

The effect of statistical fluctuation, finite size error, and digital resolution on the phase percolation and transport properties of the NIST cement hydration model

by

E.J. Garboczi and D.P. Bentz
Building and Fire Research Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899 USA

Reprinted from Cement and Concrete Research, Vol. 31, No. 10, 1501-1514, October 2001.

NOTE: This paper is a contribution of the National Institute of Standards and Technology and is not subject to copyright.



NIST
National Institute of Standards and Technology
Technology Administration, U.S. Department of Commerce



The effect of statistical fluctuation, finite size error, and digital resolution on the phase percolation and transport properties of the NIST cement hydration model

E.J. Garboczi*, D.P. Bentz

*Building Materials Division, Building and Fire Research Laboratory, National Institute of Standards and Technology, 100 Bureau Drive,
Stop 8621, Gaithersburg, MD 20899-8621, USA*

Received 26 October 2000; accepted 25 June 2001

Abstract

The National Institute of Standards and Technology (NIST) cement hydration model starts with a three-dimensional digital image of cement particles, and then uses cellular automaton rules to simulate the reaction of cement with water and the development of cement paste microstructure. This stochastic model uses a digital image with a certain resolution to represent the reality of continuum shapes in a fairly small periodic unit cell. As such, it is potentially subject to statistical fluctuation, finite size error, and the effects of digital resolution. This paper evaluates the model in light of these potential sources of error, along with the effects of cement particle size distribution (PSD), water/cement (w/c) ratio, and cement chemistry, by focusing on the phase percolation and transport property predictions of the model. Statistical fluctuation and finite size error are shown to be of minor importance in the model. The effect of digital resolution is significant, however, but different from the case of a finite element or finite difference model. Unlike these cases, the “best” resolution is not necessarily the finest resolution, but must be chosen based on comparison with experiment and various physical length scales present in cement paste. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Hydration; Hydration products; Electrical properties; Transport properties; Modeling

1. Introduction

Concrete is a composite material made up of the primary phases of cement paste, pores, and aggregates of various sizes. However, it is not a simple composite material. It can be classified as a porous, multiscale [1–4], interactive [2], composite material. It is a porous material because the cement paste matrix is a porous material, where the pores are produced mainly by the chemical reactions between cement and water and the original size and spatial distribution of the particles. It is a multiscale material because there are morphologically different microstructures of importance at length scales ranging from nanometers to millimeters [4]. The distribution of pores in the pore space is modified by the presence of aggregates, which is why we call concrete an

interactive composite. As more aggregates are added, the additional, highly porous interfacial transition zone regions force the bulk matrix to be denser in order to conserve the overall water/cement (w/c) ratio [2].

This paper focuses on the cement paste matrix of concrete, which controls the transport properties of concrete. In this paper, “transport properties” means diffusion of ions and flow of water through the pore space of the cement paste. Any porous material allows transport through its pores if the following conditions are fulfilled: (a) the pores themselves allow transport within them and (b) these pores are continuous or percolated [5] through the material. An illustration of the first condition is whether or not a pore is saturated with water in a cement paste. If the pore is dry, obviously, no ions will diffuse through the pore space but gases can diffuse through the empty pore. If the pore is full of water, then ions can diffuse through it but now gases cannot (unless they first become dissolved in the liquid water). Only fully saturated cement paste will be considered in this paper. To understand the second condition, just

* Corresponding author. Tel.: +1-301-975-6707; fax: +1-301-990-6891.

E-mail address: edward.garboczi@nist.gov (E.J. Garboczi).

imagine that the entrained air voids were the only pores in a concrete. Since these are in general isolated from each other, there will be no diffusive or fluid transport through these pores since they do not form a connected network.

The fact that the pores that allow transport are percolated does not completely determine the rate of transport. The tortuosity of the pore network and the size of the pores also influence this rate [6]. Different transport processes depend differently on these pore space parameters, however. For example, fluid permeability depends directly on the square of a pore size, while ionic diffusivity depends much less sensitively on pore size [6–8].

Like concrete, cement paste is also a random composite material. The original cement grains are themselves made up of many chemically distinct phases and the chemical reactions with the mixing water produce many other chemically distinct phases. In any random material, one may ask the following question about any one of the phases: is it percolated or isolated? If it is the pore phase that is being considered, one can also ask what is the ionic diffusivity or fluid permeability through these pores, if they are connected?

The current version of the model considered in this present work has been thoroughly described in a recent review article [9]. The model consists of a 3-D digital image with cubic pixels where each pixel is identified by a phase label that tells which phase inhabits that pixel. Images of cement particles are superimposed on the lattice using various methods [9] and hydration is simulated by dissolving the cement pixels and allowing the dissolved species to diffuse in the pore space and form new products according to cellular automaton rules that mimic the known chemistry of the original cement particles. At any given point in time, the algorithm can be interrupted and various properties measured, like degree of hydration, phase percolation, diffusivity, elastic moduli, heat signature, or chemical shrinkage. Ref. [9] describes how full portland cement chemistry is now handled by the model, in contrast to earlier models that were for C_3S only [10].

Any model of a random particulate material has sources of error inherent in its representation. For example, the hydration model only represents a piece of cement paste of submillimeter dimensions. There can be errors that occur because the small size of the model, relative to usual cement paste sample dimensions, cannot completely capture the statistics of how real cement particles are spatially arranged. Since the cement particles are randomly mixed, two different arrangements of the particles in the unit cell of the model can, in principle, give different results. A third source of error comes about because the model uses a digital image to capture the real shapes of hydrated and unhydrated phases. Inadequate resolution, which is the real spatial length per pixel, can cause errors in the outputs of the model.

The focus of this paper is to evaluate the effect of the above three potential sources of error on the predictions of the model for percolation and transport properties. In doing

this, we will also consider the effect of changes in chemistry, computer power, and new understanding of the parameters of the model. The coupling of model statistics and real microstructure will be considered and discussed. This paper can be considered to be an update of several earlier works [10,11], as well as a critical analysis of the “nuts and bolts” of the hydration model.

2. Review of percolation concepts

The concept of percolation is important for many aspects of science and technology [5]. For materials, percolation really only applies to a *random* material, made up of two or more phases.

Through some process, one or more phases are being randomly built up, for example, through being the products of chemical reaction, or randomly consumed via being the reactant in a chemical reaction. At any instant of time, a phase is either topologically connected throughout the microstructure or topologically disconnected. If connected, then a hypothetical ant could walk from one side of the material to the other along this phase. If disconnected, then the ant would have to “jump” some gaps between parts of the phase in order to proceed along it. When a phase goes from being disconnected to connected, or vice versa, so that there is a change in topology, then we call this point a percolation threshold, and denote the volume fraction of the phase at this point with a subscript “c.”

Percolation obviously has a strong influence on transport. When a material only has one phase that is active in transport, then when that phase is percolated, the material has a nonzero transport property. When that phase is disconnected, then the material has a zero transport property. When two or more phases are active in transport, for example, electrical conductivity, then the contrast between the conductivity of each kind of material matters, as well as their relative connectedness.

Since the hydration model is digitally based, it is easy to check whether any given phase or combination of phases is percolated using a modified burning algorithm [12]. This algorithm computes whether or not a phase is percolated. If the phase is percolated, the algorithm also computes how much of the phase is percolated and how much is isolated. In this paper, the “fraction connected” is reported in many of the percolation figures, which is numerically equal to the fraction of the phase in question that is percolated.

3. Sources of error when simulating random materials

Any model is only an approximation of reality. There are always shortcomings built into any model, which must be quantitatively addressed when assessing the accuracy of model results. Digital models, like the National Institute of Standards and Technology (NIST) cement hydration

model, are no different. Here, the term “model” refers to a fundamental computer simulation model [13], where the microstructure is directly represented in the computer.

When one builds this kind of model of a random system, there are immediately two main sources of error built into the process. These two error sources are [14,15]: (1) statistical fluctuation and (2) finite size effect. As mentioned in Section 1, statistical fluctuation errors come about simply because the system modelled is random, while finite size effect errors arise because the model is usually not as large as the real material. Digital models have a third important source of error, digital resolution. Here, the term “digital model” means a direct representation by a 2-D or 3-D digital image, where each pixel bears a label identifying it as a phase of the material. Each of these three sources of error will be discussed in turn for the percolation and transport property aspects of the cement paste microstructure model.

Any model of a random material can only represent a piece of the material. Finite size error comes about if the piece of the material is not big enough to be “typical” of a representative size sample. The terminology that is used in the composite material field is “representative elementary volume” [16,17]. Uchikawa [18] has examined this concept for cement paste, mortar, and concrete and has observed qualitatively in 2-D images that an area of $100 \times 100 \mu\text{m}$ was representative for cement paste. An example would be the usual rule for concrete samples, that they should be three to five times as big as the largest aggregate for them to be considered “typical.” Empirical testing has determined that this is “big enough” for concrete. This concept has been dealt with very carefully for a limited class of problems, but the formalism is general [14,19,20]. For any model of a random material, the only way to tell if the model is “big enough” is by testing the property of interest for different size models. In our case, one uses bigger (smaller) unit cells with more (fewer) particles. The unit cells have periodic boundary conditions, which help greatly in reducing the effect of the surface by essentially removing the surface. This has been done for the cement hydration model. Models that are $100 \times 100 \times 100 \mu\text{m}$ in size, containing typically 2000 particles, have been found to be large enough to be representative. Larger models containing more particles at the same digital resolution and with periodic boundary conditions, have not been found necessary for accuracy.

Since cement paste is a random material, the digital representation must also be random. This means that different realizations of the model will have some differences between them. Random numbers determine where the cement particles are originally placed in the unit cell and how the hydration takes place [9]. As will be seen below, averaging a given percolation curve over different random numbers does not make much of a difference. The averaged curves become smoother than any one curve, but the qualitative behavior of the curve, i.e., where its percolation threshold is, does not change significantly. Other properties

of the hydration process, like heat release curves, have also been shown to vary little with statistical fluctuation [21].

The final source of error is digital resolution, which is unique to digital models [12,14,15]. For example, the model tries to represent a cement particle, which has an order of 10^{13} atoms for a typical $15\text{-}\mu\text{m}$ diameter C_3S particle, by a fairly small number of pixels. In the model, the pixels dissolve and move around, resulting in digitally fragmented remainders that have digitally fragmented surfaces. Another way of thinking about digital resolution error is to compare the pixel size to an intrinsic length scale in the problem. For example, if the smallest CH crystal of importance is $0.1 \mu\text{m}$ and the pixel size is $0.4 \mu\text{m}$, then obviously the smallest CH crystals cannot be represented as single units and the model will fail to resolve these particles. This may not induce much error in predicted properties, however, if these small particles do not play a significant role in determining the properties of interest. The main purpose of this paper is to determine exactly how digital resolution error matters. Of the three sources of error discussed in this section, it turns out that the digital resolution error is the only important one for the cement hydration model. We investigate this effect by keeping the same physical size model, a cube $100 \mu\text{m}$ on a side, but with a different length per pixel so that more pixels are used at finer resolution.

4. Previous work

There have been two main previous model studies of the percolation characteristic of cement paste phases using two different models [10,22].

The first of these was on an early version of the present model, with fixed cement chemistry (C_3S only), fixed cement particle size distribution (PSD), and fixed resolution (100^3 pixels). Computational limitations at that time prohibited consideration of higher resolutions, but it was determined that finite size errors and statistical errors were negligible. As will be seen later, the cement PSD does not matter very much as long as it is broad. Unfortunately, the PSD that was used in this early work was for only four different sizes of particles. This in itself was not so much of a problem, but the four different sizes had equal numbers, which meant that the distribution was heavily skewed in terms of particle mass towards the larger particles. The effect of the cement PSD is studied in Section 5, along with the effect of cement chemistry and other changes to the basic model made since the time of the first study (1991). The effect of digital resolution is left to Section 6.

The second study of percolation in cement paste focused on the percolation of the capillary pore space, which is defined as the original water-filled space around the cement particles and their hydration products. The model used was a version of an older model [23], where the particles were represented by collections of continuum spheres, not digital images. Hydration was simulated by the spherical growth of

cement particles along with the generation of CH particles. The particles were allowed to overlap, with some correction made for the common volume between two overlapping particles as they both grew.

Percolation of the capillary pore space was evaluated by digitizing the resulting continuum image at varying resolutions and then studying the percolation properties as a function of resolution. The resolutions studied with this model ranged from 10 to 1 $\mu\text{m}/\text{pixel}$ [22] (the resolutions studied in the present paper range from 1 to 0.125 $\mu\text{m}/\text{pixel}$, so that the ranges are complementary).

The percolation threshold of the capillary pore space was found to depend on the resolution used, which is not surprising [22]. The reason is that the microstructure was generated using continuum objects, so one had to go to fine-enough resolution until the continuum pore space was properly sampled. In particular, some fraction of pore throats that were narrower than the size of the pixel used would be lumped into the solid phase and not treated as a pore at all. This would tend to decrease the connectivity of the pore space, as connections via these kind of pores would not be counted until the resolution was fine enough. The results seem to agree with this explanation. This model is similar in many respects to a simple overlapping solid sphere model, where the pore space is just the space around a collection of randomly centered and sized solid spheres. A recent study of the elastic properties of this model found that the effect of digital resolution was quite pronounced [14].

5. Effects of w/c, PSD, and cement chemistry

Changes have been made to the hydration model since the first percolation study was carried out [9,24]. The changes involve continuous dissolution and reaction and cement chemistry. The original model dissolved some cement, diffused and reacted all dissolved cement, then dissolved some more, in a cyclic fashion. The current version of the hydration model operates in a more continuous dissolution/reaction mode while still retaining the concept of cycles. By changes in cement chemistry is meant that now the model can handle full portland cement chemistry vs. only being able previously to handle pure C_3S pastes. Because of these changes, it was important to go back and investigate the effect of PSD and recheck the effect of w/c ratio. The models considered in this section are all 100^3 -size models, with a resolution of 1 $\mu\text{m}/\text{pixel}$.

A realistic PSD for the original cement particles was used and the effect of w/c ratio and cement chemistry studied. The PSD for the C_3S models to be considered comes from the thesis of Boumiz [25] and the PSD for the portland cement models comes from the Cement and Concrete Reference Laboratory (CCRL) standard cement 115 [21] (details of the cement are in Ref. [21]).

Fig. 1 shows the “fraction connected” for the capillary pore space for the C_3S model plotted vs. degree of hydration

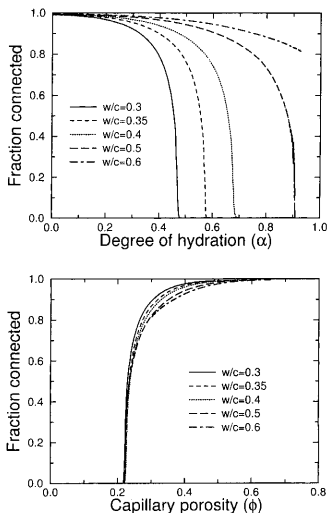


Fig. 1. Fraction of the capillary pore space that is part of a connected pathway for various w/c ratio C_3S pastes as a function of the degree of hydration (top) and the capillary porosity (bottom) using the PSD of Boumiz [25].

(top) and capillary porosity (bottom) for values of w/c ratio varying between 0.3 and 0.6. This version of the model includes continuous dissolution of particles. Notice, in the bottom graph of Fig. 1, how all the curves collapse to approximately a single curve with a common percolation threshold of about 0.22. That means, for this range of w/c ratios, that the capillary pore space becomes topologically disconnected when its total volume fraction falls below 22% of the total cement paste volume. This value (22%) is different than the number reported in the earlier work [10] (18%). We have done careful simulations to show that this difference is not due to any of the small changes in the model carried out since 1991, like continuous vs. cyclic dissolution, but rather due to differences in cement PSD, as will be shown below.

Fig. 2 shows results for capillary porosity percolation for two different cement systems at two different w/c ratios. In both cases, only four sizes of particles were used, with diameters of 3, 9, 13, and 19 pixels. In one case, equal

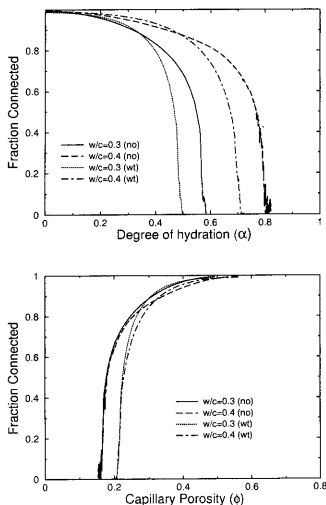


Fig. 2. Fraction of the capillary pore space that is part of a connected pathway for two w/c ratio C_3S cement pastes as a function of the degree of hydration (top) and the capillary porosity (bottom). Two discrete four-size PSD were used: one with equal numbers of particles in each class (no), and one with equal volumes of particles in each class (wt).

numbers of each size particle were used, while in the other case, equal volumes of particles were used. The equal number case is almost identical to that used in the earlier work [10]. In Fig. 2, we see pairs of results that look very similar, except that the percolation thresholds, as seen in the bottom parts of the figures, are significantly different between the two cases. The equal number case, which has an average particle diameter heavily weighted towards the larger particles, has a percolation threshold of about 16%, close to that found in the earlier work. The equal volume case, which has a much broader distribution, has a percolation threshold of about 21%, much closer to the results found in Fig. 1. Other works on portland cement particles have shown that for this resolution system, $1 \mu\text{m}/\text{pixel}$, as long as the cement PSD is reasonably broad, the capillary percolation threshold will be about 20–22% [26].

A question that has arisen in other work, and has now become important for the model work, is: how close are C_3S

results to results on portland cement? The model has incorporated full portland cement chemistry for some time now [9] and the percolation results should be compared to the C_3S -only model. Fig. 3 shows the same kind of results as was shown in Fig. 1, but for a portland cement, the NIST 115 standard cement. The resemblance to Fig. 1 is striking, with almost exactly the same percolation threshold of about 21%. Therefore, one can safely say that there is no difference in the capillary pore space percolation characteristics between portland cement paste and C_3S cement paste, as both models distribute hydration products on the cement (C_3S) particle surfaces and within the available pore space in a similar manner.

The percolation of other phases can also be considered, and comparisons made between C_3S and portland cement pastes. Fig. 4 shows the percolation data for the C–S–H phase for C_3S pastes, and Fig. 5 shows the equivalent for an NIST 115 portland cement paste. The qualitative shape and behavior of the two systems are identical, with a percolation

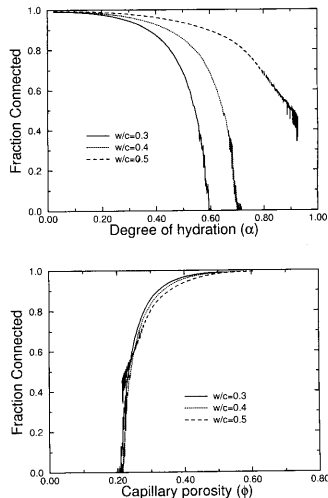


Fig. 3. Fraction of the capillary pore space that is part of a connected pathway for three w/c ratio portland cement pastes as a function of the degree of hydration (top) and the capillary porosity (bottom) using the NIST 115 PSD.




Fig. 4. Fraction of the C–S–H phase that is part of a connected pathway for three w/c ratio C_3S cement pastes as a function of the degree of hydration (top) and the volume fraction of the C–S–H phase (bottom) using the Boumiz [25] PSD.

threshold for the C_3S paste of about 0.225, while the portland cement paste had a percolation threshold of about 0.18. In this case, there is a quantitative difference between the percolation aspects of the two systems. If we consider a C_3S paste and an NIST 115 paste at the same w/c ratio, and the same degree of hydration, we find that the ratio of C–S–H volume fractions between the two pastes is about the same as the ratio between the C–S–H percolation thresholds. Both pastes are producing about the same total volume of hydration products. In the NIST 115 case, there are other products that help to cover the surfaces of the cement particles, allowing the C–S–H produced to be more efficient at bridging between particles, thus, allowing a lower volume fraction to still percolate.

Figs. 6 and 7 show the CH percolation information for C_3S and portland cement pastes, respectively. These two systems both have a CH percolation threshold of about 0.155. The noise seen in Fig. 7 is due to the fact that a mechanism for CH dissolution and recrystallization was

allowed in the portland cement system. The noise was caused by dissolution and recrystallization of key small crystals, which did not change the volume fraction of CH very much, but changed the connectivity a lot. One might imagine small crystals that completed a key “bridge” that just connected two large masses of CH. If these small crystals should dissolve, a large amount of connectivity would be lost, but not a large amount of CH. The CH phase forms in the available pore space between cement particles, so its percolation threshold is less sensitive to cement chemistry.

6. Effects of digital resolution on phase percolation

Increasing the fineness of the digital resolution of the hydration model, while keeping the physical size fixed, will increase the computational size of the models. Even having access to the computational facilities of NIST, when the




Fig. 5. Fraction of the C–S–H phase that is part of a connected pathway for three w/c ratio portland cement pastes as a function of the degree of hydration (top) and the volume fraction of the C–S–H phase (bottom) using the NIST 115 PSD.

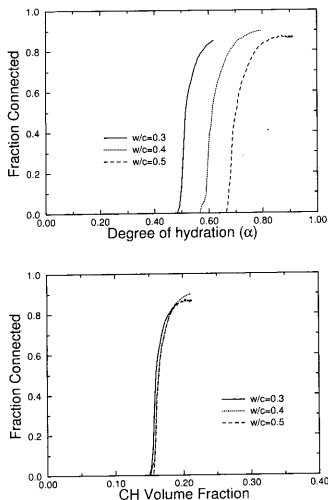


Fig. 6. Fraction of the CH phase that is part of a connected pathway for three w/c ratio C_3S cement pastes as a function of the degree of hydration (top) and the volume fraction of the CH phase (bottom) using the Boumiz [25] PSD.

hydration models become large, the computer time for hydrating a model becomes large as well. Having more continuous dissolution/reaction also slows down the model, compared to the older versions. We have used grids as large as 800^3 pixels in size for the C–S–H results. Since most of the percolation results for other phases required longer hydration times, only a maximum of 400^3 -size systems were able to be used for these results.

Larger sizes are possible with the memory on machines available to us, but the slow turnaround time for systems above about 1000^3 is prohibitive, even for the C–S–H results. Memory requirements, with the stripped-down version of the model used in this paper, which does not include the set point computation, is about $1.7 n^3$ bytes, where the model is $n \times n \times n$ pixels in size. Since monitoring the setting process takes much more memory (to store the original particle label in the appropriate pixel [9]), this was not carried out for the larger sizes, and so no digital resolution comparisons for the set point are available.

However, insight into the effect of digital resolution on setting can be obtained from the results for the C–S–H phase percolation, and will be discussed later.

One aspect of the model must be addressed before the digital resolution effect can be discussed. A key part of the model is the dissolution of cement particles. This has been done in the past by considering all cement pixels that are touching water. For such a pixel, one of the six nearest neighbors of the cement pixel is selected randomly. If it is a water-filled pore pixel, then the cement pixel can dissolve into it. If it already occupied by a solid phase, then the cement pixel cannot dissolve.

When simulating systems with coarser PSD and, therefore, larger particles, it has been found that hydration using this dissolution rule stops too soon. To correct this deficiency, and allow the use of larger particles (on the order of $40 \mu m$), the second and third nearest neighbor pixels were also allowed to be queried as possible dissolution sites [26,27]. There are a total of 26 neighbors when first, second,




Fig. 7. Fraction of the CH phase that is part of a connected pathway for three w/c ratio portland cement pastes as a function of the degree of hydration (top) and the volume fraction of the CH phase (bottom) using the NIST 115 PSD.

and third nearest neighbors are considered (all the pixels in a $3 \times 3 \times 3$ cube, with the central pixel being the cement pixel being investigated). This change allows further hydration to be carried out on larger particles, and provided excellent agreement between experiment and model data [27].

Physically, one can think of a "dissolution length," a length over which the material can dissolve. By going from 6 to 26 neighbors, we are effectively going from a dissolution length of 1–1.4 pixel. Of course, as the resolution increases, and the physical length per pixel decreases, the physical length of this pixel size decreases. This will be discussed further in Section 8.

The question then arises, what is the effect of digital resolution on the percolation properties of the cement paste model? In the results to be discussed next, the effect of having 6 or 26 neighbor dissolution will be considered simultaneously with the effect of digital resolution on the percolation properties of the model.

Section 5 showed that different w/c ratio pastes have percolation quantities ("fraction connected") that collapse onto a single curve when plotted against the volume fraction of the phase being considered. Therefore, the following results will be for a single w/c ratio of 0.3, plotted only against the volume fraction of interest, and not the degree of hydration. This lower value of w/c ratio was chosen, as it was computationally easier to drive it through its percolation thresholds, since less hydration was needed. In addition, only C_3S cement particles were used because they were simpler and quicker to hydrate, yet had the same qualitative behavior as did the portland cement systems.

Each system was created with the same number of randomly placed particles, but with each particle appropriately larger in pixels according to the resolution. In continuum terms, the same size particles were used at each resolution, but with a different number of pixels per unit length. All systems were the same physical size, $100 \times 100 \times 100 \mu\text{m}$. The resolutions were: 1 $\mu\text{m}/\text{pixel}$ for the 100³-pixel system, 0.5 $\mu\text{m}/\text{pixel}$ for the 200³-pixel system, 0.25 $\mu\text{m}/\text{pixel}$ for the 400³-pixel system, and 0.125 $\mu\text{m}/\text{pixel}$ for the 800³-pixel system. For comparison, the 100³ and 200³ sizes were investigated using both 6 and 26 neighbor dissolution. At the 400³ and 800³ sizes, the 26 neighbor dissolution was necessary to be able to hydrate the cement to the point of capillary porosity depercolation and CH percolation. Both the 6- and 26-neighbor dissolution schemes were able to hydrate to the point of C–S–H percolation, which comes fairly early in the hydration process (see Fig. 4).

Fig. 8 shows a comparison, at the same w/c ratio and the same porosity of 0.32, between the 100³-, the 200³-, and the 400³-resolutions. Note that each system had the same physical size, $100 \times 100 \times 100 \mu\text{m}$. The images shown in Fig. 8 are not full slices, but show an area of approximately $50 \times 50 \mu\text{m}$ in size. The capillary pore space is in white, and all solid phases have been turned black, in order to focus on the capillary pore space.

Because each particle is the same physical size, but is made up of more pixels at higher resolutions, the amount of hydration per dissolution cycle will be smaller as the resolution increases. The amount of hydration per cycle will also be affected by the number of neighbors allowed to be considered for dissolution. Fig. 9 shows a plot of the degree of hydration vs. number of dissolution cycles for three different resolutions at $w/c=0.3$, with both the 6 and 26 neighbor dissolution. Note that decreasing resolution and increasing number of neighbors considered for dissolution both tended to increase the amount of hydration per cycle.

Fig. 10 shows the capillary pore space percolation for this paste, at three resolutions: 100³, 200³, and 400³. For the two lower resolutions, the 6- and 26-neighbor dissolution curves are identical in shape and displaced slightly in percolation threshold. Therefore, there is only a small quantitative, and no qualitative, effect, of using a different number of neighbors for dissolution on capillary pore percolation. More interesting is the progressive lowering of the capillary pore percolation threshold as the resolution is increased from 0.22 at 100³ to 0.12 at 400³. Essentially, this happens because smaller pores can be resolved at the greater resolutions, and a pathway that would seem to be closed off at low resolution is seen to be narrowly open at higher resolutions.

This resolution effect is also seen in simpler systems, for example in 2-D percolation of digital circles [28]. When the circles are made up of only 1 pixel and are randomly placed on a square digital lattice, this process is equivalent to site percolation, which has a percolation threshold of 0.59 for the circle phase. As the resolution is increased, the continuum circle limit of 0.68 area fraction of circles is achieved. In 3-D, the same process for spheres leads from site percolation (1 pixel/sphere), where the percolation threshold is 0.249, to continuum sphere percolation, where the threshold is 0.29 [29].

The same story holds true for CH phase percolation. Fig. 11 shows the percolation curves for CH. The difference between the 6- and 26-neighbor dissolution now seems a little more pronounced than in Fig. 10, but still small compared to the effect of digital resolution. The movement of the CH threshold with resolution seems a bit different from the capillary percolation threshold case. The 800³ result seems close to an equilibrium point, with the infinite resolution value of the CH percolation threshold at about 0.1. Only a small part of this curve could be generated with the available computer capacity. CH phase percolation is topologically that of shapes growing at random locations, which gradually impinge enough so that a connected backbone is achieved. Certainly the shape of the crystals, which is a function of their growth habit, sensitively determines their percolation threshold. For the case of simply overlapping, equal size and shape ellipsoids, it was found that the percolation threshold was sensitively dependent on the aspect ratio of the ellipsoids of revolution [29]. It is certainly possible that the way the model grows CH does not agree well with real cement paste. For example, at times CH tends




Fig. 8. 2-D slices through the microstructure of 0.3 w/c cement C_3S paste at three different resolutions: 100^3 (top left), 200^3 (top right), and 400^3 pixels (bottom), and the same capillary porosity-0.32. Black is solid, white is the capillary pore space. The physical size of the slices is about $50\text{ }\mu\text{m} \times 50\text{ }\mu\text{m}$.

to form flat, hexagonal crystals, at least when there is room in the pore space for them to grow. The model grows them as

roughly isotropic masses. CH morphology is also strongly dependent on temperature and silica fume content [30].




Fig. 9. Degree of hydration attained vs. the number of model cycles run for three different resolutions, and 6 vs. 26 neighbor dissolution for 0.3 w/c C_3S cement paste.




Fig. 10. Fraction of the capillary pore space that is part of a connected pathway as a function of the capillary porosity for 0.3 w/c at different resolutions for a C_3S cement paste.




Fig. 11. Fraction of the CH phase that is part of a connected pathway as a function of the total CH volume fraction for 0.3 w/c at different resolutions for a C_3S cement paste.

The “fraction connected” curves are shown in Fig. 12 for the C–S–H phase. Here it is seen that there are only very small differences between 6- and 26-neighbor dissolution. However, the C–S–H phase percolation threshold is probably the most sensitive to the digital resolution of the three phases considered. The value of this percolation threshold is seen to move regularly downward as resolution is increased. In fact, the asymptotic value of percolation threshold should scale like $1/L$, where L is the size of the system in pixels for the following reason. It is conjectured that the C–S–H phase becomes percolated when a thin sheet partially covers each cement particle and also touches a neighboring cement particle. This sheet can be made up of material that is dissolved and reacted from the first layer of the cement particles. If this is true, then the degree of hydration at C–S–H percolation should scale like $1/L$, because the surface-to-volume ratio of the particles also scales this




Fig. 12. Fraction of the C–S–H phase that is part of a connected pathway as a function of the total C–S–H volume fraction for 0.3 w/c at different resolutions for a C_3S cement paste.

way, and essentially all that is happening is that the surface of the cement particles are dissolving and reacting to percolate the C–S–H layer. The volume fraction of the C–S–H should also scale the same way, since it goes like the total surface area of the cement divided by the volume of the system. Again, this ratio goes like $1/L$. At the two highest values of L , 400 and 800, this scaling is approximately true, as the percolation threshold for the 800^3 system is about one-half that of the 400^3 system. The comparison of this result to real cement paste is made in Section 8.

7. Resolution effects on computed diffusivity and permeability

One of the biggest effects of any change in capillary pore space percolation properties is on the diffusivity of the cement paste. Clearly, if one size system is still percolated when a smaller size is not at the same overall porosity, then there could be a significant difference between the diffusivities of the two, even though the larger system will be connected through smaller pores.

To test this possible effect on diffusivity, various microstructures of sizes 100^3 – 400^3 were saved at several different porosities between 0.12 and 0.32. The C–S–H phase was allowed to have a diffusivity between 0.0 and 0.01, where the pore space had a diffusivity of unity. In previous work, where an equation was developed relating porosity and diffusivity, a value of 0.0025 was used for the C–S–H phase [11]. Since computing the diffusivity for the 400^3 systems consumed much CPU time, runs were made at a smaller number of porosities.

The data is summarized in Fig. 13, which shows the computed diffusivity for all systems using only $D_{C-S-H} = 0.0$ and 0.005. It is clear that as the capillary porosity percola-




Fig. 13. Normalized diffusivities for different size systems vs. capillary porosity for two different values of $D_{C-S-H} = 0.0$ and 0.005.

tion threshold decreases, the diffusivity at a given porosity increases because the pore space is better connected farther away from its percolation point. As the capillary porosity decreases, the effect of the capillary porosity also decreases, with the overall diffusivity being controlled by the amount and diffusivity of the C–S–H phase. Since at any given capillary porosity, the phase fractions of all the different size systems are the same, the diffusivities should all become similar once the porosity has passed the percolation threshold for all systems. The graphs have almost reached that point in Fig. 13.

The permeability of each system was computed using a lattice Boltzmann algorithm modified to run on parallel computers [31]. Fig. 14 shows the computed fluid permeability curves. This transport property is much more sensitive to pore size than is diffusivity. The extra connectivity of the pore space of the larger models is through smaller pores, therefore, it is interesting to compare permeability and diffusivity, which is not very sensitive to pore size, on the same systems. In Fig. 14, it is clear that the permeability becomes very small for porosities smaller than the percolation threshold for the 100^3 system. This is because the biggest pores, which dominate the permeability, are disconnected by the 100^3 threshold. The permeability for the larger systems is dominated by much smaller pores for porosities lower than this. This figure broadly agrees with the experimental results of Powers et al. [32] on where the permeability of various w/c ratio pastes sharply changes, indicating a change in the size of pore that dominates the permeability.

Fig. 8 shows the effect of resolution on the capillary pore size very well. Since permeability is so much more sensitive to pore size than is diffusivity [6], digital resolution affects it much more than diffusivity. Chemical shrinkage results have shown that there is a pronounced flattening in the degree of hydration vs. time curve when the capillary porosity reached values near those predicted for the percola-

tion threshold predicted by 100^3 models [21]. This was interpreted as evidence that the capillary porosity was closing off near this porosity, thus, restricting the inflow of external water and causing a decrease in hydration. Since the flow of water is controlled by permeability, not diffusivity, the results of Fig. 14 are instructive in this case. This figure shows that, even though the higher resolution models still have some permeability for porosities lower than the 100^3 threshold, this permeability is very small. For most purposes, the porosity at the 100^3 threshold is the controlling porosity, thus, explaining the rough agreement with the experimental evidence.

One conclusion is clear from the above diffusivity work. Extracting an effective conductivity/diffusivity for the C–S–H phase by tuning the C–S–H property so that the bulk property agrees with experiment is not necessarily an accurate process. The interaction between C–S–H diffusivity and digital resolution of the model was clearly seen in Fig. 13.

Now this does not mean that the equation for diffusivity of cement paste as a function of capillary porosity (ϕ) that was derived in earlier work (Eq. (1)) [11],

$$D/D_i = 0.001 + 0.07\phi^2 + 1.7H(\phi - \phi_c)(\phi - \phi_c)^2 \quad (1)$$

is inaccurate. The parameters in this equation are $H(x)$, which equals 0 for negative values of the argument and 1 for positive values of the argument, and ϕ_c , which is the percolation threshold of the capillary porosity, as computed several times in earlier parts of this paper. This equation is accurate, as it was fit to experimental measurements, by using a relative diffusivity of 0.0025 for the C–S–H phase, where the value of 1 is used for the capillary pore space. It is just that using this value other than in the context of this equation, or assuming it to be an accurate measurement of a C–S–H property, is not necessarily valid. Therefore, while this equation, or an improved version [26], can be used for predicting cement paste transport, one must be cautious about extracting direct C–S–H information from it.

8. Further discussion of digital resolution

As is the usual case for random geometry models, the percolation properties of the cement hydration model do depend fairly strongly on digital resolution. They are robust, however, with respect to other factors like w/c ratio, and details of the model like cyclic vs. continuous operation, dissolution, etc. The question then arises, is there a “right” resolution?

This is not a trivial question. When solving differential equations using, say, a finite difference technique, the error gets smaller as the mesh size becomes smaller. In this case, full accuracy is only achieved in the limit of infinitely fine resolution, although usually there is a resolution at which the errors are “small enough.” If the cement hydration




Fig. 14. Permeabilities for different size systems (in units of μm^2) vs. capillary porosity.

model is like this case, then up to the limit of resolution tested, $0.125\text{ }\mu\text{m/pixel}$, the resolution was not small enough, as it appeared that two out of three percolation thresholds were still changing (the capillary pore space percolation threshold was not computed for the 800^3 case). However, the cement hydration model is designed to simulate a real material, cement paste, so we may turn to the real material to find a length scale that can give us clues as to the “right” resolution. It cannot be too strongly emphasized that a cellular automaton model is not the same mathematically as a finite difference model, and so will respond to digital resolution differently.

One note must be made before the discussion proceeds. If the notion of a dissolution length in the model is accepted, then the pixel representation of this length scale did not scale with resolution as did the number of pixels per particle. The reason for this is that going beyond the third nearest neighbor for dissolution brings in unphysical aspects of the dissolution. If dissolution can go beyond this point, then a pixel that is underneath solid pixels can “jump out” over its neighbor, dissolving before a solid pixel that is closer to the pore space. This does not happen if the model stays at a $3 \times 3 \times 3$ dissolution box.

The first length scale of interest to the model comes from the resolution of the original cement image generated by scanning electron microscopy that is used to generate the starting particles [9,24]. This resolution is only $0.5\text{ }\mu\text{m/pixel}$, therefore, in one way it does not make much sense to go beyond this resolution in hydration, as the original materials are only known to this resolution. This resolution would correspond to the 200^3 models considered in this paper. However, this does not give an absolute limit and could be improved in the future.

Second, the structure of the hydration model itself sets a lower limit to how many micrometers per pixel make sense. In this model, the C–S–H phase is represented as a featureless, uniform solid. Actually, there are nanometer-scale pores in this material, which are not resolved in this model. To begin to see these pores would require a resolution of about 50 nm/pixel or $0.05\text{ }\mu\text{m/pixel}$, which would correspond to a model size of about 2000^3 , for a fixed unit cell length of $100\text{ }\mu\text{m}$. Therefore, if the model is designed to not see the C–S–H pores, there is no reason to go larger than this size model. A size of 800^3 , corresponding to a resolution of $0.125\text{ }\mu\text{m/pixel}$, is approaching this limit.

Another size scale is set by the dissolution kinetics. The first few hours of hydration time corresponds to the first few cycles of dissolution/hydration in the model [9]. It is known that small cement particles, $1\text{ }\mu\text{m}$ or less in diameter, dissolve and hydrate in the early hours of hydration. In the same way, 1-pixel particles dissolve very rapidly in the model, with increasing slowness as the number of pixels in the particle increases. Probably, in terms of matching this dissolution behavior, the lower limit on the length per pixel is $0.2\text{ }\mu\text{m/pixel}$, which would correspond to a 500^3 -pixel model in this paper. Having

more pixels in a $1\text{-}\mu\text{m}$ diameter particle would make it dissolve too slowly to match reality.

Another length scale in cement paste might be given by the thickness of the initial C–S–H layer that forms on the surfaces of the cement grains, which determines the initial percolation of the C–S–H phase. If this layer was on average, say, $0.25\text{ }\mu\text{m}$ thick, then it would make sense to have a pixel length scale of no less than $0.25\text{ }\mu\text{m/pixel}$, or a 400^3 -size model. Then there would be no chance of forming an unphysically thin C–S–H layer of less than this thickness, as the 1-pixel coating that percolates would be the right thickness. The thickness of this layer might be able to be determined using a technique like atomic force microscopy (AFM), which has the capability of imaging the topography of a cement particle through water, so that the thickness of a thin surface layer of C–S–H might be able to be measured.

The set point can also be used to guess at an appropriate resolution for the model. Note that the set point of the larger models was not directly measured, since extra memory is required for this determination, which made the larger models too large for our study. However, the C–S–H percolation results can be used to show how the set point scales with resolution.

The set point in the portland cement model is determined by the percolation of the joint C–S–H/ettringite cement phase. Physically, the set point is defined as the point at which the growing C–S–H/ettringite first “glues” the cement particles together [33]. Cement particle flocculation is ignored in the model, as particles only count as being connected if a C–S–H or ettringite bridge connects them. To make a connected phase out of the joint cement/C–S–H/ettringite phase is easier than percolating the C–S–H phase separately. Clearly, the cement particles can be “glued” together into a percolated phase without the C–S–H part of the “glue” being itself percolated. Therefore, the degree of hydration at which the C–S–H percolation point occurs must be an upper bound for the degree of hydration at which the set point occurs. Since the set point is still dependent on C–S–H production, the set point must then scale similarly with model resolution as did the C–S–H percolation threshold, as seen in Fig. 12.

Set point results of the 100^3 model have been found to be in fairly good agreement with experimental measurements. The differences tended to be that the predicted degree of hydration was a bit too low [33]. The higher resolution C–S–H percolation data in the present study would only further reduce the degree of hydration at set, implying that the model would agree less well with experiment. Therefore, the model resolution cannot go too high above 100^3 , if at all, if agreement with set point experiments with this version of the model is to be preserved. More good measurements of the degree of hydration at set are needed to help calibrate the model's resolution. Of course, one must remember that the model predicts geometrical percolation and the experiments predict set point via

some kind of physical measurement. Geometrical percolation almost certainly takes place before the threshold for physical property measurement is reached, therefore, the experimentally measured degree of hydration at set is most likely above that determined by the model. Therefore, the lower values of degree of hydration at set that would be produced by a somewhat higher resolution model might still be in good agreement with reality.

One piece of evidence for a resolution near that of 1 μm /pixel is found in electrical conductivity measurements of frozen cement paste using impedance spectroscopy [34]. In this work, a low enough temperature was achieved (-40°C) so that the only conduction paths were considered to be through the C–S–H phase, and not the capillary pores. There was very good agreement between predictions of the model and experimental measurements, and the percolation threshold of the C–S–H appeared to be about 0.12–0.16. Since the experimental measurements were on portland cement paste, the model values of the C–S–H percolation threshold for this material should be used to compare. This value was found to be about 0.18, which argues for a resolution near to that used in Ref. [34] — 1 μm /pixel.

9. Summary

The effects of chemistry, w/c ratio, and cement PSD on the percolation aspects of the NIST cement hydration model were studied. The role of cement PSD was significant. The effect of w/c ratio was the same as that found before: all percolation curves for different w/c ratios for a given cement collapsed on a single curve when plotted against the volume fraction of the phase of interest. Portland cement chemistry beyond simple C_3S cements did not have a significant effect on percolation quantities, except for those of the C–S–H phase. This was explained from the relative production of C–S–H in the portland cement hydration vs. that of the C_3S .

The three major sources of error present for every digital, random model — finite size error, statistical fluctuation, and digital resolution — were studied. It was found that finite size error and statistical fluctuation were not important for the cement hydration model as it is usually used (1 μm /pixel, 100^3 size, about 2000 particles). Digital resolution did have an important effect on the percolation quantities and the transport properties of diffusivity and permeability. However, intrinsic length scales present in cement paste, which include the typical size of cement particles that very quickly dissolve and hydrate, the thickness of the first C–S–H layer, and the size of the largest C–S–H pores, along with the values of the set point for real materials, restrict the resolutions that will give realistic values of experimental quantities for the model.

Further detailed experimental work on cement paste should aid in determining the “right” resolution to use in the model. It does seem that this “right” resolution must be

near 1 μm /pixel, as comparison with numerous experimental results appear to point to that resolution. This is much different than the case of finite difference or finite element computations, where the “right” resolution is as fine as possible [14,15,35].

Acknowledgments

We would like to thank the Partnership for High Performance Concrete Technology program at NIST for support of this work. We would also like to thank John Hagedorn (NIST Information Technology Laboratory) and Nicos Martys (NIST Building and Fire Research Laboratory) for computing the permeability of the various size systems.

References

- [1] D.P. Bentz, E.J. Garboczi, E.S. Lagergren, Multi-scale microstructural modelling of concrete diffusivity: Identification of significant variables, *Cem. Concr. Aggregates* 20 (1998) 129–139.
- [2] E.J. Garboczi, D.P. Bentz, Multi-scale analytical/numerical theory of the diffusivity of concrete, *Adv. Cem. Based Mater.* 8 (1998) 77–88.
- [3] E.J. Garboczi, D.P. Bentz, Analytical formulas for interfacial transition zone properties, *Adv. Cem. Based Mater.* 6 (1997) 99–108.
- [4] D.P. Bentz, D.A. Quenard, V. Baroghel-Bouny, E.J. Garboczi, H.M. Jennings, Modelling drying shrinkage of cement paste and mortar: Part I. Structural models from nanometers to millimeters, *Mater. Struct.* 28 (1995) 450–458.
- [5] D. Stauffer, A. Aharony, *Introduction to Percolation Theory*, second ed., Taylor & Francis, London, 1992.
- [6] E.J. Garboczi, Permeability, diffusivity, and microstructural parameters: A critical review, *Cem. Concr. Res.* 20 (1990) 591–601.
- [7] A.J. Katz, A.H. Thompson, Quantitative prediction of permeability in porous rock, *Phys. Rev. B* 34 (1986) 8179–8181.
- [8] A.J. Katz, A.H. Thompson, Prediction of rock electrical conductivity from mercury injection measurements, *J. Geophys. Res.* 92 (1987) 599–607.
- [9] D.P. Bentz, Three-dimensional computer simulation of portland cement hydration and microstructure development, *J. Am. Ceram. Soc.* 80 (1997) 3–21.
- [10] D.P. Bentz, E.J. Garboczi, Percolation of phases in a three-dimensional cement paste microstructural model, *Cem. Concr. Res.* 21 (1991) 325–344.
- [11] E.J. Garboczi, D.P. Bentz, Computer simulation of the diffusivity of cement-based materials, *J. Mater. Sci.* 27 (1992) 2083–2092.
- [12] E.J. Garboczi, D.P. Bentz, N.S. Martys, Digital imaging and pore morphology, in: P. Wong (Ed.), *Methods in the Physics of Porous Media*, Academic Press, San Diego, 1999, pp. 1–41.
- [13] E.J. Garboczi, D.P. Bentz, Fundamental computer simulation models for cement-based materials, in: J. Skalny (Ed.), *Materials Science of Concrete vol. II*, American Ceramics Society, Westerville, OH, 1991, pp. 249–277.
- [14] A.P. Roberts, E.J. Garboczi, Elastic properties of porous ceramic models, *J. Am. Ceram. Soc.* 83 (2000) 3041–3048.
- [15] A.P. Roberts, E.J. Garboczi, Elastic properties of model open-cell cellular solid materials, *J. Mech. Phys. Solids*, 2001 (in press).
- [16] Z. Hashin, Analysis of composite materials: A survey, *Appl. Mech.* 50 (1983) 481–505.
- [17] S. Torquato, *Theory of Composite Materials*, 2001 (Oxford, in press).
- [18] H. Uchikawa, Similarities and discrepancies of hardened cement paste, mortar, and concrete from the standpoints of composition and

- structure, in: E. Gartner (Ed.), *Advances in Cement Manufacture and Use*, Engineering Foundation, New York, 1989, pp. 271–294.
- [19] B. Lu, S. Torquato, Local volume fraction fluctuations in heterogeneous media, *J. Chem. Phys.* 93 (1990) 3452–3459.
 - [20] B. Lu, S. Torquato, Photographic granularity: Mathematical formulation and effect of impenetrability of grains, *J. Opt. Soc. Am.* 7 (1990) 717–724.
 - [21] D.P. Bentz, A three-dimensional cement hydration and microstructure program: I. Hydration rate, heat of hydration, and chemical shrinkage, NIST Intern. Rep. 5756 (See Part I, Chapter 4, Section 2, <http://ciks.cbt.nist.gov/monograph>).
 - [22] P. Navi, C. Pignat, Simulation of cement hydration and the connectivity of the capillary pore space, *Adv. Cem. Based Mater.* 4 (1996) 58–67.
 - [23] H.M. Jennings, S.K. Johnson, Simulation of microstructure development during the hydration of a cement compound, *J. Am. Ceram. Soc.* 69 (1986) 790–795.
 - [24] D.P. Bentz, CEMHYD3D: A three-dimensional cement hydration and microstructure development modelling package, Version 2.0, NIST Intern. Rep. 6485 (April, see also <http://ciks.cbt.nist.gov/monograph>, Part I, Chapter 4, Section 2).
 - [25] A. Boumiz, Etude comparee des evolutions mecaniques et chimiques des pates de ciment et mortiers a tres jeune age, PhD thesis, Universite Paris 7 (Denis Diderot), 1995.
 - [26] D.P. Bentz, E.J. Garboczi, C.J. Haecker, O.M. Jensen, Effects of cement particle size distribution on performance properties of portland cement-based materials, *Cem. Concr. Res.* 29 (1999) 1663–1671.
 - [27] D.P. Bentz, C.J. Haecker, An argument for using coarse cements in high performance concrete, *Cem. Concr. Res.* 29 (1999) 615–618.
 - [28] E.J. Garboczi, M.F. Thorpe, M. DeVries, A.R. Day, Universal conductivity curve for a plane containing random holes, *Phys. Rev. A* 43 (1991) 6473–6482.
 - [29] E.J. Garboczi, K.A. Snyder, J.F. Douglas, M.F. Thorpe, Geometrical percolation threshold of overlapping ellipsoids, *Phys. Rev. E* 52 (1995) 819–828.
 - [30] P.E. Stutzman, J.R. Clifton, Microstructural features of some low water/solids, silica fume mortars cured at different temperatures, NIST Intern. Rep. 4790 (April, see also <http://ciks.cbt.nist.gov/monograph>, Part I, Chapter 3).
 - [31] N.S. Martys, J.G. Hagedorn, D. Goujon, J.E. Devaney, Large scale simulations of single and multi-component flow in porous media, *Proc. SPIE* 3772 (Also available at <http://ciks.cbt.nist.gov/monograph>, Part II, Chapter 6, Section 3).
 - [32] T.C. Powers, L.E. Copeland, H.M. Mann, Capillary continuity or discontinuity in cement pastes, *Portland Cem. Bull.* 110 (1959) 3–12.
 - [33] D.P. Bentz, E.J. Garboczi, N.S. Martys, Application of digital-image-based models to microstructure, transport properties, and degradation of cement-based materials, in: H.M. Jennings, J. Kropp, K. Scrivener (Eds.), *The Modelling of Microstructure and Its Potential for Studying Transport Properties and Durability*, Kluwer, Dordrecht, 1996, pp. 167–185.
 - [34] R.A. Olson, B.J. Christensen, R.T. Coverdale, S.J. Ford, G.M. Moss, H.M. Jennings, T.O. Mason, E.J. Garboczi, Interpretation of the impedance spectroscopy of cement paste via computer modelling: III. Microstructural analysis of frozen cement paste, *J. Mater. Sci.* 30 (1995) 5078–5086.
 - [35] E.J. Garboczi, Finite element and finite difference programs for computing the linear electric and elastic properties of digital images of random materials, NIST Intern. Rep. 6269 (See Part II, Chapter 2, <http://ciks.cbt.nist.gov/monograph>).