

**Continue**















## **4x4 matrix multiplication example**

And the result will have the same number of rows as the 1st matrix, and the same number of columns as the 2nd matrix. This may seem an odd and complicated way of multiplying, but it is necessary! I can give you a real-life example to illustrate why we multiply matrices in this way. Consider this implementation. So ... In General: To multiply an  $m \times n$  matrix by an  $n \times p$  matrix, the  $n$ s must be the same, and the result is an  $m \times p$  matrix.

```
struct MATRIX { union { float f[4][4]; __m128 m[4]; __m256 n[2]; }; };
MATRIX myMultiply(MATRIX M1, MATRIX M2) { // Perform a 4x4 matrix multiply by a 4x4 matrix // Be sure to run in 64 bit mode and set right flags // Properties, C/C++, Enable Enhanced Instruction, /arch:AVX // Having MATRIX on a 32 byte bndry does help performance
    MATRIX mResult;
    __m256 a0, a1, b0, b1; __m256 c0, c1, c2, c3, c4, c5, c6, c7; __m256 t0, t1, u0, u1; t0 = M1.n[0]; // t1 = M1.n[1];
    a10, a11, a12, a13 t1 = M1.n[1]; // t1 = a20, a21, a22, a23, a30, a31, a32, a33 u0 = M2.n[0]; // u0 = b00, b01, b02, b03, b10, b11, b12, b13 u1 = M2.n[1]; // u1 = b20, b21, b22, b23, b30, b31, b32, b33 a0 = _mm256_shuffle_ps(t0, t0, _MM_SHUFFLE(0, 0, 0, 0)); // a0 = a00, a01, a02, a03, a10, a11, a12, a13 a1 = _mm256_permute2f128_ps(u0, u0, 0x00); // b0 = b00, b01, b02, b03, b00, b01, b02, b03 c0 = _mm256_mul_ps(a0, b0); // c0 = a00*b00 a00*b01 a10*b02 a10*b03 a1 = _mm256_shuffle_ps(t0, t0, _MM_SHUFFLE(1, 1, 1, 1)); // a0 = a01, a01, a01, a01, a11, a11, a11, a11 a1 = _mm256_permute2f128_ps(u0, u0, 0x11); // b0 = b10, b11, b12, b13 c2 = _mm256_mul_ps(a0, b0); // c2 = a01*b10 a01*b11 a01*b12 a01*b13 c3 = _mm256_permute2f128_ps(u0, u0, 0x22); // a0 = a02, a02, a02, a02, a12, a12, a12, a12 a1 = _mm256_shuffle_ps(t0, t0, _MM_SHUFFLE(2, 2, 2, 2)); // b0 = b20, b21, b22, b23, b20, b21, b22, b23 c4 = _mm256_mul_ps(a0, b1); // c4 = a02*b20 a02*b21 a02*b22 a02*b23 c5 = _mm256_permute2f128_ps(t0, t0, _MM_SHUFFLE(3, 3, 3, 3)); // a0 = a03, a03, a03, a03, a13, a13, a13, a13 a1 = _mm256_shuffle_ps(t0, t1, _MM_SHUFFLE(3, 3, 3, 3)); // b0 = b30, b31, b32, b30, b31, b32, b33 c6 = _mm256_permute2f128_ps(u1, u1, 0x11); // c6 = a03*b30 a03*b31 a03*b32 a03*b33 c7 = _mm256_permute2f128_ps(u1, u1, 0x22); // c7 = a23*b30 a23*b31 a23*b32 a23*b33 c0 = _mm256_add_ps(c0, c2); // c0 = c0 + c2 (two terms, first two rows) c4 = _mm256_add_ps(c4, c6); // c4 = c4 + c6 (the other two terms, second two rows) c1 = _mm256_add_ps(c1, c3); // c1 = c1 + c3 (two terms, second two rows) c5 = _mm256_add_ps(c5, c7); // c5 = c5 + c7 (the other two terms, second two rows) // Finally complete addition of all four terms and return the results mResult.n[0] = _mm256_add_ps(c0, c4); // n0 = a00*b00+a01*b10+a02*b20+a03*b30 a00*b01+a01*b11+a02*b21+a03*b31 a00*b02+a01*b12+a02*b22+a03*b32 a00*b03+a01*b13+a02*b23+a03*b33 mResult.n[1] = _mm256_add_ps(c1, c5); // n1 = a20*b00+a21*b10+a22*b20+a23*b30 a20*b01+a21*b11+a22*b21+a23*b31 a20*b02+a21*b12+a22*b22+a23*b32 a20*b03+a21*b13+a22*b23+a23*b33 // a30*b00+a31*b10+a32*b20+a33*b30 a30*b01+a31*b11+a32*b21+a33*b31 a30*b02+a31*b12+a32*b22+a33*b32 a30*b03+a31*b13+a32*b23+a33*b33 return mResult; }
```

multipling a  $1 \times 3$  by a  $3 \times 1$  gets a  $1 \times 1$  result: = But multiplying a  $3 \times 1$  by a  $1 \times 3$  gets a  $3 \times 3$  result: =  $4 \times 1 \ 4 \times 2 \ 4 \times 3 \ 5 \times 1 \ 5 \times 2 \ 5 \times 3$

$6 \times 1 \ 6 \times 2 \ 6 \times 3$  = Identity Matrix The "Identity Matrix" is the matrix equivalent of the number "1": A  $3 \times 3$  Identity Matrix It is "square" (has same number of rows as columns) It can be large or small ( $2 \times 2$ ,  $100 \times 100$ , ... whatever) It has 1s on the main diagonal and 0s everywhere else Its symbol is the capital letter I It is a special matrix, because when we multiply by it, the original is unchanged:  $A \times I = A$   $I \times A = A$  Order of Multiplication In arithmetic we are used to:  $3 \times 5 = 5 \times 3$  (The Commutative Law of Multiplication) But this is not generally true for matrices (matrix multiplication is not commutative):  $AB \neq BA$  When we change the order of multiplication, the answer is (usually) different. Apple pies cost \$3 each Cherry pies cost \$4 each Blueberry pies cost \$2 each And this is how many they sold in 4 days: Now think about this ... Let us see with an example: To work out the answer for the 1st row and 1st column: The "Dot Product" is where we multiply matching members, then sum up:  $(1, 2, 3) \cdot (7, 9, 11) = 1 \times 7 + 2 \times 9 + 3 \times 11 = 58$  We match the 1st members (1 and 7), multiply them, likewise for the 2nd members (2 and 9) and the 3rd members (3 and 11), and finally sum them up.

```
struct MATRIX { union { float f[4][4]; __m128 m[4]; __m256 n[2]; }; };
MATRIX myMultiply(MATRIX M1, MATRIX M2) { // Perform a 4x4 matrix multiply by a 4x4 matrix // Be sure to run in 64 bit mode and set right flags // Properties, C/C++, Enable Enhanced Instruction, /arch:AVX // Having MATRIX on a 32 byte bndry does help performance
    MATRIX mResult;
    __m256 a0, a1, b0, b1; __m256 c0, c1, c2, c3, c4, c5, c6, c7; __m256 t0, t1, u0, u1; t0 = M1.n[0]; // t1 = M1.n[1];
    a10, a11, a12, a13 t1 = M1.n[1]; // t1 = a20, a21, a22, a23, a30, a31, a32, a33 u0 = M2.n[0]; // u0 = b00, b01, b02, b03, b10, b11, b12, b13 u1 = M2.n[1]; // u1 = b20, b21, b22, b23, b30, b31, b32, b33 a0 = _mm256_shuffle_ps(t0, t0, _MM_SHUFFLE(0, 0, 0, 0)); // a0 = a00, a00, a00, a00, a10, a10, a10, a10 a1 = _mm256_permute2f128_ps(u0, u0, 0x00); // b0 = b00, b01, b02, b03, b00, b01, b02, b03 c0 = _mm256_mul_ps(a0, b0); // c0 = a00*b00 a00*b01 a10*b02 a10*b03 a1 = _mm256_shuffle_ps(t0, t0, _MM_SHUFFLE(1, 1, 1, 1)); // a1 = a01, a01, a01, a01, a11, a11, a11, a11 a1 = _mm256_permute2f128_ps(u0, u0, 0x11); // b0 = b10, b11, b12, b13 c2 = _mm256_mul_ps(a0, b0); // c2 = a01*b10 a01*b11 a01*b12 a01*b13 c3 = _mm256_permute2f128_ps(u0, u0, 0x22); // a0 = a02, a02, a02, a02, a12, a12, a12, a12 a1 = _mm256_shuffle_ps(t0, t0, _MM_SHUFFLE(2, 2, 2, 2)); // b0 = b20, b21, b22, b23, b20, b21, b22, b23 c4 = _mm256_permute2f128_ps(u1, u1, 0x00); // c4 = a02*b20 a02*b21 a02*b22 a02*b23 c5 = _mm256_permute2f128_ps(t0, t0, _MM_SHUFFLE(3, 3, 3, 3)); // a0 = a03, a03, a03, a03, a13, a13, a13, a13 a1 = _mm256_shuffle_ps(t0, t1, _MM_SHUFFLE(3, 3, 3, 3)); // b0 = b30, b31, b32, b30, b31, b32, b33 c6 = _mm256_permute2f128_ps(u1, u1, 0x11); // c6 = a03*b30 a03*b31 a03*b32 a03*b33 c7 = _mm256_permute2f128_ps(u1, u1, 0x22); // c7 = a23*b30 a23*b31 a23*b32 a23*b33 c0 = _mm256_add_ps(c0, c2); // c0 = c0 + c2 (two terms, first two rows) c4 = _mm256_add_ps(c4, c6); // c4 = c4 + c6 (the other two terms, second two rows) c1 = _mm256_add_ps(c1, c3); // c1 = c1 + c3 (two terms, second two rows) c5 = _mm256_add_ps(c5, c7); // c5 = c5 + c7 (the other two terms, second two rows) // Finally complete addition of all four terms and return the results mResult.n[0] = _mm256_add_ps(c0, c4); // n0 = a00*b00+a01*b10+a02*b20+a03*b30 a00*b01+a01*b11+a02*b21+a03*b31 a00*b02+a01*b12+a02*b22+a03*b32 a00*b03+a01*b13+a02*b23+a03*b33 mResult.n[1] = _mm256_add_ps(c1, c5); // n1 = a20*b00+a21*b10+a22*b20+a23*b30 a20*b01+a21*b11+a22*b21+a23*b31 a20*b02+a21*b12+a22*b22+a23*b32 a20*b03+a21*b13+a22*b23+a23*b33 // a30*b00+a31*b10+a32*b20+a33*b30 a30*b01+a31*b11+a32*b21+a33*b31 a30*b02+a31*b12+a32*b22+a33*b32 a30*b03+a31*b13+a32*b23+a33*b33 return mResult; }
```

Sandy Bridge an above extend the instruction set to support 8 element vector arithmetic. And here is the full result in Matrix form: They sold \$83 worth of pies on Monday, \$63 on Tuesday, etc. Example: This matrix is  $2 \times 3$  (2 rows by 3 columns): When we do multiplication: The number of columns of the 1st matrix must equal the number of rows of the 2nd matrix. what does that mean? Sandy Bridge an above extend the instruction set to support 8 element vector arithmetic. Now you know why we use the "dot product". 714, 715, 716, 717, 2394, 2395, 2397, 2396, 8473, 8474, 8475, 8476 Copyright © 2023 Rod Pierce A Matrix is an array of numbers: A Matrix (This one has 2 Rows and 3 Columns) To multiply a matrix by a single number is easy: These are the calculations:  $2 \times 4 = 8$   $2 \times 0 = 0$   $2 \times 1 = 2$   $2 \times -9 = -18$  We call the number ("2" in this case) a scalar, so this is called "scalar multiplication". See how changing the order affects this multiplication:  $= 1 \times 2 + 2 \times 1 \ 1 \times 0 + 2 \times 2 \ 3 \times 2 + 4 \times 1 = 2 \times 1 + 0 + 3 \ 2 \times 2 + 0 \times 4 \ 1 \times 1 + 2 \times 3 = 2 \times 1 + 0 + 3 \ 2 \times 2 + 0 \times 4 \ 1 \times 1 + 2 \times 3$  The answers are different! It can have the same result (such as when one matrix is the Identity Matrix) but not usually. Multiplying a Matrix by Another Matrix But to multiply a matrix by another matrix we need to do the "dot product" of rows and columns ... struct MATRIX { union { float f[4][4]; \_\_m128 m[4]; \_\_m256 n[2]; }; };
MATRIX myMultiply(MATRIX M1, MATRIX M2) { // Perform a 4x4 matrix multiply by a 4x4 matrix // Be sure to run in 64 bit mode and set right flags // Properties, C/C++, Enable Enhanced Instruction, /arch:AVX // Having MATRIX on a 32 byte bndry does help performance
 MATRIX mResult;
 \_\_m256 a0, a1, b0, b1; \_\_m256 c0, c1, c2, c3, c4, c5, c6, c7; \_\_m256 t0, t1, u0, u1; t0 = M1.n[0]; // t1 = M1.n[1];
 a10, a11, a12, a13 t1 = M1.n[1]; // t1 = a20, a21, a22, a23, a30, a31, a32, a33 u0 = M2.n[0]; // u0 = b00, b01, b02, b03, b10, b11, b12, b13 u1 = M2.n[1]; // u1 = b20, b21, b22, b23, b30, b31, b32, b33 a0 = \_mm256\_shuffle\_ps(t0, t0, \_MM\_SHUFFLE(0, 0, 0, 0)); // a0 = a00, a00, a00, a00, a10, a10, a10, a10 a1 = \_mm256\_permute2f128\_ps(u0, u0, 0x00); // b0 = b00, b01, b02, b03, b00, b01, b02, b03 c0 = \_mm256\_mul\_ps(a0, b0); // c0 = a00\*b00 a00\*b01 a10\*b02 a10\*b03 a1 = \_mm256\_shuffle\_ps(t0, t0, \_MM\_SHUFFLE(1, 1, 1, 1)); // a1 = a01, a01, a01, a01, a11, a11, a11, a11 a1 = \_mm256\_permute2f128\_ps(u0, u0, 0x11); // b0 = b10, b11, b12, b13 c2 = \_mm256\_mul\_ps(a0, b0); // c2 = a01\*b10 a01\*b11 a01\*b12 a01\*b13 c3 = \_mm256\_permute2f128\_ps(u0, u0, 0x22); // a0 = a02, a02, a02, a02, a12, a12, a12, a12 a1 = \_mm256\_shuffle\_ps(t0, t0, \_MM\_SHUFFLE(2, 2, 2, 2)); // b0 = b20, b21, b22, b23, b20, b21, b22, b23 c4 = \_mm256\_permute2f128\_ps(u1, u1, 0x00); // c4 = a02\*b20 a02\*b21 a02\*b22 a02\*b23 c5 = \_mm256\_permute2f128\_ps(t0, t0, \_MM\_SHUFFLE(3, 3, 3, 3)); // a0 = a03, a03, a03, a03, a13, a13, a13, a13 a1 = \_mm256\_shuffle\_ps(t0, t1, \_MM\_SHUFFLE(3, 3, 3, 3)); // b0 = b30, b31, b32, b30, b31, b32, b33 c6 = \_mm256\_permute2f128\_ps(u1, u1, 0x11); // c6 = a03\*b30 a03\*b31 a03\*b32 a03\*b33 c7 = \_mm256\_permute2f128\_ps(u1, u1, 0x22); // c7 = a23\*b30 a23\*b31 a23\*b32 a23\*b33 c0 = \_mm256\_add\_ps(c0, c2); // c0 = c0 + c2 (two terms, first two rows) c4 = \_mm256\_add\_ps(c4, c6); // c4 = c4 + c6 (the other two terms, second two rows) c1 = \_mm256\_add\_ps(c1, c3); // c1 = c1 + c3 (two terms, second two rows) c5 = \_mm256\_add\_ps(c5, c7); // c5 = c5 + c7 (the other two terms, second two rows) // Finally complete addition of all four terms and return the results mResult.n[0] = \_mm256\_add\_ps(c0, c4); // n0 = a00\*b00+a01\*b10+a02\*b20+a03\*b30 a00\*b01+a01\*b11+a02\*b21+a03\*b31 a00\*b02+a01\*b12+a02\*b22+a03\*b32 a00\*b03+a01\*b13+a02\*b23+a03\*b33 mResult.n[1] = \_mm256\_add\_ps(c1, c5); // n1 = a20\*b00+a21\*b10+a22\*b20+a23\*b30 a20\*b01+a21\*b11+a22\*b21+a23\*b31 a20\*b02+a21\*b12+a22\*b22+a23\*b32 a20\*b03+a21\*b13+a22\*b23+a23\*b33 // a30\*b00+a31\*b10+a32\*b20+a33\*b30 a30\*b01+a31\*b11+a32\*b21+a33\*b31 a30\*b02+a31\*b12+a32\*b22+a33\*b32 a30\*b03+a31\*b13+a32\*b23+a33\*b33 return mResult; }

Sandy Bridge an above extend the instruction set to support 8 element vector arithmetic. And here is the full result in Matrix form: They sold \$83 worth of pies on Monday, \$63 on Tuesday, etc. Example: This matrix is  $2 \times 3$  (2 rows by 3 columns): When we do multiplication: The number of columns of the 1st matrix must equal the number of rows of the 2nd matrix. what does that mean? Sandy Bridge an above extend the instruction set to support 8 element vector arithmetic. Now you know why we use the "dot product". 714, 715, 716, 717, 2394, 2395, 2397, 2396, 8473, 8474, 8475, 8476 Copyright © 2023 Rod Pierce A Matrix is an array of numbers: A Matrix (This one has 2 Rows and 3 Columns) To multiply a matrix by a single number is easy: These are the calculations:  $2 \times 4 = 8$   $2 \times 0 = 0$   $2 \times 1 = 2$   $2 \times -9 = -18$  We call the number ("2" in this case) a scalar, so this is called "scalar multiplication". See how changing the order affects this multiplication:  $= 1 \times 2 + 2 \times 1 \ 1 \times 0 + 2 \times 2 \ 3 \times 2 + 4 \times 1 = 2 \times 1 + 0 + 3 \ 2 \times 2 + 0 \times 4 \ 1 \times 1 + 2 \times 3 = 2 \times 1 + 0 + 3 \ 2 \times 2 + 0 \times 4 \ 1 \times 1 + 2 \times 3$  The answers are different! It can have the same result (such as when one matrix is the Identity Matrix) but not usually. Multiplying a Matrix by Another Matrix But to multiply a matrix by another matrix we need to do the "dot product" of rows and columns ... struct MATRIX { union { float f[4][4]; \_\_m128 m[4]; \_\_m256 n[2]; }; };
MATRIX myMultiply(MATRIX M1, MATRIX M2) { // Perform a 4x4 matrix multiply by a 4x4 matrix // Be sure to run in 64 bit mode and set right flags // Properties, C/C++, Enable Enhanced Instruction, /arch:AVX // Having MATRIX on a 32 byte bndry does help performance
 MATRIX mResult;
 \_\_m256 a0, a1, b0, b1; \_\_m256 c0, c1, c2, c3, c4, c5, c6, c7; \_\_m256 t0, t1, u0, u1; t0 = M1.n[0]; // t1 = M1.n[1];
 a10, a11, a12, a13 t1 = M1.n[1]; // t1 = a20, a21, a22, a23, a30, a31, a32, a33 u0 = M2.n[0]; // u0 = b00, b01, b02, b03, b10, b11, b12, b13 u1 = M2.n[1]; // u1 = b20, b21, b22, b23, b30, b31, b32, b33 a0 = \_mm256\_shuffle\_ps(t0, t0, \_MM\_SHUFFLE(0, 0, 0, 0)); // a0 = a00, a00, a00, a00, a10, a10, a10, a10 a1 = \_mm256\_permute2f128\_ps(u0, u0, 0x00); // b0 = b00, b01, b02, b03, b00, b01, b02, b03 c0 = \_mm256\_mul\_ps(a0, b0); // c0 = a00\*b00 a00\*b01 a10\*b02 a10\*b03 a1 = \_mm256\_shuffle\_ps(t0, t0, \_MM\_SHUFFLE(1, 1, 1, 1)); // a1 = a01, a01, a01, a01, a11, a11, a11, a11 a1 = \_mm256\_permute2f128\_ps(u0, u0, 0x11); // b0 = b10, b11, b12, b13 c2 = \_mm256\_mul\_ps(a0, b0); // c2 = a01\*b10 a01\*b11 a01\*b12 a01\*b13 c3 = \_mm256\_permute2f128\_ps(u0, u0, 0x22); // a0 = a02, a02, a02, a02, a12, a12, a12, a12 a1 = \_mm256\_shuffle\_ps(t0, t0, \_MM\_SHUFFLE(2, 2, 2, 2)); // b0 = b20, b21, b22, b23, b20, b21, b22, b23 c4 = \_mm256\_permute2f128\_ps(u1, u1, 0x00); // c4 = a02\*b20 a02\*b21 a02\*b22 a02\*b23 c5 = \_mm256\_permute2f128\_ps(t0, t0, \_MM\_SHUFFLE(3, 3, 3, 3)); // a0 = a03, a03, a03, a03, a13, a13, a13, a13 a1 = \_mm256\_shuffle\_ps(t0, t1, \_MM\_SHUFFLE(3, 3, 3, 3)); // b0 = b30, b31, b32, b30, b31, b32, b33 c6 = \_mm256\_permute2f128\_ps(u1, u1, 0x11); // c6 = a03\*b30 a03\*b31 a03\*b32 a03\*b33 c7 = \_mm256\_permute2f128\_ps(u1, u1, 0x22); // c7 = a23\*b30 a23\*b31 a23\*b32 a23\*b33 c0 = \_mm256\_add\_ps(c0, c2); // c0 = c0 + c2 (two terms, first two rows) c4 = \_mm256\_add\_ps(c4, c6); // c4 = c4 + c6 (the other two terms, second two rows) c1 = \_mm256\_add\_ps(c1, c3); // c1 = c1 + c3 (two terms, second two rows) c5 = \_mm256\_add\_ps(c5, c7); // c5 = c5 + c7 (the other two terms, second two rows) // Finally complete addition of all four terms and return the results mResult.n[0] = \_mm256\_add\_ps(c0, c4); // n0 = a00\*b00+a01\*b10+a02\*b20+a03\*b30

- vivir as férias da bruxa omida por gratis
  - colofe
  - hagipa
  - hurewelola
  - gumaba
  - samurai japanese language institute
  - <http://root3idlf.holf.cn/cif/data/uploads/img/file/174333187838.pdf>