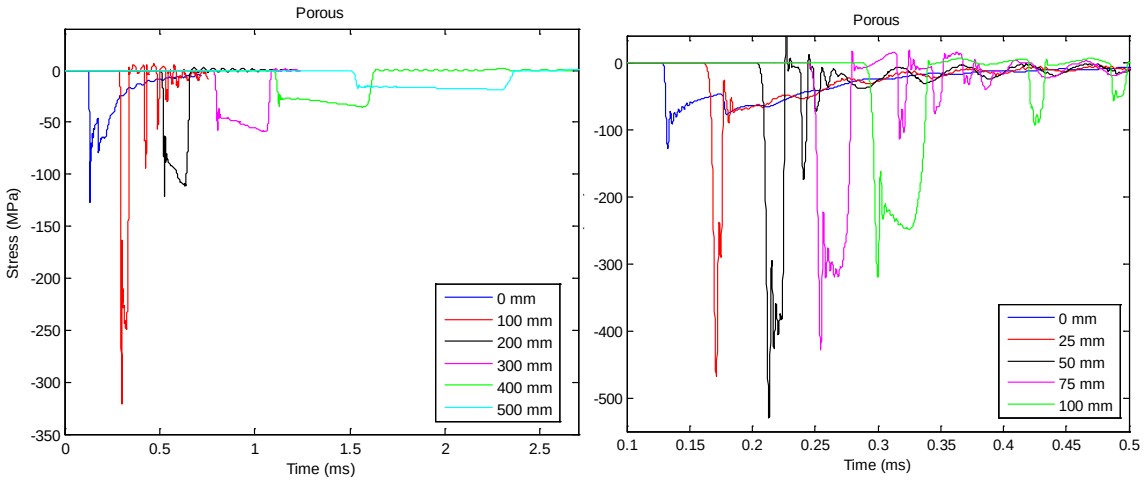

Figure 4.7 Damping material near bar – 1D results (Porous sand)

However, the results from wet sand show a totally different tendency with no damping for any thickness. Wet sand is the only non-porous material, so this seems to indicate that there is something special about the porous materials.

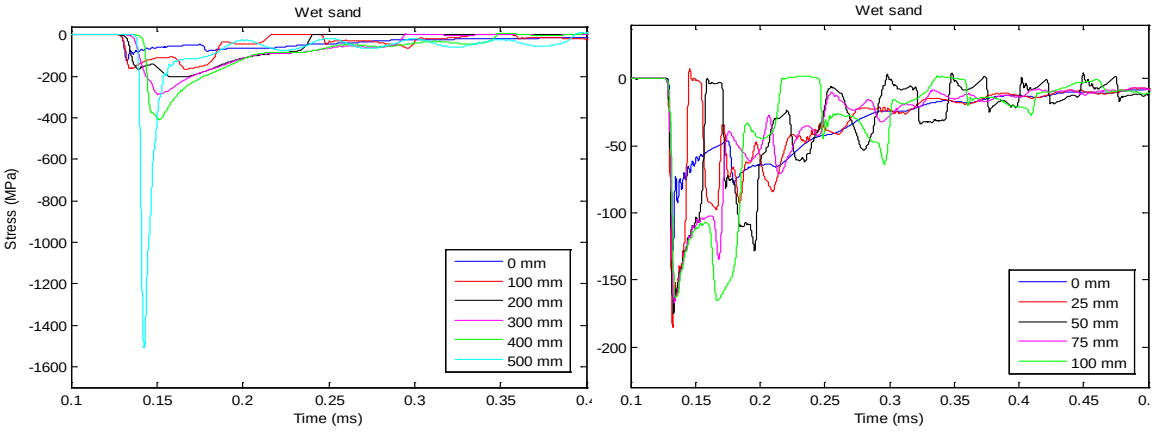

Figure 4.8 Damping material near bar – 1D results (Wet sand)

To visiualize things easier, let us define a factor “relative damping” as the ratio of the maximum stress for a given thickness of damping material compared with the maximum stress when no material is present. Thus, if the “relative damping” is more than 1.0, we have shock enhancement and not damping. Figure 4.9 shows the results for each material as a function of thickness of the damping material (notice that the scales on the axes are different):

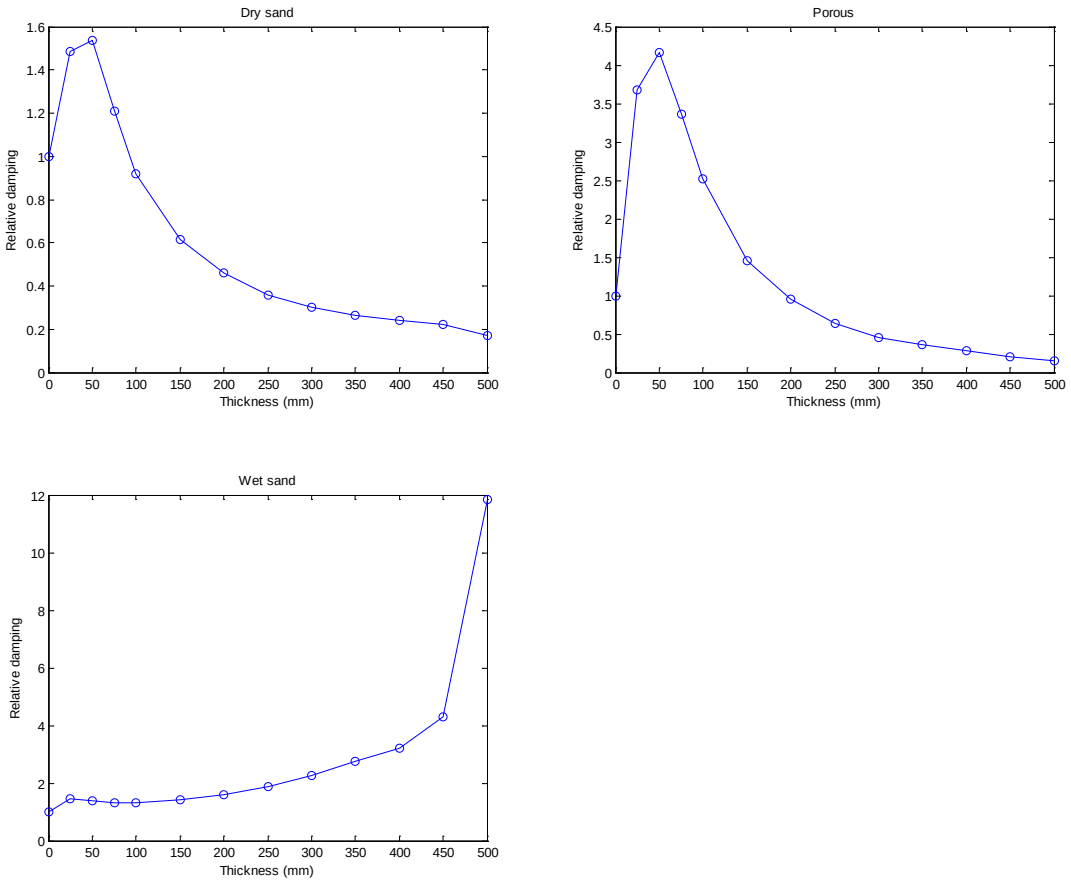

Figure 4.9 Relative maximum stress as a function of material thickness for all three materials.

Let us briefly sum up what we have learned so far from our 1D-simulations with the damping material located close to the steel bar:

- The two porous materials are able to attenuate the shock amplitude if the thickness is large enough. This leads to longer duration of the wave at a lower amplitude. However, if the thickness is insufficient, these materials may actually increase the amplitude of the shock wave compared to the case with no damping material!

- The non-porous wet sand behaves completely different from the porous materials. In none of the situations investigated was it able to attenuate the shock wave. The trend is also different from the two other materials in that the amplitude increases for thick layers of wet sand.

Later we will try to explain these points physically. However, let us first look at Setup 4.1b in where the damping material is close to the charge instead of close to the steel bar and see whether the behaviour is roughly similar. This might give us further clues about what is going on.

4.3.2 Damping material near charge

The numerical results for dry sand in Setup 4.1b are shown in Figure 4.10. It is clear that the stress wave measured in the bar completely changes character, except, of course, for the cases of 0 mm and 500 mm thickness, which are obviously the same as in the previous setup.

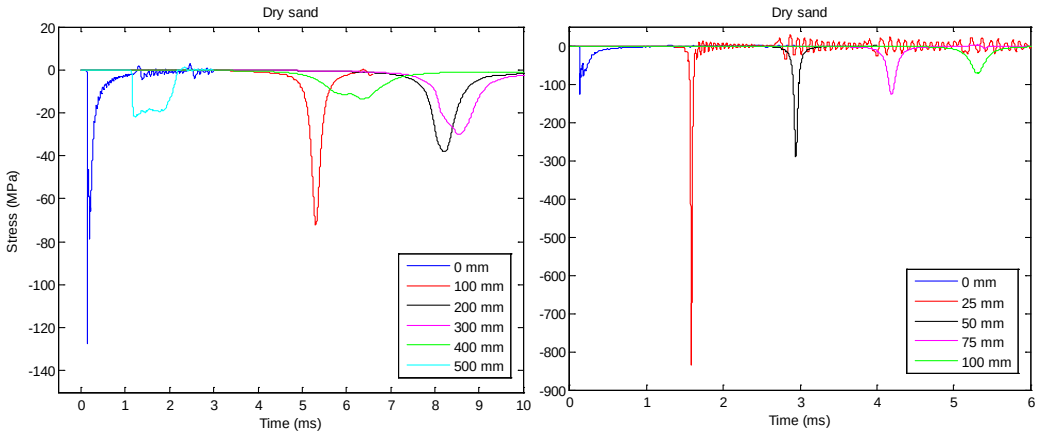

Figure 4.10 Results for dry sand when the material is close to the charge.

Notice, in particular, the behaviour of the arrival time of the wave. It varies a lot more as a function of layer thickness than when the damping material was close to the bar. This is due to the whole process being very different from the previous setup. Instead of the shock wave propagating through the damping material, the material is accelerated by the shock wave and moved until it impacts the bar. For thin layers, there is not much material to accelerate, which speeds up the process, but on the other hand, the material is then initially located further from the bar. So, there are two "competing" factors, contributing to the relatively complex behaviour of the arrival time.

In Figure 4.11 we have summarised the results for relative damping for all three materials.

Porous sand has much the same characteristics as in the previous setup and behaves similar to dry sand, although in a more extreme fashion. However, wet sand again seems to behave peculiarly. There is a slight amplification for 25 mm, but for the other thicknesses the maximum amplitudes are substantially reduced, until around 400-500 mm thickness when the amplification starts increasing dramatically. We will get back to explaining this behaviour, but first we shall look at the final possible 1D-setup.

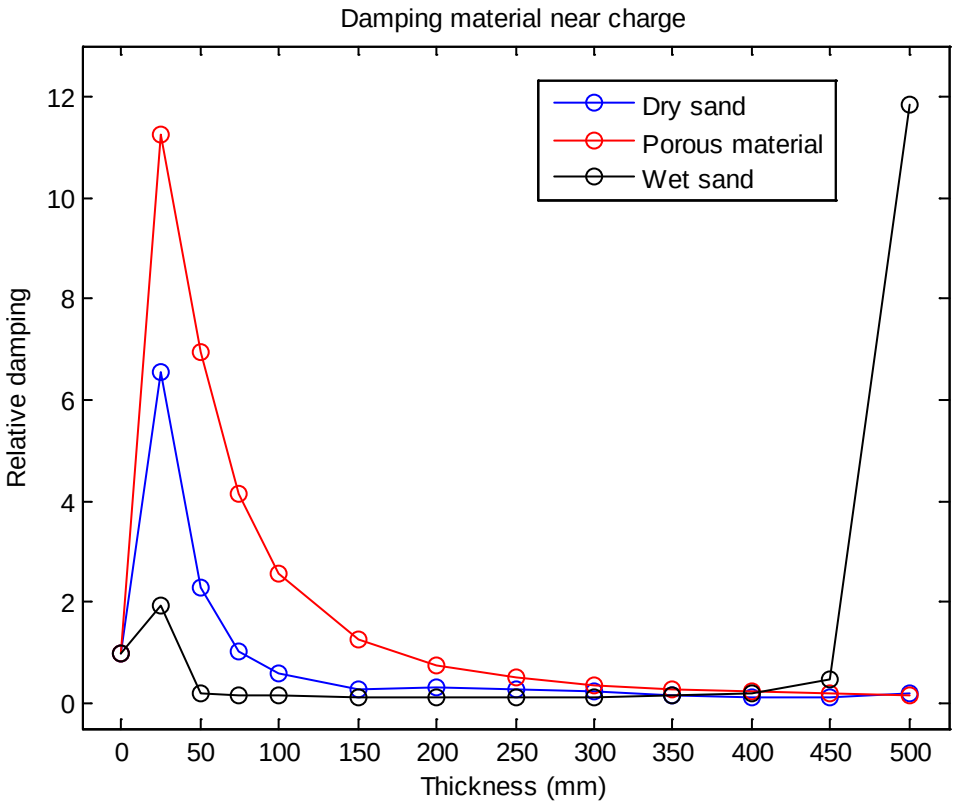

Figure 4.11 Relative damping for all materials when the damping material is near the charge.

4.3.3 Damping material at fixed distance from charge

In Setup 4.1c the charge is at a fixed distance (500 mm) from the damping material. This means that for thicker layers, the charge is moved further away from the steel bar. For all thicknesses, exactly the same shock wave will arrive at the damping material.

The relative damping results for this case are shown in Figure 4.12 for all materials.

The overall picture is relatively similar to the other cases for dry sand and porous material. Again wet sand differs quite a bit. The relative damping is almost constant (giving increased stress by a factor of roughly 2.0), but unlike in the other two set-ups, it does not “blow up” for thicknesses close to 500 mm.

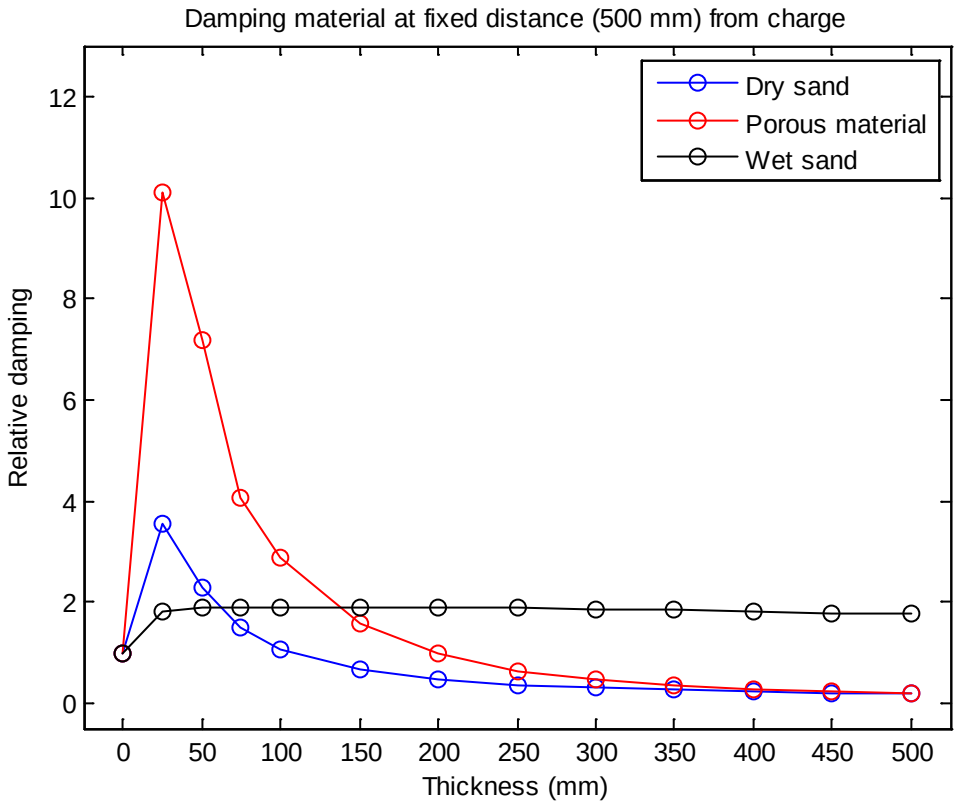

Figure 4.12 Relative damping for all materials when the damping material is at a fixed distance from the charge.

4.4 Summary of 1D results so far

The confined 1D-simulations have shown that shock attenuation (i.e. reduced amplitude) is certainly possible in some cases, but it has also raised several questions that need answering in order to fully understand the process:

- Why do we get shock enhancement instead of attenuation for thin layers of the porous materials?
- What is the criterion to get attenuation and enhancement? Both in terms of material thickness and material properties.
- Why is the non-porous material (wet sand) behaving so totally differently, especially why this huge enhancement for thick materials?
- Why are the results so different for the various setups?

To better answer these questions, it will be useful to vary some other parameters to see what happens and obtain more information that may be helpful in understanding what is going on. Let us first look at the effect of increasing the amount of TNT.

4.5 Effect of charge size

To study the effect of charge size, we ran simulations with the TNT thickness increased to 20 mm, i.e. 4 times more TNT than previously. Setup 4.1a (damping material near steel bar) was used.

In Figure 4.13 we compare the relative damping results for 5 mm and 20 mm TNT.

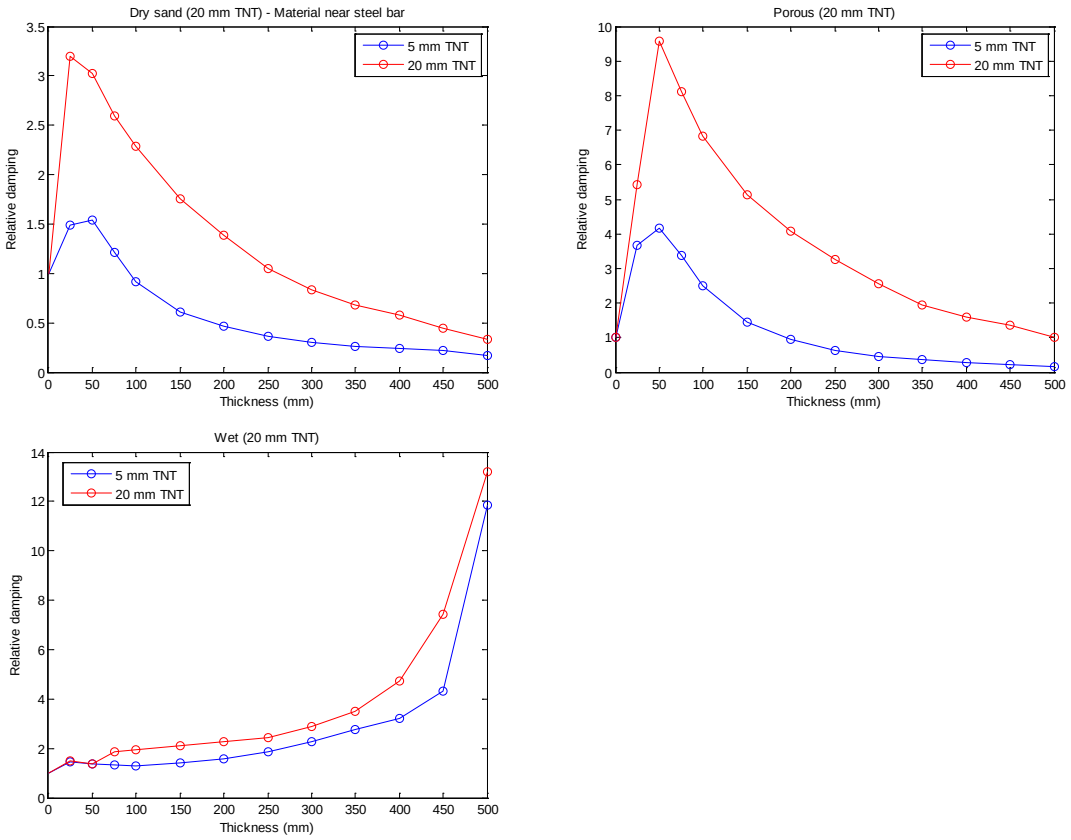

Figure 4.13 Relative damping for 5 mm and 20 mm TNT when damping material is near the bar.

Looking first at the results for dry sand, we see that there generally is less attenuation of the shock wave amplitude for the larger charge. Now a much thicker layer of dry sand is needed to decrease the relative damping. (This was not obvious – remember that the figure shows relative damping, not absolute damping!) In fact, about 250 mm dry sand is needed to reduce the amplitude to the same level as without any damping material. Anything less will just lead to an increased maximum stress amplitude.

For porous sand the tendency is much the same as for the dry sand, but, as usual, more extreme. Most thicknesses increases the stress amplitude for the big charge and only for around 500 mm is the amplitude down to the same as without any damping material.

The wet sand behaves in a similar fashion to the 5 mm situation, except for slightly more amplification for most thicknesses. Again the behaviour of (non-porous) wet sand remains dramatically different from the two porous materials.

It is a potentially important observation that for porous materials a thicker layer of damping material is required to achieve attenuation for larger explosive charges. It could mean that in a real situation, where an object is to be protected, the required amount of damping material depends on what kind of incident shock wave is expected.

4.6 Attenuation as a function of mass

In each plot so far, we have looked at attenuation of the shock amplitude as a function of the material thickness for each damping material. However, we could also look at damping as a function of material mass. The porous sand has a much smaller density than the other materials and a given layer of it will have much less mass than a similar layer of dry or wet sand. In some practical situations, one might want to use as little mass as possible to achieve the desired attenuation. If we plot the results from Figure 4.12 (charge at fixed distance from explosive) as a function of mass/area instead of as thickness, we obtain Figure 4.14.

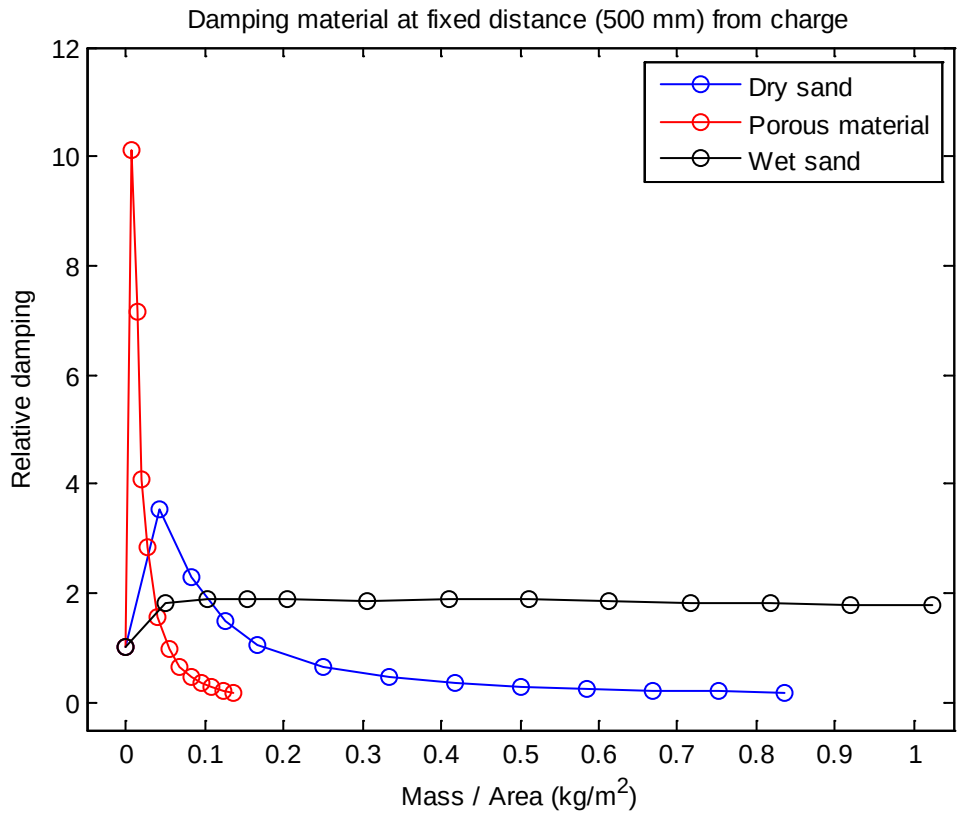

Figure 4.14 Relative damping for each material as a function of total mass of the damping material.

Basically the plot is the same as Figure 4.12 but with the graphs having been scaled a little bit. However, the new plot gives us a totally different impression. It is now immediately clear that only a little mass of porous sand (which will obviously take up quite a bit of volume because of the low density) will lead to attenuation of the shock amplitude. If mass is the defining parameter in a practical situation, the porous sand seems to do well compared to dry and wet sand, as long as there is enough of it.

4.7 Other ways of measuring attenuation

What if we had measured attenuation a different way, for example by looking at the impulse transferred to the bar, instead of the maximum stress amplitude. Would the results have been roughly the same? Let us examine this in more detail, using Setup 4.1a (damping material close to steel bar) as an example.

4.7.1 Impulse in steel bar

AUTODYN stores the total impulse of the steel bar as a function of time, making this parameter easy to obtain. Integration of the stress curves in Figure 4.5 would also have given quite similar results for the impulse. In Figure 4.15 we have plotted the impulse for different layer thicknesses (5 mm TNT – damping material close to the bar).

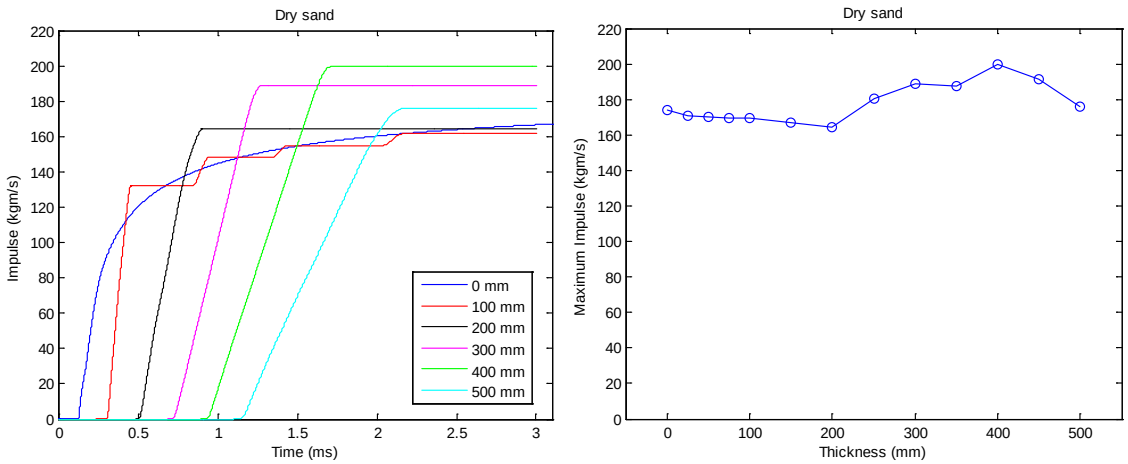

Figure 4.15 Impulse in steel bar for different layer thickness of dry sand (damping material close to the bar).

We note that the impulse of the "no material" case takes a very long time to converge (if it converges at all). Further, only for 100 mm and 200 mm thickness does there seem to be any damping of the impulse. Finally, note how the impulse seems to rise in steps for the 100 mm case.

Figure 4.15 also shows the maximum impulse as a function of layer thickness. As we see there is no clear tendency, except that the impulse is not much affected by the damping material. From measurements of the momentum instead of the maximum stress, one might be tempted to conclude quite differently about the shock attenuation capabilities of dry sand.

4.7.2 Impulse in short bar (projectile)

What if we had measured the impulse in a short steel bar (effectively a projectile) instead of our very long bar that does not move much at all? This was tested by making the bar 40 mm long and then running the same setups as in the previous chapter. The results are given in Figure 4.16. We see much of the same pattern as for the long bar, except that the impulse decreases substantially when we have more attenuation material.

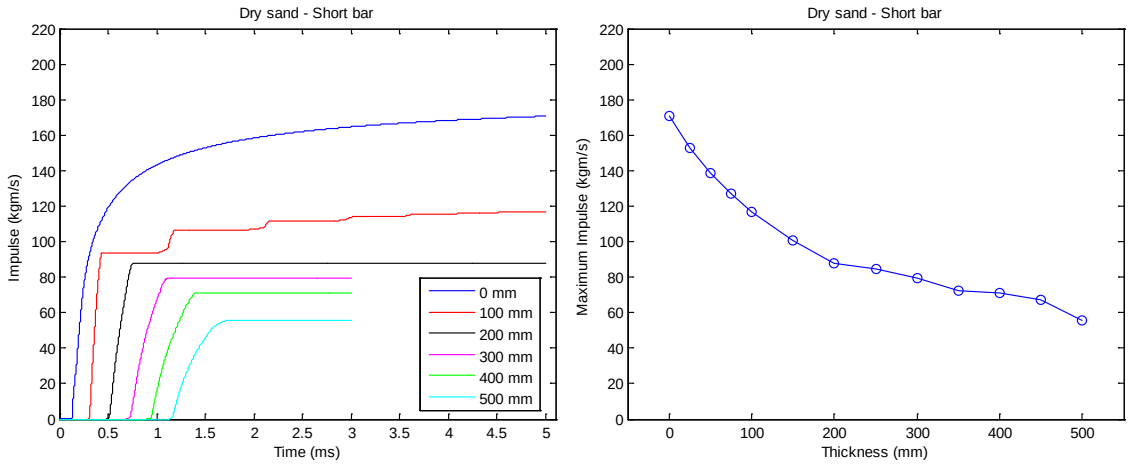

Figure 4.16 Impulse in short steel bar for different layer thickness of dry sand (damping material close to the bar).

In contrast to the long bar simulations, Figure 4.16 shows that the maximum impulse decreases with increasing layer thickness. Thus, from this experiment, the dry sand seems to be able to attenuate the impulse. However, the damping material is exactly the same both in the long and short bar experiment. So, how can the results be so different?

The reason is that the short bar now suddenly has a mass that is comparable to the damping material. After interaction with the shock wave, the damping material will be moving in the same direction and at the same velocity as the bar. It therefore carries the "missing momentum", just as we saw in our single-point object example in Chapter 3.2.2. For a long bar, we don't notice this effect because the mass of the bar is so large compared to the damping material, making the momentum in the damping material negligible.

It is starting to look like our measuring method seems to decide what kind of results we obtain. This may have something to do with DSTL, NAVAIR, NTNU and FFI obtaining so different results on attenuation.

5 Shock wave propagation in the damping material

To learn more about shock attenuation and perhaps obtain further clues about what is going on, we will study how shock waves actually propagate inside the various damping materials in our 1D-setup. As an example we will look at the cases of 5 mm TNT in Setup 4.1c (charge at fixed distance from the damping material) for dry and wet sand.

5.1 Shock wave propagation properties in dry and wet sand

Let us start with dry sand and see how the velocity and pressure waves travel through the first 50 mm of attenuation material. The velocity and pressure profiles are depicted at different points in time in Figure 5.1, showing how the wave evolves as it propagates through the damping material.

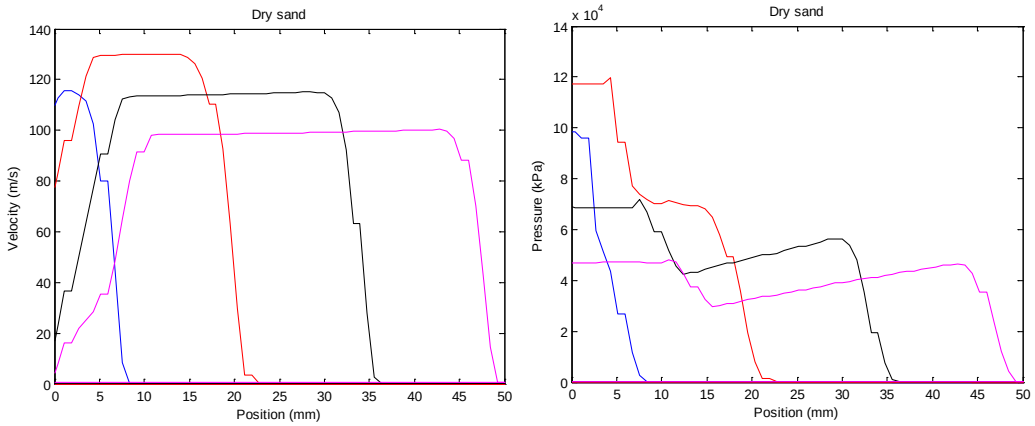

Figure 5.1 Shock wave propagation in dry sand

We see that the maximum velocity first increases a little bit, but then starts decreasing. In the beginning, the dry sand obtains a velocity of more than 100 m/s. The pressure increases slightly first, but then also falls off. Note that the position of the gas/dry sand interface, initially at *x = 0 mm*, moves slightly since the sand acquires such a high velocity.

So, how does the wet sand velocity and pressure profiles compare? This is shown in Figure 5.2.

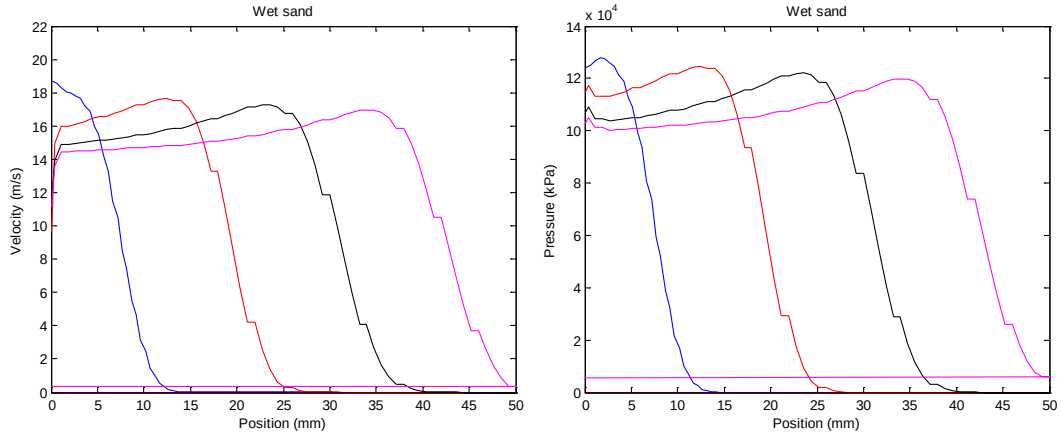

Figure 5.2 Shock wave propagation in wet sand

It is clear that there is a huge difference in the behaviour of the shock waves in dry and wet sand. In wet sand, both the pressure and velocity wave have an almost constant amplitude and the wet sand acquires a much lower velocity than the dry sand, so the gas/sand interface does not move much.

If the 50 mm of attenuation material is followed by a steel bar (as in our example in Chapter 4), the properties of the shock wave at the steel bar interface will depend strongly on what kind of material it has passed through. It is not immediately obvious whether the dry or wet sand will transfer the highest stress pulse to the steel bar. However, by going back to Figure 4.12 we see that the low pressure dry sand impacting at high velocity gives a slightly higher stress amplitude in the bar than the high pressure wet sand impacting at low velocity.

But, what if the damping material thickness had been 500 mm? How do the waves continue to propagate in both materials? Let us examine that as well. The results for dry sand are shown in Figure 5.3 and in Figure 5.4 for wet sand.

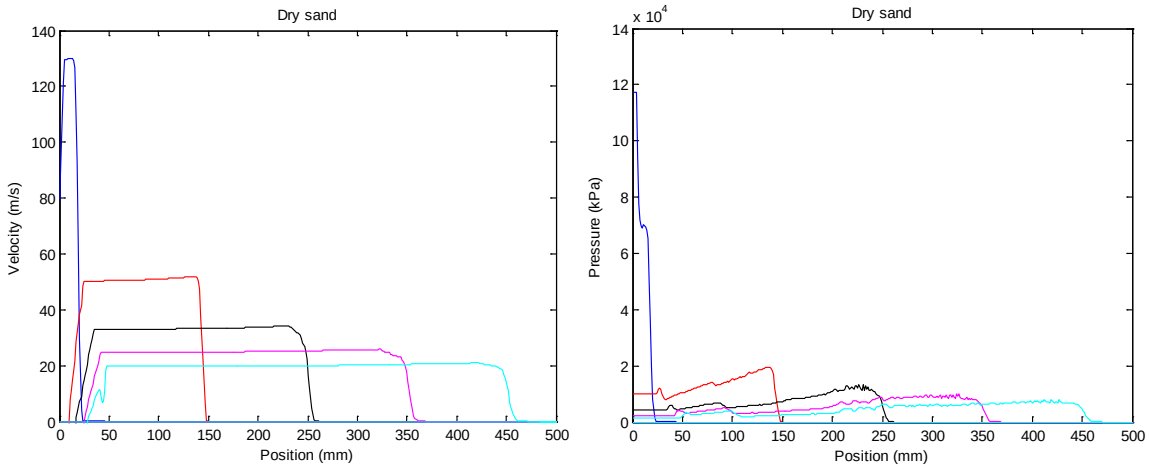

Figure 5.3 Shock wave propagation in dry sand (500 mm thickness)

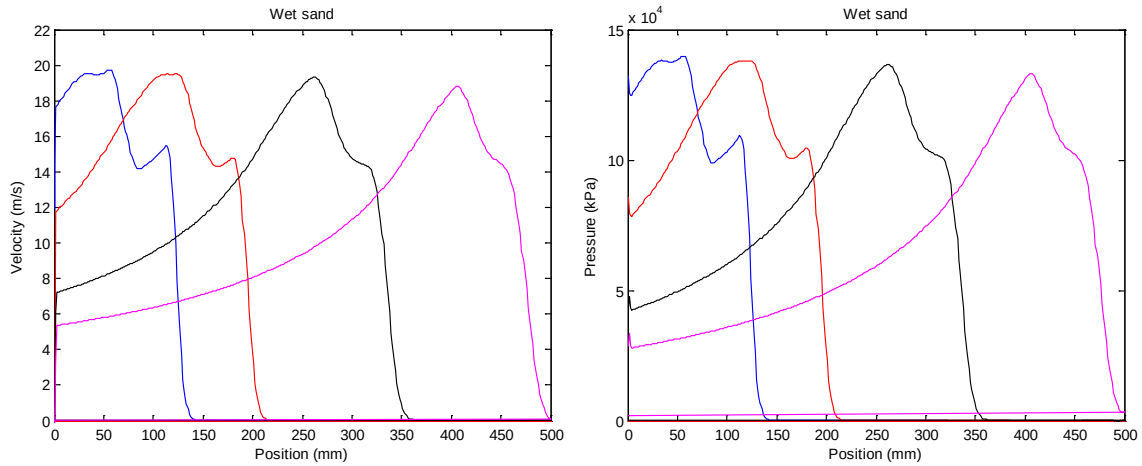

Figure 5.4 Shock wave propagation in dry sand (500 mm thickness)

We see the same tendency. For dry sand, the velocity and pressure decreases, whereas for wet sand the velocity and pressure are roughly constant as the wave propagates. We also notice that for 500 mm damping material, the velocity of the dry sand has dropped to roughly the same as the wet sand wave. However, the pressure in the dry sand is much smaller.

What happens now to the shock wave if the 500 mm attenuation material is followed by a steel bar? Again, this is not obvious, but from Figure 4.12 we see that now the dry sand gives a much lower amplitude in the steel bar than the wet sand.

The propagation properties of the shock waves in dry and wet sand are summed up in Figure 5.5, which shows the velocity and pressure amplitude as a function of how far the waves have travelled.

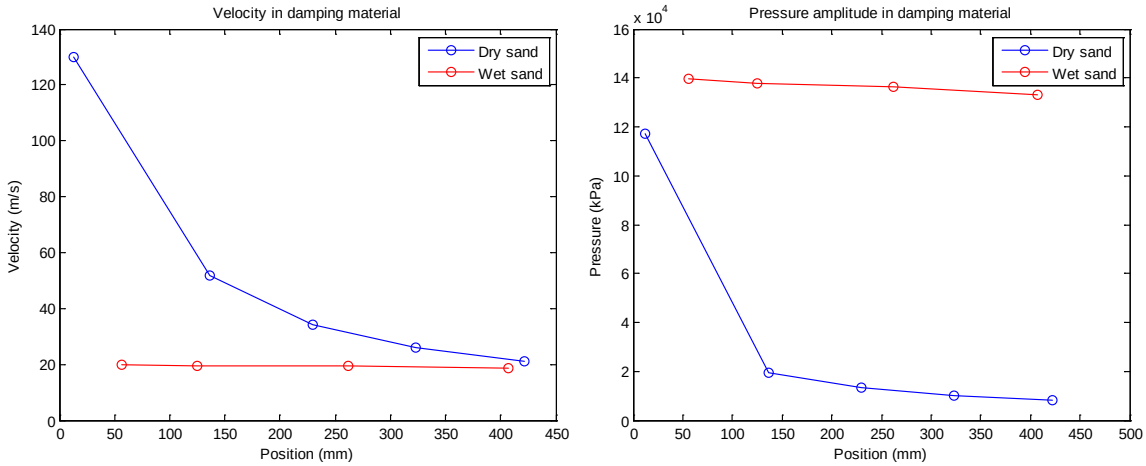

Figure 5.5 Shock wave properties in dry and wet sand.

So, a picture of what is going on is starting to emerge. A porous material is quite easily compacted and when impacted by a shock wave it therefore acquires a high velocity (even more so if the initial density is low). However, as the shock wave travels through the porous medium, both the pressure and velocity amplitudes fall quite quickly. This is very different from the situation in a non-porous medium, where the wave will travel at roughly constant amplitude. (This is of course standard wave propagation theory).

If the damping thickness is small, a porous material acquires a much higher velocity than a non-porous material because the velocity does not “have time to fall”. This answers the earlier question of why we get higher stress for thin layers. If the damping material obtains such a high velocity that the subsequent impact generates a higher stress than the blast wave would have done without the damping material, then we have this situation. (It is not obvious how to calculate analytically the condition for this to happen). However, for thick layers, both velocity and pressure have fallen off considerably before the wave has propagated through the porous damping material, leading to low amplitude stress waves in the steel bar compared with a non-porous material.

5.2 Analysis of porous shock wave propagation

One question that remains is why the porous material behaves so differently from the non-porous materials. Why does both the pressure amplitude and velocity amplitude fall as the wave propagates through the material? This is an important question and we will therefore look into it in some detail.

The behaviour of a shock wave in a porous material is best understood by looking at a simpler wave than what is generated by an explosive. To illustrate things we will therefore examine the propagation of a square pulse through our dry sand. This can easily be set up in a 1D-situation using a velocity boundary condition in AUTODYN. The behaviour of such a wave is shown in Figure 5.6, where the pressure and velocity profiles are shown at different points in time, illustrating how they propagate through the porous dry sand.

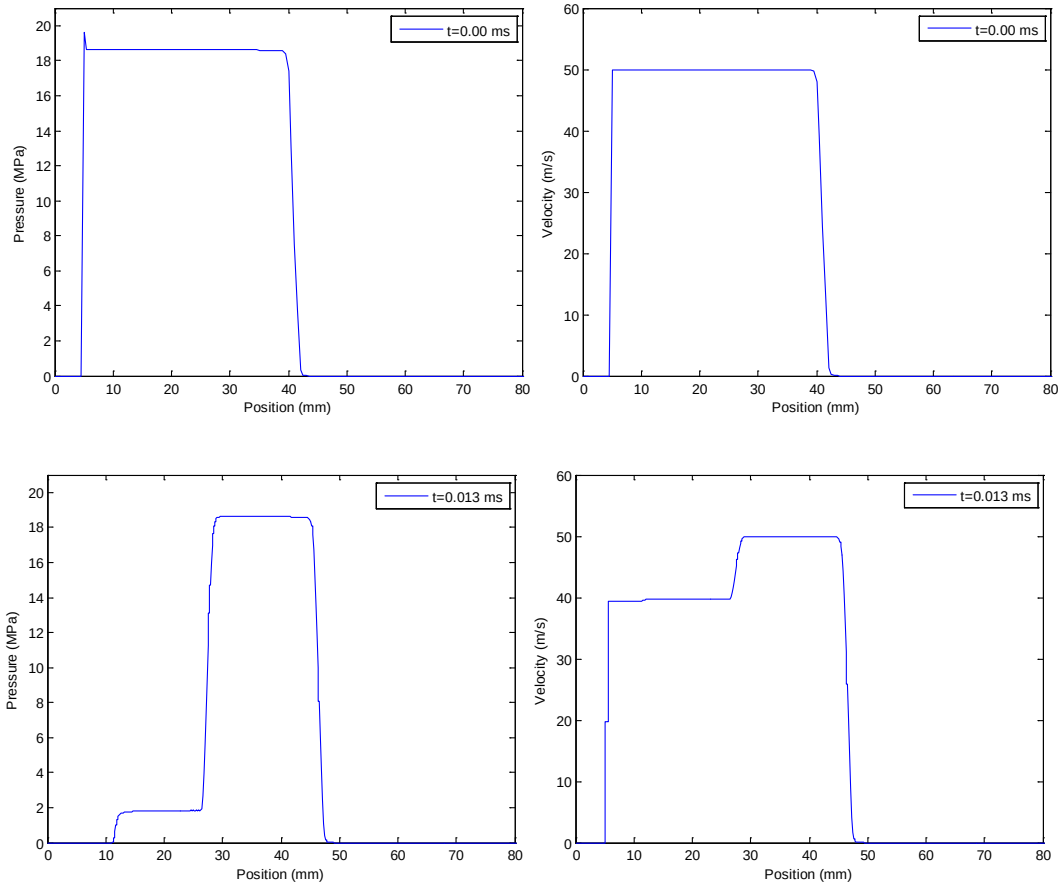

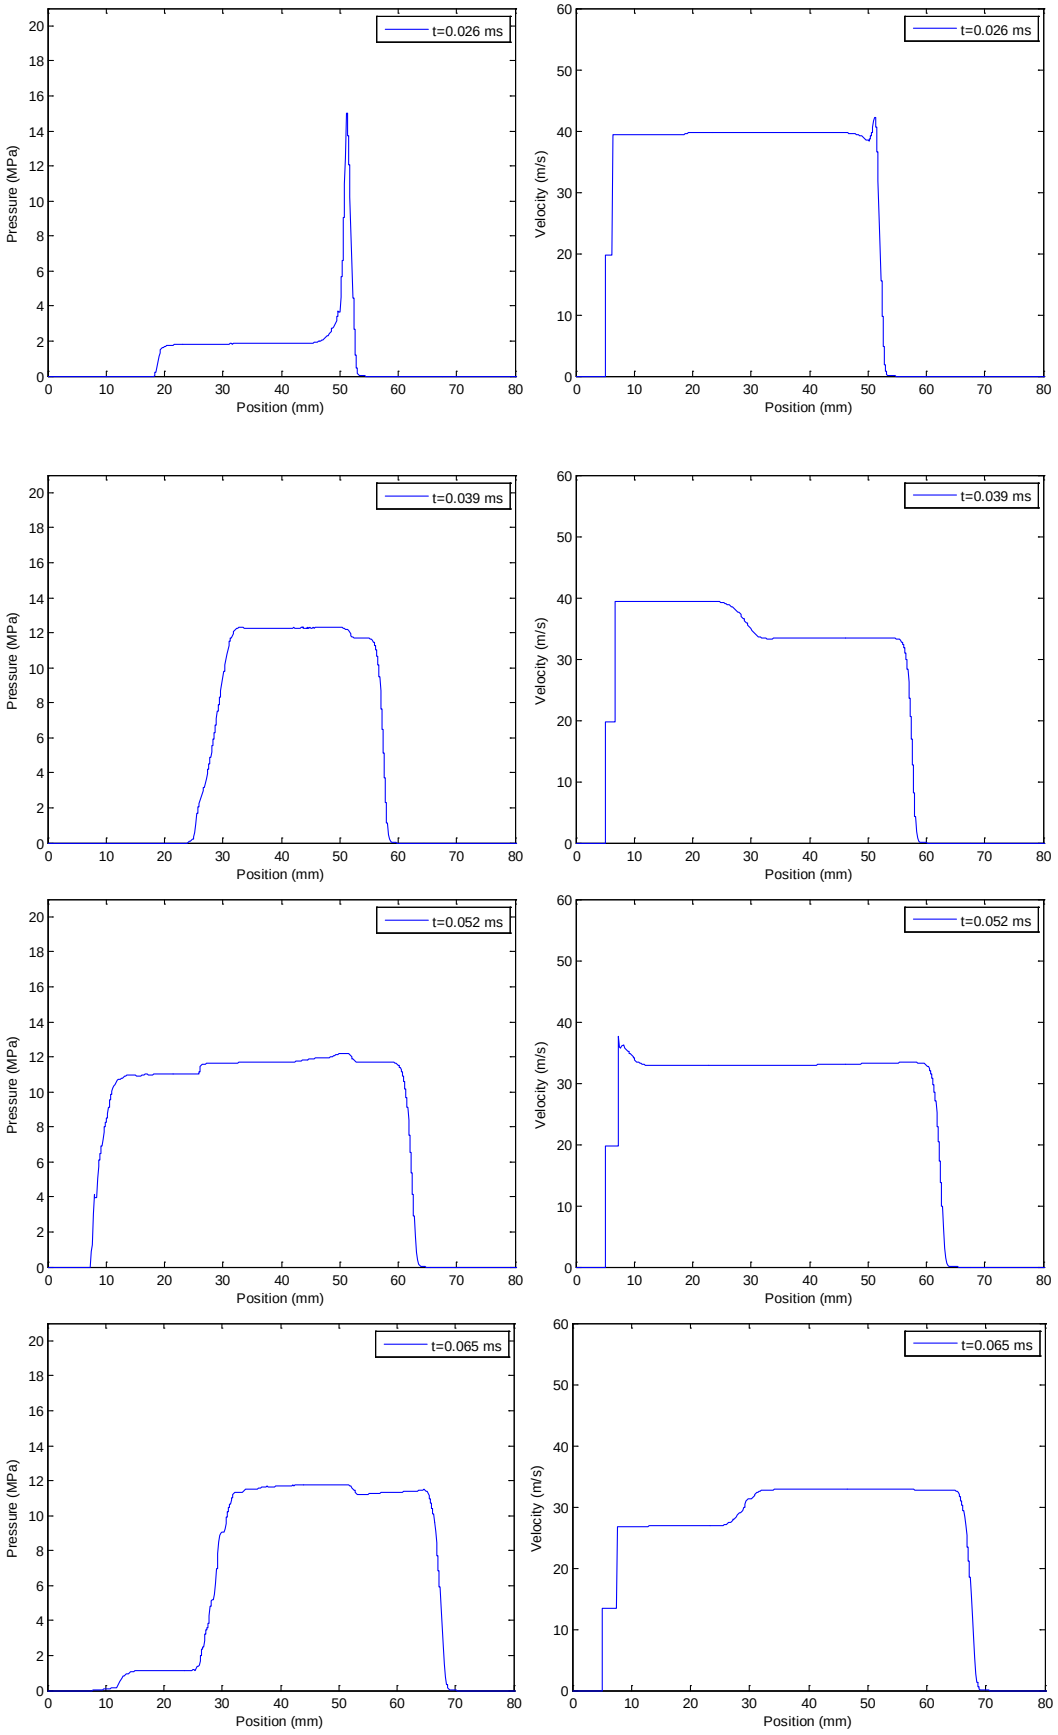

Figure 5.6 Square shock wave propagating through dry sand.

We see that both the pressure and velocity have a square shape initially. Now, look at the situation a little later. The important thing is that after the loading has stopped, an unloading wave propagates through the material, reducing the pressure quite significantly and also the velocity slightly. Further, the propagation velocity of this unloading wave is higher than the velocity of the shock front. Thus, the unloading wave will eventually catch up with the shock front. The illustrations show this happening as well as what happens next. The unloading wave is reflected as a loading wave, going back through the material increasing the pressure and further reducing the velocity. The whole process then repeats once this wave reaches the back of the rear end of the material and reflects as an unloading wave (last two plots in Figure 5.6).

The end result, after several reflections, is a wave of long duration and relatively low velocity and pressure. As we saw in Figure 5.5, this is typical for a shock wave in a porous material. It all comes down to the difference in loading and unloading properties for porous materials. As mentioned in Chapter 4.2.2, when a porous material unloads it follows a linear path in a pressure-density diagram instead of going back to the initial state. Since the wave propagation velocity depends on the slope of the loading/unloading in such a diagram, it follows that unloading waves will travel much faster than loading waves in a porous material. This is the big difference with a non-porous material, where loading and unloading have the same slope and loading and unloading waves therefore travel at the same velocity.

But, what happens if the loading has such a long duration that the wave has already propagated through the whole porous material before unloading begins? An alternative way of saying this is that the damping thickness is too small compared with duration of the loading. This is exactly the situation we have already encountered in Chapter 4, where the whole material moves at a very high velocity and we receive shock enhancement instead of damping.

5.3 Wet sand behaviour

We can also explain why the stress amplitude “blows up” for wet sand near 500 mm thickness in the two set-ups where the damping material is close to the bar (Figure 4.9) and to the explosive (Figure 4.11). This is actually a confinement effect of the explosive. Since the wet sand hardly compacts, large pressures build up as the detonation wave is confined. If there is some gap between the charge and material, the detonation wave amplitude will have attenuated slightly before impacting the damping material and the effect weakens. In Setup 4.1c the charge is always at a distance from the material and there is therefore no confinement effect and “blow up” for wet sand (Figure 4.12). Note that this confinement effect also happens for dry sand and porous sand, but since these are much lighter and easier to compact and accelerate, the confinement is not as strong. Further, the pressure is still considerably attenuated as it moves through the material.

Having more or less understood how things work in a 1D confined situation, we shall now look at more complicated setups, i.e. the experiments described in Chapter 2. We will perform numerical simulations of all these setups, using our three “test materials” to gain further insight into the

physics. With this insight and applying our obtained knowledge from the 1D confined situation, we will try to explain why the various experiments seem to give different conclusions.

6 NTNU experiment

Let us start by reviewing the NTNU pendulum experiment (1).

6.1 Experimental set-up and results

As we remember from Chapter 2, the objective of the NTNU experiments was to investigate whether aluminium foam panels could be used for protection against blast loading. Using a ballistic pendulum setup they measured the transferred impulse from a detonated charge with and without aluminium foam covering the pendulum.

Two different PE4 charges were used, 1 kg and 2.5 kg and these were placed at a distance of 500 mm from foam panels attached to the ballistic pendulum. The panels had dimensions 68.4 cm x 70 cm x 6 cm. Two different foam densities were used and also the addition of an aluminium cover plate was tried. The experiment is illustrated in Figure 6.1.

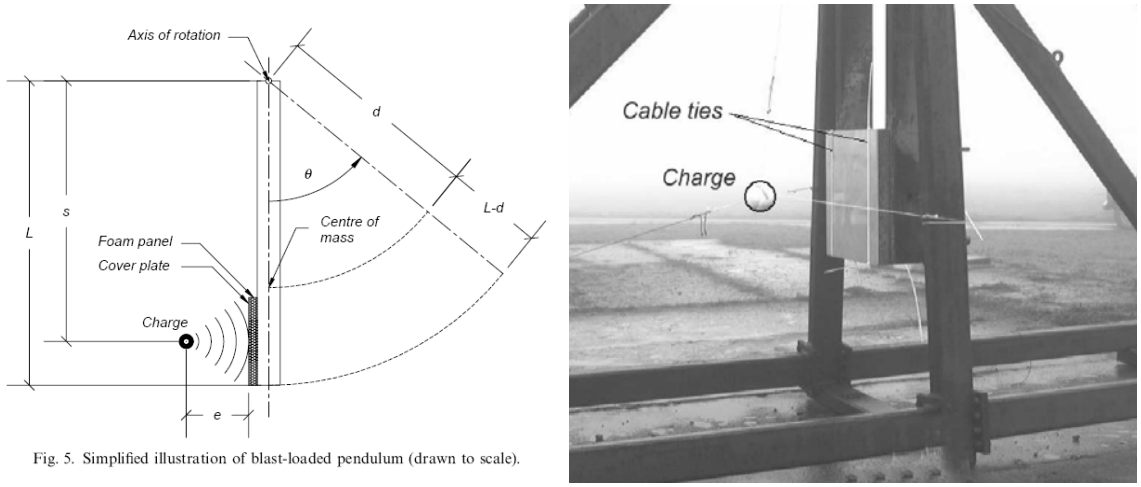

Figure 6.1 NTNU experiment

The results of the experiment came as a surprise to NTNU. Instead of a reduction of impulse in the pendulum, more impulse was actually transferred with the damping material present. However, armed with the knowledge we have obtained from basic impact physics and the 1D-simulations, this does not sound unnatural, given that momentum is a conserved quantity. If the aluminium foam leads to increased reflection of the incoming wave, more impulse must be transferred to the pendulum to compensate.

On comparison with our simple numerical 1D-setups, we note that the NTNU experiment bears some similarities to Setup 4.1a, where the damping material is near the steel bar (Chapter 4.3.1).

However, there are also several complications compared with that setup:

- 3D effects
    - The shock wave will be geometrically damped as it travels outward
    - The shock wave will not hit the pendulum uniformly
- The pendulum is “thin” (i.e. more similar to the projectile case,as in Chapter 4.7.2.)

Let us now perform numerical simulations of the experiment using our three test materials to see if we can understand in detail how these complications affect the results (if at all).

6.2 Numerical simulations

The numerical simulations were performed using AUTODYN. The pendulum was modelled in Lagrange and the rest in Euler. A graded Euler grid was used with square cells of side length 2.50 mm in the relevant part. The Lagrange cells were chosen to have slightly coarser mesh than the Euler cells to avoid problems with interaction. For validation some simulations were also performed in which the damping material was modelled in Lagrange, and this did not make any difference to the results.

The pendulum material was modelled using 4340 steel from the AUTODYN material library and for the explosive we used C4 from the same library (PE4 and C4 are almost the same explosive).

The pendulum is a complicated structure, so instead of modelling everything in 3D, we approximated the pendulum in 2D as a huge cylindrical disk (diameter 700 mm) behind the damping material with the same mass as the actual pendulum. Since the plate covers the material, this should not make any difference. Instead of looking at the swing of the pendulum to estimate the transferred impulse, we obtained the momentum of the “pendulum disc” directly from AUTODYN. (This method also removes any ambiguity in the impulse measurement. NTNU obtained different answers depending on how they calculated the impulse from the swing of the pendulum.)

In the initial stage, before the shock wave reaches the aluminium panel, the situation is spherically symmetric. This was exploited by running the initial stage in 1D (with spherical symmetry) and then remapping to a 2D grid once the shock wave reached the damping material. The state after remapping is shown in Figure 6.2.

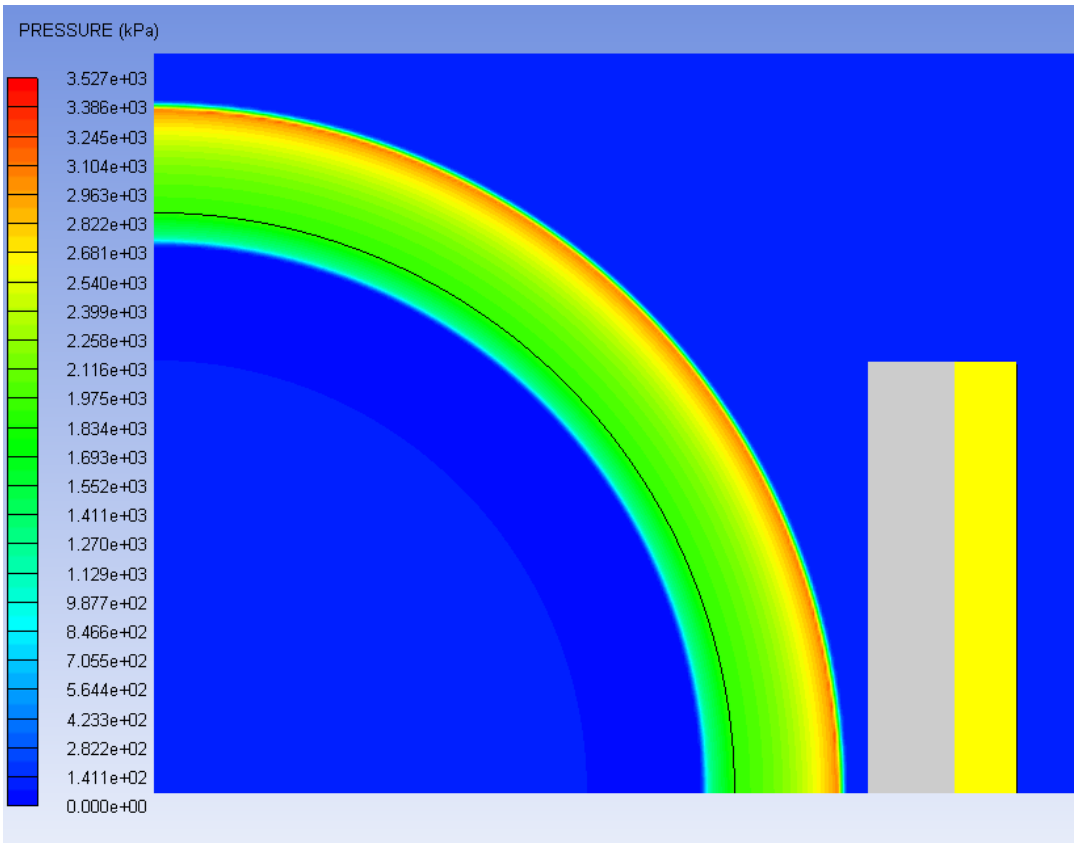

Figure 6.2 Initial state of our 2D simulation of the NTNU-experiment.

The NTNU experiments were performed using aluminium foam as damping material, for which we do not have an exact material model. In any case, our objective is not to replicate the NTNU results, but to obtain further insight into the physics of shock attenuation. We will therefore perform these simulations using our three test materials to see how they behave, compare with the 1D results and try to draw some conclusions from what we observe. The porous sand is probably most similar to the aluminium foam actually used in the experiment.

In Figure 6.3 we have plotted the impulse transferred to the pendulum as a function of time for the three different damping materials in the 2.5 kg explosive case.

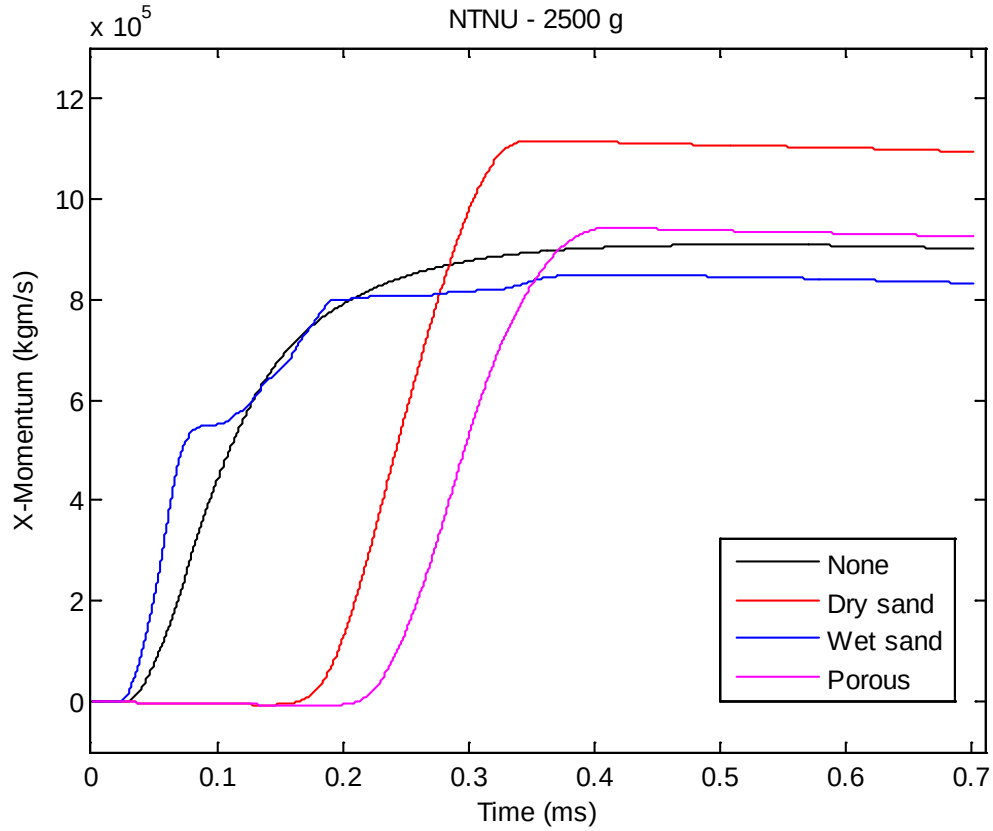

Figure 6.3 Momentum transferred to the pendulum as a function of time, for the different damping materials.

In Table 6.1 we sum up the results for both charges, compared with the case of no damping material.

	Dry sand	Wet sand	Porous	NTNU Average experimental result (aluminium foam)
1 kg	+26.4%	-2.9%	+8.8%	+1.3%
2.5 kg	+21.6%	-7.6%	+2.8%	+14.8%

Table 6.1 Results for impulse from the simulation of the NTNU experiments compared with the case of no damping material.

The overall results are in reasonable agreement with the NTNU experiments. Exact agreement was not to be expected since we did not use the same materials. The porous materials increase the impulse but the effect is not very big, just as we found in our 1D-confined simulations. Again the wet sand behaves slightly differently from the two porous materials.

It seems that despite the added complications, the 3D effects do not introduce anything new. The pendulum is "thin", but the total mass is so large that it moves very little, so the situation is actually quite similar to the "long bar" case studied in Chapter 4. Also, the radius of the

pendulum disc and damping material are so large, that the problem can be considered almost confined during the loading duration. Thus the NTNU results are not really surprising at all. In fact, they are just what one should expect from such an experiment.

6.3 Further analysis

Let us proceed a little further. In Chapter 4 we saw that although the damping materials could not attenuate the impulse, they could attenuate the maximum amplitude of the shock wave. The stress inside the pendulum was not measured in the experiments, but in the simulations this can easily be done by putting numerical gauges inside the "pendulum".

In addition, we will also extend the "pendulum" so that reflections from the rear end will not reflect immediately. This will make it simpler to see how the initial wave is influenced by the damping material as the results are not obscured by reflected waves. By doing this we are obviously creating a heavier pendulum, which, as we saw in Chapter 4, should increase the impulse slightly. This should not make any fundamental difference to the results, but will hopefully make the physics slightly clearer.

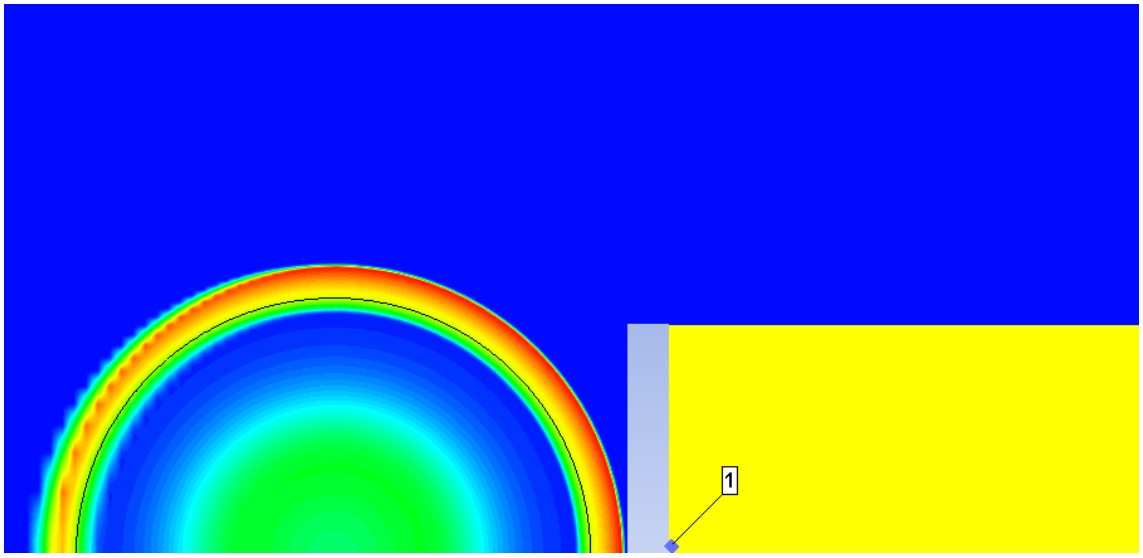

Figure 6.4 Modified NTNU-setup with thicker pendulum disc (2D – axial symmetry)

The stress is measured in gauge #1, near the damping material. In Figure 6.5 we have compared the stress for dry sand and no material for the two charges (1.0 and 2.5 kg) used in the experiments. We see exactly the same phenomenon as in the 1D-simulations. A short pulse with a high amplitude is replaced by a long pulse with lower amplitude.

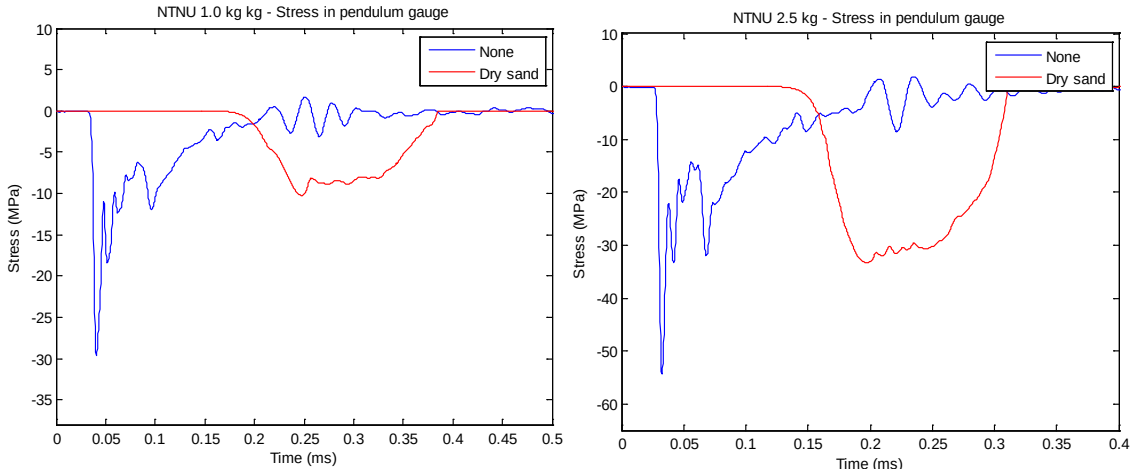

Figure 6.5 Stress in pendulum in modified NTNU-experiments

Another interesting observation is that for the 2.5 kg case there is less difference between the dry sand amplitude and no material amplitude. We saw in the 1D case that if the charge is sufficiently large (or the damping material thickness sufficiently small), the damping material may increase the shock amplitude. Could the same phenomenon be at work here since the amplitude is much larger for the 2.5 kg charge than for the 1.0 kg charge?

To investigate, let us increase the charge size even more. By trying 10 kg and 50 kg (not used in experiments, but in simulations we can easily use any charge size we want) charges, we obtained the results shown in Figure 6.6.

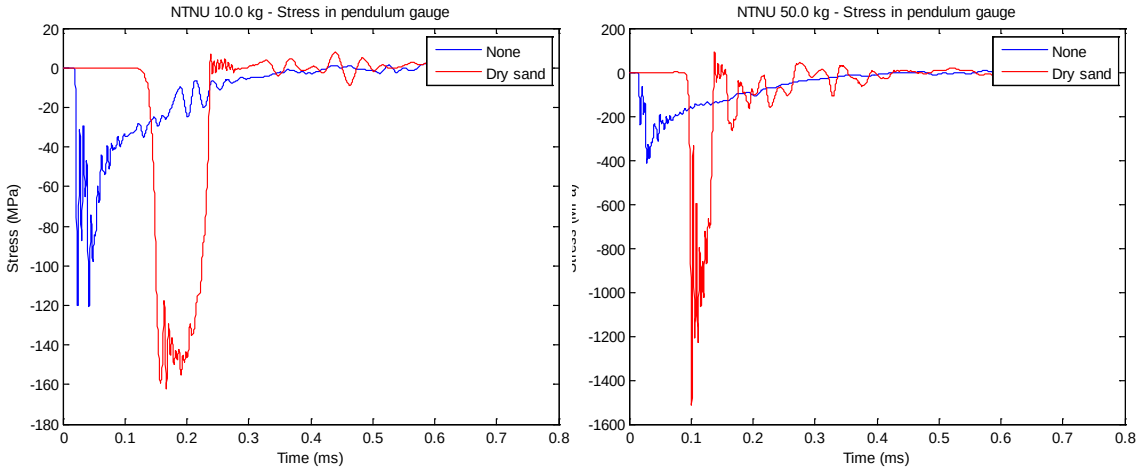

Figure 6.6 Stress in pendulum in modified NTNU-experiments for larger charges (10 kg , 50 kg).

Indeed, we see exactly the same phenomenon as in the 1D-simulations. Now the charge has grown so large that we have reached the regime where a thickness of 10 mm damping material does not give attenuation of the amplitude at all. If we increased the damping material thickness

for the large charges, we would again get reduced amplitude, though we would have to increase the thickness considerably for the 50 kg case.

By running the setup for wet and porous sand with the new charge sizes, we can compare the maximum stress ratios (maximum stress in pendulum protected by damping material divided by maximum stress without damping material). This is done in Figure 6.7.

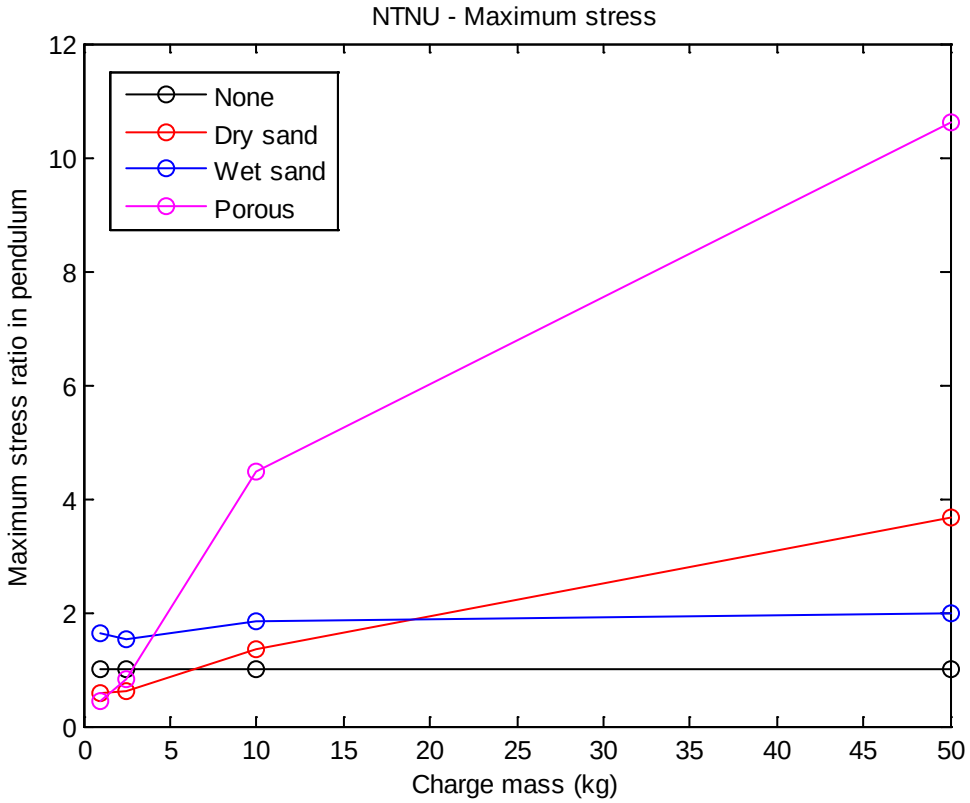

Figure 6.7 Stress amplitude ratios as a function of charge mass for different damping materials in the modified NTNU-experiment.

We see that for the actual charge masses used in the NTNU experiments, porous and dry sand reduces the amplitude, but for increased charge masses they increase the amplitude, in particular the porous sand. Wet sand would not have given damping in any case, and the amplitude seems largely independent of the charge size.

While we are at changing the charge mass, it could also be interesting to see how this affects the impulse transferred to the pendulum. The impulse ratio (i.e. for damping and without damping material) is shown in Figure 6.8 Again the results are similar to what we obtained in the 1D case, i.e. the damping material has little effect on momentum transfer. Interestingly, it seems that the charge masses actually used by NTNU should be expected to result in most impulse transfer to the pendulum. Although the difference is not that large as a function of charge mass, it's no wonder they were confused by the experimental results.

The simulation results obtained in this chapter strengthens our suspicion that, in general, impulse experiments on attenuation should be avoided because impulse is a conserved quantity and the experiments will say nothing about attenuation.

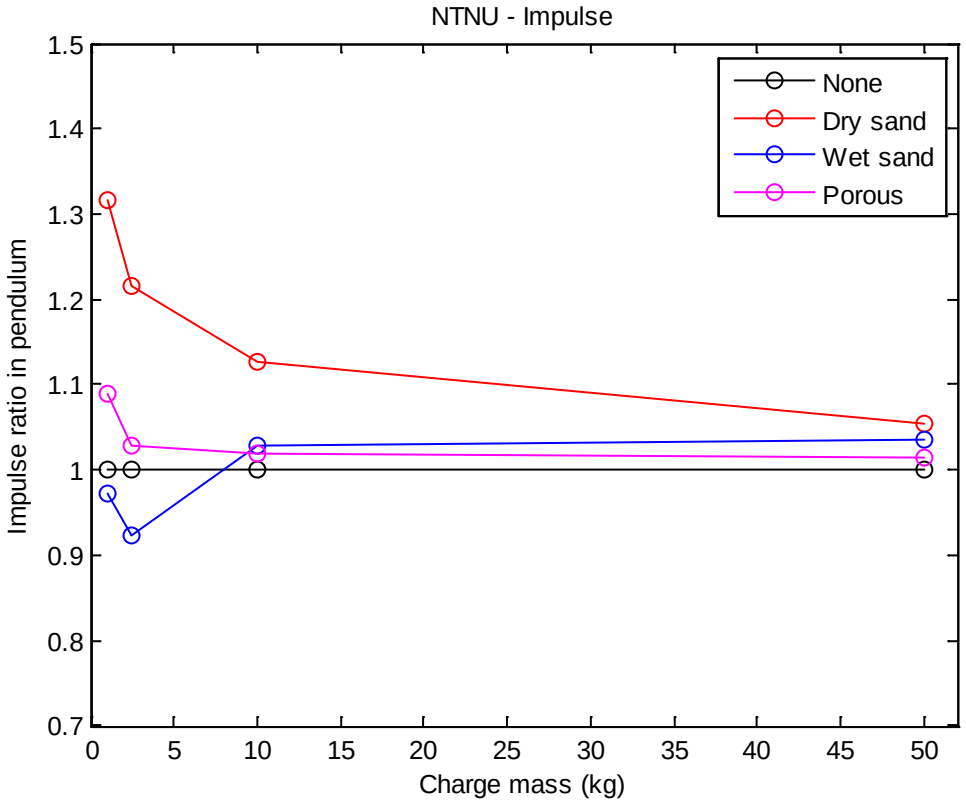

Figure 6.8 Stress amplitude ratios as a function of charge mass for different damping materials in the modified NTNU-experiment.

7 DSTL experiment

We now turn our attention to the DSTL experiments, a large range of free-field blast tests using spherical PE4 charges of different mass (20 g – 5 kg). The charges were generally surrounded by various damping materials of different thickness.

Figure 7.1 Illustration of DSTL-experiment: Charge in red, damping material in blue, gauges in green.

In one test series the resulting pressure wave was measured using pressure gauges at various distances, as illustrated in Figure 7.1. In a later series, a ballistic pendulum was used to measure the transferred impulse. Results were compared with tests without damping material.

7.1 Pressure measurements

Let us first examine the test series where pressure gauges were used to measure the pressure wave. In one case where the pressure history is available to us, a 40 g spherical PE4 charge was surrounded by 200 mm Perlite P25. The corresponding pressure results compared to a bare charge are shown in Figure 7.2. We see that the amplitude is mitigated considerably at the various pressure points at different distances from the charge.

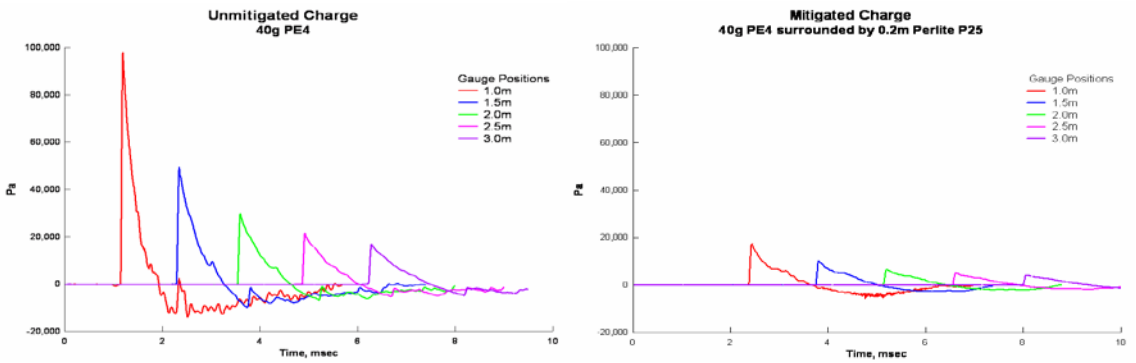

Figure 7.2 DSTL experimental pressure results for 40g PE4 charge surrounded by Perlite P25. Reproduced from (4).

We do not have an exact material model for Perlite P25, so instead we will model this test series using our three "test materials". Again this should give us some insight into what is happening physically and how the material properties affect the results.

The experimental setup is almost spherically symmetric, except for the damping materials having a cubical shape. For simplicity, we used spheres of the same diameter in the simulations, as the spherical symmetry enabled all AUTODYN simulations to be performed in 1D. An Euler grid with a very fine resolution of 1 mm was used. The explosive PE4 was modelled using the C4 material model from the AUTODYN material library.

The AUTODYN pressure results for the bare charge and for the three damping materials are shown in Figure 7.3.

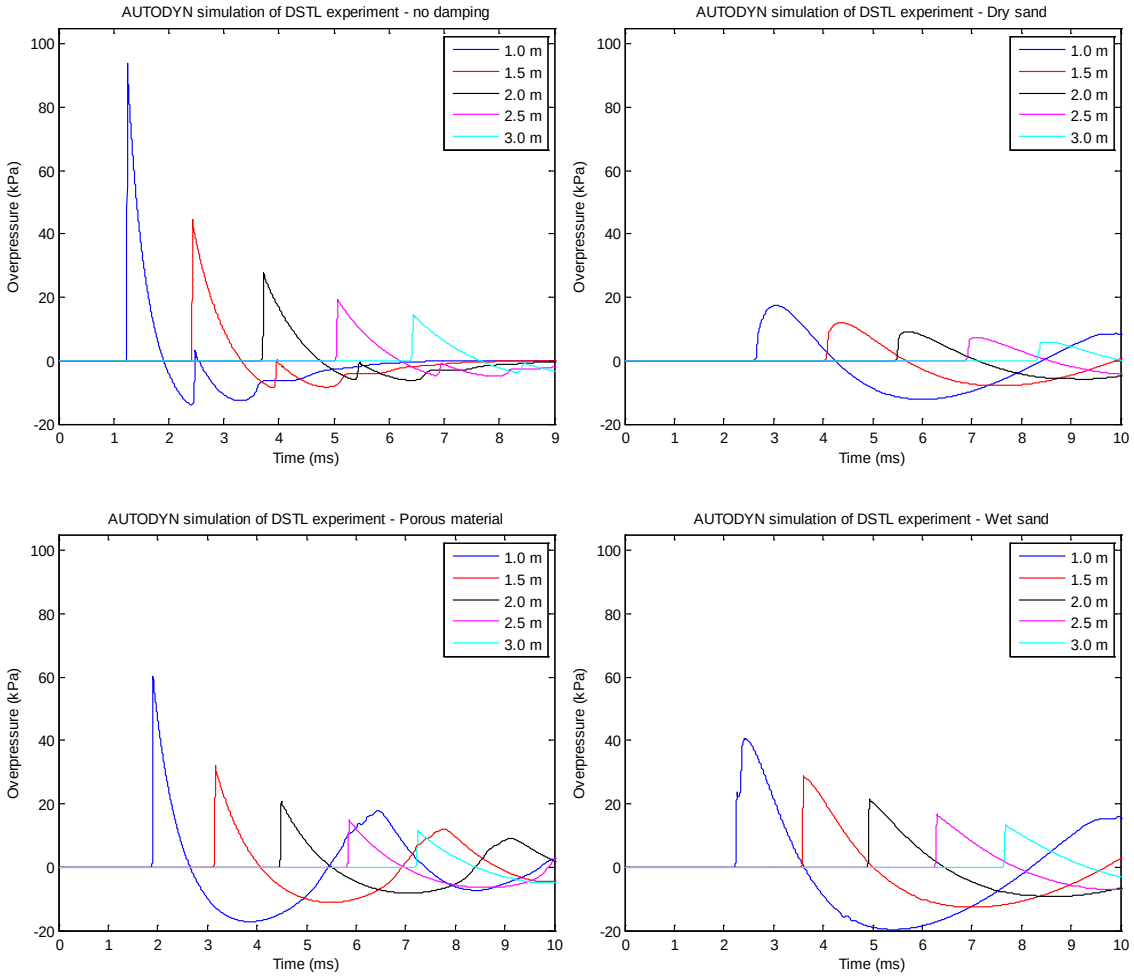

Figure 7.3 AUTODYN pressure results at different distances for 40g C4 charge, both bare and surrounded by different damping materials.

We see that there appears to be almost perfect agreement between experiment and simulation for the unmitigated charge. This should, of course, be expected, but it is still nice to see.

The simulations for dry sand give very similar results to the experiment with perlite with regards to maximum amplitude. However, the blast wave seems to more "rounded" than in the experiments. The porous sand does not give much damping and is the "worst" damping material, but the shock wave is "sharper" than for dry sand. Wet sand falls somewhere in between dry sand and and porous.

7.2 Analysis

Is it possible for us to explain these results physically based on the understanding we have developed so far? For example, why is dry sand so good and porous relatively bad in this case?

A good start is to review exactly what happens after detonation of the charge. First the blast wave propagates through the explosive and to the boundary with the attenuation material. At this interface, there will be both transmission and reflection depending on the relative properties of the

explosive and the damping material. The transmitted wave then continues through the damping material until the interface with air, where again some is transmitted and some reflected, depending on the relative properties of the two materials. So, there are two effects involved in determining how the blast wave reaching the air turns out.

Let us first consider the effect at the interface between attenuation material and air. For an elastic wave arriving here (which the current blast wave is not, but assuming elasticity probably gives a rough idea of what can be expected), the transmission and reflection would be totally dependent on the relative impedance of air and the damping material.

Impedance is density multiplied with sound velocity. This means that high density materials with high sound velocity (like steel etc.) have high impedance, while low density materials with low density (like air etc.) have low impedance. Materials with similar impedance at an interface will have more transmission and less reflection (so two identical materials will have no reflection). Comparing our three materials, we find that porous sand has impedance closest to air, followed by dry sand and finally wet sand. Thus, we would expect the porous material to give the highest transmission, followed by dry sand and wet sand. Except for wet and dry sand changing places, this is exactly what we see in Figure 7.2.

But, why would wet sand give a higher amplitude than dry sand? This must be due to the other factor, the ongoings at the interface between explosive and damping material. Since wet sand has the highest density of our three attenuation materials, it is harder to accelerate which leads to greater confinement of the detonation products and consequently more build up of pressure. Thus, the transmitted wave at the first interface has a higher amplitude for wet sand than for dry sand. In total the effects at the two interfaces add up to exactly what we see.

7.3 Pendulum experiments

DSTL also used a pendulum to measure the transferred momentum for different damping materials. Their setup was more or less exactly the same as in the NTNU experiments, except that the damping material surrounded the charge instead of being located near the pendulum. Further, the charges were placed at various distances from the pendulum. Their results for a variety of materials are reproduced in Figure 7.4.

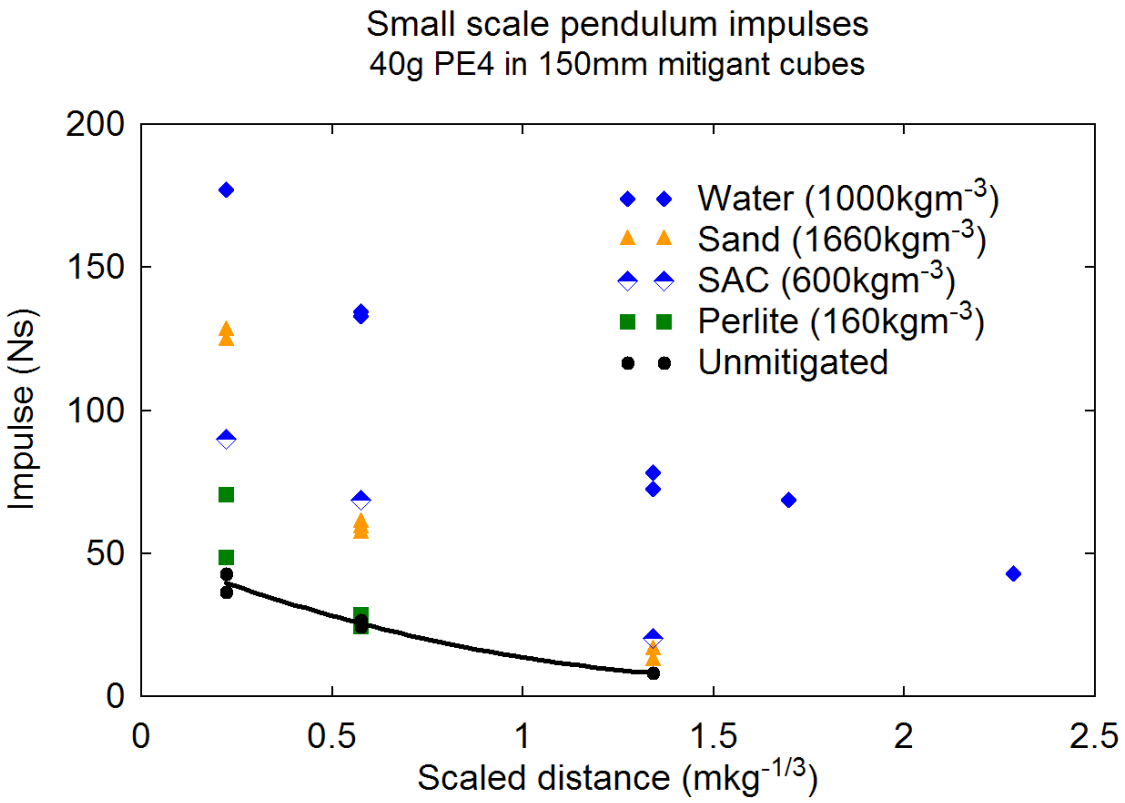

Figure 7.4 Results from DSTL pendulum experiments (Reproduced from (4)). Distance is scaled by the cube root of the charge mass.

As usual, let us perform numerical simulations with our three "test" materials to see if we obtain roughly the same behaviour. Similarly to the NTNU simulations, we simplified the pendulum to become an accelerated wall to avoid having to model in 3D. However, in the NTNU simulations we knew the exact dimensions of the aluminium foam plate in front of the pendulum. Here, there is no plate in front of the pendulum, so the dimensions of the pendulum, which are not known to us, become more important.

We solved this by calibrating the simulations without damping material to the experiments. A "pendulum plate" of radius 25 mm was found to reproduce the experimental result. Apart from this change, exactly the same grid as in the NTNU experiments was used. Simulations were first performed in 1D with spherical symmetry and remapped to 2D when the blast wave reached the pendulum. As in the previous section, we used spheres with radius of 75 mm for the damping material instead of cubes.

Our results are shown in Table 7.1 and Figure 7.5. We see that the numerical simulations closely resemble the experimental results, even though the materials are not the same. The wet sand behaves similarly to water, except for the impulse falling off slightly faster, whereas the porous material has a similar behaviour to Perlite (which it might resemble). Clearly at short distance, the materials all give enhancement of the momentum, as is expected.

Distance	None	Dry	Wet	Porous
86 mm	54.0	129.1	228.2	69.5
205 mm	44.0	74.2	140.0	43.9
462 mm	26.5	14.3	56.0	21.9
581 mm	20.9	10.7	24.9	14.5
786 mm	14.2	6.7	13.5	10.1

Table 7.1 Momentum transferred to "pendulum" in simulations with different damping materials surrounding the charge. Pendulum at different distances. Momentum given in Ns.

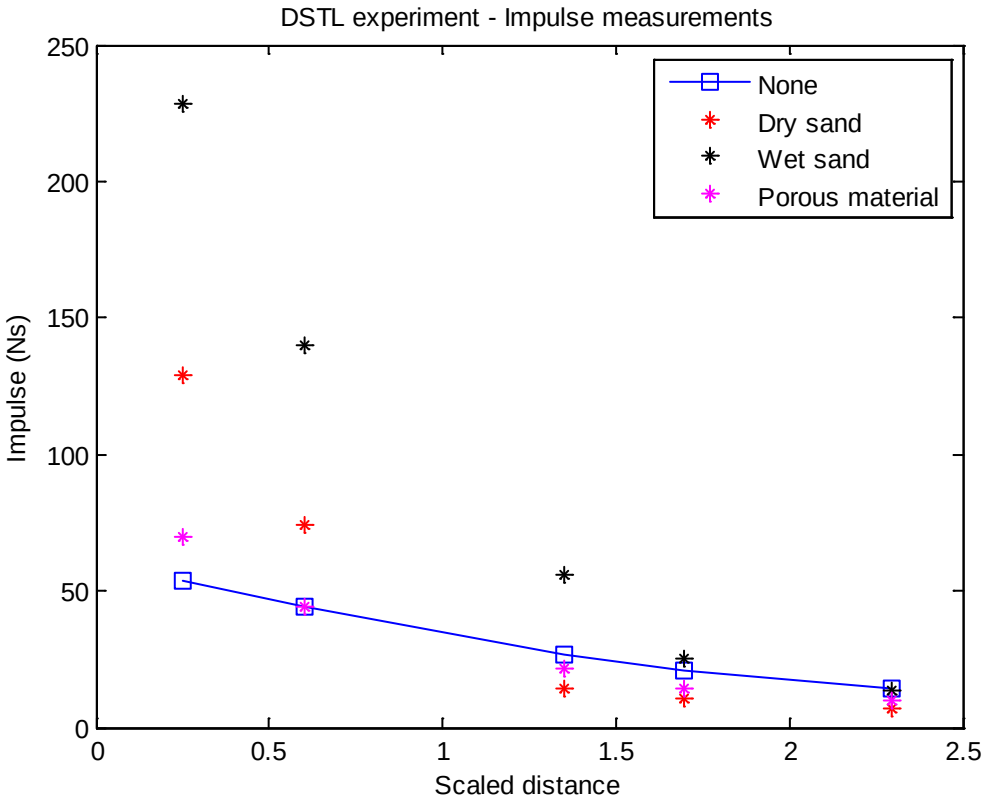

Figure 7.5 Results from simulations of DSTL pendulum experiments. Distance is scaled by the cube root of the charge mass.

These results can also be explained theoretically. This time only the effect at the interface between explosive and damping material is important. Heavier materials give greater confinement which means more impulse is reflected. However, due to momentum conservation, this also means that more impulse must be transmitted into the damping material. Thus, a heavy damping material will acquire more momentum (basically same effect as discussed in Chapter 3.1, where a heavy wall impacted by a wall will obtain a considerable momentum despite having no kinetic energy). And, this is exactly what we see in Figure 7.5.

The experimental results have exactly the same tendency, except that water increases the impulse more than dry sand, even though dry sand is heavier. How can we explain this? It is probably due to the material properties of water, which is incompressible, whereas dry sand can be compacted. This increases the confinement, even though water is not as heavy as sand. A numerical simulation using water as damping material was run and confirmed this experimental result on comparison with dry sand.

At larger distances, we note that the different materials transfer about the same amount of impulse. This is because the pendulum is further away and less material actually impacts the pendulum. Thus, there is, in the long range case, a kind of "scatter effect" which complicates the situation, depending on how much the various damping materials are scattered away from the pendulum.

8 NAVAIR

We now turn our attention to the setup used by NAVAIR to study the shock attenuation properties of porous materials.

8.1 Experimental setup

In the NAVAIR experiments, a Pentolite charge with a mass of 175 g was placed directly on top of a damping material and detonated. Epoxy was used to glue the pieces of damping material together. Several attenuation materials were tested, each having a thickness of 1, 2 or 3 inches (25.4 mm, 50.8 mm or 76.2 mm). The pressure was measured using PVF gauges on a PMMA block placed directly under the attenuation material. A photo of their setup is shown in Figure 8.1.

On comparing with our 1D uniaxial setup, we immediately notice the following:

- The damping material layer is quite thin, especially compared with the charge size. In the uniaxial setup, for dry sand and porous material this geometry would clearly give shock enhancement of the amplitude.
- However, the charge is not confined so some of the energy will be directed away from the damping material. This will probably tend to decrease the "effective" charge size as compared with our 1D uniaxial situation.
- The damping material is not confined either. This should decrease the amplitude of the shock wave inside the damping material as release waves can propagate from the edge.

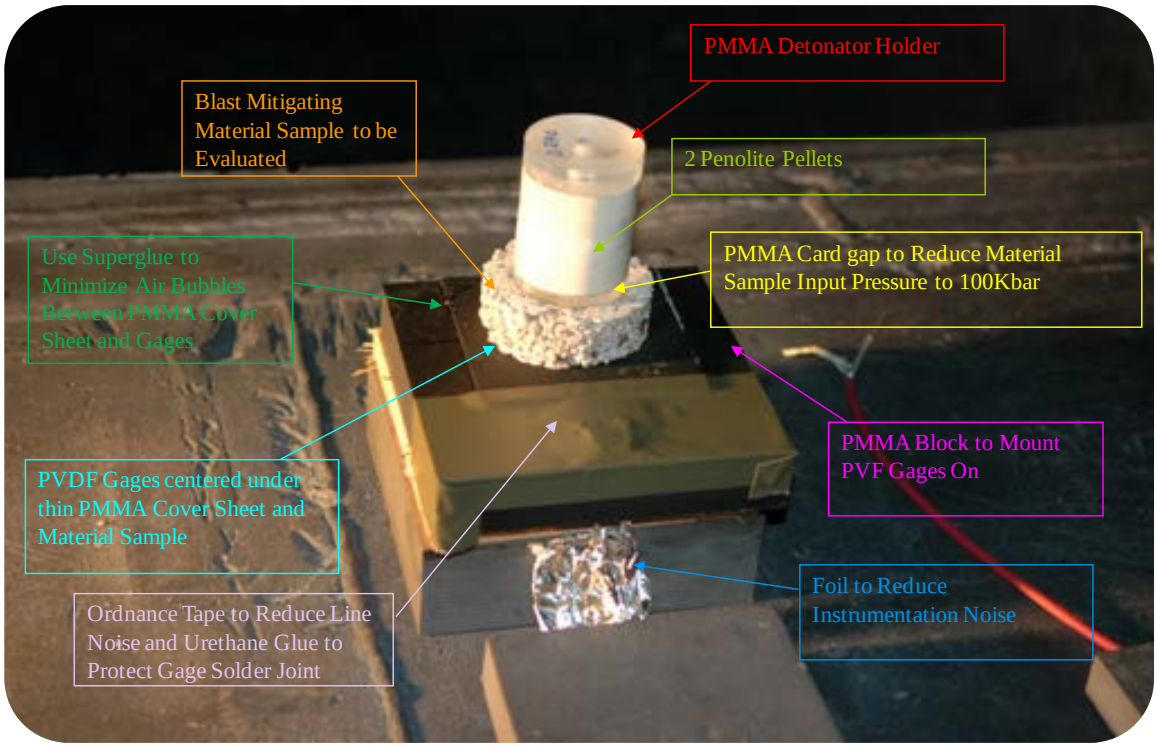

Figure 8.1 NAVAIR setup (Reproduced from (5)).

It is important to note that the explosive was placed directly on the mitigation material (except for a 7 mm thin PMMA disc). This means that the distance from the explosive to the pressure gauge varied with the thickness of the damping material.

NAVAIR presented the results as the pressure measured by the PVDF gauges. They had some trouble with calibration of their sensors but it was thought that the trend for different thicknesses should be about right.

8.2 Numerical simulations

As usual AUTODYN simulations of the experiments were run for our three “test materials”. Cylindrical symmetry in the configuration allowed us to model everything in 2D. Each setup was modelled using the Euler processor with a uniform grid of 0.25 mm cells. The AUTODYN setup for 3 inches of dry sand as damping material is shown in Figure 8.2. A “flow out” boundary condition was used below the PMMA.

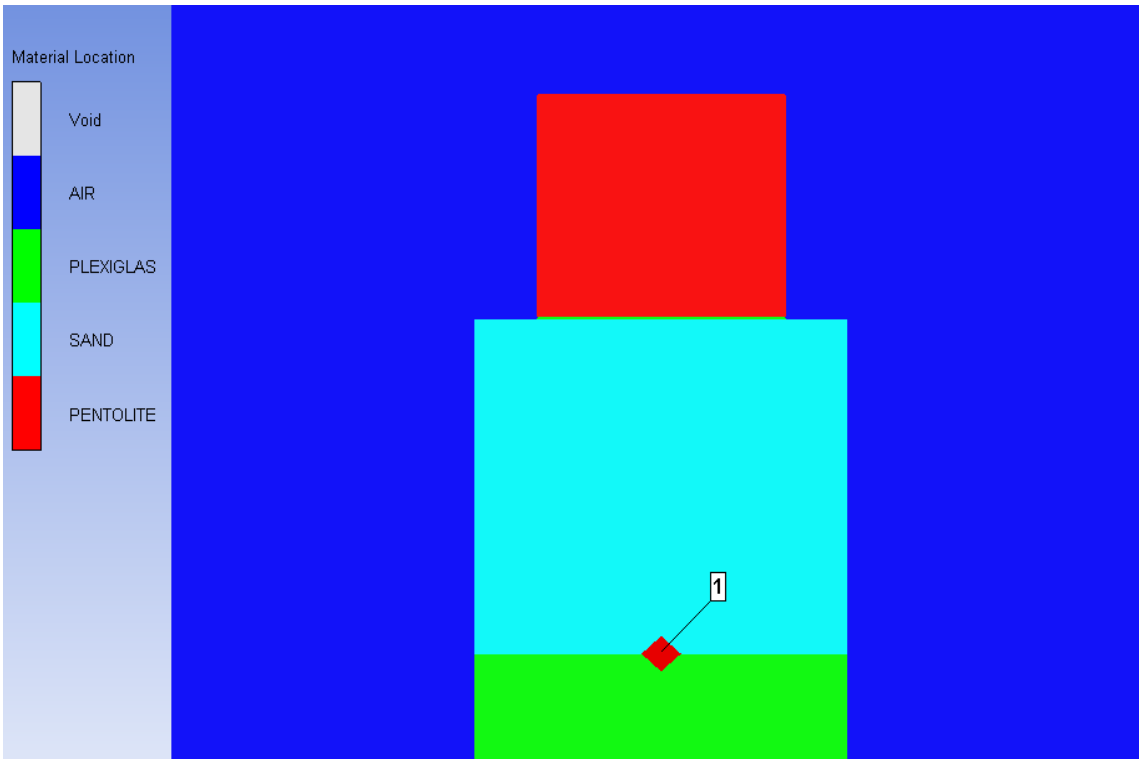

Figure 8.2 AUTODYN setup for 3 inches dry sand in the NAVAIR experiment. Location of the gauge point is shown.

Simulations were run for 1, 2 and 3 inches thickness of the damping material (as in the experiment), as well as 10 mm to investigate very thin layers of attenuation material.

The results for maximum pressure are shown in Table 8.1 and in Figure 8.3.

Thickness	**Dry**	**Wet**	**Porous**	**None**
10 mm	9.00	10.51	9.96	4.50
1 inch (25.4 mm)	3.06	7.87	8.69	3.32
2 inches	0.42	5.59	4.65	1.99
3 inches	0.09	4.13	1.65	1.24

Table 8.1 Simulation results for maximum pressure in GPa.

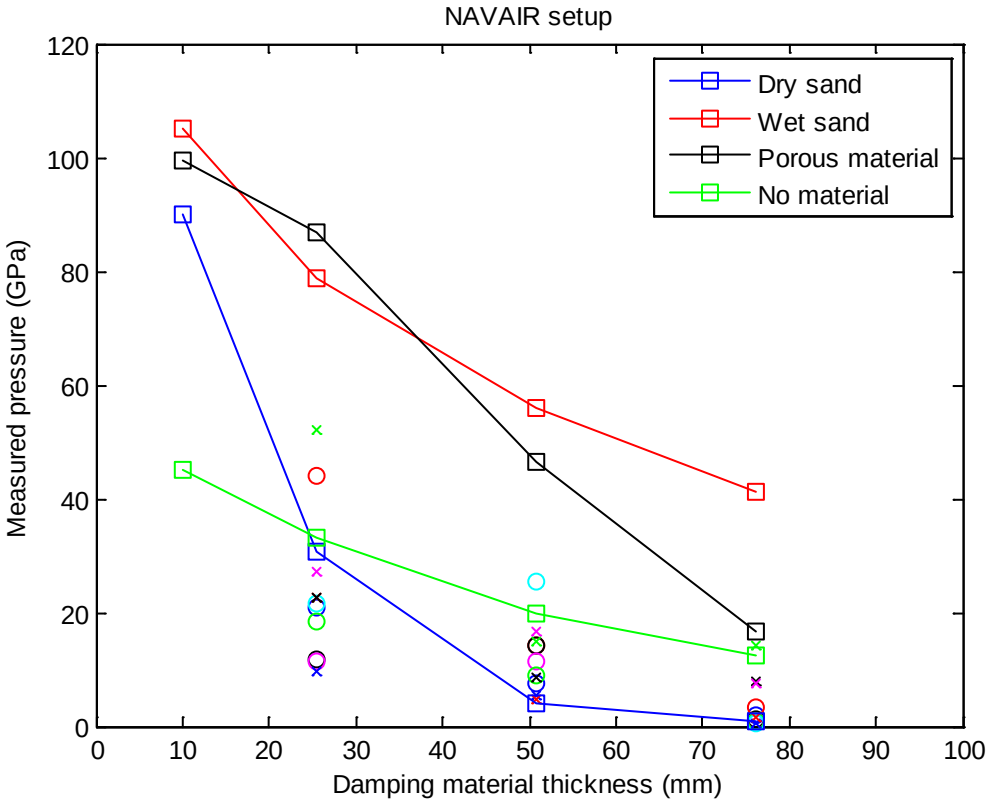

Figure 8.3 Pressure as a function of material thickness in the NAVAIR setup for our three damping materials.

Another way to present the results is to normalize with respect to the case where no damping material is present. This is done in Table 8.2 and Figure 8.4, where we have also included normalized experimental results. Note that the experiments are not for the same materials and to avoid cluttering up the figure, we have not drawn lines between the experimental points.

	Dry	**Wet**	**Porous**
10 mm	2.00	2.33	2.21
1 inch	0.92	2.37	2.62
2 inch	0.21	2.81	2.33
3 inch	0.07	3.33	1.33

Table 8.2 Normalized simulation results

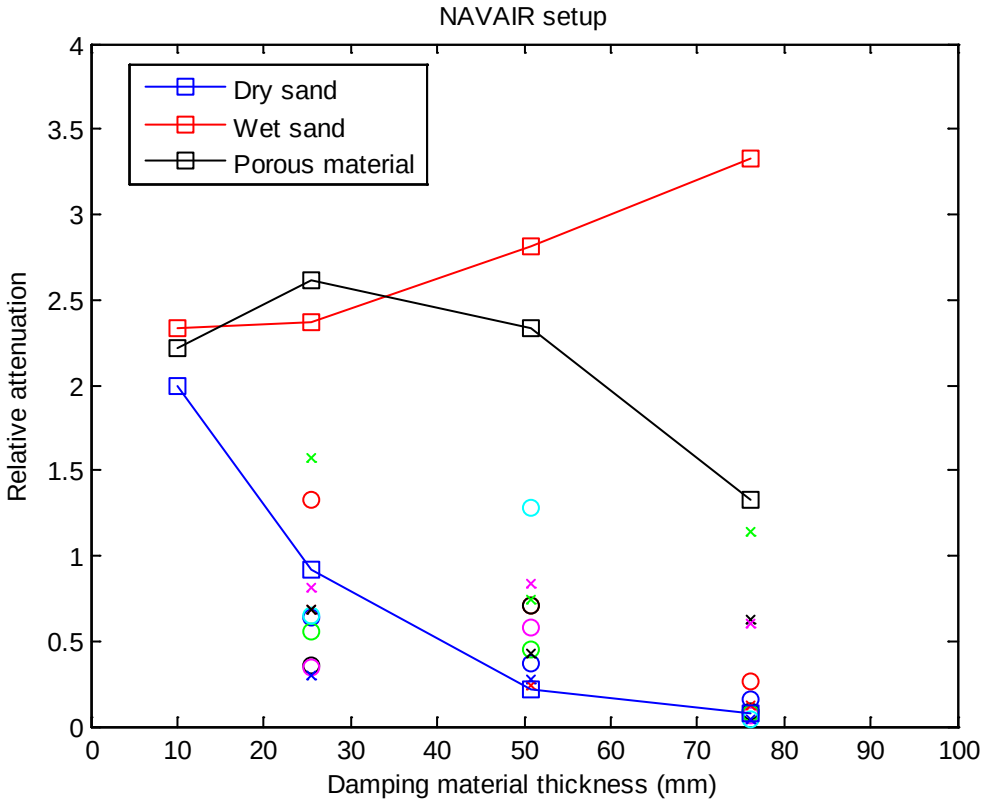

Figure 8.4 Simulation results and experimental results by NAVAIR

Again the simulation results are consistent with our earlier findings. We see that a too thin layer of porous material will give enhancement whereas it tends towards attenuation for thicker layers. In the NAVAIR setup, even though there is no confinement, still the porous material layer is so thin (or the charge so large) that it gives shock enhancement for all thicknesses. Dry sand shows the same tendency as porous material, pressure amplitude enhancement for thin layers and shock mitigation for larger layers, again consistent with our 1D uniaxial results.

On comparing with the experiments we see that dry sand behaves similarly to most of the materials. Porous sand behaves somewhat like Italian pumice (cyan circles) in having a maximum before falling off (although it has worse attenuation than Italian pumice). Again wet sand behaves totally different from the two other ones. NAVAIR did not test any similar materials to wet sand, so this is not inconsistent with their findings.

Finally, it is also interesting to see the relative damping as a function of mass of the attenuation material (instead of thickness). This is shown in Figure 8.5, which particularly illustrates how differently the two porous materials behave from the non-porous wet sand.

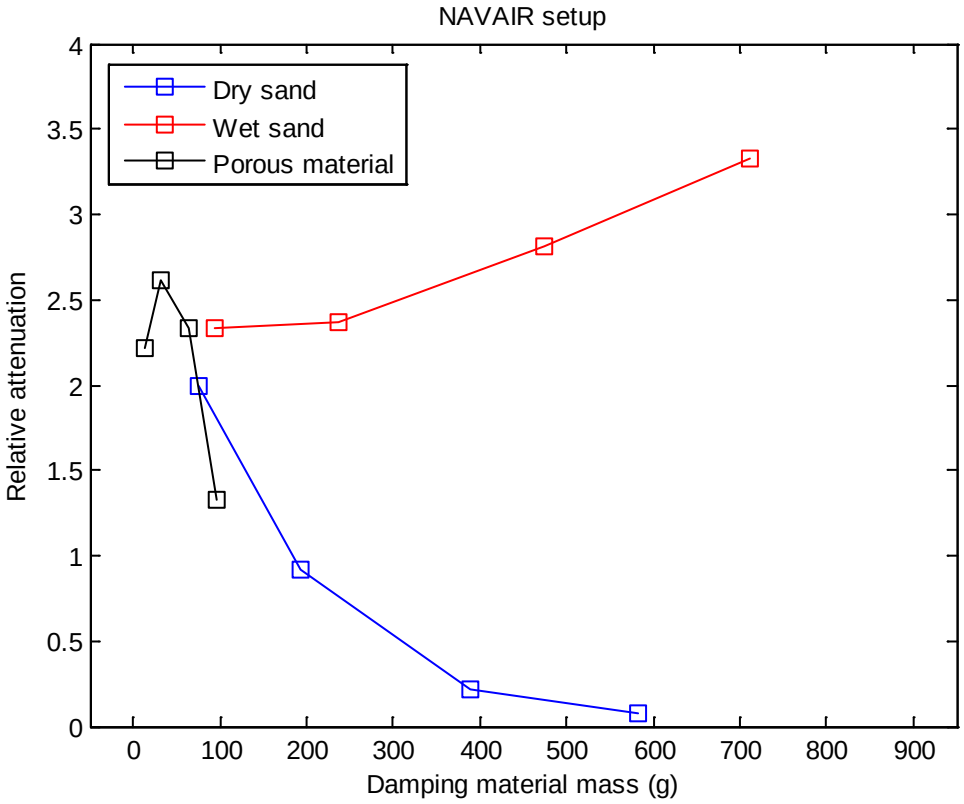

Figure 8.5 Attenuation in simulations of NAVAIR experiment as a function of damping material mass for all three materials.

9 FFI Hopkinson bar

We now turn our attention to the FFI Hopkinson bar experiments. In these tests a cylindrical TNT charge (radius 25 mm, 129 g) was detonated at 80-100 mm distance from a Hopkinson steel bar (radius 25 mm, length 3000 mm). Attenuation materials of different thicknesses were placed between the charge and the steel bar (either close to the charge or close to the bar). Strain gauges were placed at the Hopkinson bar to measure the strain (and calculate the stress) transferred from the explosive. In this way, different damping materials could be compared with each other and with the case of no damping material. The setup is illustrated in Figure 9.1.

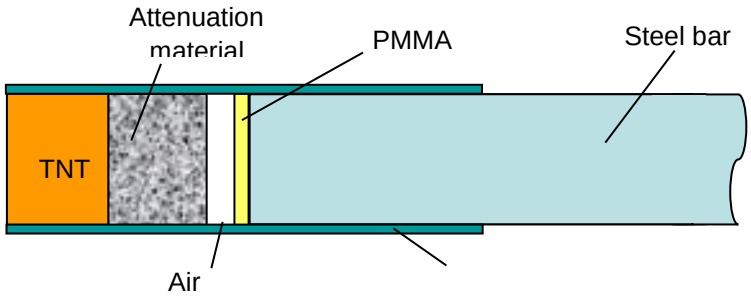

Figure 9.1 FFI set-up

A whole range of different materials were tested, including pumice, LECA (coarse and fine), aluminium foam, rubber granules, gravel, wood shavings, saw dust, Glasopor, Siporex and brick. At least two tests were performed for each material and the repeatability of the experiments seemed to be excellent.

The FFI tests are quite similar to the 1D-simulations in Chapter 4, though with some important differences. Most importantly the situation is not pure 1D since the movement is not contained in any way. Much of the blast wave is not directed towards the steel bar at all, while some of the damping material may scatter away from the bar and not interact with it.

There were two experimental series performed: one with the damping material close to the charge and another with the material close to the bar. There were some problems with the strain gauges in the first series, so the results are quite uncertain and will not be discussed further here. However, in the second series, optical strain gauges were used which gave very reliable and repeatable results. Also note that the different material samples had the same mass, implying that the thickness varied from material to material.

9.1 Numerical simulations

The simulations were performed in 2D using cylindrical symmetry. A graded Euler grid was used for the explosive and damping material, having a resolution of 2 mm x 2 mm in the important part. The Hopkinson Bar was modelled using a Lagrange grid of 8 mm x 8 mm and 4340 Steel from AUTODYN material library. To reduce computation time, we removed the Euler grid after the main shock wave had entered the bar. In Figure 9.2 we show a typical setup, where the explosive is detonated, a shock wave created which then enters the damping material and propagates towards the bar.

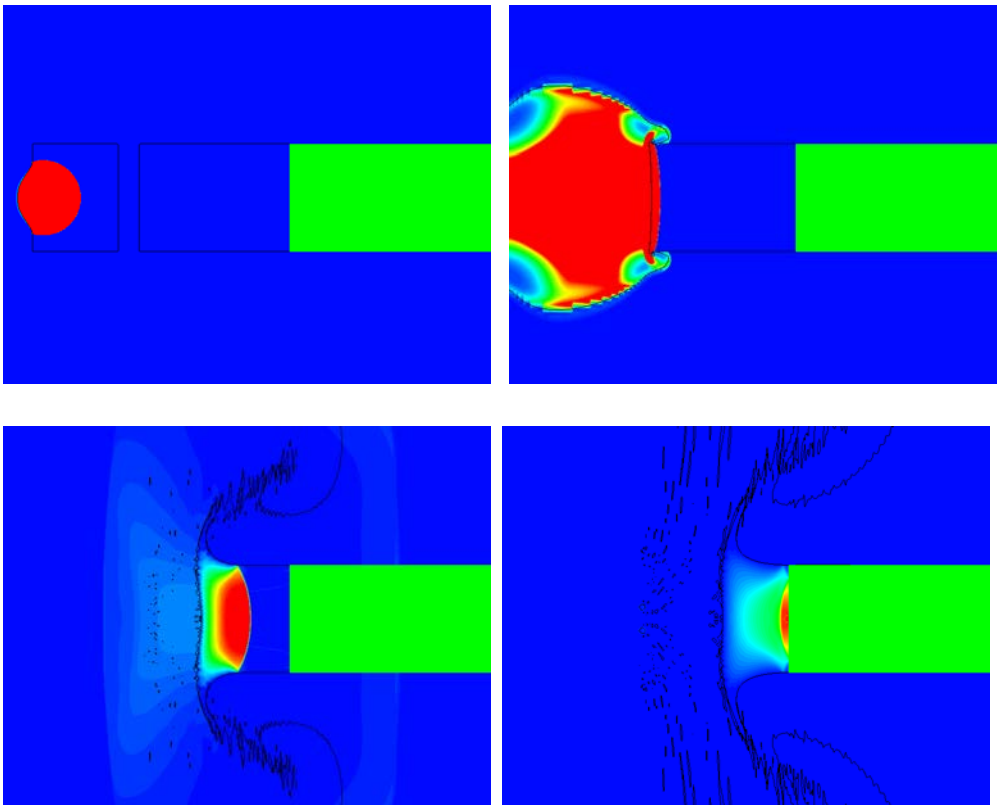

Figure 9.2 Contour plot from 70 mm dry sand near bar. Notice how some of the material is scattered and the shock wave attenuated as it propagates.

For validation, a simulation was performed without damping material and compared with the experiments. The agreement for the measured strain was excellent as is shown in Figure 9.3.

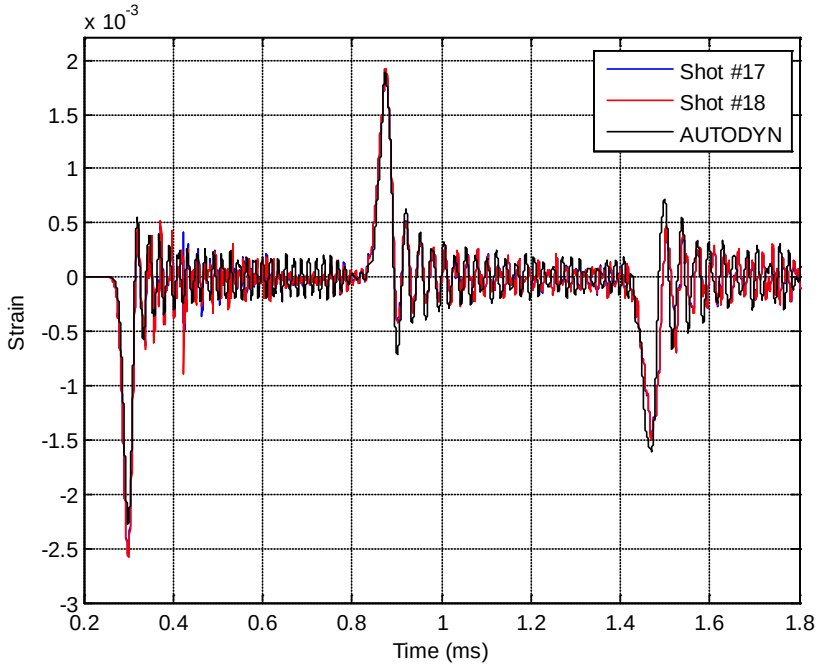

Figure 9.3 AUTODYN simulations of FFI Hopkingson bar setup compared with experiments without damping material.

9.2 Damping material near bar

Let us now do the usual routine of performing numerical simulations of the setup with all our three test materials, starting with the case where the damping material is close to the steel bar. The stress and strain was measured in the middle of the bar, just as in the tests. In Figure 9.4 we show the results for stress as a function of time for dry sand at various thicknesses (including zero thickness, i.e. no damping material).

We see that as more material is added, the stress becomes gradually smaller indicating an attenuation effect of the material. In Figure 9.5 we show the stress ratio (i.e. comparison with the case of no damping material) for all three materials as a function of thickness. Notice how wet sand suddenly seems to be by far the best damping material according to this test. Also, note that the amplitude is reduced in all cases, except for 10 mm porous material.

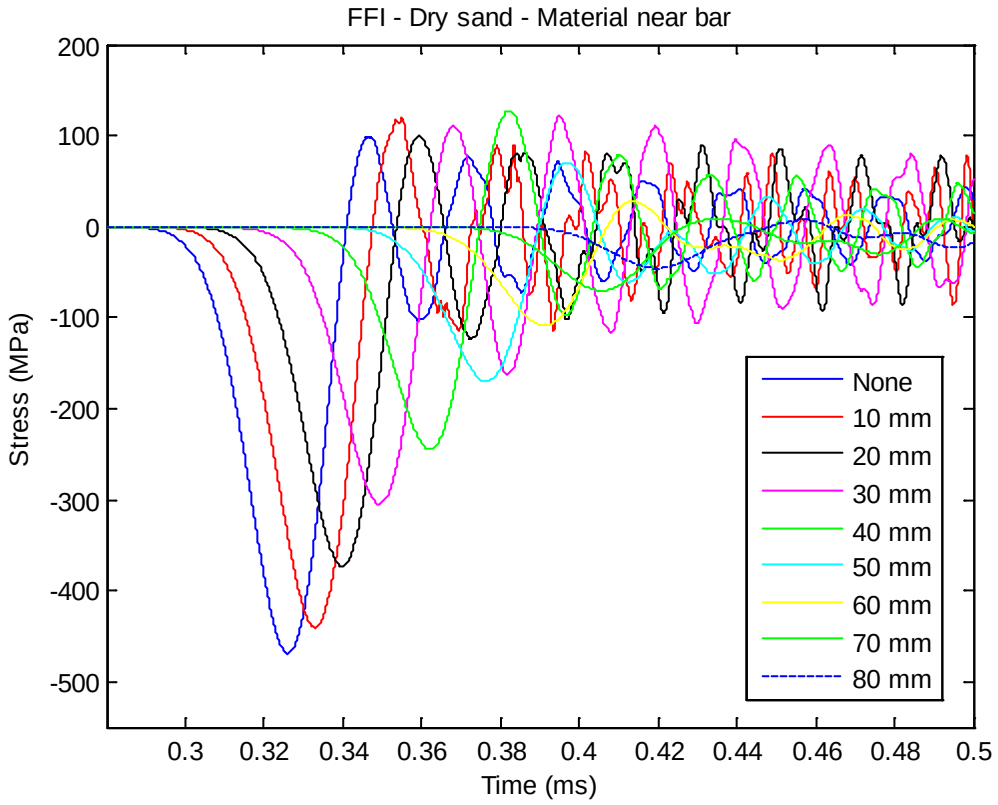

Figure 9.4 AUTODYN results for dry sand near bar in FFI setup.

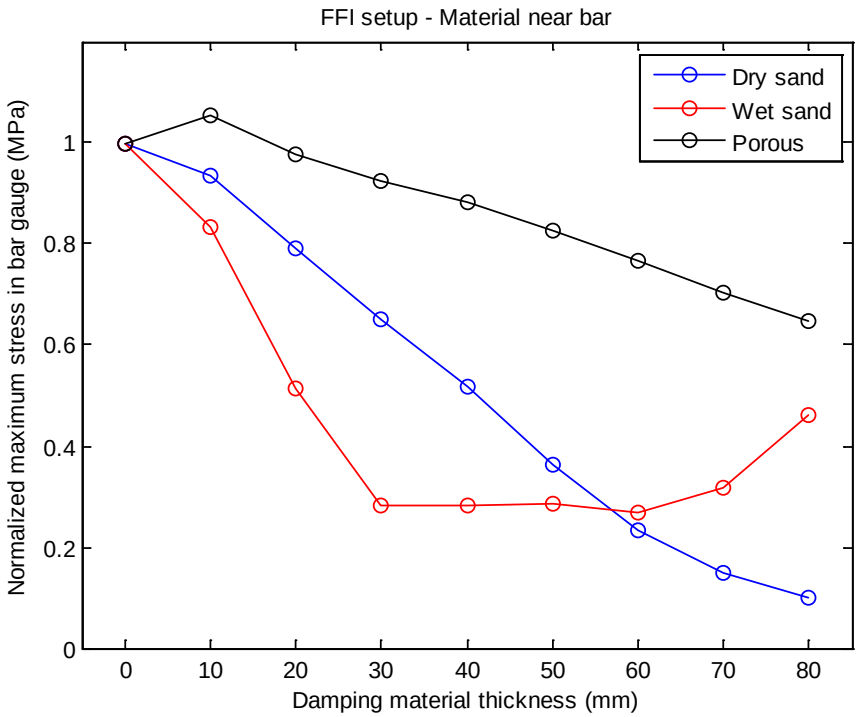

Figure 9.5 AUTODYN stress ratio for all three damping materials in FFI Hopkinson bar set-up.

If we ignore material type we can plot the actual experimental data together with the simulation results in Figure 9.6.

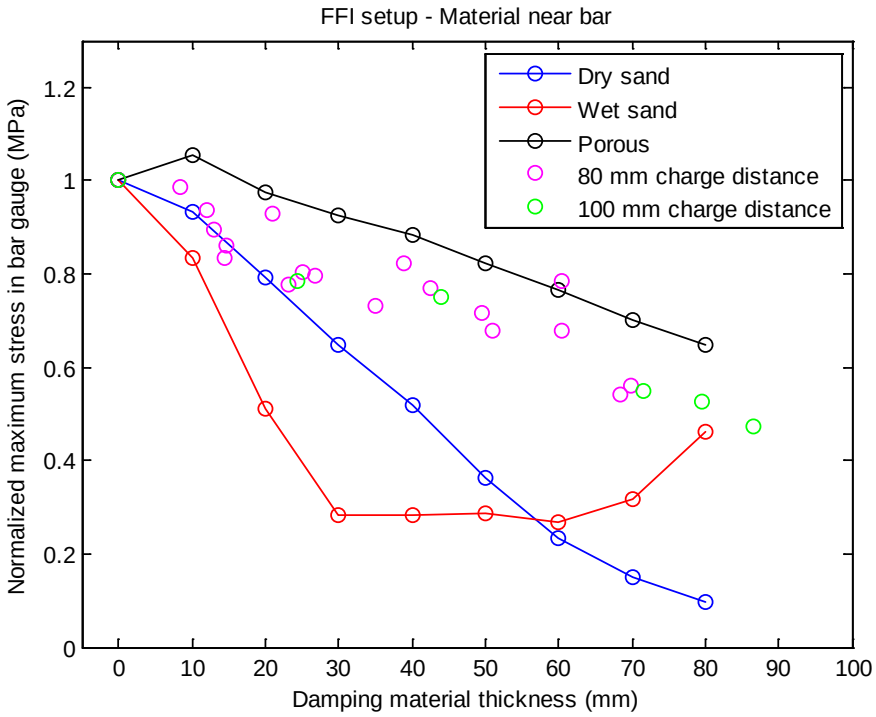

Figure 9.6 Numerical simulation of the FFI Hopkinson Bar setup together with experimental results (not for the same materials).

The damping materials used in the experiments were mostly porous so it may be quite natural that they mostly fit somewhere between simulations of the “extreme” porous sand and dry sand.

In Figure 9.7 we have plotted everything in terms of material mass instead of thickness.

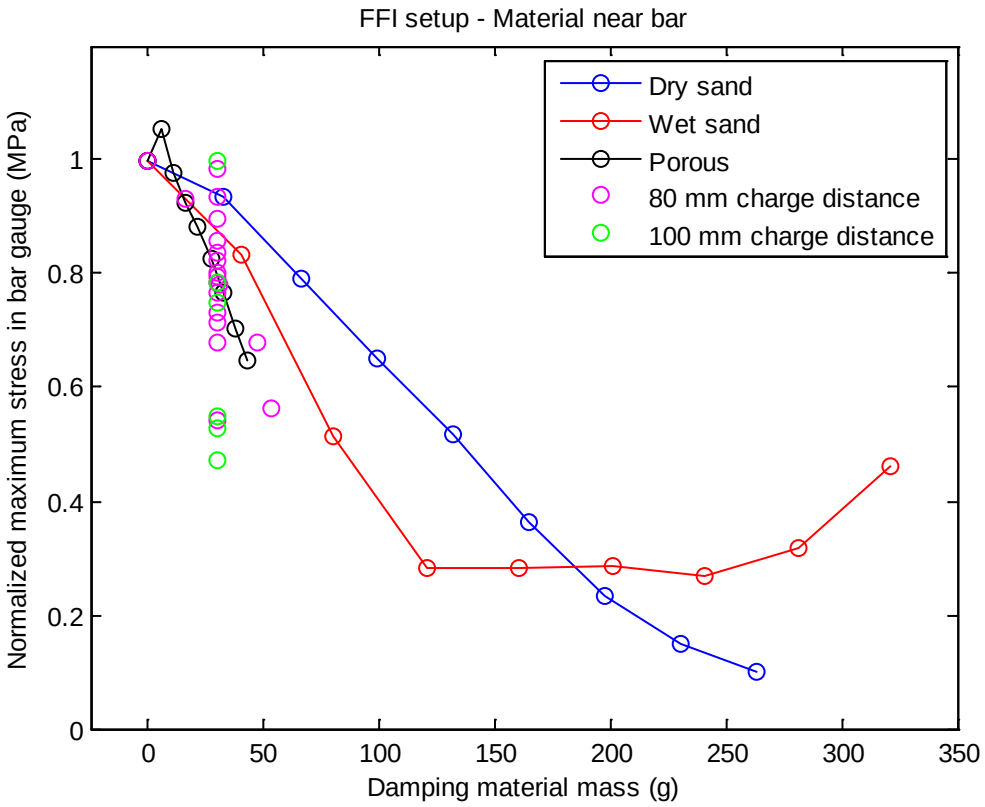

Figure 9.7 Numerical simulation of the FFI Hopkinson Bar setup together with experimental results (not for the same materials) as a function of damping material mass.

It is interesting to compare these results with the uniaxial 1D results of Chapter 4.3.1. Basically, the only difference is the lack of confinement. This means that some of the damping material will scatter and never interact with the steel bar and thus transfer no momentum. Observing the stress as a function of time, we see that the width of the pulse is more or less independent of the damping material thickness, only the amplitude changes. The momentum is the integral of the pulse, which means that less momentum is transferred when the amplitude is reduced. Since total momentum is necessarily conserved, this means that lots of momentum in the damping material does not interact with the bar. If we had made the steel bar much wider, there would not be any “missing” momentum. In fact, by doing so we would be turning the setup into something similar to the NTNU setup.

In the 1D uniaxial setup in Chapter 4 we saw that both enhancement and mitigation of the shock amplitude was possible, depending on the charge size and damping material layer thickness. Except for one case of thin porous material (10 mm) we always get damping in the FFI

Hopkinson bar setup. The reason that it is difficult to get shock enhancement in this setup is the "scattering effect" just discussed, which is due to the lack of confinement.

Looking back to the 1D results, we noticed that porous sand gave the highest amplitude for thin material layers. This seems to be quite consistent with the results (Figure 9.6) in the current setup, since porous sand always gives less damping then the other materials. Remember that the material layer here is always quite thin (though the necessary thickness depends on charge size, which is difficult to compare with the 1D-setup since the charge here is much larger but not confined). There is some enhancement of the amplitude for 10 mm, but for larger thicknesses the scattering effect makes sure that the total result is attenuation, although less than for the other materials.

Dry sand generally gave much more damping than porous sand in the 1D uniaxial case and we see that it performs better here as well. Although the scattering effect may also depend on material properties, it does not seem to change that dry sand is better than porous material. Interestingly, wet sand, for once is the best damping material, at least for thin material layers. This is also quite consistent with the 1D uniaxial results where wet sand performed well for thin layers compared to the other materials. Here the scattering effect means that we get damping all the time instead of the slight enhancement. Also note that wet sand does not "blow up" for the case when there is no free space between charge and steel bar. This is also due to lack of confinement and the scattering effect.

9.3 Damping material near charge

Let us now perform the same simulations for the situation when the damping material is close to the explosive, as was the case in the first FFI Hopkinson bar experiments. Results for dry sand as a function of time are shown in Figure 9.8. In Figures 9.9-9.10 we show the results for maximum amplitude for all materials. We also include the one available experimental data point.

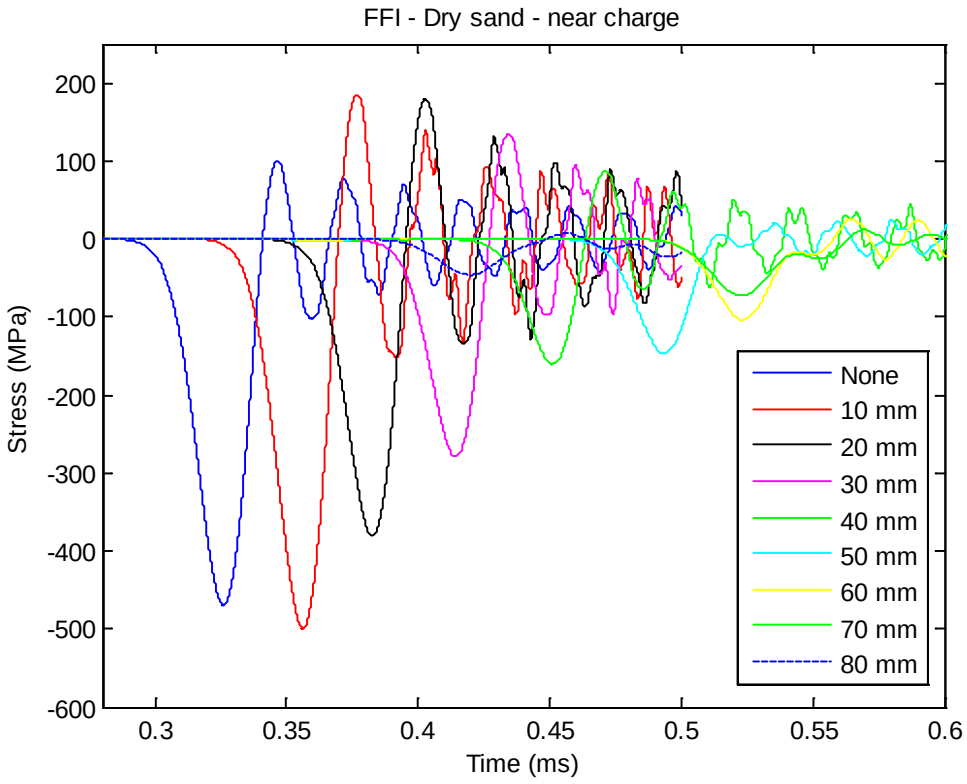

Figure 9.8 AUTODYN results for dry sand near charge in FFI setup.

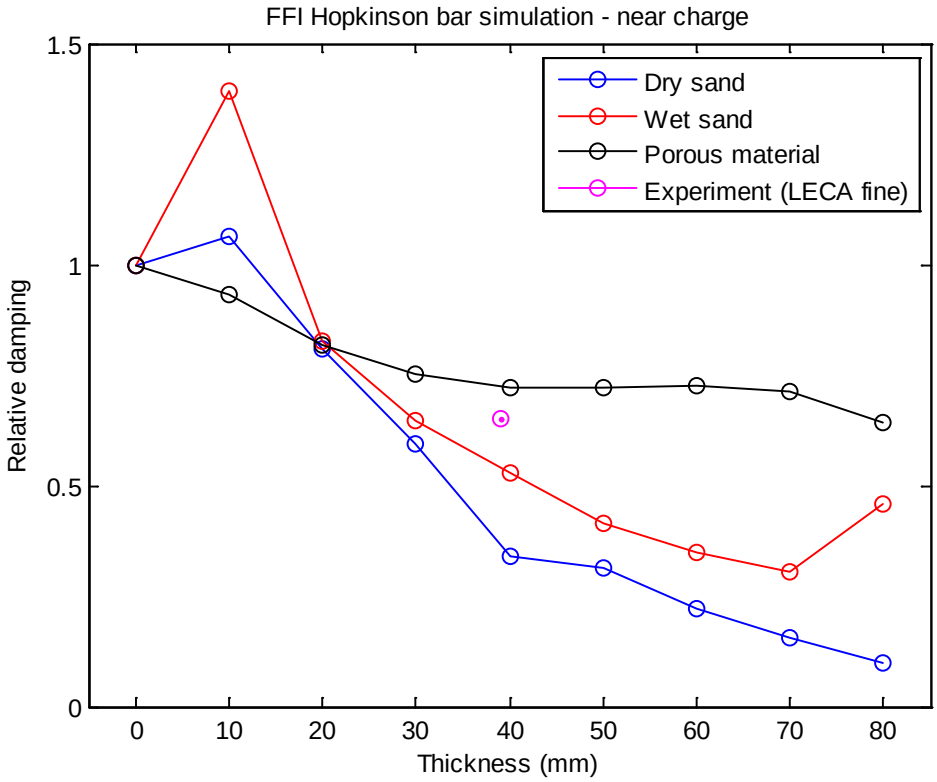

Figure 9.9 AUTODYN stress ratio for all three damping materials in FFI Hopkinsor bar set-up.

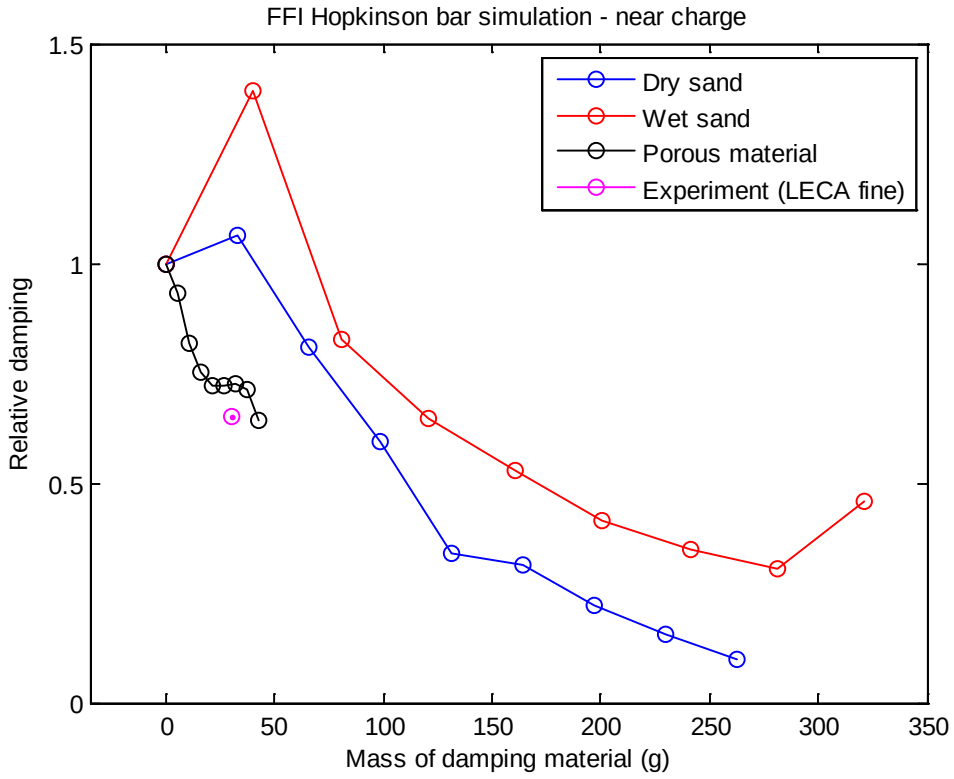

Figure 9.10 AUTODYN stress ratio for all three damping materials in FFI Hopkinsor bar set-up as a function of damping material mass.

Compared with the setup where the damping material is near the bar, it is clear that wet sand now performs much worse, giving quite a bit of shock enhancement for very thin layers. The same happens for dry sand, while the porous sand now is the best material for thin layers, quite the contrary to the other setup. This is also different from the 1D uniaxial scenario where wet sand was best, followed by dry sand and where porous sand was not good.

To explain this, we have to realize that the mechanisms are now quite different from when the material was near the bar. First the damping material is compacted and accelerated by the explosive before eventually clashing into the steel bar. In addition there is the scattering effect. Porous material is very light, so when only a thin layer is near the charge, it is easily scattered and does not impact the bar. This explains why it performs so much better in this setup. On the other hand, wet sand is much heavier, so the scattering effect is smaller than for the other materials, which is the reason why it performs worse in this case.

10 Application of theory to charge buried under vehicle

Having learned about the potential of shock attenuation from various simulations and experiments, let us know apply the theory to a practical situation. One of the worst threats for personnel in Afghanistan and Iraq have been IED attack on a vehicle.

10.1 Buried charge

IEDs come in many varieties, but let us here consider a buried charge detonated below a vehicle. Before doing any simulations, let us use our knowledge developed in the previous chapters to estimate the effect of a buried charge on the vehicle, compared to a bare charge at the same distance.

We note that the buried charge scenario is quite similar to the DSTL experiment, in that there is a charge surrounded by "damping" material. In that setup we saw, both experimentally and numerically, an increase in the momentum transferred when the attenuation material was present. This is in agreement with "common knowledge" since "everybody" knows that a buried charge is much more dangerous than a charge positioned on the ground. Further, from the DSTL setup, we would also expect a charge in wet sand to transfer more impulse than a charge in dry sand.

Let us now see whether our intuition is correct. It would have been possible to set up a numerical simulation involving a complete vehicle, but for simplicity we will only look at a simpler setup with a buried charge accelerating a plate. In addition to this being a much simpler geometry to model, such experiments have actually been performed, enabling us to compare numerical results with actual data.

The setup to be studied was investigated experimentally in (13) and numerically in (10). It is a small scale experiment where cylindrical charges are covered by a thin layer of sand, as outlined in Figure 10.1. A cylindrical Composition-B high explosive (HE) charge of mass 0.625 kg was buried in a cylindrical sand container with a depth of burial (DOB) of 50 mm. Diameter to height ratio of the charge was approximately 3. A rectangular steel plate, 80 x 80 x 6 cm^3 and mass 300 kg, was positioned, using four wooden legs, with a standoff distance (SOD) of 0.2 m or 0.3 m measured from the sand surface to the bottom side of the plate. The sand container rested on a steel plate.

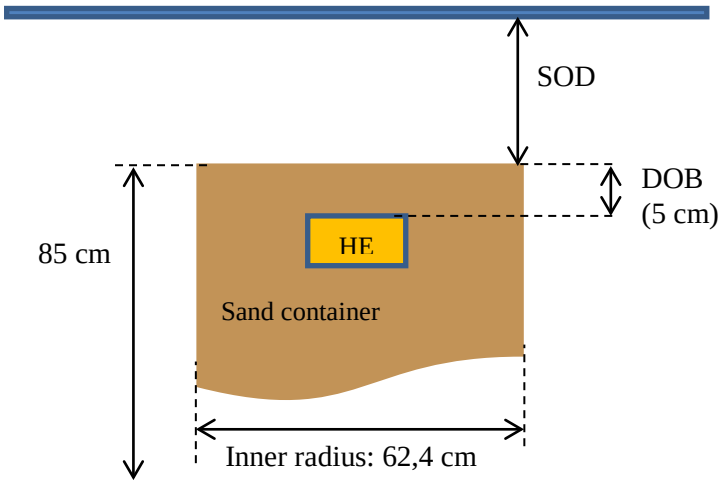

Figure 10.1 Buried charge setup used in (13).

The transferred impulse to the steel plate was measured and several different plate geometries were studied, in particular flat and V-shaped plates. Here we will only look at the flat plate case with an SOD of 200 mm. The other geometries do not add anything new.

As usual we will perform numerical simulations using our three test materials. To enable the problem to be run in 2D using AUTODYN, the rectangular plate was replaced by a circular plate of same area. Although we have not performed comparison simulations, this should not have significant impact on the results. Furthermore, we are interested in relative differences, not exact nominal values in the following discussion. We used the same materials as earlier in the report, except that this time the dry sand density was decreased by 0.3 g/cm^3 to be in approximate accordance with the experimental sand. The results are shown in Figure 10.2.

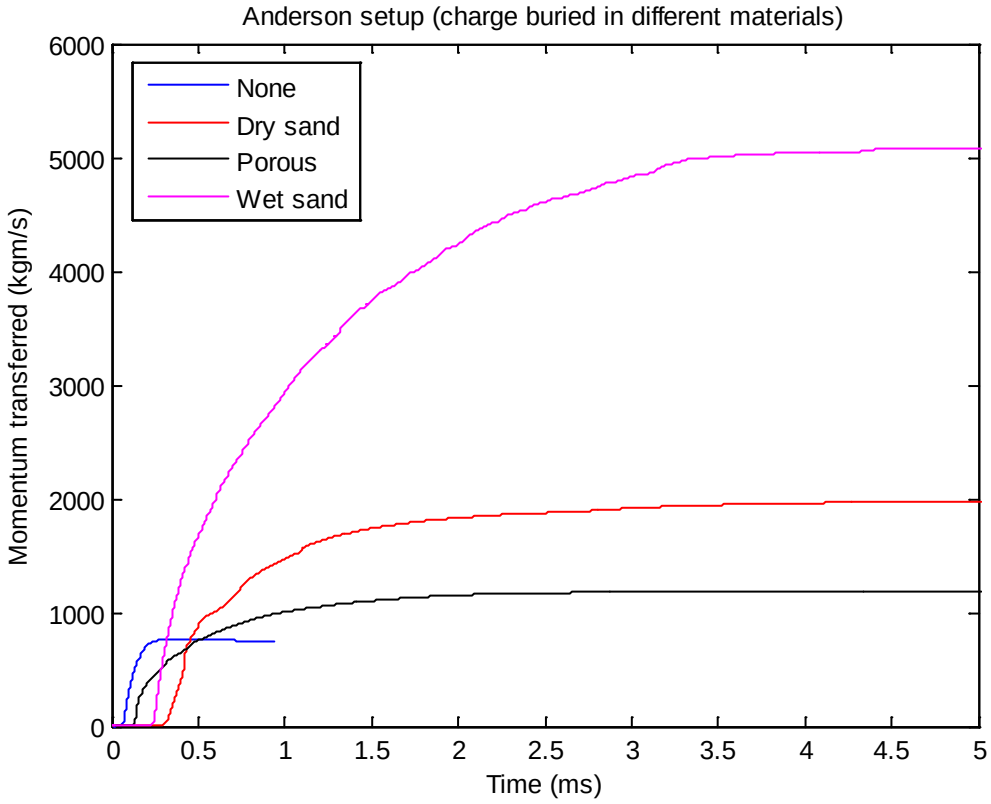

Figure 10.2 Transferred momentum for charge buried in different materials.

We see that our intuition was indeed correct. The transferred impulse varies enormously with the material the charge is buried in. For wet sand, the situation is particularly bad for the plate, with a transferred impulse of 5078 Ns, compared with 769 Ns for a bare charge. The result for dry sand is 1999 Ns, which is in good agreement with the experimental result of 1979 Ns. However, the dry sand used here may not be exactly the same as used in the experiments, so this may be coincidental. Experiments were only carried out using one type of dry sand.

10.2 Non-buried charge and protected plate

Now, let us consider possible protection for the plate. Let us assume that the plate is protected by a 200 mm layer of our three different materials (charge is at the same distance from the plate as in the unprotected case). For our non-buried charge (no sand), this situation is similar to the NTNU experiment. Thus, we should not expect much difference in transferred impulse, though the layer may or may not reduce the maximum amplitude. The AUTODYN setup is shown in Figure 10.3.

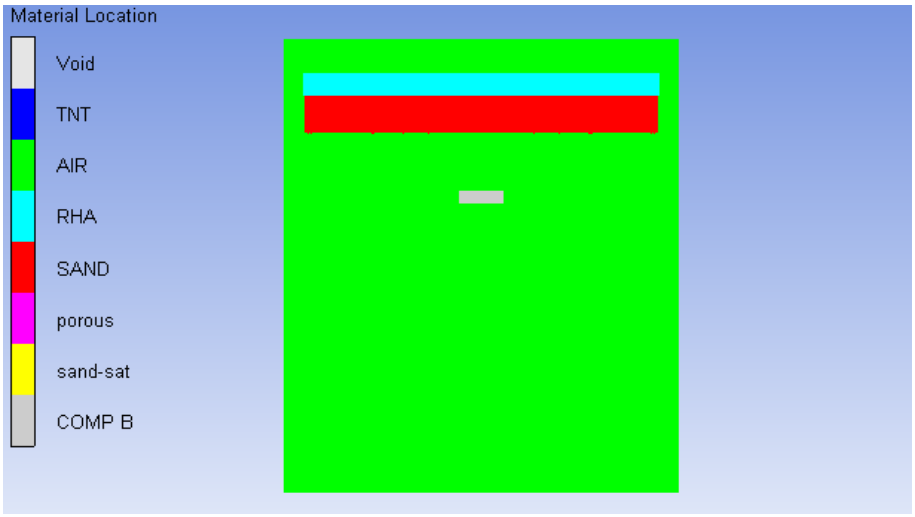

Figure 10.3 Non-buried charge and plate protected by dry sand.

On running these simulations for all damping materials, we obtain the results in Figure 10.4.

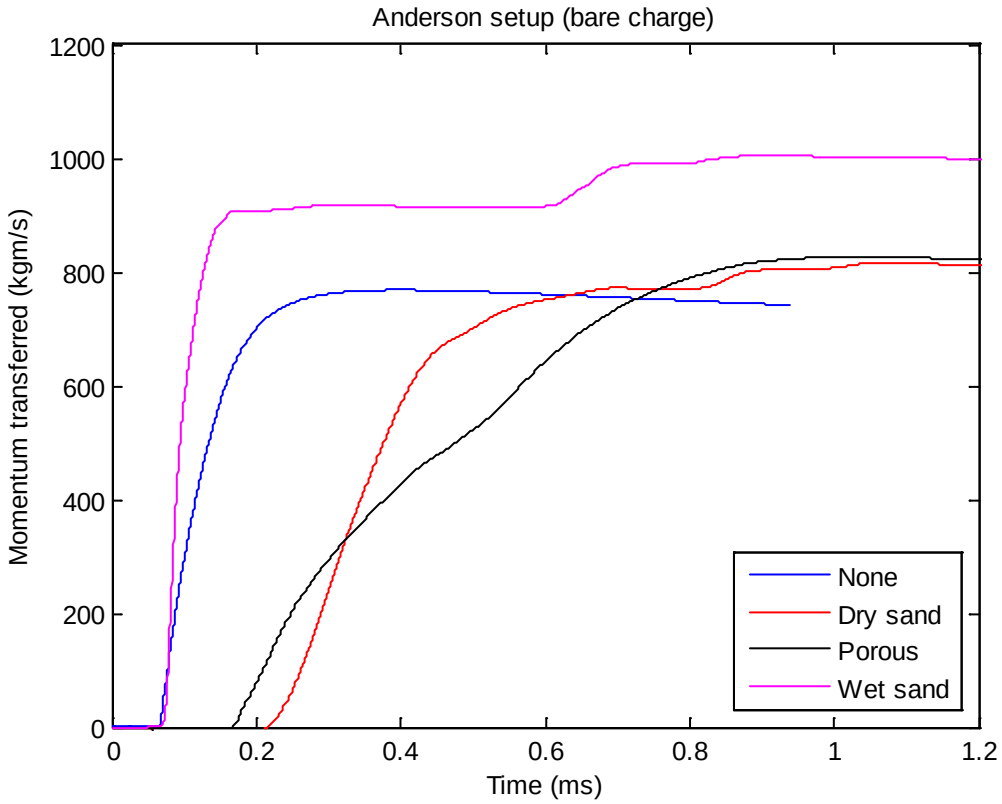

Figure 10.4 AUTODYN results for a plate protected by different materials.

Just as expected, we see that there is little change in the total transferred impulse. However, let us look at the average force on the plate (i.e. derivative of transferred momentum), as this could be important for the occupants of a vehicle. This is shown in Figure 10.5.

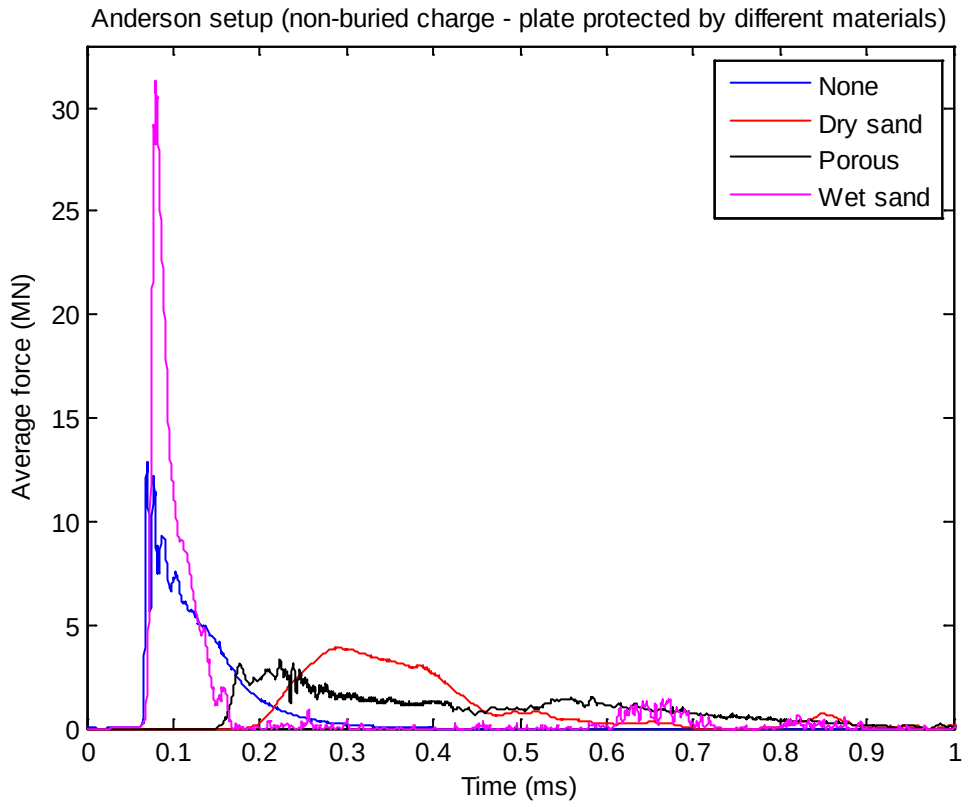

Figure 10.5 Average force on plate from non-buried charge. Plate is protected by different materials.

It is clear that both the porous materials (dry and porous sand) give longer lasting forces with lower amplitude than a non-protected plate. The non-porous wet sand gives a much higher peak force than even without the protection. Again this is consistent with everything we have learned before.

11 Summary

In this report we have examined shock attenuation both analytically, numerically and experimentally. We first reviewed several experiments on shock damping and noted that they seemed to give conclusions that were in contradiction to each other.

To understand how this could come about we first reviewed some basic impact physics and then performed a number of very simple 1D numerical simulations. From this approach we demonstrated the futility of trying to attenuate the impulse from a blast wave, since impulse is a conserved quantity. Despite this, shock attenuation actually makes sense because the manner in which momentum is transferred also matters.

We saw that it was indeed possible for a damping material to ensure that the momentum is delivered over a longer period of time, thereby giving a smaller force on the object to be protected. However, worryingly, we also noticed that under some circumstances the damping material may lead to an enhanced shock amplitude instead of the opposite.

Armed with the knowledge obtained from analytical theory and simplified simulations, we numerically analysed all the experiments that were described earlier. From this approach we were able to explain all the results obtained and see that they were not contradictory after all, but a consequence of using different measurement setups.

Finally we applied the theory of shock attenuation to buried charges under a vehicle and examined how it might be possible to protect the vehicles.

In summary, it is clear that shock attenuation is a very complex topic and a thorough understanding (as we have developed in this report) is necessary in order to design protection systems for a given target. Without such an understanding of the phenomenon one risks developing something that actually increases the damage to the target.

References

(1) Hanssen A G, Enstock L, Langseth M, Close-range blast loading of aluminium foam panels, International Journal of Impact Engineering 27 (2002), pp. 593-618

(2) Allen R M, Kirkpatrick D J, Longbottom A W, Milne A M, Bourne N K, Experimental and numerical study of freefield blast mitigation, AIP Conf. Proc. 706, 823 (2004)

(3) Kirkpatrick D, Argyle A, Harrison K, Leggett J, A comparison of blast & fragment mitigation performance of several structurally weak materials, AIR Conf. Proc. 955, pp. 951-954 (2007)

(4) Leggett J, Mitigation of explosive blast and fragmentation, Powerpoint presentation at ANNC WGIII meeting 2007

(5) Grover J, Evaluation of blast-mitigating materials worldwide, NAWCWD TM 8631, 2011

(6) Teland J A, Svinsås E, Frøyland Ø, Porous geological material for shock attenuation, Proceedings of 3rd European Survivability Workshop, Toulouse, Frankrike 16-19 mai 2006

(7) Teland J A, Skriudalen S, Sagvolden G, Svinsås E, Experimental investigation of shock attenuation properties of various protective materials, Proceedings of the 4th European Survivability Workshop, Great Malvern, England, 15-17 april 2008

(8) Teland J A, Skriudalen S, Shock attenuation experiments using a Hopkinson bar method – summary of experimental results, FFI/RAPPORT-2009/00582

(9) Fairlie G, Bergeron D, Numerical simulation of mine blast loading on structures, Proceedings of the 17th MABS, Las Vegas, Nevada, USA, June 2002

(10) Skriudalen S, Teland J A, Sand models in IMPETUS Afea and AUTODYN Comparing experimental results to simulations, FFI/RAPPORT-2013/01588

(11) Grujicic M, Bel V C, A computational analysis of survivability of a pick-up truck subjected to mine detonation loads, Multidisipline modelling in materials and structures, vol. 7, no. 4, pp. 386-423, 2011.

(12) Glanville J, Private e-mail message, April 28, 2009

(13) C. E. Anderson Jr, T. Behner, and C. E. Weiss, Mine blast loading experiments, International Journal of Impact Engineering, vol. 38, pp. 697-706, 2011.

Appendix A Material models

A.1 Dry sand

Equation of State	Compaction
Reference density	2.6410 g/cm^3

Loading	
Density (g/cm^3)	**Pressure (MPa)**
1.6740	0.00
1.7395	4.577
1.8738	14.98
1.9970	29.151
2.1438	59.175
2.2500	98.098
2.3800	179.443
2.4850	289.443
2.5850	450.198
2.6713	650.660
Unloading (linear method)	
Density (g/cm^3)	**Soundspeed (m/s)**
1.6740	265.200
1.7456	852.100
2.0863	1721.70
2.1468	1875.50
2.3000	2264.80
2.57200	2956.10
2.59800	3112.20
2.63500	4600.00
2.64100	4634.00
2.80000	4634.00

Strength model (MO-Granular) Pressure part	
Pressure (MPa)	**Yield stress (MPa)**
0.00	0.00
3.401	4.235
34.898	44.695
101.324	124.035
184.650	226.000
500.000	226.000

Strength model (MO-Granular) Density part	
Density (g/cm^3)	**Yield stress (MPa)**
2.00	0.00
10.00	0.00

Strength model (MO-Granular) Shear modulus G	
Density (g/cm^3)	**G (MPa)**
2.03000	76.9000
2.10100	869.400
2.44200	4030.00
2.50200	4910.00
2.65600	7769.00
2.92800	14800.9
2.95400	16571.0
2.99100	36718.0
2.99700	37347.0

Failure model	**Hydro (Pmin)**
Hydro tensile limit	-1.0 kPa

A.2 Wet sand

Equation of State	**Compaction**
Reference density	2.6410 g/cm^3

Loading	
Density (g/cm^3)	**Pressure (MPa)**
2.046	0.00
2.052	72.28
2.058	144.58
2.063	216.80
2.069	289.10
2.075	361.40
2.081	433.70
2.086	506.30
2.092	578.30
2.156	1510.00
Unloading (linear method)	
Density (g/cm^3)	**Soundspeed (m/s)**
2.046	3812.0
2.156	3812.0

Strength model (MO-Granular) Pressure part	
Pressure (MPa)	**Yield stress (MPa)**
0.00	0.00
30.77	9.952
61.53	32.81
92.30	68.60
123.10	117.3
153.80	178.9
184.60	253.50
215.37	253.50
246.13	253.50

Strength model (MO-Granular) Density part	
Density (g/cm^3)	**Yield stress (MPa)**
2.00	0.00
10.00	0.00

Strength model (MO-Granular) Shear modulus G	
Density (g/cm^3)	**G (MPa)**
2.03000	76.9000
2.10100	869.400
2.44200	4030.00
2.50200	4910.00
2.65600	7769.00
2.92800	14800.9
2.95400	16571.0
2.99100	36718.0
2.99700	37347.0

A.3 Porous sand

Equation of State	Compaction
Reference density	2.10 g/cm^3

Loading	
Density (g/cm^3)	**Pressure (MPa)**
0.275	0.00
0.4395	4.577
0.6738	14.98
0.897	29.151
1.1438	59.175
1.3500	98.098
1.5800	179.443
1.7800	289.443
1.9850	450.198
2.1700	650.660
Unloading (linear method)	
Density (g/cm^3)	**Soundspeed (m/s)**
0.274	265.200
0.4456	852.100
0.8863	1721.70
1.0468	1875.50
1.3000	2264.80
1.6700	2956.10
1.7900	3112.20
2.0350	4600.00
2.1400	4634.00
2.4000	4634.00

Strength model (MO-Granular) Pressure part	
Pressure (MPa)	**Yield stress (MPa)**
0.00	0.00
3.401	4.235
34.898	44.695
101.324	124.035
184.650	226.000
500.000	226.000

Strength model (MO-Granular) Density part	
Density (g/cm^3)	**Yield stress (MPa)**
0.01	0.00
10.00	0.00

Strength model (MO-Granular) Shear modulus G	
Density (g/cm^3)	**G (MPa)**
1.674	76.9000
1.7457	869.400
2.0863	4030.00
2.1468	4910.00
2.3000	7769.00
2.5720	14800.9
2.5980	16571.0
2.6350	36718.0
2.6410	37347.0

Failure model	**Hydro (Pmin)**
Hydro tensile limit	-1.0 kPa